Thesis report

Global CO Emission Estimation for Steel Plants Using TROPOMI Data and Enhanced Machine Learning for Pollution Detection

AE5822: Thesis Space

Arthur Bronstring

Delft University of Technology & Netherlands Institute for Space Research

TU Delft

SRON
Netherlands Institute for Space Research

Thesis report

Global CO Emission Estimation for Steel Plants Using TROPOMI Data and Enhanced Machine Learning for Pollution Detection

by

Arthur Bronstring

Student Name	Student Number
Arthur Bronstring	4780744

Commitee chair:	Dr. L.L.A. Vermeersen
Supervisors:	T. Borsdorff & Dr. S.M. Cazaux
External examiner:	Dr. J.P. Veefkind
Project Duration:	April 2024 - February 2025
Faculty:	Faculty of Aerospace Engineering, Delft

Cover:	The Sentinel-5P satellite watching over Earth, the image credits belong to ESA/ATG medialab
Style:	TU Delft Report Style, with modifications by Daan Zwaneveld

Netherlands Institute for Space Research

Preface

Submitting this thesis means the end of my studies at TU Delft.The research presented herein was conducted under the guidance and supervision of Tobias Borsdorff and Dr. Stéphanie Cazaux. I thank my supervisors for their guidance and support during the project.

I would also like to thank Peter Sterk and Thomas Plewa for their contributions to this project. I must also express my heartfelt gratitude to my friends and family for their constant support. I dedicate this work to everyone who has inspired and encouraged me throughout this time.

Arthur Bronstring
Zoetermeer, January 2025

Abstract

The TROPOspheric Monitoring Instrument (TROPOMI) provides high-resolution, global measurements of carbon monoxide (CO) for environmental pollution monitoring. The Automated Plume Detection and Emission Estimation algorithm (APE), developed by SRON, identifies pollution plumes and estimates emissions based on the satellite data. This study implemented four machine learning algorithms to enhance APE and applied them to 180 steel plant locations for 6 years to estimate the average emissions from detected plumes using the divergence method. The models identified up to 136.1% more plumes than APE. Comparing the estimated emissions with the European Pollutant Release and Transfer Register (E-PRTR) dataset shows the ResNet-44 model achieves the lowest bias (1.20 kg/s) and standard deviation (2.36 kg/s) compared to APE, which had a bias of 3.41 kg/s and a standard deviation of 2.40 kg/s. This demonstrates the potential of machine learning to improve plume detection and emission estimation for remote sensing of pollution from space.

Contents

Nomenclature

List of abbreviations

Abbreviation	Definition
APE	Automated Plume detection and Emission estimation algorithm
ATBD	Algorithm Theoretical Basis Document
CFM	Cross-sectional Flux Method
CNN	Convolutional Neural Network
CNSA	China National Space Administration
CrIS	Cross-track Infrared Sounder
ECMWF	European Centre for Medium-Range Weather Forecasts
EDGAR	Emissions Database for Global Atmospheric Research
Envisat	Environmental Satellite
E-PRTR	European Pollutant Release and Transfer Register
ESA	European Space Agency
EU	European Union
GEM	Global Emission Monitor
Geo-AQ	Geostationary Air Quality mission
GHG	Greenhouse gases
IODD	Input Output Data Definition
IPCC	The Intergovernmental Panel on Climate Change
LEO	Low Earth Orbit
MARS	Methane Alert and Response System
ML	Machine learning
NaN	Not a Number
NASA	National Aeronautics and Space Administration
NDC	Nationally Determined Contributions
NOAA	National Oceanic and Atmospheric Administration
PRF	Product Readme File
PUM	Product User Manual
RBF	Radial Base Function
RFC	Random Forest Classifier
Std	Standard deviation
Suomi-NPP	Suomi National Polar-orbiting Partnership
SVM	Support Vector Machine
TANGO	Twin ANthropogenic Greenhouse Gas Observers
TROPOMI	TROPOspheric Monitoring Instrument
UN	United Nations
VIIRS	Visible Infrared Imaging Radiometer Suite

List of symbols

Symbol	Definition	Unit
A	Amplitude	[-]
C	CO column measured by TROPOMI.	[mol/m^2]
D	Source of emission	[kg/s]
E	Emission	[kg/s]
f_i	Feature i in the data	[-]
f_{norm}	Normal distribution	[-]
F_{norm}	Cumulative normal distribution	[-]
$f_{\text{skewed}}(x)$	Skewed Gaussian distribution	[-]
L	Split value for a feature in a tree classifier	[-]
n	Total number of transaction lines	[-]
N_{skewed}	Amplitude of the skewed Gaussian distribution	[-]
$p[f]$	feature in data point	[-]
Q_i	CO flux through transaction line i	[kg/s]
s	Position in the cross-section	
t$_0$	Time of measurement	[s]
$\vec{u}$	U-component of wind	[m/s]
$\vec{v}$	V-component of wind	[m/s]
$\vec{w}$	Wind vector	[m/s]
x	X position in the grid	[m]
x_0	X position of the steel plant	[m]
y	Y position in the grid	[m]
y_0	Y position of the steel plant	[m]
z$_i$	Cross-section i	[m]
α	Skewness factor	[-]
Δx	Distance between the grid points	[m]
δC^i_{co}	Area of the CO enhancement in the cross-section	[kg/m^2]
θ	Direction of the pixel compared to the source	
θ_0	Direction of the wind	
μ	Average value	[-]
μ_E	Average emission estimate	[kg/s]
σ_a	Standard deviation of parameter a	[-]

1 Introduction

This thesis project was done in cooperation with SRON. SRON is a non-profit organization that does research in the field of space science. Their main research fields are Earth observation, exoplanets and astrophysics. This project takes place in the field of Earth observations at SRON and is about the analysis of CO plume emissions from industry with the help of the APE algorithm. APE stands for Automated Plume detection and Emission estimate algorithm and is designed to work on TROPOMI CO concentration data. TROPOMI is an instrument onboard the Sentinel-5P satellite. SRON has contributed to the creation of the instrument. The Sentinel-5P satellite is part of the ESA's Copernicus program. This program is focused on Earth observation [1]. With the TROPOMI instrument trace gasses like CO can be measured. CO plays a role in climate change but is not a greenhouse gas itself. Nevertheless, CO is harmful and therefore the emissions of this gas should be monitored. SRON is responsible for the CO data product of TROPOMI [2].

The thesis project aims to further the development of the APE algorithm. This project focuses on improving industrial sources' plume detection and emission estimation. For this project, the main focus was on researching machine learning methods as a replacement or addition to the current plume detection implemented in APE. A smaller focus was put on developing a new emission estimation method using the divergence method. Ultimately with the development of the APE algorithm, the aim is to accurately estimate the emissions from industry to help with the reduction of CO emission worldwide. This thesis project is part of this pursuit by SRON.

The thesis is structured as follows, chapter 2 explains the background information for this project. This is followed by the research questions and research plan for the thesis in chapter 3. In chapter 4, the implementation methods for the machine learning approaches are discussed as well as the augmentation algorithm. The creation of three datasets for the training and testing of the machine learning methods is denoted in chapter 5. The development of the new emission estimation methods is explained in chapter 6. In this chapter, a method to calculate a time series of emissions using the individual divergence of the single overpass data is explained. This is followed by an explanation of how to use an entire time series of emission estimations to estimate an average emission of the plume data for the entire time series without the use of a plume detection method. The different plume detections are compared with each other in chapter 7. The answers to the research questions of the project are written in chapter 8. The chapter ends the thesis by giving a future outlook on the development of APE.

2
Background

This chapter covers the background information for the thesis project. This information comes from the internship completed before the thesis and the literature study done in the early stages of the thesis [3]. In section 2.1 a brief description of the climate change and air pollution problem is given. Then an overview of space-based monitoring systems is given insection 2.2. This is followed by section 2.3, providing a description of TROPOMI onboard Sentinel-5P. In section 2.4 dataset products and retrieval are discussed. This is followed by a description of the APE algorithm in section 2.5. The final section discussed different improvements for the plume detection of APE and can be found in section 2.6.

2.1. Climate change and air pollution

Earth's atmosphere houses a variety of different chemical compounds. However, 99% of the atmosphere consists of only three constituents. These are nitrogen, argon and oxygen gas. The remaining one percent of the atmospheric contents are known as trace gases. Trace gases include carbon dioxide (CO_2), carbon monoxide (CO), methane (CH_4), nitrogen dioxide (NO_2), Ozone (O_3) and many more. Of which some are considered to be greenhouse gases. Large accumulations of greenhouse gases enhance the greenhouse effect that warms Earth's average temperature [4].

In recent years, the awareness of the need for policy and climate action has significantly increased. Since 2015, 195 out of the 198 members of the United Nations (UN) agreed to limit the increase of Earth's average temperature to 1.5°C by signing the Paris Agreement, with Earth's average temperature measured from 1850-1900 as benchmark average value [5]. According to the most recent Intergovernmental Panel on Climate Change (IPCC) report the average temperature of Earth has increased by 1.1°C in 2010-2020 since 1850-1900 [6]. According to the Climate Action Tracker, in December 2023 the world has already warmed by an average temperature of 1.3°C[1]. These temperature increases show an accelerated need for solutions. In January 2025, it was reported by Copernicus that Earth's average temperature in 2024 had already increased by 1.6 °C [7].

The IPCC report also points out that the greenhouse gases (GHG) that cause climate change are not significantly decreasing as stated before these GHGs enhance the greenhouse effect, which means that these gases absorb longwave Earth radiation and emit thermal heat as a result. Without GHG, the energy due to Earth's radiation would have escaped into space instead of trapping them in the atmosphere. The greenhouse effect is good as it allows life to exist on Earth. However, as the amount of GHG increases, so does the thermal heat emitted from these gases [8]. In Figure 2.1, it is shown that the current Nationally Determined Contributions (NDC) of governments around the world to decrease global GHG emissions are not enough. Even worse, the figure shows an increase in GHG emissions. The report also outlines the risks of the increase in temperature. Effects include extreme heat, droughts, rising sea levels and an increase in extreme weather events. To avoid this future, a large decrease in GHG emissions is necessary [9]. The IPCC figure shows that the current GHG emissions will have to be decreased by 43% by 2030 to ensure the maximum average temperature increase of 1.5°C.

[1] Retrieved on 1-5-2024, `https://climateactiontracker.org/publications/no-change-to-warming-as-fossil-fuel-endgame-brings-focus-onto-false-solutions/`

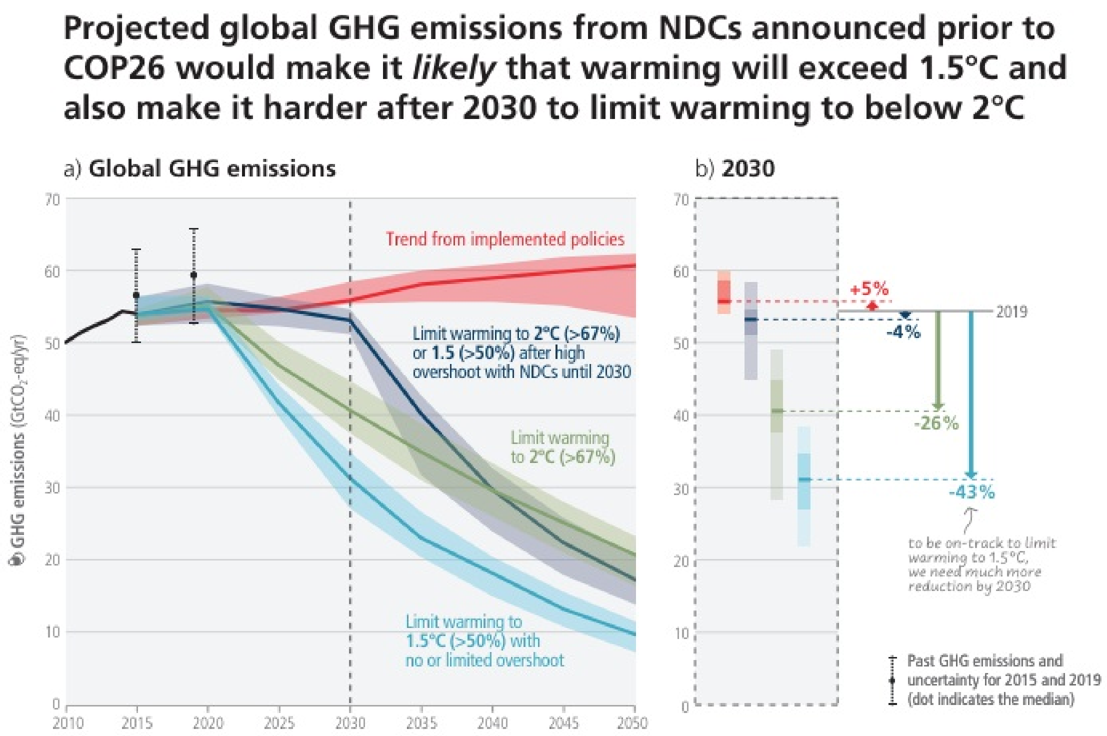

Figure 2.1: Figure 2.5 from the IPCC report showing that currently implemented climate policies are not enough [6]. Subfigure a) shows the prediction of the emissions over time based on the actions of governments and subfigure b) shows the needed reduction of the emissions by 2030 to reach certain targets.

The emissions that cause climate change also affect the quality of the air. It is well known that living near industrial sources increases health risks [10]. The quality of the air is determined by increased pollutants, such as trace gases and particulate matter, that constitute the atmosphere in the area. Thus there is a need to monitor the concentrations of trace gases to ensure that health risks are minimized. Most governmental bodies around the world have regulations on air quality [11]. These include standards for the concentration of gases that people are allowed to be exposed to for specific amounts of time, the banning of the use of technology that is known to cause excessive pollution, taxes on the emissions of certain gases and cap and trade systems for certain emission pollutants [12].

To ensure compliance with climate goals, it is necessary to monitor the amount of emissions emitted globally. To this end, institutions and companies have started to publish emission inventories. The issue with these is that they are often not specific enough to verify with measurements or are not based on observations [13]. Therefore, an independent way to estimate the emissions of industry is needed. There have already been some successes with independent monitoring of certain GHGs. In particular, the Methane Alert and Response System (MARS) is a monitoring system that uses satellite data to detect methane leaks. After these are found the relevant authorities are contacted to ask them to deal with the problem [2].

2.2. Space-based monitoring of trace gases

Possible ways to measure trace gases are using in-situ, airborne and satellite measurements. All of these methods are ways to monitor trace gases on various scales. However, due to the global scale of the problem, most of the methods are not appealing as they focus on the regional scale. The task of trace gas monitoring needs to be done on a global scale, especially when considering the global targets set to fight climate change. It is important to note that GHGs mix in the atmosphere and stay there for

[2]Retrieved on 1-5-2024, `https://www.unep.org/topics/energy/methane/international-methane-emissions-observatory/methane-alert-and-response-system`

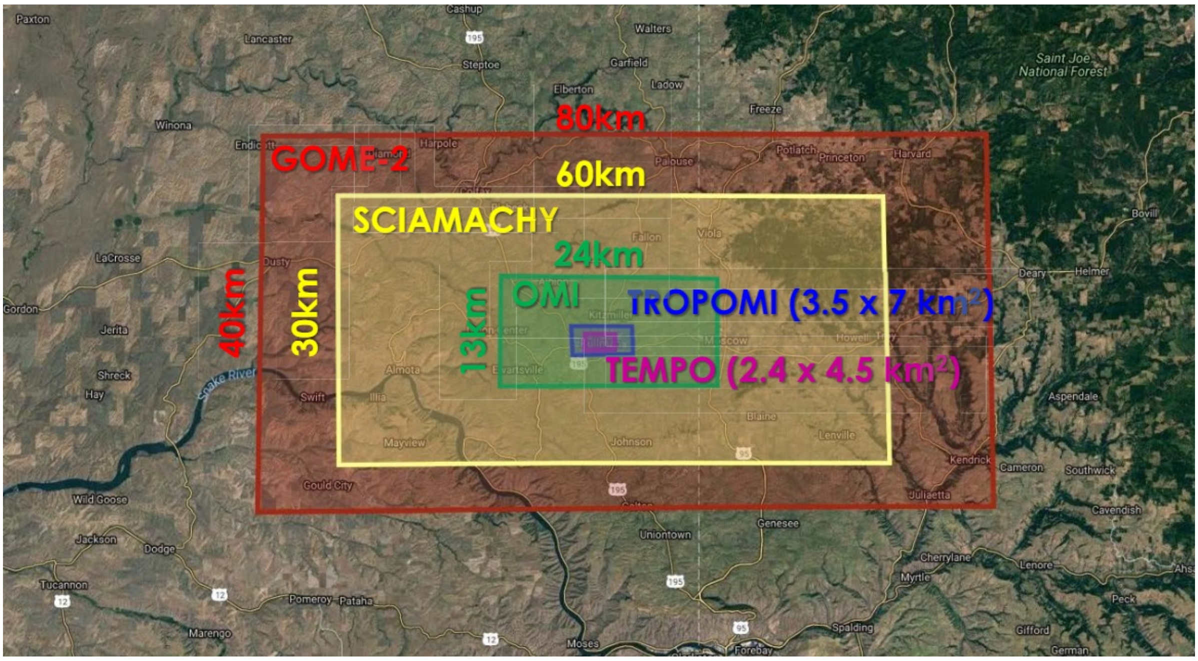

Figure 2.2: Improvement of resolution of trace gas satellite measurements over time[17]

a long time, making it hard to trace the origin of the pollutants. According to the National Aeronautics and Space Administration (NASA), CO_2 stays in the atmosphere of Earth for 300 to 1000 years [3].

The global coverage that satellites provide is preferable over regional measurement strategies, however, it is still useful to use regional measurements for certain areas to improve existing models and to make sure that the global models are working correctly. Satellites have long been used to measure trace gases for decades. Earlier missions include the Aura mission from NASA and the Environmental Satellite (Envisat) mission from the European Space Agency (ESA). These missions paved the way for further understanding the effects of trace gas pollution [14] [15]. With increasing demands for more accurate data in the fight against climate change, the European Union has significantly increased the number of Earth observation missions by working on the Sentinel program as a part of the Copernicus Earth observation system. The program aims to supply timely, accurate and accessible information to manage the environment, study the effects of climate change and provide civil security. Ultimately the program should aid the European Union in sustainable policy-making[4]. The China National Space Administration (CNSA) launched the Gaofen-5 satellite. The purpose of this satellite is to measure CO_2 emissions [16]. There are also commercial satellite operations that measure trace gases in the atmosphere. For example, the company GHGsat can be contacted to perform monitoring operations[5].

Over the years the technology of the different satellites has been improving. In Figure 2.2 the spatial resolution of different missions is shown with respect to instruments from various missions. The GOME-2 instrument is from the MetOp mission, the SCIAMACHY instrument was developed for the Envisat mission, OMI was created for the Aura mission, the TROPOMI instrument is onboard the Sentinel-5P satellite and TEMPO is mounted on top of an Intelsat 40e satellite.

Particularly interesting Sentinel missions are Sentinel-4, Sentinel-5, Sentinel-5P and Sentinel-CO2M. These missions have been designed to monitor changes in the atmospheric composition. Sentinel-4 focuses on the air quality over Europe and has an hourly data rate[6]. Sentinel-5 and Sentinel-5P focus on monitoring the concentration of trace gas all over the world[7][8]. The Sentinel-CO2M mission will

[3] Retrieved on 7-5-2024, `https://science.nasa.gov/earth/climate-change/greenhouse-gases/the-atmosphere-getting-a-handle-on-carbon-dioxide/`

[4] Retrieved on 29-4-2024, `https://www.esa.int/About_Us/Ministerial_Council_2012/Global_Monitoring_for_Environment_and_Security_GMES`

[5] Retrieved on 7-5-2024, `https://www.ghgsat.com/en/`

[6] Retrieved on 30-4-2024, `https://sentinels.copernicus.eu/web/sentinel/missions/sentinel-4`

[7] Retrieved on 30-4-2024, `https://sentinels.copernicus.eu/web/sentinel/missions/sentinel-5`

[8] Retrieved on 30-4-2024, `https://sentinels.copernicus.eu/web/sentinel/missions/sentinel-5p`

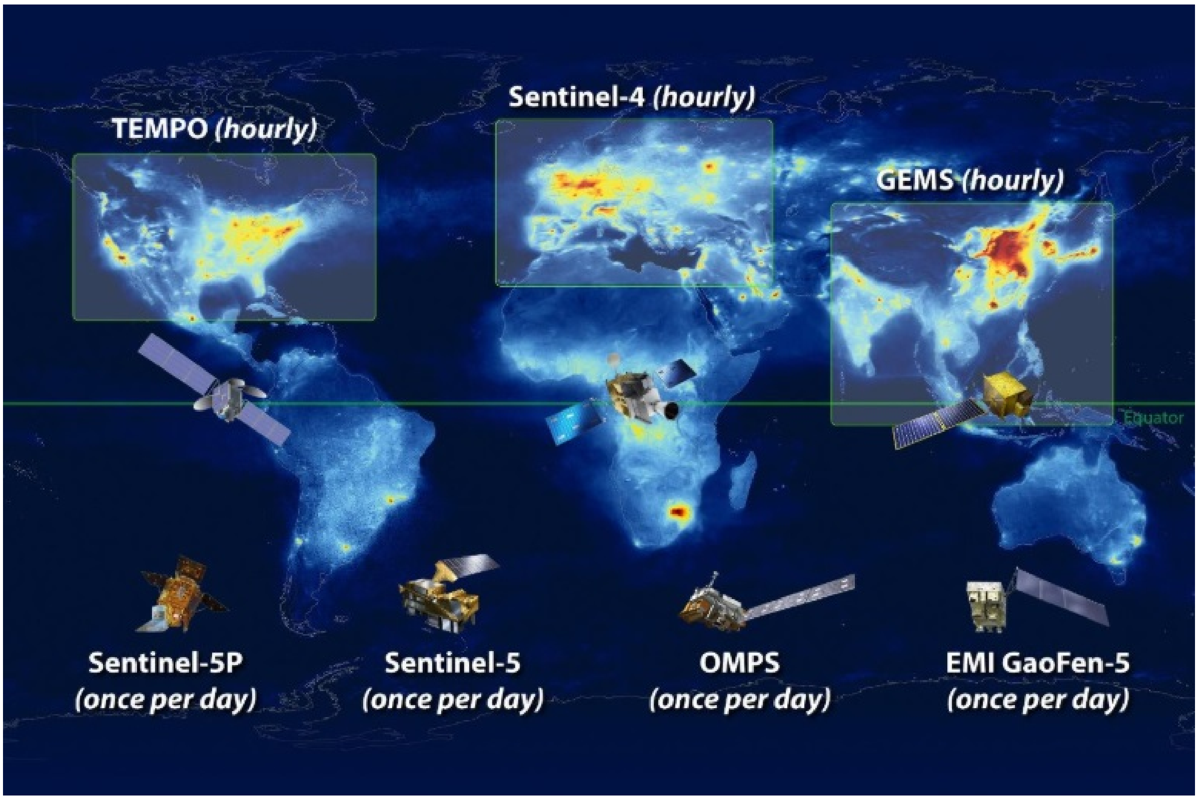

Figure 2.3: Coverage of different Earth observation missions [18]

focus on anthropogenic climate change caused by the emission of CO_2[9]. The Twin ANthropogenic Greenhouse Gas Observers (TANGO) mission, not a part of the Sentinel program, will use the Sentinel data to support its operation. The satellites of TANGO have a spatial resolution of 300 meters by 300 meters compared to the resolution of Sentinel-5P of 5.5 km by 7 km. TANGO will use it to focus on specific areas and does not have global coverage. To identify interesting regions to focus on, other satellite data, such as data from the Sentinel missions, will be used[10].

The distinction of Sentinel-4 compared to the other Sentinel missions is its focus on Europe. This satellite is part of the Geostationary Air Quality mission (Geo-AQ). These satellites focus on a particular area of the world and are able to make hourly measurements. These missions sacrifice their global coverage for more data from a particular area. Sentinel-5P for example has global coverage but only measures a certain place once per day while Sentinel-4 will have coverage of Europe but can make measurements once per hour. The purpose of Geo-AQ is to make the data products of the satellite consistent with each other [18]. The satellites in Figure 2.3 show that TEMPO, Sentinel-4 and GEMS are part of Geo-AQ. The other satellites in the figure are seen as complementary low-Earth orbit missions [18].

It has to be noted that most of the missions mentioned have not yet been launched. Currently, out of the four mentioned Sentinel missions, only Sentinel-5P is in operation. The other satellites are still awaiting their launches. In the near future, there will be more data available but currently, the data is limited to the Sentinel-5P data.

2.3. Sentinel-5P satellite mission

The Sentinel missions are a part of the Copernicus program, which is part of the European Union Space program. The program should facilitate effective policy-making by providing data related to a large

[9]Retrieved on 29-4-2024, `https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Copernicus_Sentinel_Expansion_missions`

[10]Retrieved 30-4-2024, `https://www.sron.nl/news/tango-missievoorstel-gepresenteerd-in-laatste-selectieronde-voor-esa-s-scout-missie/`

Figure 2.4: The Sentinel-5P satellite in space with TROPOMI pointed to Earth, image credit to ESA/ATG medialab

range of different issues[11]. The Sentinel program replaces earlier Earth observation missions that are set to retire or have since retired[12]. The program aims to cover different aspects of Earth observation that were done by the previous satellites individually, to improve upon previous capabilities and to increase the time coverage of the data. Each of the six Sentinel satellites has a different objective. The overall objective of the Sentinel program is to further our knowledge of the climate and to significantly improve the existing climate models by increasing the measured data [19].

The launch of the Sentinel-5P occurred on October 13, 2017. The Sentinel-5P satellite is meant to be the precursor mission to Sentinel-5. The mission is supposed to act as a bridge between the completed Envisat and Aura missions and the future Sentinel-4 and Sentinel-5 missions. The purpose of the mission is to measure trace gases and aerosols to monitor their concentration and study their effect on the climate on a global scale [20]. The minimum mission duration is 7 years[13]. For illustration purposes, the Sentinel-5P satellite is shown in Figure 2.4.

The satellite is in low Earth orbit (LEO) and flies in a near-polar sun-synchronous orbit with an inclination of 98.7° at a height of 824 km. At 13:30 Mean local solar time, the satellite crosses the equatorial ascending node[14]. This is presumably for trace gas measurements as it would enable the satellite to take measurement of Europe in the afternoon when emissions would be greater. Sentinel-5P flies in loose formation with NASA's Suomi National Polar-orbiting Partnership (Suomi-NPP) satellite. This is to have synergy between the measurements of Sentinel-5P's TROPOMI instrument and Suomi-NPP's Visible Infrared Imaging Radiometer Suite (VIIRS) and Cross-track Infrared Sounder (CrIS) instruments [18]. Examples of this synergy are VIIRS's cloud measurement data that is used for destriping and the VIIRS active fire data that could be used with TROPOMI trace gas data to estimate emissions from wildfires [21] [22].

[11] Retrieved on 29-4-2024, https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Introducing_Copernicus

[12] Retrieved on 29-4-2024, https://sentinels.copernicus.eu/web/sentinel/missions

[13] Retrieved on 6-5-2024, https://www.tropomi.eu/

[14] Retrieved on 7-5-2024, https://sentinel.esa.int/web/sentinel/missions/sentinel-5p/orbit

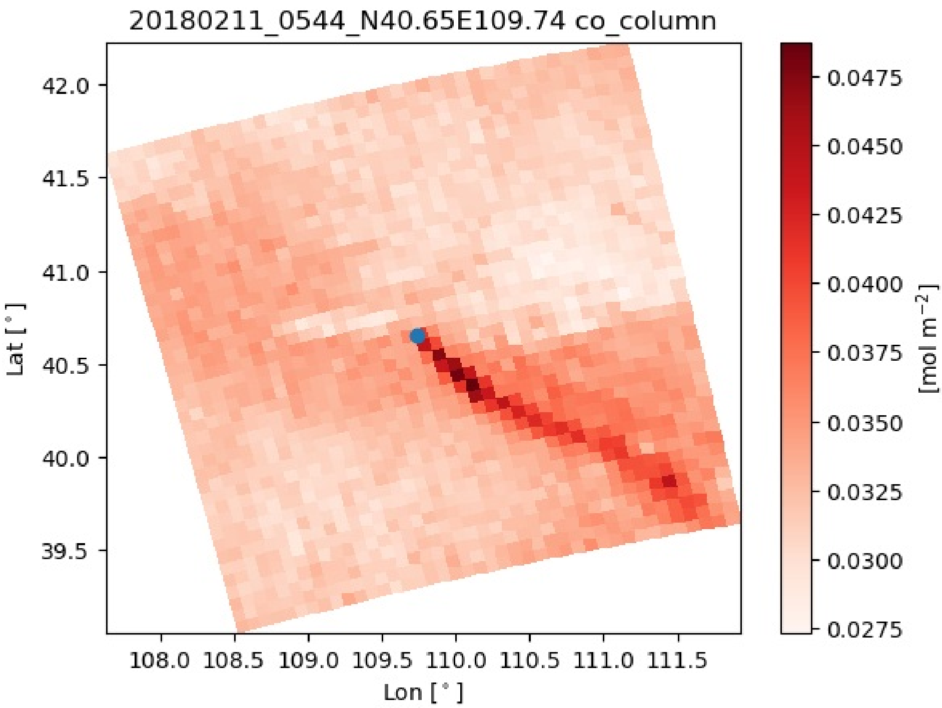

Figure 2.5: An example of a CO plume in TROPOMI data, the blue dot represents the location of a steel plant from the internship report [3]

The Sentinel-5P satellite holds the TROPOspheric Monitoring Instrument (TROPOMI) as its only payload. TROPOMI was developed by various institutes and organizations, including KNMI and SRON. The instrument was commissioned by ESA and the Netherlands Space Office[15]. In terms of coverage, TROPOMI has a swath width of approximately 2600 km, the instrument has daily coverage over the area of below -7° and above 7° latitude for radiance and reflective measurements. It can provide more than 95% daily area coverage for the latitudes between 7° and -7°[16]. The spatial resolution of the instruments is approximately 5.5 km by 3.5 km or 7 km by 5.5 km. The data products were updated in 2019. Before that update all data had a resolution of 7 km in the flight direction [23].

The instrument can measure the concentrations of trace gases like SO_2, CH_4, CO, HCHO, NO_2 and aerosols among others [20]. A CO plume from TROPOMI can be seen in Figure 2.5. The TROPOMI instrument consists of four spectrometers with different spectral ranges. These include SWIR, UV, UVIS and NIR. The last three spectrometers are included in the UVN module. The specifics of each spectrometer can be found in Table 2.1. The trace gases that TROPOMI can measure at each wavelength compared to previous missions are shown in Figure 2.6. It has to be noted that even though TROPOMI measures less wavelengths than other missions, it has a much better spatial resolution which can be seen in Figure 2.2. With these measurements, in addition to the quantification of emissions, a large amount of atmospheric phenomena can be studied. These include the effect of the COVID lockdowns and the Australian Black Summer on the atmosphere among others[24][25]. The MARS project, mentioned in section 2.1, also uses Sentinel-5P data.

2.4. TROPOMI L2 data products

The TROPOMI datasets discussed in this section are kept to the ones that are available for downloading, which are the L1B and the L2 data products of the instruments[17]. These dataset levels represent the processing done to the instrument output data. L0 data represents the raw data measured by the instrument. For TROPOMI this data is already time-ordered and archived but unavailable for use. L1B

[15] Retrieved on 14-5-2024, `https://www.tropomi.eu/`

[16] Retrieved on 8-5-2024, `https://sentinel.esa.int/web/sentinel/missions/sentinel-5p/geographical-coverage`

[17] Retrieved on 16-5-2024, `https://sentinels.copernicus.eu/web/sentinel/technical-guides/sentinel-5p/products-algorithms`

Table 2.1: Overview of the spectrometers of TROPOMI from (Reshi et al., 2024[23])

Spectrometer	UV		UVIS		NIR		SWIR	
Band ID & parameters	1	2	3	4	5	6	7	8
Spectal range [nm]	267-300	300-320	305-400	400-499	661-725	725-786	2300-2343	2343-2389
Spectral resolution [nm]	0.45-0.50		0.45-0.65		0.34-0.35		0.227	0.225

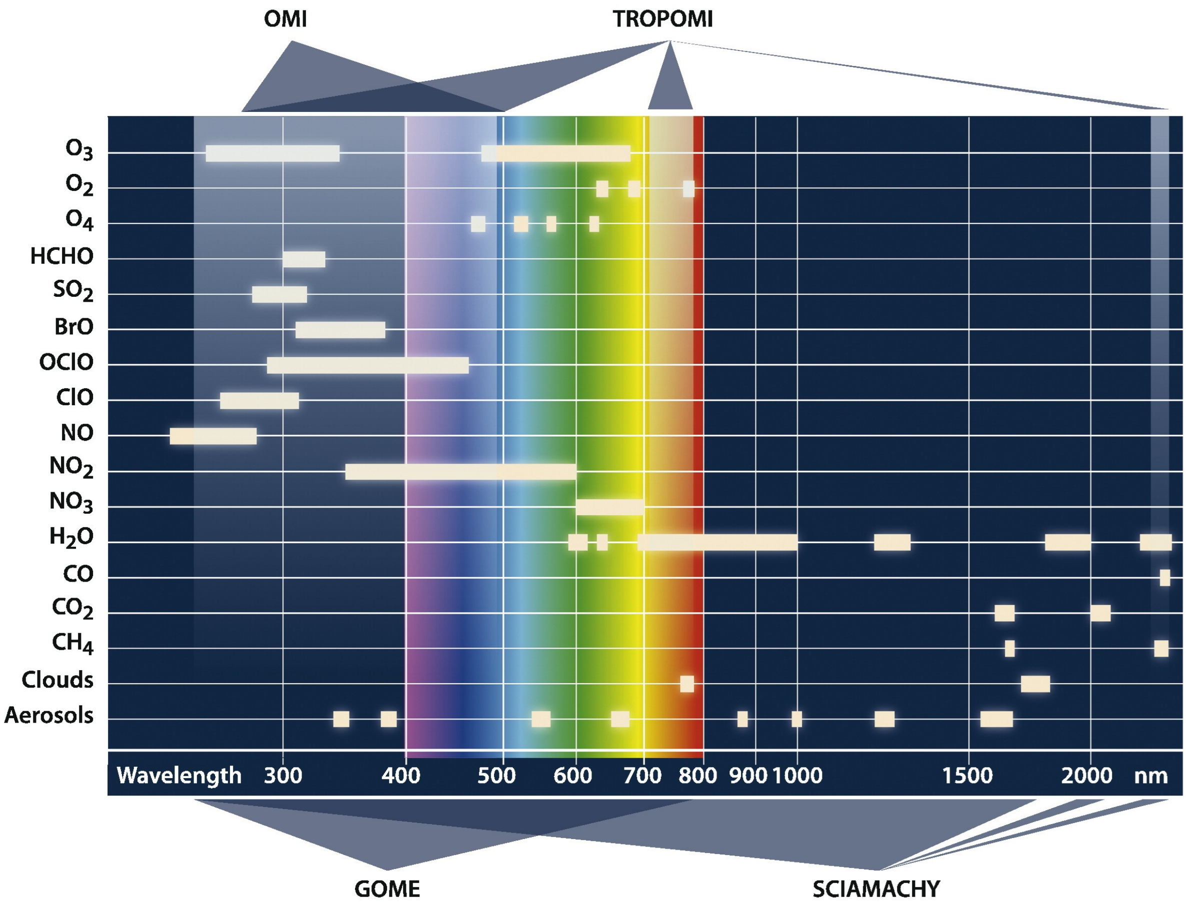

Figure 2.6: Comparing the capabilities of TROPOMI with previous mission instruments from (Veefkind et al., 2012 [20])

Table 2.2: Available L2 TROPOMI data products

Sentinel-5/UVNS Level-2 product	Parameter(s)	Distribution
O3	Ozone (O3) total column, tropospheric column, stratospheric vertical profile	To all users
NO2	Nitrogen dioxide (NO2) total column, tropospheric column	To all users
SO2	Sulfur dioxide (SO2) total column, layer height (TBC)	To all users
HCHO	Fomaldehyde (HCHO) total column	To all users
CHOCHO	Glyoxal (CHOCHO) total column	To all users
CH4	Methane (CH4) total column	To all users
CO	Carbon monoxide (CO) total column	To all users
Cloud	Cloud effective fraction, effective height, cloud mask	To all users (TBC)
Aerosol	Aerosol UV absorption index, layer height, optical depth (TBC)	To all users
Surface	Surface effective albedo, scene heterogeneity	To all users (TBC)
UV	UV spectrally resolved irradiance at surface, UV index	To all users

is the radiance and irradiance data and L2 is the trace gas and aerosols concentration data. The L1B datasets are used to generate the L2 concentration data. For different data products a retrieval is used to convert the radiance and irradiance data into the intended product. The available L2 datasets can be found in Table 2.2. This table is taken from the Sentinel Online website[18]. The TBCs, which probably stand for To Be Continued but this is not mentioned, in the table represent datasets being worked on, have not fully been processed and will possibly not be updated in the future. These imperfect datasets can be found on the S5P-PAL Data Portal[19]. The L2 datasets can be found on the Copernicus Dataspace website[20].

The files are in NetCDF-4 format and the general format of level 2 data can be seen in Figure 2.7. In this figure, the root level, first level group, second level group and third level group represent the data structure, not actual variables or values. The ... in the figure presumably denotes the product-specific variables and data that are universal for every data product[26].

A guide on each L2 dataset can be found on the Sentinel Online website. There are several types of guides to inform potential users of the capabilities of the data for each available dataset. These guides are a Product User Manual (PUM), an Algorithm Theoretical Basis Document (ATBD), an Input Output Data Definition (IODD) and a Product Readme File (PRF). The PUM is supposed to give users the technical information of the products, the ATBD gives information on the retrieval algorithm used to create the products, the IODD describes the input and the output data used to create the products and the PRF is a shorter document that describes the changes between product versions and the overall quality of the product[21].

As stated before the datasets are of particular interest due to their use in scientific research for the monitoring of emissions of trace gases. An algorithm that studies the trace gas CO using L2 TROPOMI data is called Automated Plume detection and Emission estimation algorithm (APE). This algorithm is discussed further in section 2.5.

[18] Retrieved on 20-5-2024, `https://sentinels.copernicus.eu/web/sentinel/missions/sentinel-5/data-products`
[19] Retrieved on 21-5-2024, `https://data-portal.s5p-pal.com/`
[20] Retrieved on 22-5-2024, `https://dataspace.copernicus.eu/`
[21] Retrieved on 21-5-2024, `https://sentiwiki.copernicus.eu/web/s5p-products`

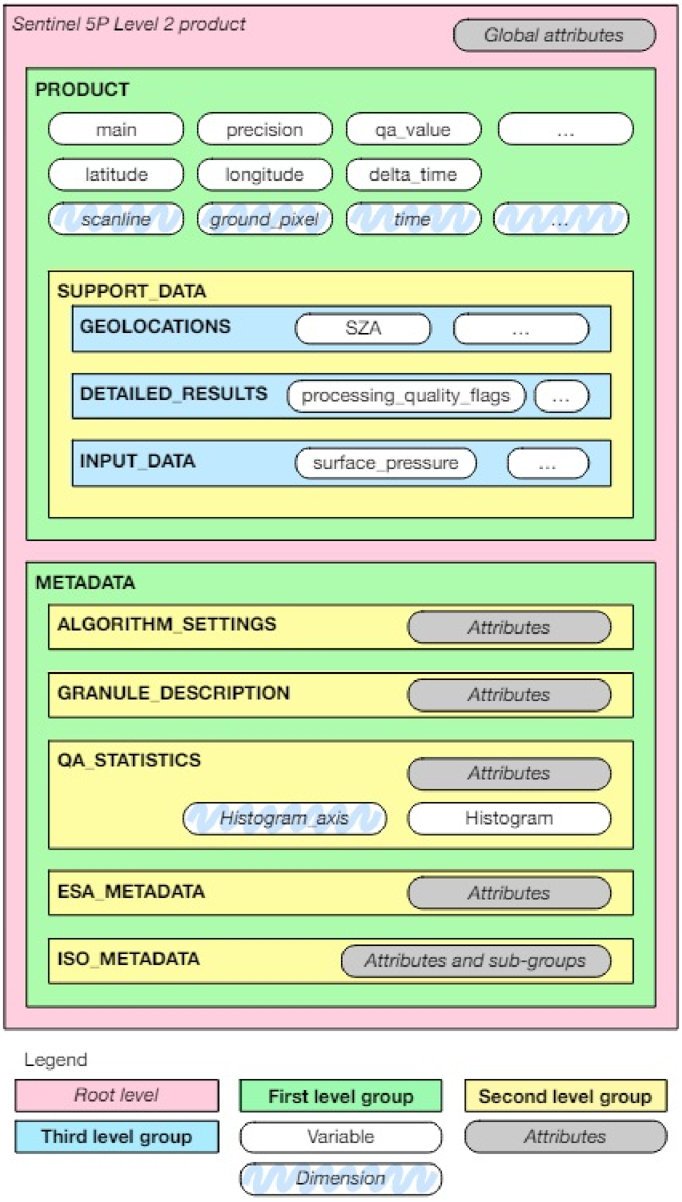

Figure 2.7: The general structure of TROPOMI level 2 data products from the CO L2 Product User Manual [26]

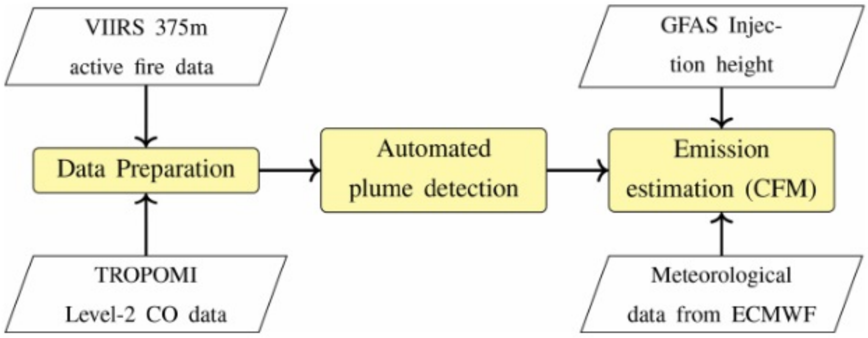

Figure 2.8: Flowchart of APE from Goudar et al. 2023 [27] showing the three parts of APE in yellow.

2.5. APE algorithm

APE is a tool created by SRON to estimate CO emissions. The tool is capable of calculating emission estimates based on the L2 TROPOMI CO concentration column data. The intended purpose of APE is to function similarly to MARS as described in section 2.1 but for CO emissions instead of CH_4. The point of the tool is thus to detect in near real-time the emissions of CO. To this effect, the tool should be able to quickly process the data from TROPOMI, detect the visible plumes and estimate the emissions of these plumes. For this purpose, APE should be able to make accurate emission estimations relatively fast. This means that the focus is not on making the most accurate emission estimations possible but on balancing the computational speed with the accuracy of the emission estimates. Thus any improvement for APE has to take these considerations into account. In subsection 2.5.1, the first two parts are explained which encompass the data downloading, processing and plume detection. The third part of APE is explained in subsection 2.5.2. This part goes over the emission estimation using CFM and the average divergence. The final subsection in subsection 2.5.3 goes over issues encountered in the process of estimating the emissions.

2.5.1. Automated plume detection

Scientific researchers at SRON have been developing APE for the quantification of CO emissions from wildfires and industry using data from the TROPOMI instrument onboard the Sentinel-5P satellite and have written one paper on its development. This subsection is largely based on the information provided by that paper by Goudar et al. 2023 [27]. Currently APE consists of three parts. A flowchart of APE is presented in Figure 2.8. This flowchart shows the three different parts of APE in yellow that quantify wildfire emissions.

The first part of APE consists of the data downloading and processing part. To start the process, APE uses two main inputs, which are date and location. Thus the latitude and longitude coordinates of the location of interest and the time period of investigation. When given this input APE will go through the TROPOMI orbits and cut out granules of 41 by 41 size pixels. To avoid large pixel sizes, APE restricts the pixel to be smaller than 12 km in swath width direction. In the 41 by 41-pixel granules, 80 % of the data must have a quality value of 0.5 or higher. The definition of the quality value parameter is explained in Figure 2.5.3. In short, it is a measure of the cloud conditions during the measurement of a CO concentration value. For the 7 by 7 pixels around the source, 85% must have a quality value of higher than 0.5 [27].

As seen in the flowchart for wildfire plumes in Figure 2.8, APE uses the level-2 CO TROPOMI data and the VIIRS 375m active fire data. When investigating CO, using CO TROPOMI data makes sense. The use of the VIIRS active fire dataset is less obvious and thus requires explanation. As explained in section 2.3, VIIRS is an instrument onboard the Suomi NPP satellite. This satellite is run by both NASA and the National Oceanic and Atmospheric Administration (NOAA). The Suomi NPP and Sentinel-5P satellites are flying in the same orbit with a temporal separation of 3.5 minutes [27]. This creates a strong synergy between the data products of the two satellites. Thus VIIRS datasets can be used almost seamlessly in conjunction with TROPOMI datasets. An example of using the active fire product on TROPOMI data can be seen in Figure 2.9. APE is also capable of processing the data of industrial sources. During the internship a list of 1115 steel plants was created [3].

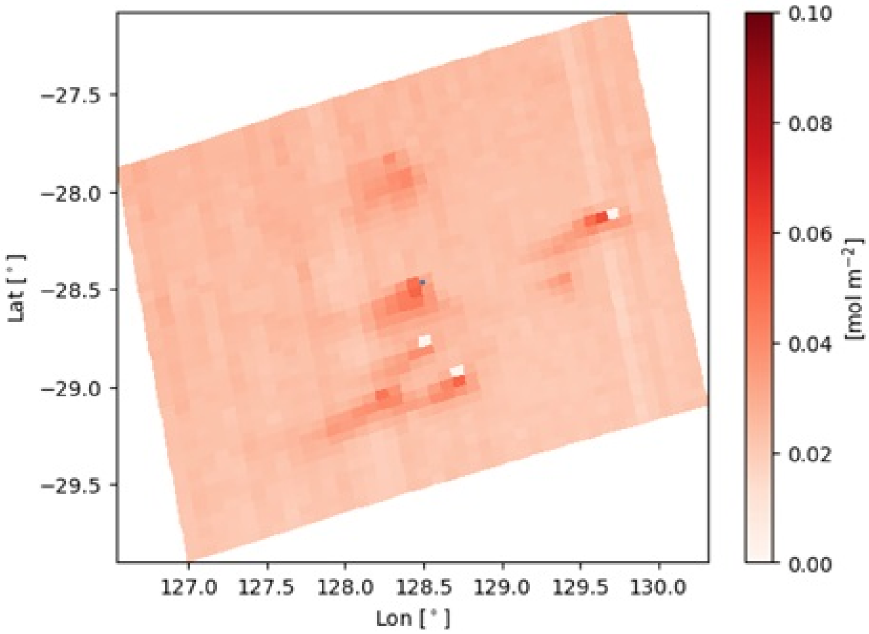

Figure 2.9: Example of using VIIRS active fire data with TROPOMI level 2 CO-data. The blue dot represents the location of a plume as detected by APE. The data is not destriped and thus shows lines in the latitudinal direction.

The second part of APE is the plume detection algorithm. It relies on the Marker-based watershed transform method. In this method, the CO concentration area is seen as a topographic map where the concentrations represent heights. A plume in the data would appear as peaks of a mountain range while low concentration would come across as valleys. An overview of the Marker-based watershed transform method can be found in Figure 2.10. The method for plume detection needs to calculate two maps before detecting the plume. These are the gradient map and the marker map. Before the plume detection can take place, these maps have to be calculated first. In Figure Figure 2.10a the CO column data processed in the first part from raw TROPOMI data is shown. The gradient map as shown in Figure Figure 2.10b is created by applying Gaussian smoothing to the CO column data, to remove extreme concentration values from affecting the gradient image and then using a Sobel operator. The next step is to create the marker image. This is done in several steps. First, the values of the CO column data of Figure Figure 2.10a that are larger than the median value of the smoothed CO column data or smaller than the mean of the center 15 by 15 pixels of the CO concentration figure are obtained. Second, the pixels that conform to these cases are clustered together by checking which pixels are next to each other. This creates the image in Figure Figure 2.10c. Here each different connected region receives a different label as seen in the color bar of the figure. The marker image can then be found focusing only on the enhancements that are located near the plume and results in Figure Figure 2.10d. The gradient map and the marker image are then inputted into the watershed algorithm which labels and segments the remaining plume candidates as seen in Figure Figure 2.10e. Using the location of the plume from the Active Fire data of VIIRS the correct plume is found. The resulting plume mask of the plume is shown in Figure Figure 2.10f. After the plume is detected, several further tests are performed to ensure Emission estimation.

Firstly, the length of the plume has to be larger than 25 km otherwise the emission is unable to be estimated. This has to do with possible issues with the plume shape. Short plumes could have strange shapes which makes automated processing hard and thus these are ignored for emission estimation. Secondly, there must not be ten or more plumes at a short distance from the source as this will affect the background value and thus the enhancement that the plume represents.

An issue that arises when using this method is that a plume needs to be present in the data for the method to work. If this is not the case, the method can confuse certain enhancements in the data for a plume, resulting in an emission estimation that is meaningless. For wildfire plumes, it is possible to rely on the VIIRS active fire data to find the locations of plumes. For industrial sources, this data cannot be

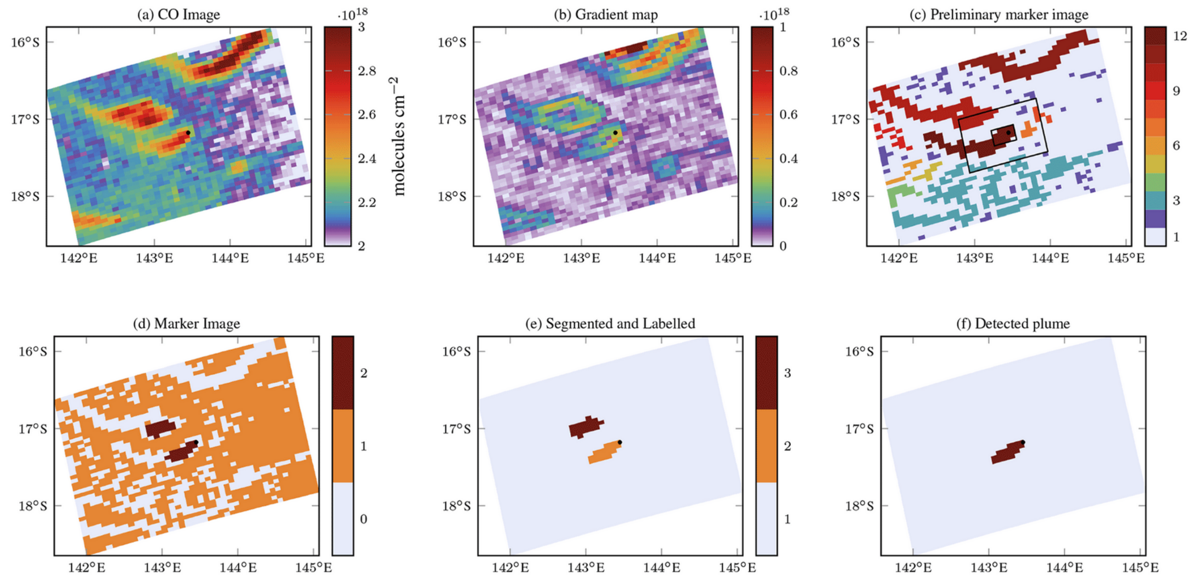

Figure 2.10: Plume detection using the marker-based watershed transform method in APE from (Goudar et al.,2023 [22])

used. It is not known for every steel plant when they are emitting pollutant. Therefore, APE has the tendency to allow for large amount of meaningless emission estimations for industrial sources.

In the Goudar et al., 2023 [27] paper, it is explained that machine learning for plume detection has not been implemented due to the lack of available training data. Therefore by creating such a dataset, the current plume detection could potentially be replaced. However, this dataset should be optimized. The optimization depends on the particular machine learning method that is implemented. A machine learning approach for plume detection could get rid of the problem of plume-less data being used for emission estimation.

2.5.2. Emission estimation

As written in Goudar et al. 2023 [27], the third and final part of APE consists of a cross-sectional flux method (CFM) that estimates the emission of a plume. An explanation of this method using the information in the paper can be found in the following subsection. What's not written in any paper is that APE is also capable of applying an average divergence method to calculate the average emission over a chosen time period. This method is explained in the second subsection.

The CFM method in APE

The basic idea behind the CFM method is to determine the emission of a point source by estimating the rate of change in the concentration of the pollutant over different locations in the plumes while using an estimate of the velocity of the wind at those locations. This way the spread of the pollutant near the source can be used to estimate the emission of pollutant.

The input for the third part of APE consists of a 41 by 41-pixel granule of CO concentration data, a plume mask array, the time of the emission and the location of the point source emitter. An example of the CFM method can be seen in Figure 2.11. The location and the time of emission are used to get the wind field near the source of the plume. Using the plume mask, the point source emission point and the wind direction, a line is drawn from the emission point to the end of the plume. This line shows the plume's downward wind direction in the mask. Perpendicular to the downward wind direction line several other lines are drawn. These lines are known as the transaction lines and they show up as dashed lines in Figure 2.11.

The emission is calculated by taking the mean of the fluxes through the transaction lines or cross-sections which is defined in Equation 2.1. Here E is the emission in kg/s while n is the total number of transaction lines and Q_i represents the CO flux through transaction line i of the plume in kg/s. This term is further defined in Equation 2.2. Here the flux term Q_i is defined as the integral over the area of the CO enhancement in the cross-section in kg/m^2 represented by δC^i_{CO} times the wind speed in m/s represented by v^i. The enhancement of the CO means the CO concentration data is subtracted by the

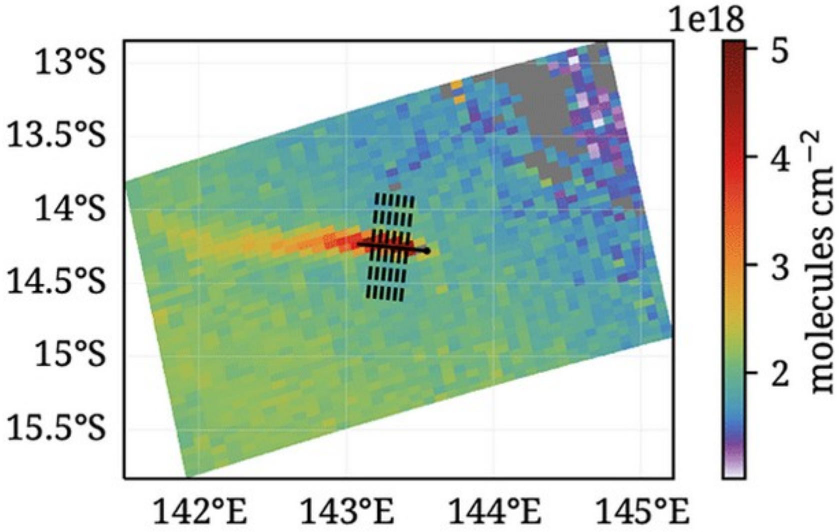

Figure 2.11: CFM method on a plume from TROPOMI. The plume line is defined as the black line in the figure. The transaction lines are shown as the dashed lines in the figure. This figure is taken from Figure 4 of Goudar et al. 2023 [27]

background CO concentration. The equation shows the dependency of the δC^i_{CO} term on the position in the cross section s and the time of measurement t_0 while the wind speed velocity is also dependent on the plume height in the cross-section z_i. The paper further explains the need for a Lagrange model to get the correct height of the plume. This is because depending on the height of the plume, the wind speeds might be significantly different. For industrial sources, the assumption is made that the height of the emission source remains the same. This is because it is not to be expected that blast furnaces of steel plants move. The plume height is very important for wildfire plumes as this height can vary significantly.

$$E = \frac{1}{n}\sum_{i=1}^{n} Q_i \tag{2.1}$$

$$Q_i = \int \delta C^i_{\text{CO}}\left(s, t_0\right) \cdot v^i\left(z_i, s, t_0\right) \cdot \mathrm{d}s \tag{2.2}$$

An issue arises from using the CFM method to calculate emissions related to the plume mask. The plume mask is the result of the plume detection of the marker-based watershed transform method. This method is explained in subsection 2.5.1 and relies on using an enhancement to find the plume in the data granule. By performing several tests using APE it was found that the masks had a huge influence on the emission estimate. By selecting a few more pixels to be part of the plume the emission estimate could be off by more than 100% of the initial estimate. This issue shows the sensitivity of the method. The problem can be made worse if the quality of certain pixels is too low or if pixels have the NaN value assigned to them.

At first, the sensitivity of the emission estimate on the plume mask might not seem like a significant issue. Therefore be explained that the CO column values could get changed easily as a result of a change in destriping. An example of this can be found in Figure 2.12 an Figure 2.13. The first case shows a not destriped array of CO Column in Figure 2.12a which results in the mask in Figure 2.12b. The second case shows the same CO concentration data after a destriping algorithm took out the stripes in Figure 2.13a. This CO column results in the plume mask in Figure 2.13b which is slightly different from Figure 2.12b. The difference in the plume mas is a simple four pixels near the end of the plume. According to APE the first case results in an emission of about 10.86 kg/s while the second case results in an emission of 27.96 kg/s. This is a significant difference of 17.10 kg/s. The four pixels in the mask between the two cases constitute a difference of nearly two times the first case emission. It is hard to see the difference between the two cases in the input CO column data. Therefore, the difference can be seen be, in Figure 2.14. Here the difference array is the stripes that were taken away in the second case by applying a destriping method on Figure 2.13a.

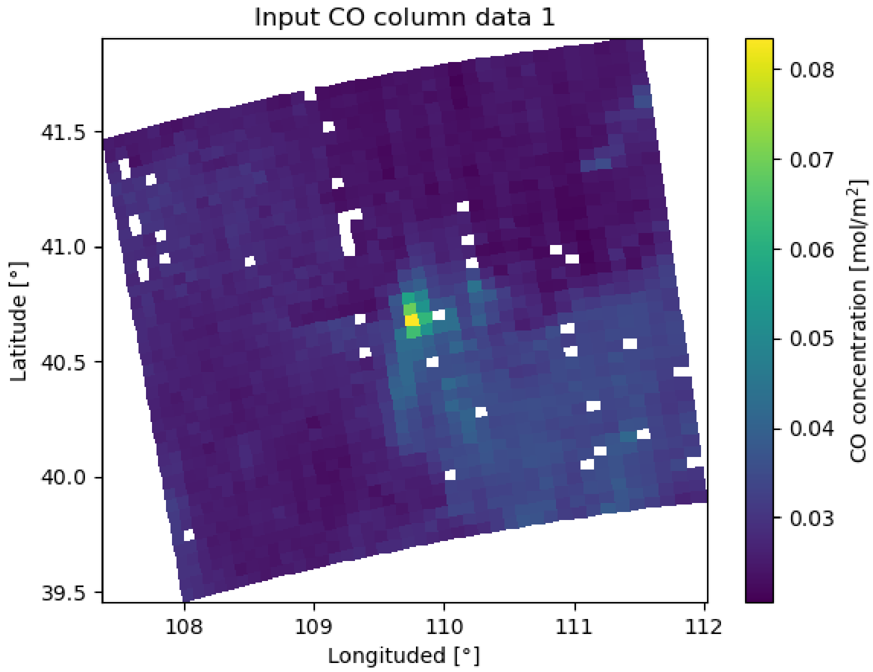

(a) Plume as detected by APE

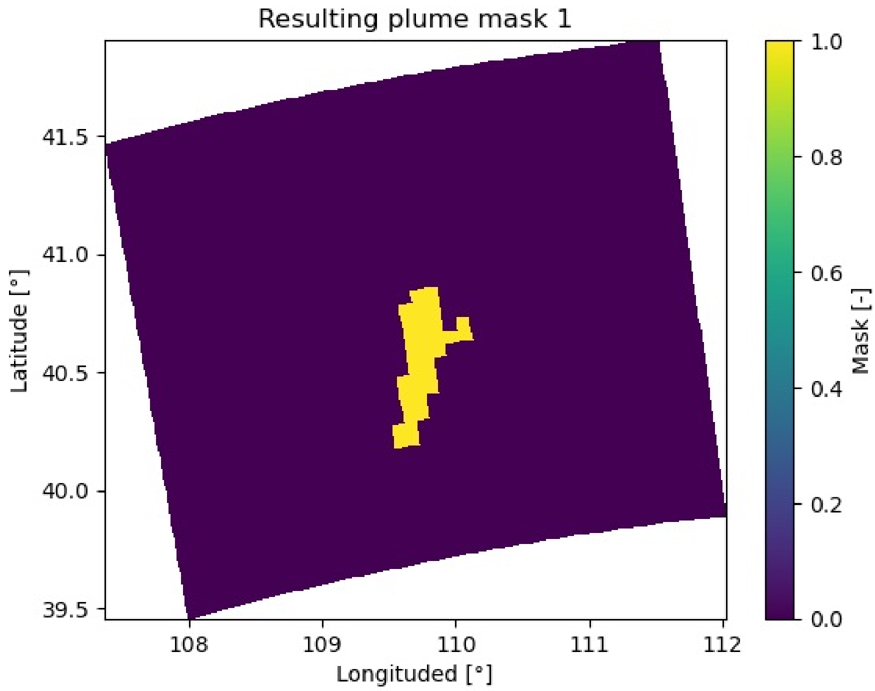

(b) Divergence of the plume

Figure 2.12: The input CO column without destriping and the resulting plume mask created by APE

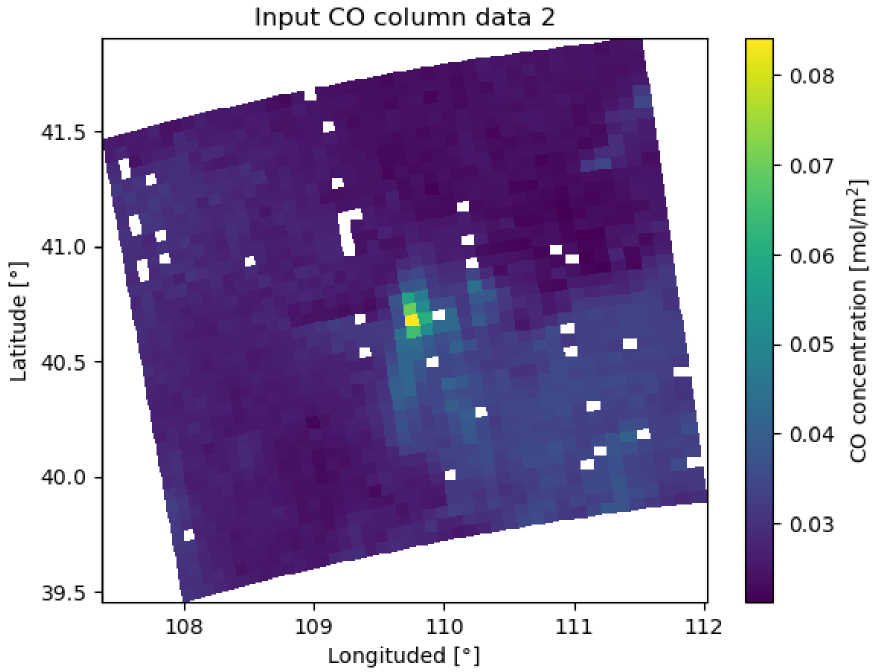

(a) Plume as detected by APE

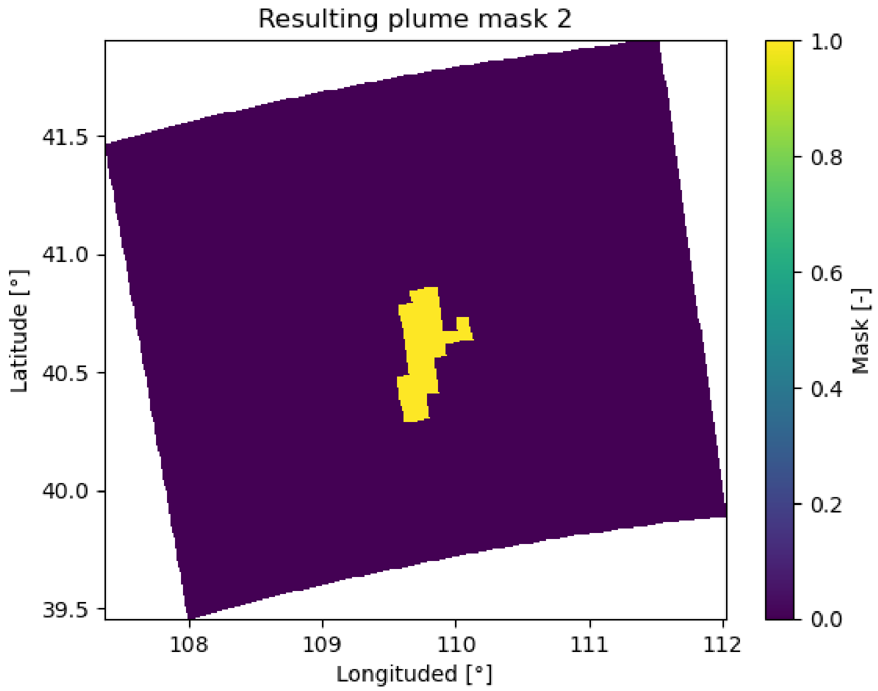

(b) Divergence of the plume

Figure 2.13: The input CO column with destriping and the resulting plume mask created by APE

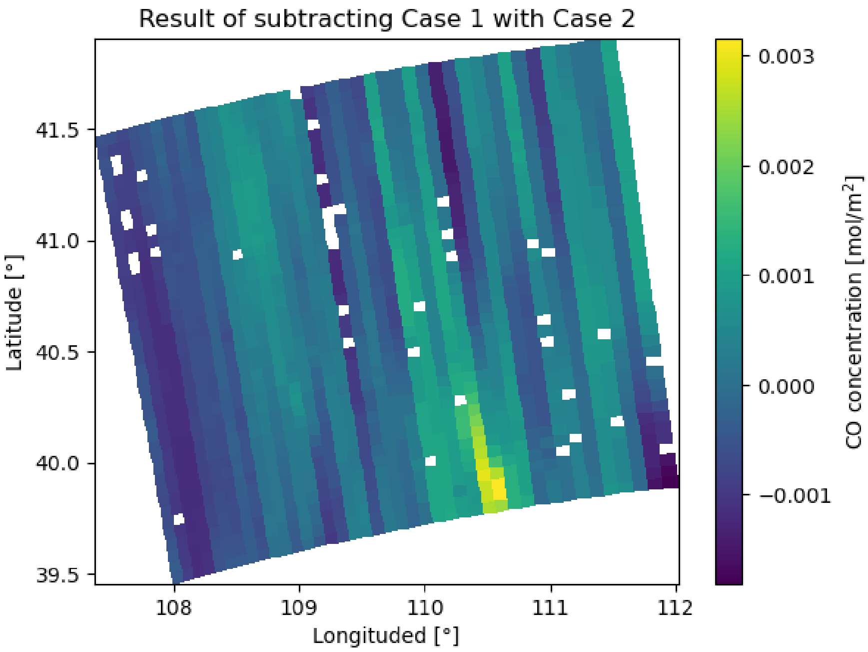

Figure 2.14: Subtracting the first case from the second case to show the effect of destriping

The accuracy of the cross-sectional flux method can be greatly increased with higher spatial resolution and thus the method should not be discarded. Especially when considering the spatial resolution of the future TANGO mission which has been announced to be 300 by 300 meters [28]. However, due to the limiting resolution of TROPOMI, which is 5.5 km by 7 kilometers, it seems clear that the cross-sectional flux method is not the most optimal emission estimation method. Another issue is the limited ERA5 spatial resolution in the ERA5 data which is only 31 km by 31 km [29]. These spatial resolutions greatly limit the accuracy of the method.

Another issue of the cross-sectional flux method is the influence of NaN data points on the results. If one or more pixels are missing in an otherwise clearly visible plume, the plume mask could be smaller resulting in a smaller plume that will be analyzed. As stated in subsection 2.5.1, the APE algorithm filters out data points based on how many NaN values exist in the 41 by 41-pixel granules as well as based on how many NaN values in the 7 by 7 pixels near the source.

The average divergence method in APE

The divergence theorem as explained in Beirle et al., 2019 [30] gives a tool to analyze the average emission of a point source. The method they provide appears simple. They define the emissions of the source as the addition of sources and sinks. Sources are locations where emissions are produced while sinks are locations where emissions dissipate. In the paper, they were able to separate different NO_x sources in the city of Riyadh, this can be seen in Figure 2.15. This method can also be utilized for CO point sources.

Since the paper focuses on NO_x rather than CO, the procedure for CO is a bit different. It takes considerably longer for CO to dissipate from the atmosphere compared to NO_x and thus the sinks of CO are not taken into consideration. Instead, the divergence term is taken to equal the emission. The equation for the divergence term can be found in Equation 2.3. In this equation, D represents the divergence in kg/sm^2, ∇ is the divergence operator, C is the column concentration of CO in kg/m^2, $\vec{w}$ is the horizontal wind in m/s, $\vec{u}$ and $\vec{v}$ are the wind directions in the latitudinal and longitudinal direction in m/s and x and y are the distance in the latitudinal and longitudinal direction respectively.

$$D = \nabla C\vec{w} = \frac{\partial\left(C \cdot \vec{u}\right)}{\partial x} + \frac{\partial\left(C \cdot \vec{v}\right)}{\partial y} \tag{2.3}$$

As can be seen in the equation, the information necessary for the calculation of the divergence is the CO column concentration and the wind field. Naturally, for the CO column concentration data,

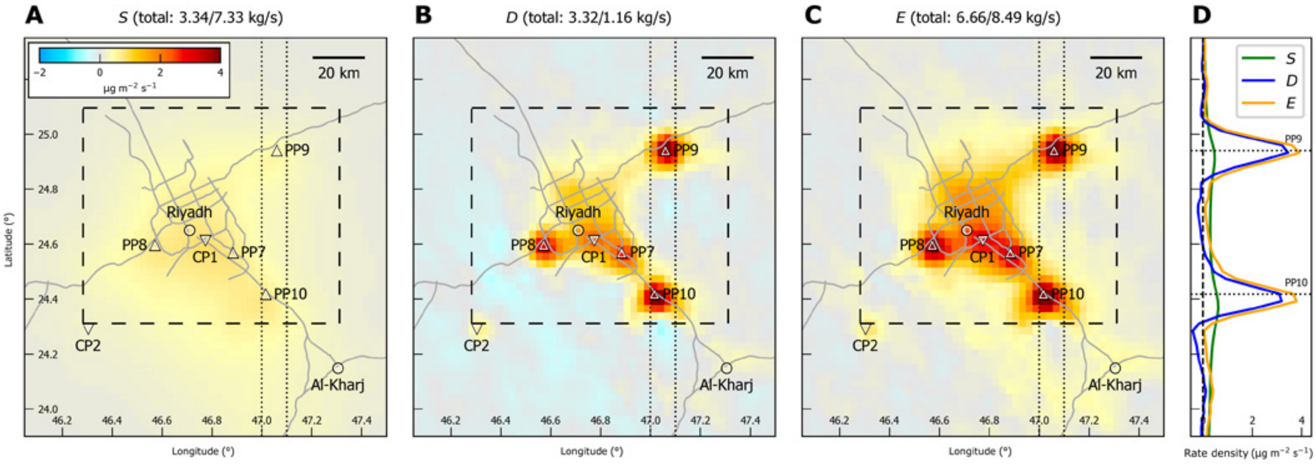

Figure 2.15: Figure **A** showing the NO_x sinks (S), **B** showing the divergence (D), **C** showing the total emission (E) over the city while **D** shows the concentration of NO_x of S, D and E at 47.05°E. The total values given on top of the figures are integrated values from the dashed line boxes and from the entire figure.

TROPOMI data was used. The wind data was taken from the European Centre for Medium-Range Weather Forecasts (ECMWF) and their dataset is called ERA5. This dataset has an hourly wind field for the entire globe with a spatial resolution of 31 km by 31 km [29].

To analyze the average divergence of a location, the divergences of a single overpass should be at the same locations. To make sure that this is the case, the individual granules are interpolated onto the same grid. This was done by creating a square 160 km by 160 km grid with the center at the emission. The spatial resolution of the divergence grid was selected to be 4 km by 4 km. The ERA5 data had to be interpolated onto the same grid. To avoid interpolation errors because the resolution of the wind data has a much higher spatial resolution than TROPOMI CO column data, it was decided to only use the wind direction value at the pixel where the point source is located.

The spatial resolution of the divergence grids was chosen to be 4 km by 4 km. This is smaller than the spatial resolution of TROPOMI, which is 5.5 km by 7 km. When averaging the values of the grid over more than 6 years of data this would not result in large integration error. For singular emission calculations from divergence, larger errors could occur. However, these are unavoidable as it is necessary to interpolate onto a new grid. This is caused by the orbit of Sentinel-5P which changes the location of TROPOMI measurement. During early work with the divergence method, a much smaller spatial resolution was chosen which resulted in artifacts showing up in the average divergence. To avoid this the resolution was updated from 1 km by 1 km to 4 km by 4 km.

The interpolation function used, cannot process NaN or masked values in the CO column and therefore a placeholder value was used for the NaN values. The placeholder value was decided to be the median value of the input 41 by 41-pixel granules. This does lead to inaccuracies in the interpolation to the new grid.

To calculate the divergence numerically a fourth order central finite scheme was used. The scheme for $\frac{\partial(C\cdot\vec{u})}{\partial x}$ can be found in Equation 2.4. A similar equation can be made for the term $\frac{\partial(C\cdot\vec{v})}{\partial y}$ by exchanging the variables.

$$\frac{(\partial C\vec{u})}{\partial x} = \frac{(C\vec{u})_{x-2} - 8\,(C\vec{u})_{x-1} + 8\,(C\vec{u})_{x+1} - (C\vec{u})_{x+2}}{12\Delta x} \tag{2.4}$$

As an example of the method, the average divergence of the Inner Mongolia BaoTou Steel Union Co Ltd was calculated and can be seen in Figure 2.16. In this figure, the average divergence of the steel plant was calculated over a time period of September 2019 to August 2024. The spatial resolution in the image is 4 km. The divergence peak shown in the image is at the location of the steel plant which proves that over the almost five-year period, a large emission source is located at the location of the steel plant. Integrating over the area of the peak gives an emission estimation of the steel plant. The emission estimation in this case would give a lower estimate of the emission as no plume detection was done on the dataset meaning that data not containing plumes is included. Thus when taking the average of the divergence data the no-plume divergence data will lower the average divergence. This

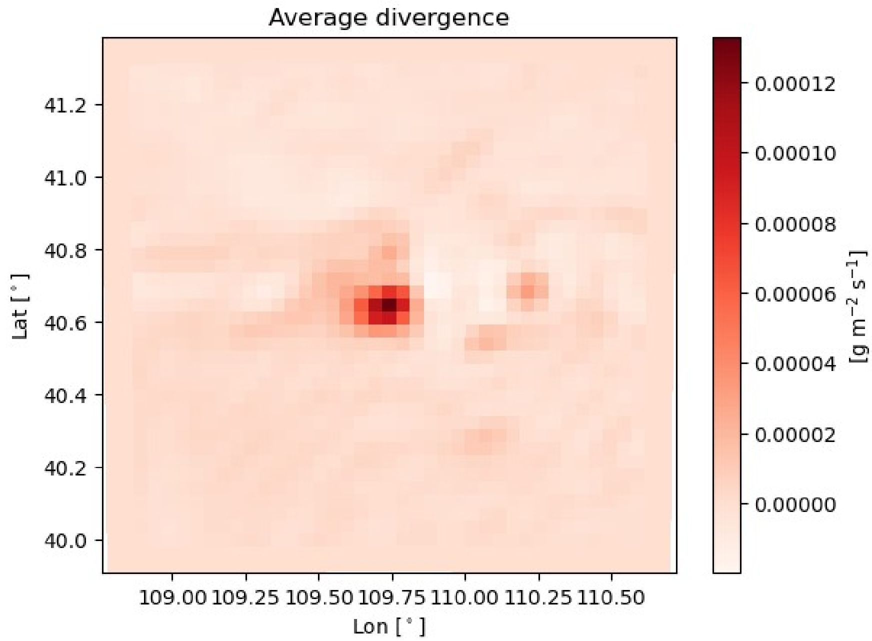

Figure 2.16: The average divergence of the Inner Mongolia BaoTou Steel Union Co Ltd from September 2019 to August 2024.

image in Figure 2.16 is still useful as it can be used to validate the steel plant location. This is because an operating steel plant should show a peak in the average divergence if it is actively emitting CO in that time period.

During the internship this version of the divergence method was added to the APE algorithm as an option to calculate the emissions for steel plant point emission sources [3]. As stated before, this update has not yet been documented in a paper. The average divergence method was used to analyze 514 out of the 1115 steel plants described earlier. Out of these 185 steel plants were found to have an average divergence peak.

2.5.3. Known issues with emission quantification of APE

When working with the CO data from TROPOMI several issues can occur. For APE to work properly these issues have to be taken into account. The following paragraphs discuss several problems that have been encountered by APE.

Limited spatial resolution

Due to the limited spatial resolution, it is not always possible to pinpoint the emissions sources. The emission sources are limited by the resolution. The CO data from TROPOMI has a spatial resolution of 7 km by 5.5 km [22]. This could mean that a relatively large industrial area will be measured as a single pixel, however, this is a hardware limitation that cannot be overcome easily. It does make it possible for sources to be simplified as a point source. Averaging the data can give larger resolutions, by utilizing that the satellite overpass will be slightly different each time, but this does come at the cost of individual time measurements.

Creating a new grid for the measurements

When using the divergence method an issue with the measurements of satellites is that depending on the orbit that the satellite takes, they might not make the measurements at the same place each time. This results in datasets where the measurements are not the same. To fix this issue, the data will have to be regrid so that the data can be compared to each other. To do this, the concentration data will have to be calculated for different locations. Depending on how this is done, the outcomes of the emissions might be affected as some methods are less accurate. Therefore, when selecting a method to regrid, the effects on the data should be taken into account.

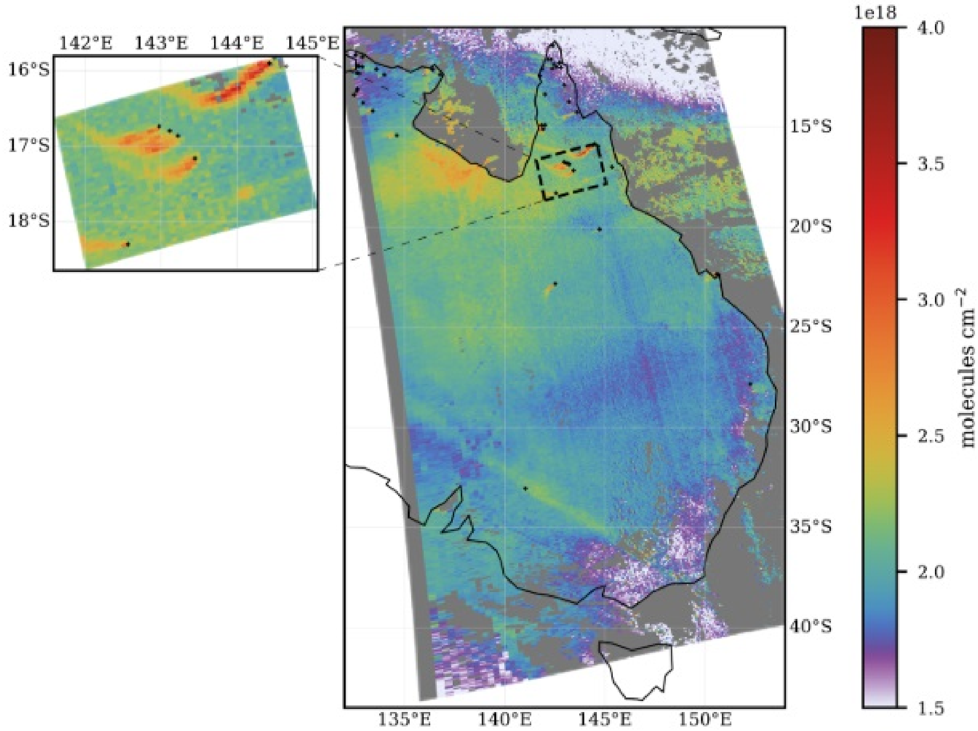

Figure 2.17: Example of overlapping plumes from Goudar et al., 2023 [22] due to wildfires in Australia in 2019

Dealing with noise in the data

Due to quality differences between the measurements, it could be that there is a noise present in certain pixels. This could lead to mistakes in the calculation errors. In Koene et al., 2021 [31], several methods to deal with this problem are listed.

Overlapping plumes

When multiple plumes overlap, it becomes very complex to model their emission. This is because many emission estimation strategies are only meant to model one distinct plume. Issues like this can happen in various situations. For example, Figure 2.17 shows three plumes that are hard to distinguish from the background due to the wildfires in Australia in 2019. The paper by Koene et al., 2019 [31] states that the issue of mixing and overlapping plumes will be left for the future and that their algorithms will not produce emissions for these.

Missing data/cloud coverage

Due to the influence of clouds, not every measurement of the trace gas will be available. The effect on the quality of data is shown in Figure 2.18. Here the quality of the CO data from TROPOMI is given a quality value based on the position of the clouds in the data scene. The Product Readme File states that quality values below 0.5 should be used with caution in analysis [32].

The result of defective data is that not all pixels hold data that can be used. This could result in pixels missing from plumes. To resolve this issue, a good interpolation method is needed. However, if large areas of pixels of the plume are missing the fidelity of the interpolation method will not be important. In these cases, the data should be discarded. The same is true for missing data due to other reasons. For example, if the data is corrupted, not everything will be available and therefore an interpolation method should be chosen or the plume should not be considered in the analysis.

Noise due to measurements over water

The measurements of certain trace gases above water are more sensitive than over land. This is due to the clouds above the oceans and rivers contaminating the measurements. So when measuring data, TROPOMI also takes data from above the clouds into account, which results in errors. This is the case

Qa_value	Condition	Remark
1.0	$\tau_{aer} < 0.5$ and $z_{cld} < 500$ m	clear-sky and clear-sky like observations
0.7	$\tau_{aer} \geq 0.5$ and $z_{cld} < 5000$ m	mid-levels cloud
0.4	($\tau_{aer} \geq 0.5$ and $z_{cld} \geq 5000$ m) or ($\tau_{aer} \leq 0.5$ and $z_{cld} \geq 500$ m)	high clouds, experimental data set
0.0	irow $\leq$1 or SZA$\geq$80° or defective product	corrupted or defective data

Figure 2.18: Quality value definition of CO data from the TROPOMI CO L2 Product Readme File [32]

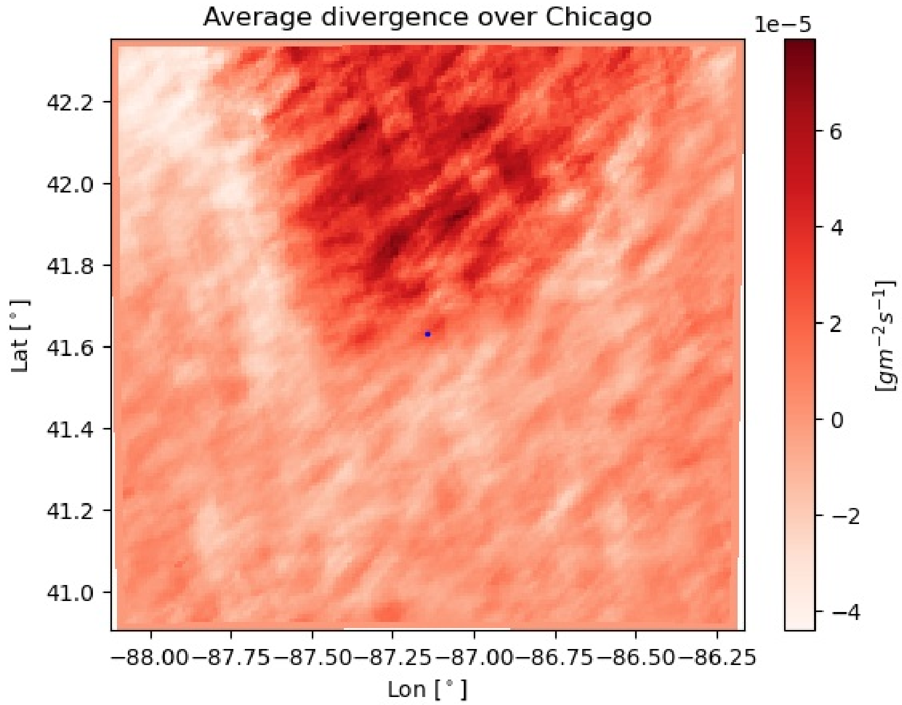

Figure 2.19: The average divergence of plume data over Chicago with the blue dot showing the Burns Harbor industrial area from the internship report [3]

for CO. For example, as can be seen in Figure 2.19 and Figure 2.20 the divergence of CO is significantly affected by these measurements. The first figure shows the divergence of the data while the second figure shows the location of Lake Michigan concerning the divergence data. As can be seen, the CO data over the lake is significantly higher. To avoid the water data affecting the emissions, a filter should be made to take this data out.

Plume height

Plumes do not stay at a constant height usually. Due to temperature and vertical wind speeds, the wind can flow upwards and downwards through the atmosphere. By assuming the height of the plume to be constant, the plume emission estimation would be influenced. A way to avoid making this assumption is by using a three-dimensional Lagrangian tracer dispersion model as is done in APE [22]. This model finds the height of the plume in a downwind direction. It does this by performing multiple Lagrangian simulations of the injection height. The first simulation is at the injection height itself, the other two are 500 meters below and 500 meters above the injection height. The injection height itself was estimated using different satellite data. The plume height is calculated by taking the average of the simulation heights. An example of the process can be found in Figure 2.21.

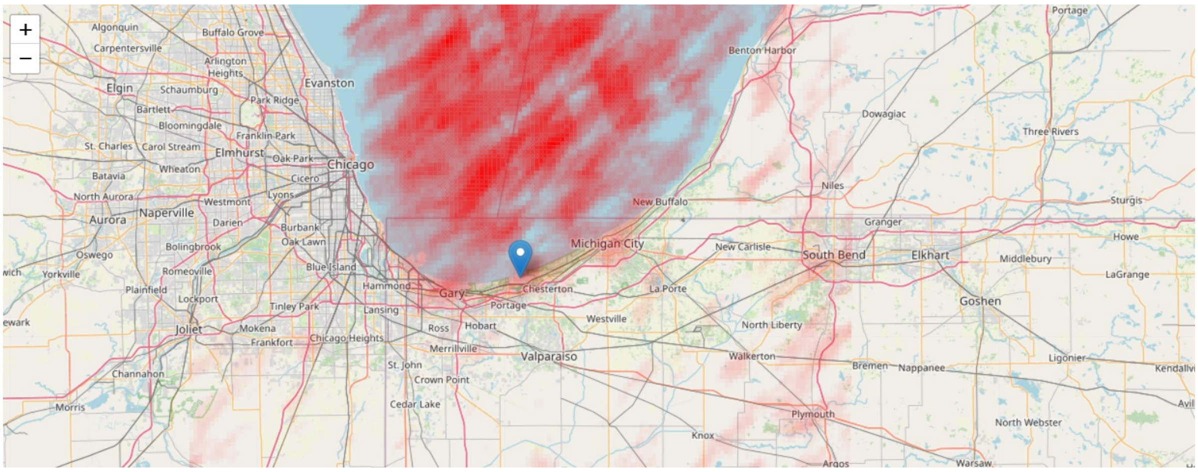

Figure 2.20: The average divergence of plume data over Chicago with the blue marker showing the Burns Harbor industrial area from the internship report [3]

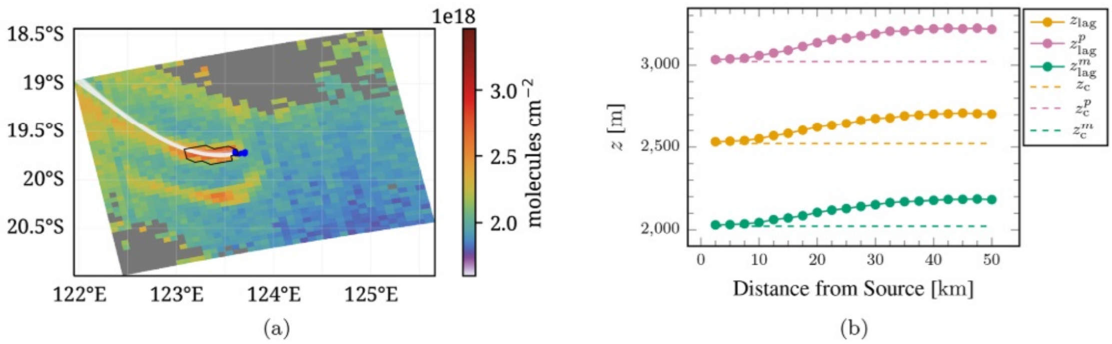

Figure 2.21: Estimation of the plume height. Subfigure (a) shows the white line of the Lagrangian model and subfigure (b) shows the different plume heights as a result. Figure from (Goudar et al., 2023 [22])

Wind speed uncertainties

The wind speed is not estimated by TROPOMI or any previously mentioned satellite. The wind speed data is taken from the European Centre for Medium-Range Weather Forecasts (ECMWF) ERA5 product. This product represents an hourly climate and weather model that includes the wind speed and direction data [22]. The spatial resolution of ERA5 wind data is known to be 0.25 ° by 0.25° for atmospheric points and 0.5 ° by 0.5 ° for ocean waves. The resolution is thus roughly 31 km by 31 km which is a worse resolution than TROPOMI CO data which has a resolution of 7 km by 5.5 km. Therefore a lot of interpolation is needed to use both data sets. In Goudar et al., 2023 [22], the uncertainty of plume emission estimation using ERA5 data is described to be less than 20% for 97% of the cases.

Stripe noise of TROPOMI CO

The raw TROPOMI data contains vertical stripes. These stripes are most likely an artifact of putting together multiple detectors or from the movement of the sensor during the measurement. To get rid of these stripes, there are several methods. In Borsdorff et al., 2023 [21], a new method using random forest classifiers is utilized. This method was developed because the current destriping method uses VIIRS cloud data. The VIIRS mission is soon to be retiring and thus this data will not be available anymore. An overview of the method can be found in Figure 2.22. The method works by first estimating a smooth background in the cross-track direction and then subtracting this from the raw data and smoothing it to remove emission peaks.

[22] Retrieved on 3-6-2024, `https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels?tab=overview`

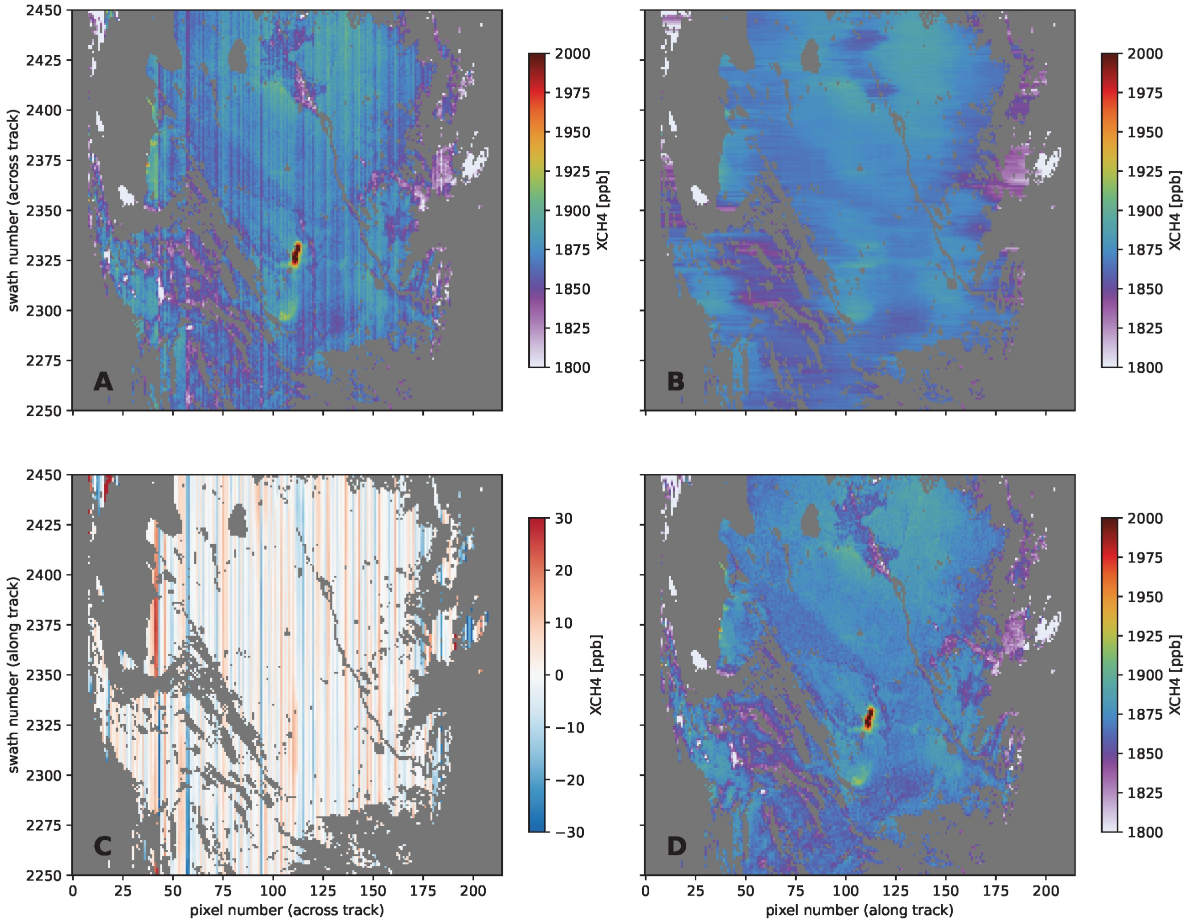

Figure 2.22: Overview of the new destriping process. Subfigure A shows the TROPOMI data. Subfigure B shows an approximation of the background. Subfigure C shows the stripes. Subfigure D shows the resulting destriped TROPOMI data. Figure from Borsdorff et al., 2023 [21]

Determining the background concentration

It can be hard to determine the background CO concentration data in the TROPOMI data. This is because the data is usually not uniform which could lead to a case where the background pixels far from the source are higher than the emission at the source. This issue can make it hard to isolate plumes in the data. A similar problem for CO_2 data is discussed in Koene et al., 2021 [31].

2.6. Alternative pollution plume detection methods for APE

As described in subsection 2.5.1, the plume detection that is currently implemented in APE is not working as well as intended. This section discusses several alternative plume detection methods that could serve as a replacement or an addition to the currently implemented plume detection method.

2.6.1. Machine learning approaches for plume detection

Machine learning has a broad range of applications in remote sensing [33]. A problem outlined in Maxwell et al., 2018 [34] was that it was not so easy to implement machine learning (ML) approaches in 2018. This was due to a lack of experience in using and optimizing ML as well as less available ML packages as is the case now. This led to under-utilization, even though the technology was seen as potentially more effective than other methods. In this section, an overview of the existing and researched ML methods are listed. Other important aspects surrounding ML are covered here as well, starting with the datasets.

For clarification purposes, in this document, a data scene or granule is defined as a combination of trace gas concentration measurements in an array with their geographical coordinates. A dataset is defined as a collection of data scenes. To make the dataset set only include useful data, a selection needs to be made. This means that in the analysis of a point, not the entire field of view of TROPOMI but a smaller selection of data points around the source is needed.

The creation of a dataset for TROPOMI is important because this dataset used for training and testing must be accurate and have little to no bias if possible. A problem that could occur is class imbalance. A class is a particular type of data. For example, if you have a dataset of pictures of fruits, you could separate the images by making each fruit a class. So pears could be a class. The class imbalance occurs when a particular class appears more than others, which could create a bias for the classifier to that class [34]. To avoid this problem, the dataset(s) should include an equal amount of data scenes from the different classes.

Another problem that can occur when using machine learning is the Hughes phenomenon. This happens when data scenes include too much information, which could lead to overfitting. Thus adding more parameters could decrease the accuracy of classification [35, 34]. To avoid this issue a smart selection in the data should be made. For plume detection purposes it makes sense to focus on the data of the plume and its surroundings which would be easily visible just by looking at the enhancement of the concentration concerning the background. Another possible solution is to use the data only in the direction of the wind, as a plume in the opposite direction of the wind would break our current understanding of science. A combination of these approaches can be made as well. During the internship, an augmentation algorithm was developed to filter out unnecessary information for plume detection.

Plume detection using support vector machines

The Support Vector Machine (SVM) approach seeks to maximize the boundaries between the points of different classes in a hyperspace. The dimension of the hyperspace is determined by the dimensions of the input vectors. The boundary is defined as a hyperplane in the hyperspace. It does this by drawing boundaries between the differently labeled data points. The SVM seeks to find the shortest distance between the data points of the different classes. These boundaries are drawn by the kernel function. The boundaries between classes are essential as they directly inform the SVM how to classify new data points. Therefore, getting the right boundaries is important thus the kernel selection is important.

To draw these boundaries, different kernels can be used. The effect of using there kernels can be seen in Figure 2.23. In this image different kernels from SCIKIT are shown [36]. By testing during the internship it became known that the radial base function (RBF) kernel produced the best results in training and testing of the plume detection data [3].

During the internship, a dataset was created to test the SVM machine learning tool and compare it with results from APE [3]. When the machine learning tool was shown to be more successful compared to APE in finding plumes, it was decided to do this follow-up project. To see how effective the

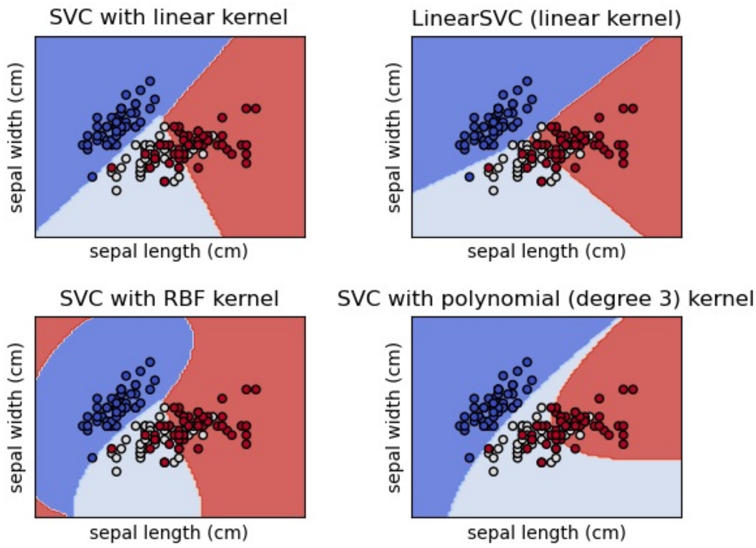

Figure 2.23: Effect of using different kernels for boundary drawing for SVM from SCIKIT [36]

new datasets are, it was decided to also run the already existing machine learning tool based on the internship dataset. The results of the machine learning tool created by the old data would serve as a benchmark.

The dataset itself consists of about 1600 images. Half of these images contain plumes and the other half contain no plume data. The dataset only includes data from the Inner Mongolia BaoTou Steel Union Co.,Ltd. plant in China. This steel plant is known to emit large plumes. According to Tian et al., 2021 [37], the emission rate of the steel plant is 34.3 ± 2.0 kg/s. The location of the plant is also ideal as it is located near the Goby desert and not close to a body of water. This makes it easy to extract plumes from the data.

The dataset was used to train the SVM tool. The use of the tool was compared with APE. When comparing the two tools it became clear that the SVM tool was able to find all 54 plumes that APE did except for one. The tool was also able to detect 125 more plumes. This showed the potential of this method. When trying to implement this method for other steel plants, an issue was encountered with the background values showing different behavior which obscured certain plumes [3].

The drawback of the SVM tool is its sensitivity to noise in the data. Therefore when using this tool an augmentation algorithm is needed to enhance the important information in the data. The augmentation algorithm reduces the information in the data to only the important features that the detection. It does this by only selecting the data points that are close to the source and data that is in the wind direction. An example of processing data with the augmentation algorithm can be found in Figure 2.24.

The tool is also not able to learn what a plume is like a human would and thus is not always able to classify very obvious plumes. An example of this happened during the internship where it was found that the direction of the plume was very important to SVM classification. To avoid issues with the directionality of a plume, every data granule is rotated three times, each time the rotation is rotated by 90°. These 3 rotations together with the original image are added to the dataset on which the SVM tool is trained. In Figure 2.25 different plume images are shown with their respective classification listed in the descriptions. As can be seen, the direction of the plume influenced the classification. The issue was solved by adding rotations of the original data scene into the dataset. This issue shows that the SVM tool is very sensitive to biases in the data.

Support vector machines were also used in Latif et al., 2023[38] on CO concentration data to predict future concentrations where they were tested among other machine learning approaches. In the paper, the approach was not considered to be the best but it was also not the worst. If the data is relatively simple, meaning without many parameters and does not include many errors, the approach produces accurate and fast results.

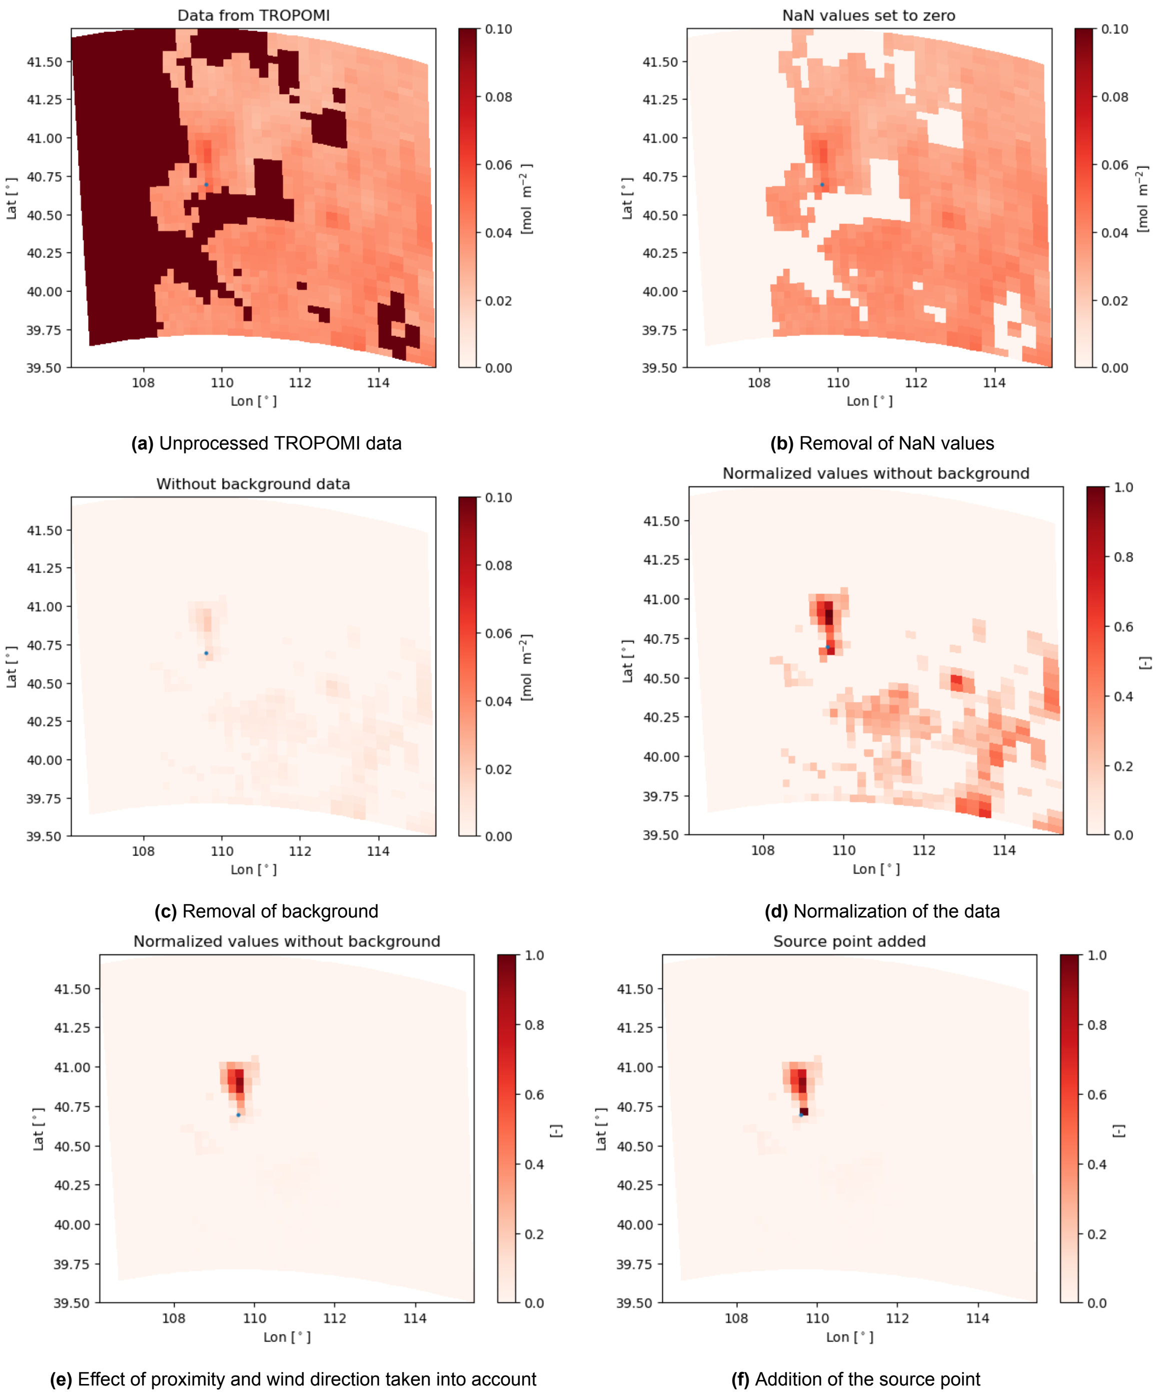

(a) Unprocessed TROPOMI data

(b) Removal of NaN values

(c) Removal of background

(d) Normalization of the data

(e) Effect of proximity and wind direction taken into account

(f) Addition of the source point

Figure 2.24: Steps taken to clean up the data for plume detection

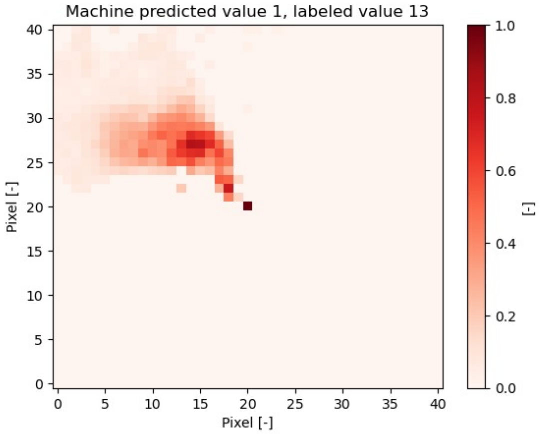

(a) Rotation 1, not detected as plume

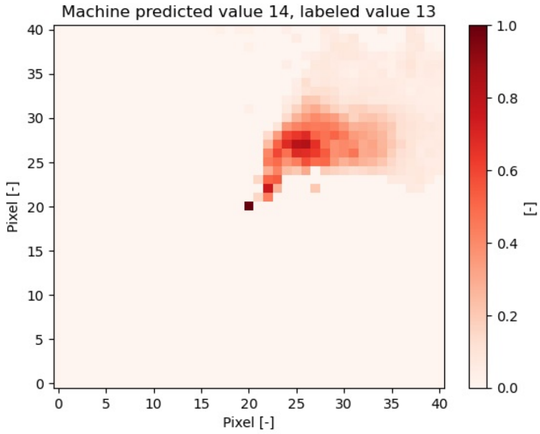

(b) Rotation 2, detected as a different kind of plume

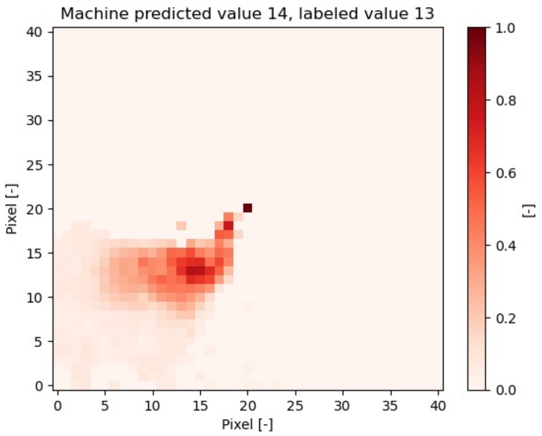

(c) Rotation 3, detected as a different kind of plume

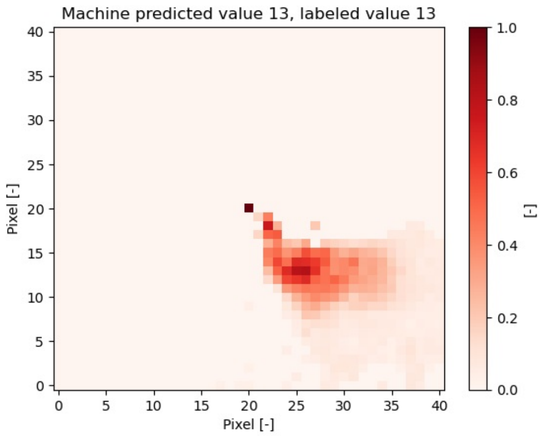

(d) Original data, correctly detected

Figure 2.25: Overview of the effect of rotations on detection by SVM. The data scenes are processed TROPOMI data using an algorithm.

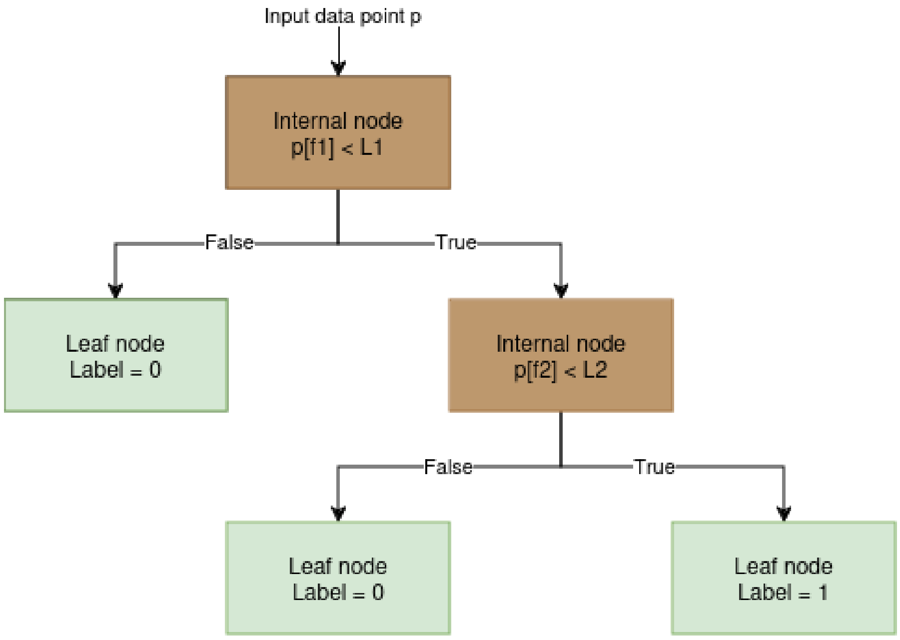

Figure 2.26: Simple overview of a small decision tree.

Plume detection using random forest classifier

For a previous project, SRON commissioned a group of software engineers to create an efficient random forest classifier (RFC). This collection of code is known as Balsa and has been made available for the project. This subsection is based on the description from the explanation given in the Balsa code as made available to the SRON Earth Group. Balsa was developed by the company Jigsaw B.V. with funding from ESA [39].

The RFC approach uses a combination of decision trees that are used for classification. The training data is used to create the decision trees. By randomly selecting parts of the training data for each decision tree, the resulting decision trees will be different from each other. For testing, each decision tree will have the same data as input. Each tree will output the label they think is the correct label for the test data. The most popular label as chosen by the trees will serve as the label used to classify the data [40] [41].

The decision tree classifiers consist of a collection of nodes. Nodes can be split into two types, leaf nodes and internal nodes. Internal nodes have two child nodes while leaf nodes have no child nodes. In an internal node, a data point is judged by checking if the value of a feature f in data point p, written as $p[f]$, is bigger or smaller than a value L. An example of a decision tree can be found in Figure 2.26. Here the basic structure of a decision tree is shown. Internal nodes use the value of features to determine whether to go to the next node. The classification process at the leaf nodes, where the class label of the leave is used to classify the data.

When data is inputted to create the decision tree classifier, the classifier splits data into two different parts that include data. This process is continued until all data features are split from each other. The split values L are stored per feature. The optimization of this tool is in finding the smallest trees while not losing too much accuracy. The larger the number of features in the data, the longer it takes to create the decision trees. To avoid the creation of huge trees, which are computationally expensive, there are several options to consider. Examples of these include limiting the depth of the tree so that not every feature is utilized or limiting the data itself by only inputting useful data. The random forest is created by randomly selecting the features of the input data for each tree. The decision trees are considered to be weak classifiers. A large combination of weak tree classifiers can result in a stronger classifier. Thus resulting in a selection of trees that are slightly different. If the training data is chosen well, the trees should generally vote for the right class resulting in good classifications overall.

The Balsa package has a maximum number of features of 256. This is a hard limit on the amount of data that can be classified. This is due to the computational complexity increasing significantly with the number of features. This is due to the system having to create more complex trees. For SVM this is not the case as the boundaries are drawn quickly and therefore there is no hard limit and the number of features to use is significantly larger. It is larger than at least 5000 features. Thus the number of features is a limiting factor for this approach, so a good selection of features is necessary to utilize this approach best. This can be done by using certain algorithms to highlight only relevant data. For example, only using the data points close to the emission source makes sense and reduces the number of data points aka features used for RFC.

According to M. Pal, 2004 [41], for remote sensing purposes, the random forest classifier has a better accuracy than SVM. They define Accuracy as the number of data scenes in testing that were classified correctly. The paper further states that RFC is able to use imbalanced data and data with unknown values.

Plume detection using deep learning

The plume detection of trace gases can be seen as an image recognition exercise. This is why it makes sense to use a proper image recognition approach. The previously mentioned approaches cannot look at the image of the data scene as they look at the data values. This causes problems such as issues with the direction of plumes as depicted in Figure 2.25. With a deep learning approach, it should be possible to avoid these kinds of problems by utilizing a more sophisticated method to recognize plumes.

Deep learning utilizes several layers of processing to detect patterns and classify data. The approach is supposed to mimic the brain by using nodes that are used to help in the classification. The nodes are called neurons and are responsible for passing through the data. Putting multiple nodes together creates a neural network. The combination of a neural network for deep learning is called a deep neural network[23]. The nodes work with weights and thresholds. If a certain threshold is met, then these particular nodes are activated. For image processing the conventional processing method is using a convolutional neural network (CNN). These networks use convolutional layers. The processing of CNN's work by first analyzing smaller features of the image and with each passing layer analyzing more and more parts of the image until analyzing the full image in the last layer. This way the CNN can fully breakdown the image and with the information gained can classify the image[24]

The CNN works in three parts, the first part is the convolution layer. In this layer, the input data is processed into feature maps which investigate different features from the input data. This is done by performing several convolutional operations to find features. During the training of the CNN, the optimal convolutional operations are found. The second layer is called the pooling layer. In this layer, the noise in the feature maps is reduced as well as the dimensions of the maps are reduced. At the end of the pooling layer, a "summary" of the important features is created. This "summary" is a vector comprised of the condensed and flattened most important features. This flattened feature map is then used as an input to the final layer which is the neural network. The neural network classifies the plume at the end. During the training process of the CNN, the weights of the convolution layer and the neural networks are tuned. Thus the CNN will learn which features are the most important for classification and create the weights for the filters. Unlike the previous methods, the CNN approach should be able to find the features in the data by itself without the need for an augmentation algorithm. The augmentation could be used to aid the network in finding the right features for classification and is thus also an asset.

An issue that arises when investigating the images as vectors as was the case for the SVM and RFC is that the features, in this case a pixel of CO concentration, are investigated by itself without investigating the features surrounding it. It is likely in an image that the surrounding pixels are in some way correlated to the center pixel. SVM and RFC take that into account implicitly. The SVM tool does not check for the correlation between the features, rather it seeks to find the optimal distance of the margin between data points. The correlation of features can be implicitly found by using non-linear kernels to create the boundary. These non-linear kernels can find the relationships by adapting to the non-linear behavior of the data. The correlation between features is also not utilized explicitly by RFC as they generate trees based on differently randomized input features. If the tree diversity is large enough the feature correlation is used implicitly for classification. This is because different trees will be able to represent the different correlations between different pixels. However, the CNN approach compared to SVM and RFC uses feature correlation explicitly. For the plume detection algorithm, this

[23] Retrieved on 4-6-2024, `https://www.ibm.com/topics/deep-learning#:~:text=Deep%20learning%20is%20a%20subset,AI)%20in%20our%20lives%20today.`

[24] Retrieved on 4-6-2024, `https://www.ibm.com/topics/convolutional-neural-networks`

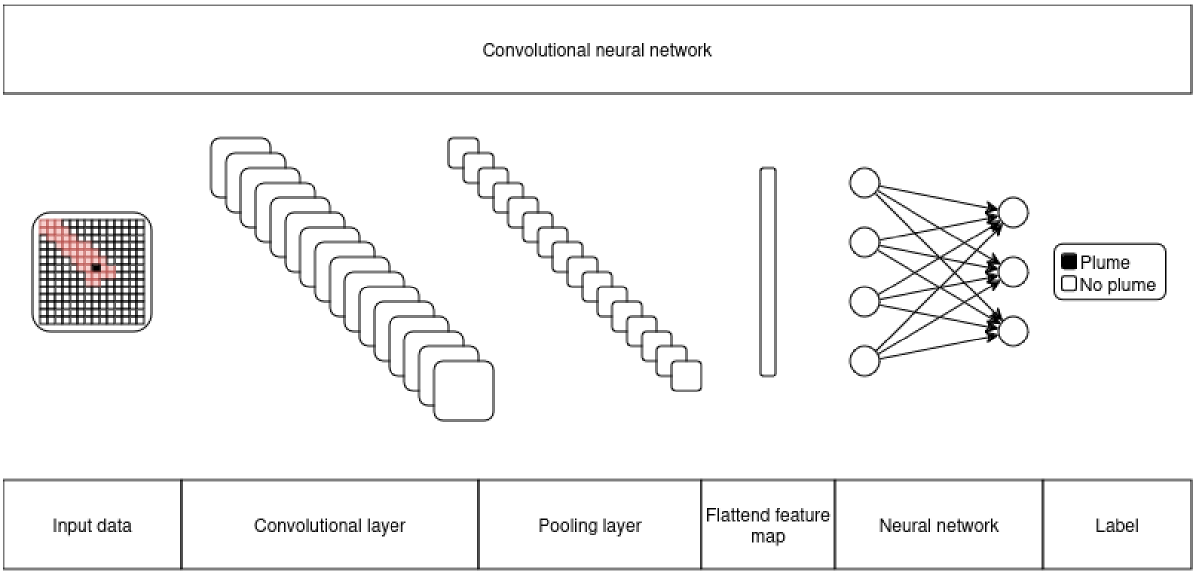

Figure 2.27: Overview of the steps taken by a CNN to classify plume data

is extremely useful. A plume is not just located in one pixel but in several pixels near the source. Thus utilizing the correlation between pixels explicitly makes more sense than hoping to approximate the effects of the correlations between the pixels in the previous methods.

In Finch et al., 2022 [42], a neural network is used to find NO_2 plumes. Using two years of data, it identified 310020 images with at least one NO_2 plume. The network was more than 90% of the time able to identify a plume correctly. The paper makes the case for using NO_2 as a tracer for CO_2 because of the much shorter lifetime of NO_2. The paper argues that by correctly identifying NO_2 plumes, it is possible to make it easier to identify CO_2plumes.

The big drawback of deep learning is that it is more computationally heavy. If the previous methods show that they are capable of producing similar accuracy in detection compared to CNN they would be more desirable to use. Thus even though this method seems to be the most promising in terms of producing accurate plume detection, it is not necessarily the best method.

2.6.2. Gaussian fitting model method

This method works by attempting to fit a Gaussian distribution through the cross-section of perceived plumes. The method models the background by fitting it with a linear function. An example of the method can be found in Figure 2.28. The figure is from Zheng et al., 2020 [43]. They use one orbit of the climate satellite OCO-2 with CO_2 data. Compared to other methods, this method works on a small part of the plume. However, with the abundance of other information available for plume in a TROPOMI dataset, this method is not the best suited to use with TROPOMI data this is because this method is able to find cross-sections of the plume whereas in TROPOMI for a lot of cases, the entire plume is visible.

2.6.3. Simple enhancement method

In this method, the enhancement of a possible plume is researched by checking the pixels around the source and the pixels further away. An example of this can be found in Figure 2.29. In Figure 2.29a, a TROPOMI CO data scene is shown with a known emission source in the middle. The middle pixels are shown inFigure 2.29b, while the background pixels are shown in Figure 2.29c. To figure out if there is an enhancement, one could average the concentration data at the source and average the background data. The enhancement ratio can be found by dividing the average of the source pixels by the average of the background pixels. If this ratio is bigger than one, there should be an emission visible in the scene. A problem with this approach is that the background of CO data is usually not uniform and thus this could lead to a case where the background pixels far from the source are higher than the emission at the source. A similar problem for CO_2 data is discussed in Koene et al., 2021 [31].

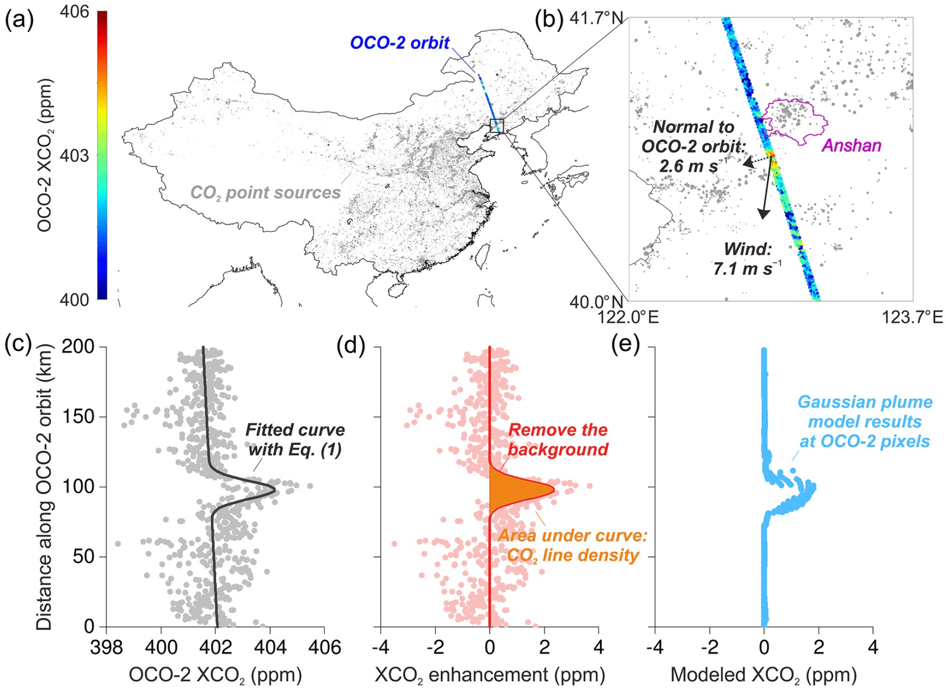

Figure 2.28: Gaussian fitting method using OCO-2 data. Subfigure (a) and (b) show an orbit with CO_2 enhancement. Subfigure (c) fits a line through the background and the enhancement. Subfigure (d) removes the background and the result is the CO_2 density. Subfigure (e) shows the modeled level of CO_2 enhancement. Figure from [43].

2.6.4. Using other trace gases such as NO

Another approach to make it easier to find plumes could by also investigating other trace gases. This is because during the combustion process other trace gases are emitted along side CO. One of these gases is NO. The issue with trying to use NO is that TROPOMI does not have a dataset for NO. It does have a dataset that measures the concentration of NO_2. This dataset can be combined with a model that can convert NO_2 to NO. This can be done using the model from Kuhlmann et al., 2021 [44]. This model would make it possible to input plumes from two trace gases to see their effect. It can also be used to refine the emission estimation method by refining the pixels that house the plume.

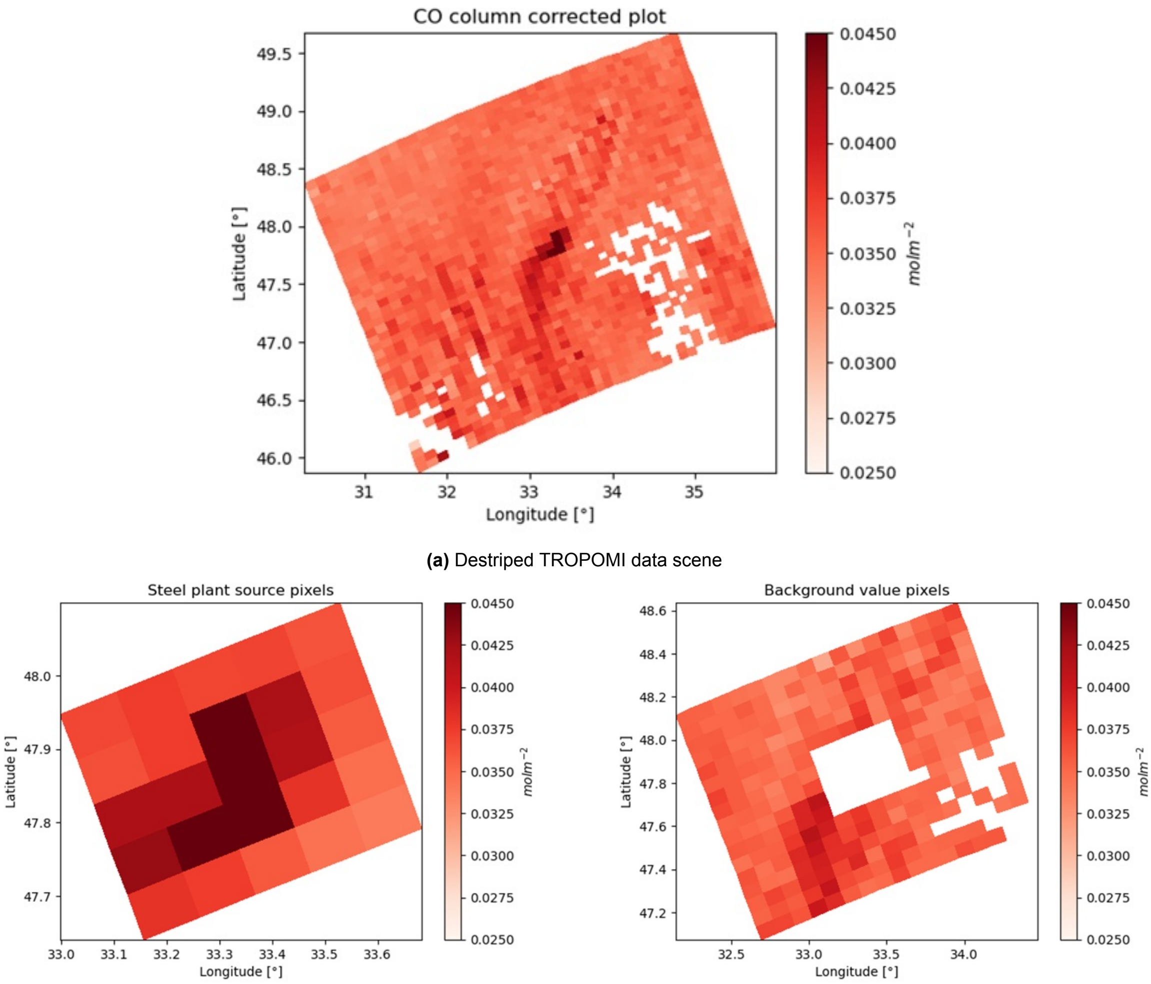

(a) Destriped TROPOMI data scene

(b) Pixels at the emission source

(c) Pixels used to determine the background

Figure 2.29: Overview of the simple enhancement method

3
Research plan

This chapter gives an overview and a justification for the thesis project. The justification takes the form of research questions, to which tasks are assigned. The project tasks will be assigned to different work packages. An estimate for how much time each work package will take is also denoted. At the end, a Gantt chart shows the thesis duration and milestones. The research questions for the project and their explanation are listed in section 3.1, followed by an overview of the tasks in section 3.2.

3.1. Research questions

Main research question:

- MQ1 - Can a machine learning approach improve the detection of pollution due to combustion at a global scale?

This question will be answered by implementing multiple machine-learning approaches using the APE algorithm. APE can calculate the emission for CO plumes but is not able to detect plumes without a predefined source. A machine learning approach should be able to find plumes even without a set location. The creation of the detection tool be done in work package 2. In work package 3, the new method will be used on a large amount of data to see exactly how successful it is.

- MQ2 - Does combining different trace gases improve the quantification of emissions?

To answer this question, in work package 1, the trace gas NO_2 data from TROPOMI will be added to the APE algorithm. This will allow the data to be used in analysis. Furthermore, the NO_2 data will be converted to NO_x data using the model from Kuhlmann et al, 2021 [44]. This will create two datasets. Both of these will be used in two ways. One is to make a more accurate plume mask to aid in emission estimation. The other purpose is to add it to the machine learning tool, to see if adding other trace gas data will be useful in finding plumes, which is the work from work package 2. In work package 3 the newly created approaches will be tested on global data that span 7 years to see exactly how it increases the accuracy.

Work package 1:

- QWP1.1 What is the effect of adding NO_2 and NO_x data for emission quantification?

In the combustion process, various trace gases are produced. These include CO, NO_2 NO and CO_2 among many others. The distribution over time of these gases gives information about the combustion process. Therefore, analyzing them together should provide us with more information than analyzing them separately. To answer the question the following steps will be taken. The NO2 data from TROPOMI will have to be included in the APE algorithm. Then the chemical model from Kuhlmann et al., 2021 [44] will be used to get the NO_x concentration data from the NO_2 data. This data can be used in different ways. For some emission estimation methods, it is necessary to find the pixels where the plume is located. This is known as the plume mask. It is expected to be easier to find the plume mask using the NO_x data than using the CO data. Furthermore, the NO_2 and NO_x datasets will also be used in work package two for plume detection.

Work package 2:

- QWP2.1 - Between the machine learning approaches for plume detection as described in section 2.6, which one detects the most plumes accurately?

To answer question WP2.1, a variety of different datasets will be created. From the literature, it is known that these datasets are quite extensive. However, to circumvent this issue the data will be augmented. It is interesting to create multiple datasets and compare results to see the range of different uses in machine learning approaches. One such data set will be just a set of raw TROPOMI images. Other datasets would use some form of augmentation to highlight specific aspects of the data to make detection easier. For example, the data that is not in the direction of the wind will be removed. The aim will be to create multiple large datasets of 10.000 scenes with worldwide coverage.

The machine learning methods that are being considered are SVM, RFC and a deep learning neural network. The SVM was already considered during the internship and is therefore already implemented. SRON has an RFC Python package made by a third party that is easily usable. During the internship, contact was established with a PhD student from the University of Heidelberg who is capable of making deep-learning models. This model thus already exists, however, I do not have it in my possession and therefore will need to discuss with the person how we continue. I will ask if I can have the code so I can run experiments and not depend on them. However, it could also be that this will not work and then I have to possibly even create a model myself, which will take away time from the rest of the project.

During the internship, two datasets were produced. One dataset of raw TROPOMI images and one that augmented the raw data to extract certain features. The algorithm removed the data from the raw images that were far from the emission source and not in the direction of the wind. The run successful SVM experiments, the augmented data was necessary. SVM can use a lot of parameters, however, RFC is limited to 256 parameters. The SVM augmentation algorithm uses 1681 paramters. Therefore, a new augmentation algorithm is necessary to make a dataset for this approach. For the deep-learning model, it will be interesting to test the different datasets to see if augmentation is necessary for it or not. This is important because augmenting the data for SVM takes 4 seconds per scene. If we have 10.000 scenes this will take 40.000 seconds which is 11 hours. Another issue in creating the datasets is bias. The direction of plumes is a sensitive issue for SVM and thus all rotations should be part of the datasets as well. Further investigations into the datasets will be necessary to take out as much data-based bias as possible.

The NO_2 and NO_x datasets created in WP1 will also be used in machine learning to see what the effect is of adding other trace gases. I am not exactly sure how I will include these in the CO data. But most likely I will simply add them to it. This will be difficult when using the RFC tool so perhaps another approach should be considered. Perhaps an algorithm to get rid of the noise in the CO data using NO_2 or NO_x data should be considered.

The training and testing of the datasets will be done iteratively to make sure that the training datasets are as optimal as possible. This means that the training data will be limited to what is strictly necessary. Thus instead of making random selections of data, the datasets will be picked one by one to see if the classification of test data increases or not. If with the introduction of a new scene into the training dataset, the classification of test scenes goes down the dataset will be taken out of the training set. This could introduce bias, so it will be done with care.

- QWP2.2 - Which plume features are the most important for detection?

When making data augmentation algorithms, certain features will be picked out to investigate which one makes the plume best visible. During the internship, two features were considered. These were the location of the source with respect to the possible plume and the direction of the wind. It could be that other features are also helpful in the process of plume detection and therefore will constructing these augmentation algorithms an open mind should be kept.

Work package 3:

- QWP3.1 What is the increase in the accuracy of the new plume detection and emission estimation?

To answer this question, the new methods will be combined and compared to the old methods of APE. This will then show if the new methods outperform the old ones. The exact method of testing

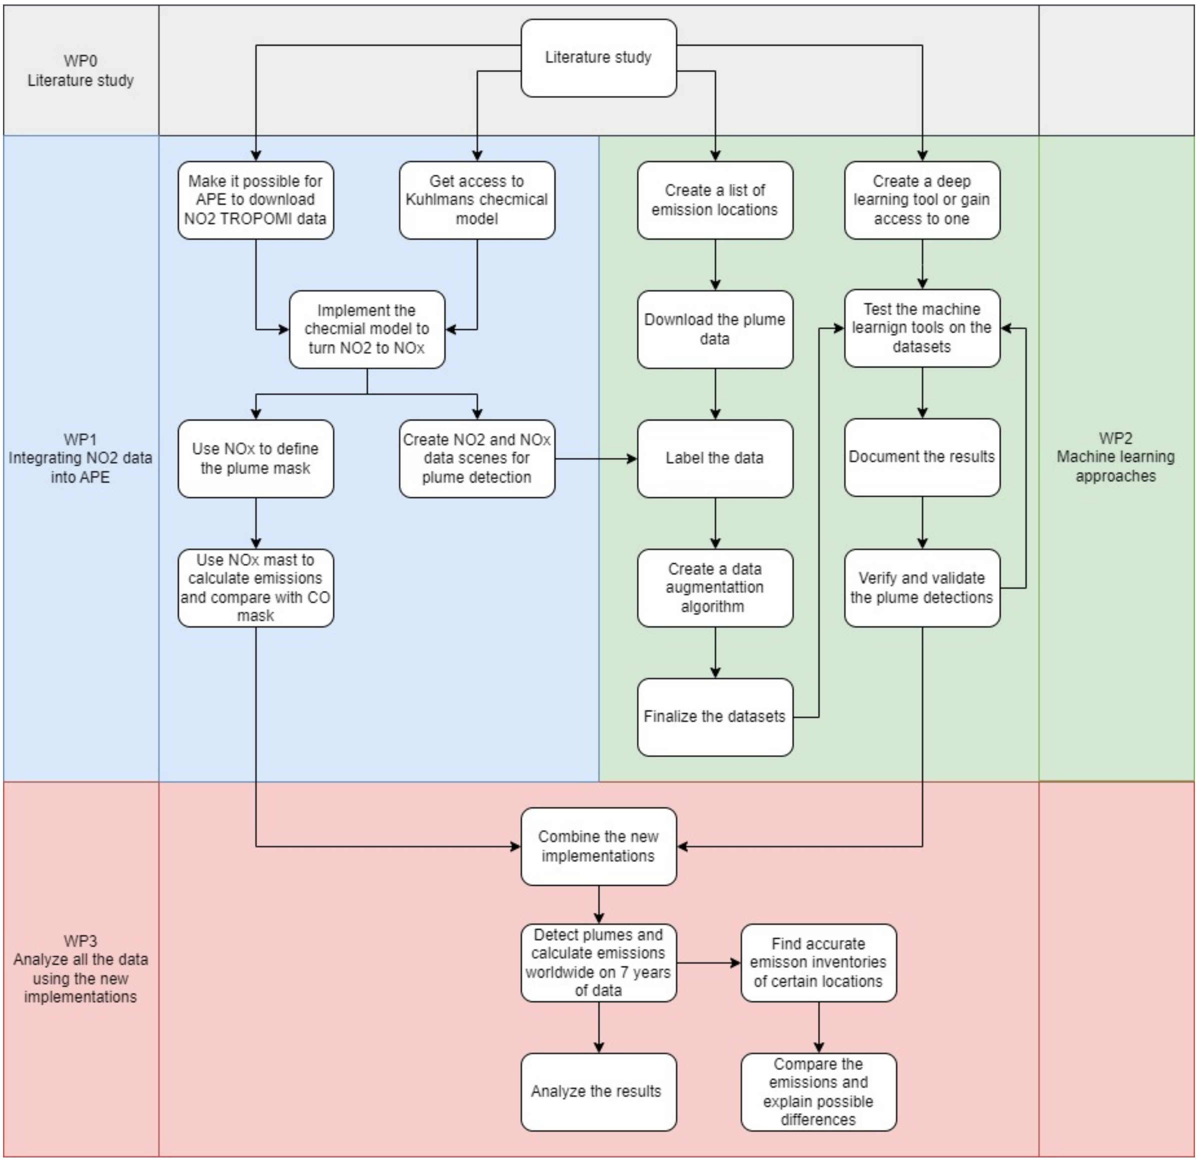

Figure 3.1: Interaction between the tasks in each work package.

will have to be thought out still but it will involve a range of different places that include plumes from industry and wildfires. It will be interesting to compare the time series of emissions to known emission inventories. I have not selected ones currently but I will look this over later.

- QWP3.2 What phenomena are visible in the data?

To answer this question, the plume detection and emission estimation will be done for several relevant locations in the data. These locations would include places with industry and wildfires. This should give us a lot of results to analyze. Going over the data could result in discoveries and validate the approach.

3.2. Thesis planning

A full overview of the work is shown in Figure 3.1. The time it takes to complete each work package is listed below. A Gantt chart showing the estimated timeline of the thesis is shown in Figure 3.2.

Planning: WP0:

- WP0: Literature study: 7 weeks, week 1 to 7

WP1:

- WP1.1 - Implement NO2 data download in APE, 2 weeks, week 8 to 9.
- WP1.2 - Get access to Kuhlmann's model, 1 week, week 8
- WP1.3 - Implement Kuhlmann's model into APE, 2 weeks, week 9 to 10.
- WP1.4 - Use NO_x data to find the plume mask, 3 weeks, weeks 11 to 14.
- WP1.5 - Estimate emission using NO_x data, 1 week, week 15.

WP2:

- WP2.1 - Create a list of emission locations, 1 week, week 8
- WP2.2 - Get access to deep learning model, 3 weeks, week 8 to 10.
- WP2.3 - Download plume data, 2 weeks, week 11 to week 12
- WP2.4 - Label the data, 2 weeks, week 12 to week 13
- WP2.5 - Create data augmentation algorithm, 2 weeks, week 14 to week 15.
- WP2.6 - Finalize the datasets, 1 week, week 16
- WP2.7 - Test the machine learning tools on the datasets, 3 weeks, week 17 to 19.
- WP2.8 - Verify and validate the plume detections, 3 weeks, week 17 to 19.
- WP2.9 - Document the results, 1 week, week 20

WP3:

- WP3.1 - Combine the new implementations, 2 weeks. week 21 to 22.
- WP3.2 - Detect plumes and calculate the emissions on 7-year data, 2 weeks, weeks 23 to 24.
- WP3.3 - Analyze the results, 2 week, week 25 to 26
- WP3.4 - Find and compare the emissions to emission inventories, 2 weeks, week 25 to 26.

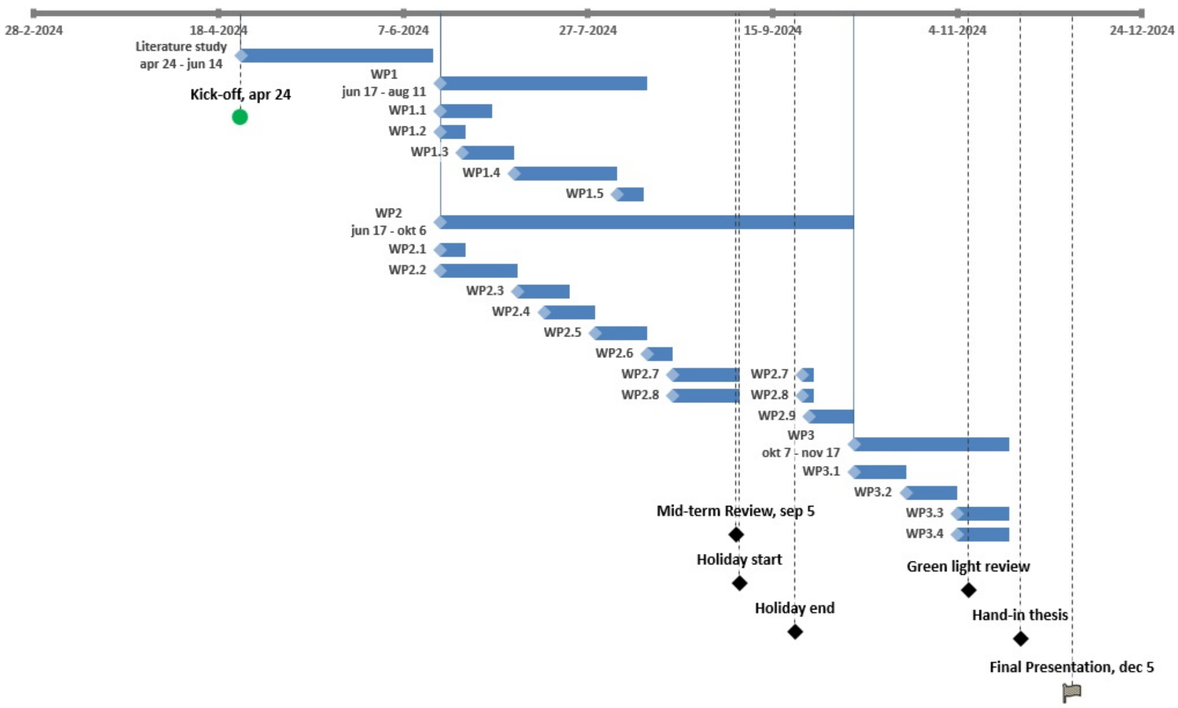

Figure 3.2: Gantt chart with the work packages shown

4

Enhancing the APE plume detection

This chapter discusses the implementation of different machine-learning approaches for plume detection. The machine learning methods should improve the APE plume detection described in subsection 2.5.1. This is done by discussing the research questions the different methods and their application in section 4.1. To better utilize the input data, the features that determine the plume detection could be enhanced. For this purpose, the input data augmentation algorithm was developed. The algorithm is described in section 4.2.

4.1. Overview of the possible improvement

To solve the first main research question of the research plan as outlined in chapter 3, which asked if a machine learning application could improve the detection of pollution on a global scale, it is important to consider the different options of machine learning that are available for the project. These methods were already discussed theoretically in subsection 2.6.1. In this section, the implementation of these methods in code is discussed.

When selecting the machine learning method, it is important to consider what problem needs to be solved. The problem that needs to be solved is that APE be expanded with another part where this happens. In subsection 2.5.1, it is explained that the emission estimation uses the plume mask provided by the Marker-based watershed transform method. The new plume detection method should thus be placed between the first part that performs the data preparation and the second part that finds the plume mask in the data. An overview of this new step can be seen in Figure 4.1. Here an updated overview of APE capabilities is shown. This figure is an updated version of Figure 2.7.

The machine learning problem of this project boils down to the extraction of plume data from a

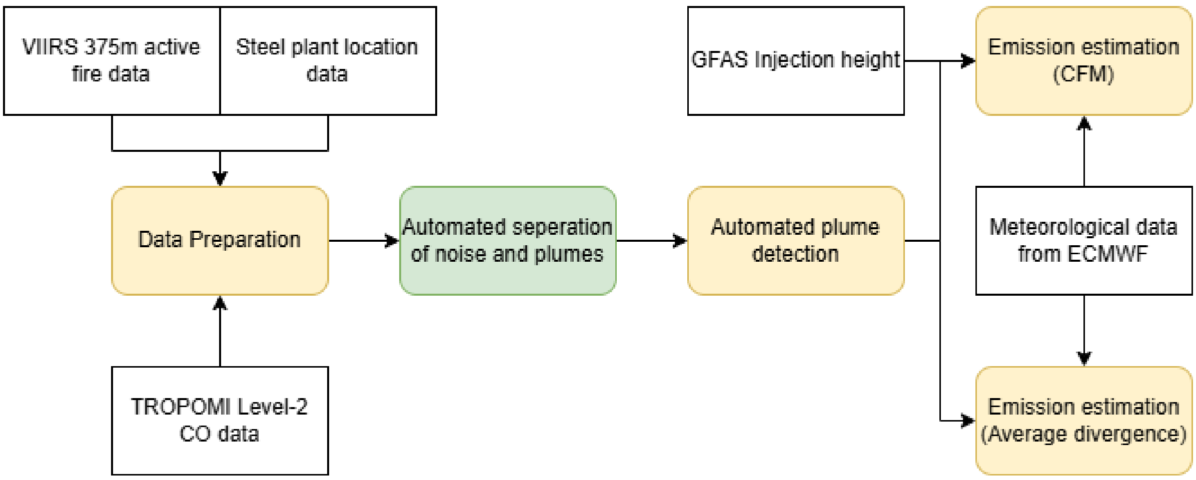

Figure 4.1: New overview of APE's capabilities with the proposed enhancement shown in green.

dataset that also includes noise data. When a granule that has a plume is found, this granule should be processed further. If no plume is found in the granule, APE should stop investigating it and move to the next data array. The machine learning tool thus acts as a filter to remove no plume-containing data from consideration. The problem here is thus a two-class classification problem. The classes are plume data and no-plume/noise data. The machine learning methods that are considered should thus be able to solve a two-class classification problem.

The plume detection methods considered are support vector machines (SVM), Random Forest Classifiers (RFC) and Convolutional Neural Networks (CNN). To save time, already existing and accessible machine-learning tools were used. The background of these machine learning methods have already been described in subsection 2.6.1. The SVM function in SCIKIT was utilized [36]. However, the SVM needed an augmentation algorithm to be used effectively. As explained in subsection 2.6.1, this algorithm was used for one particular industrial source. The augmentation algorithm did not work for other sources and thus had to be updated. This is written in section 4.2. Rotated plumes were also added to the data to avoid bias with plume directions.

For the RFC an in-house Python package called Balsa was used. Balsa was developed for other purposes, such as utilizing Balsa's RFC as an alternative to the cloud-clearing for the TROPOMI XCH_4product [21]. When utilizing the tool for the project, it was quickly realized that the 41 by 41-pixel granule was too large for Balsa to handle. The creation of the decision trees kept crashing the computer it ran on. To solve this issue it was decided to prioritize the data near the emission source. Thus the granule was cut down to the 15 by 15 center pixels as this was the maximum granule size that would not result in a crash. This is a significant reduction in terms of features compared to SVM. Just like for SVM, the data augmentation was utilized as well as the rotated data when using the RFC classifier. The reduction of features could mean that the tool works in a worse manner as less information will be available for classification. However, this problem should not be significant as the pixels closest to the source are the most important. This is because the plume will be most visible near the source and the augmentation algorithm prioritizes pixels closer to the source. This means that the pixels further away have an increasingly higher likelihood of being put to zero anyway further. Thus it is expected that this reduction in information does not have a significant impact on the accuracy of the tool.

To test the CNN approach, two methods have been utilized. Both are based on the Residual Network (ResNet) as developed by Microsoft [45]. The first is a ResNet-44 that follows the design specifications as described in He et al. 2016 [45] made available to the project by PhD student Thomas Plewa from Heidelberg University. The second one is a ResNet-26 that has had small changes in its architecture and thus differs a bit from the design as written in the previously mentioned paper made available to the project by Peter Sterk. He works as a scientist at SRON who works on machine learning applications in the Earth science group. The number next to the ResNet name indicates the layer depth. Thus ResNet-44 utilizes a layer depth of 44 and ResNet-26 utilizes a depth of 26.

The utilized Residual Network (ResNet) for this project are the ResNet-26 and the ResNet-44. These networks differ slightly from the general CNN description from subsection 2.6.1. A big drawback in the regular CNN architecture is an issue with vanishing gradients. These gradients are used to update the weights of the filters during training. When the gradients get too small the weights are not meaningfully updated. This happens as a result of the backpropagation of the error to the initial layers. This error can then be used to update the weights of all layers. With more layers, the backpropagation gets longer and can make the gradients smaller resulting in the aforementioned problem. This issue limits the depth of the networks. To avoid this, a ResNet uses the addition of the input data to the output between layers. This allows the network to get the difference between the input and the output, the residual. This has led the ResNet to improve accuracy at higher layer depths compared to the regular CNN. As stated before, the number next to the ResNet name denotes the layer depth. The higher the layer depth, the more accurate the model is expected to be. The ResNet also outputs the confidence of the tool in the given classifier. For tuning the ResNet-26 a confidence of 60% was used while for the ResNet-44 a confidence of 80% was used. This means that results produced by ResNet-44 will have a higher confidence and as the machine learning tool utilizes more layers.

For further implementation, the accuracy of the plume detection of the different methods should be determined and compared with APE. This is done in the chapter 7. Only the machine learning methods that outperform APE should be considered to be integrated into the algorithm. Before this comparison can take place, the data to train the machine learning approaches needs to be considered first. Obviously, these have to be CO plumes. The input data can be further optimized which is discussed in the next section.

4.2. Data augmentation algorithm

As written in section 3.1, question QWP2.2 asks which features are the most important for detection. These features should make it possible for the various plume detection methods to filter the data correctly. The original L1B data from TROPOMI, however, includes a lot of different other information that could cloud the judgment of these methods. Due to the inclusion of NaN values in the TROPOMI data, it is in most cases not even possible to process the unaltered TROPOMI data. Therefore some augmentation of the data is needed to perform the plume detection.

For a large portion of the project, the SVM method was the only available method for plume detection. This method was thus used for the optimization of the augmentation algorithm. It was previously discovered that the SVM tool is very sensitive and therefore the input data needed to be augmented significantly to get a good dataset to serve as input data. Thus the aim of the data augmentation algorithm is to remove the information from the input data that is not relevant and to enhance the features of the plume. This should make the plumes easier to detect.

The TROPOMI CO concentration data is used as input data. This data should be 41 by 41 pixels and the data should be destriped. The granule size of the data is determined by the APE algorithm. An example of destriping is shown in Figure 2.22. The center of the granule should be the point source emitter. It is possible to either increase or decrease the input array of data but this could come at the expense of the background data value, as it would affect the median background value. Adding more pixels would also increase the number of data points that are set to zero which would therefore not be useful in the detection. An example of the use of the algorithm can be seen in Figure 4.2. An example of input data can be seen in Figure 4.2a.

The steps in the algorithm are as follows. First, the NaN values in the data are removed from the original input data. These values represent positions where, for a variety of reasons, the TROPOMI instrument was unable to get accurate CO concentration values. The NaN values in the data are set to zero. These values are not usable in calculations and will also cause issues with the machine-learning approaches. The end goal is to normalize the data between 1 and 0, where 0 denotes uninteresting features in the data, so putting them to zero is not a completely arbitrary choice. The step can be seen in Figure 4.2b. This data is then normalized between 0 and 1.

The second step is to remove the background data and only keep the data that is significantly higher than the background concentrations which would be the enhancement. It is assumed that for a 41 by 41 pixel granule, the noise in the background is Gaussian distributed while the CO concentration of plume data is assumed to be higher. By taking the median value in the array to be the mean of the noise distribution, the standard deviation in the data can be calculated by subtracting the mean of the distribution by the value that is the 15.9$^{\text{th}}$ percentile. This value comes from calculating the area on the right side of the graph of a normal distribution that is higher than 2 standard deviations. Subtracting all the values in the array from the median value results in the enhanced data as can be seen in Figure 4.2c. Diving this by the standard deviation calculates the number of standard deviations these values are higher or lower than the median. An example of this result can be seen in Figure 4.2d. If these values are at least 2 standard deviations higher, it can be said with 95% certainty that the values are statistically significant and with 97.1% certainty that they are not part of the background noise. The values that are not higher than 2 standard deviations are set to zero as they are expected to be background data. The result of this can be found in Figure 4.2e.

To get to Figure 4.2f, two other important effects need to be taken into account. These can be seen in Figure 4.3. These include the wind direction and the distance to the source. Due to the spatial resolution of the measurement instrument, the distance of each pixel with the center pixel is variable. Especially when the data is used near the edge of the swath-widths of the instrument, the data will be stretched in the longitudinal direction. Since the plume originates near the source it should be visible near the source pixel which is set in the center, the pixels that are much further from the center are far less relevant. This is because if the plume is not visible at the center, there is no reason to investigate the data further away. Therefore the data near the center is much more important. Using the Gaussian function as denoted in Equation 4.1, the distance to the pixel was used to set the pixels of data located further away to zero. The terms σ_x^2 and σ_y^2 were both found to be 10.000 by tuning. The distance in both the latitudinal and longitudinal are equally important. The term A is set to 1 as the resulting data should be between 0 and 1, with 1 being the largest possible value. The result of taking the proximity into account can be found in Figure 4.3a.

$$f(x,y) = A \cdot exp\left(-\left(\frac{(x-x_0)^2}{2\sigma_x^2} + \frac{(y-y_0)^2}{2\sigma_y^2}\right)\right) \tag{4.1}$$

The wind direction was also taken into account. This was done by calculating the angle between the line of the source and the locations of the grid and the wind direction line from the source. The assumption is made that the plume should be in the direction of the wind at the source pixel. To avoid removing too much data in the event that the wind direction changed a lot before the measurement or because of a possible error in the wind data, the algorithm should not punish data that is not entirely in the direction of the wind. This is done by using another Gaussian function which is shown in Equation 4.2. The term A is set to 1, as the data is normalized between 0 and 1. The maximum value should be 1. The term σ_θ^2 was set to 15.000 by trial and error. This value could be further optimized. The effect of the wind direction can be found in Figure 4.3b. The combined effect of both the proximity and the wind direction can be found in Figure 4.3c.

$$f(\theta) = A \cdot exp\left(-\left(\frac{(\theta-\theta_0)^2}{2\sigma_\theta^2}\right)\right) \tag{4.2}$$

The final step is to set the value of the source pixel to one. This makes it easy to see if the plume is close to the center. It also serves as a verification. If the process of the algorithm worked as it was supposed to, in the resulting dataset the value of the middle pixel should be 1. The final result of the augmentation can be seen in Figure 4.2f. Examples of the augmentation algorithm used on other steel plants' data can be seen in Figure 4.4. The figure shows the ability of the algorithm to filter data for different atmospheric backgrounds by being able to separate enhancements from different atmospheric backgrounds of steel plants in China, India and Germany.

A plume from the Angang Steel Company Limited Anshan production base in China can be seen in Figure 4.4a. The concentration of CO that comprises this plume is much higher than the background and thus easy to separate. The result can be seen in Figure 4.4b. The plume from the Hüttenwerke Krupp Mannesmann (HKM) steel plant from Germany in Figure 4.4c is not as easy to isolate as the plume enhancement is only slightly higher than the background. The result in Figure 4.4d shows a plume but the plume is much less visible than the plume in Figure 4.4b. The final plume from the Kalyani Steels Hospet plant in India can be seen in Figure 4.4e. In this array, it is clear that the pixels are stretched in the longitudinal direction. This does not affect the result in Figure 4.4f due to filtering from the algorithm.

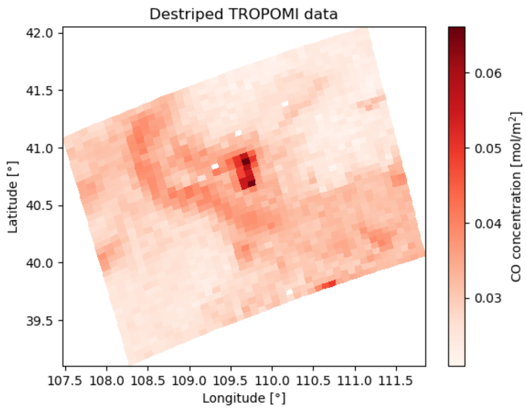

(a) Destriped TROPOMI data used as input for augmentation algorthm.

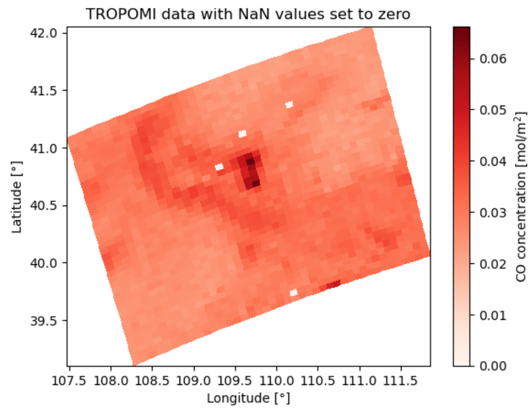

(b) The NaN values are set to zero to avoid issues with Python functions.

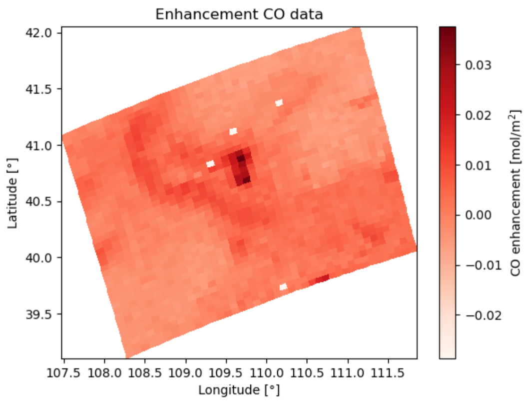

(c) The enhancement data as a result of subtracting the data from Figure 4.2b from the median value of the data in Figure 4.2a.

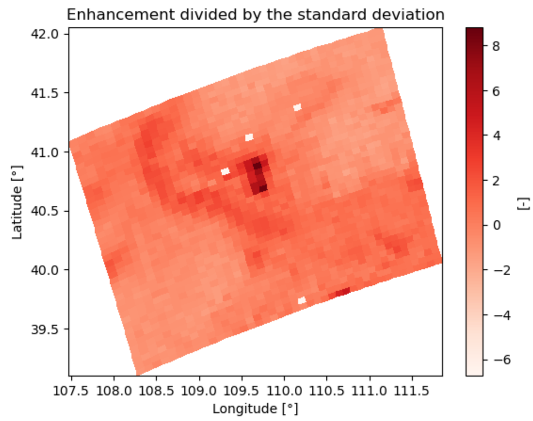

(d) Result of dividing the enhancement values from Figure 4.2c by the calculated standard deviation.

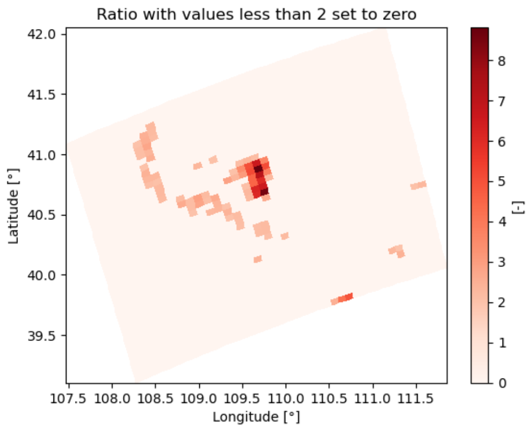

(e) Only the values higher than 2 standard deviations remain.

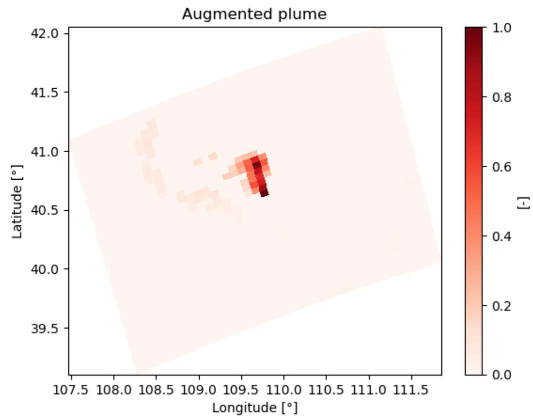

(f) Effect of Figure 4.3 taken into account as well as marking the source spot.

Figure 4.2: Overview of the steps in the augmentation algorithm of a plume over the Inner Mongolia BaoTou Steel Union Co., Ltd. plant on 3 November 2019.

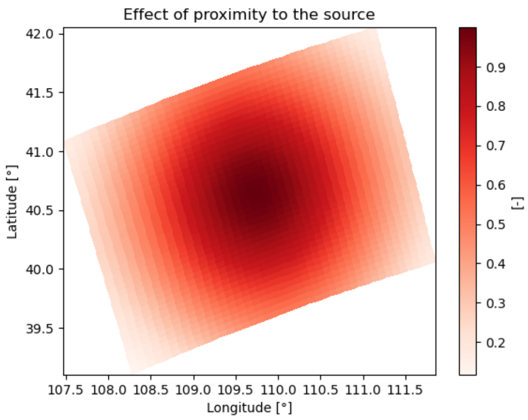

(a) The effect of proximity of the source. The closer the pixel is located to the source pixel, the higher the weight of the data.

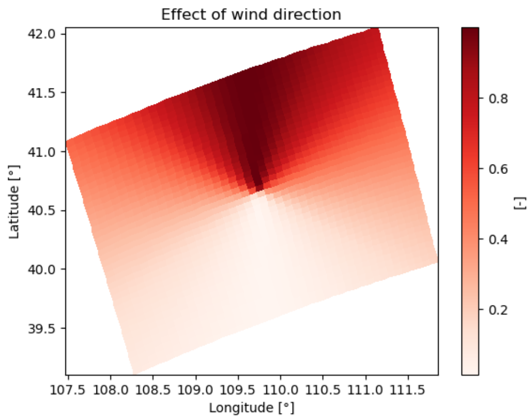

(b) The effect of the wind direction. The closer the pixel is located to the source pixel, the higher the weight of the data.

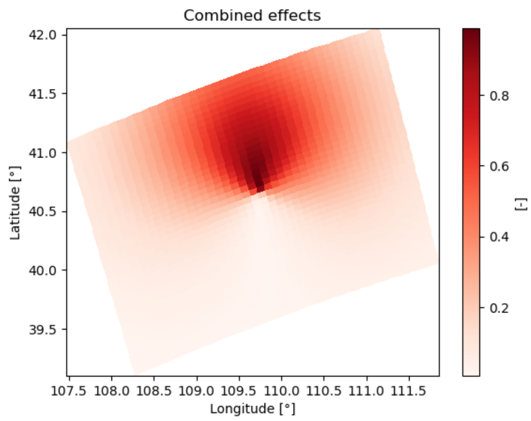

(c) The combined effects of the wind direction and the wind direction.

Figure 4.3: Overview of the effects of proximity and the effect of the wind data on the effectivity on the data.

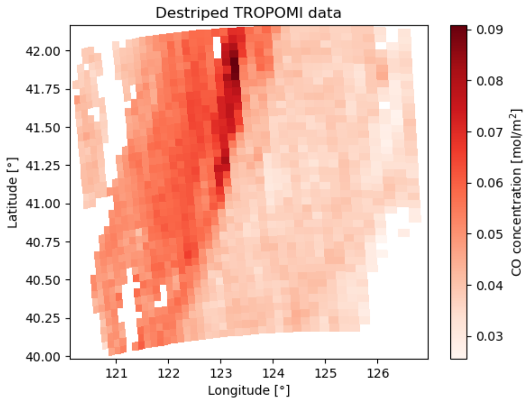

(a) A plume over the Angang Steel Company Limited Anshan production base from 24 October 2019.

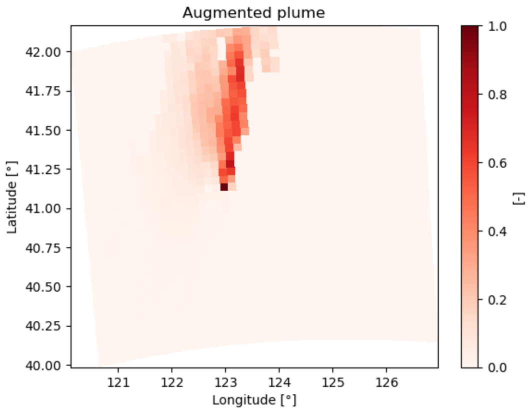

(b) Result of the plume augmentation.

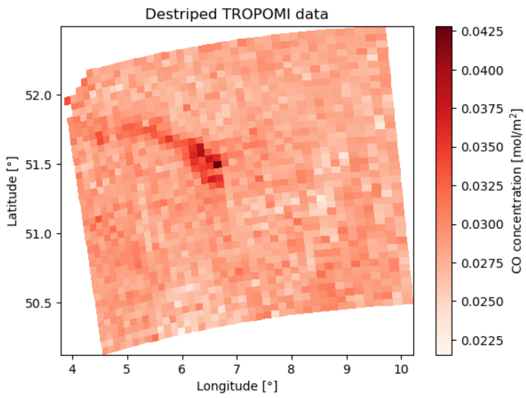

(c) A plume over the Hüttenwerke Krupp Mannesmann (HKM) steel plant from 20 October 2019.

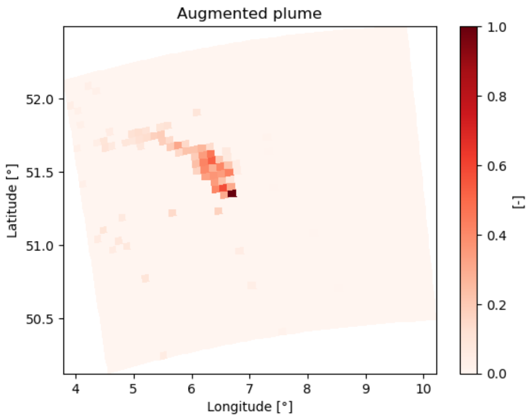

(d) Result of the plume augmentation.

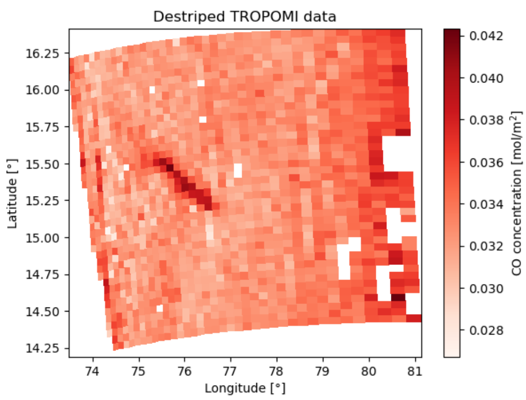

(e) A plume over the Kalyani Steels Hospet plant from 22 January 2020.

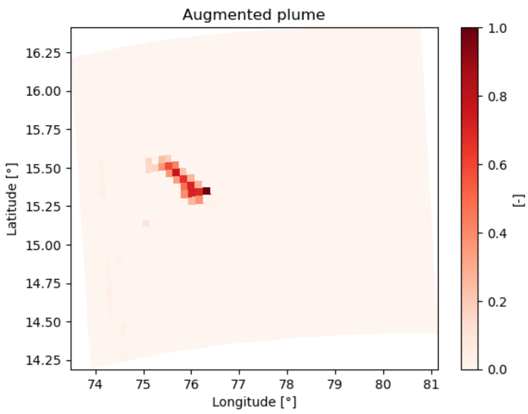

(f) Result of the plume augmentation.Result of the plume augmentation.

Figure 4.4: Examples of the augmentation algorithm used on data with different atmospheric conditions.

5

Novel training dataset for plume detection using machine learning

In this thesis, a large dataset for training AI plume detection algorithms that can be applied to pollution events from industrial and biomass-burning sources has been developed. This chapter denotes the process of locating and selecting the right data to be used for the training of the plume detection algorithms. The general approach for creating the dataset is explained in section 5.1. The sections following the general approach describe the data that was used in the datasets. In section 5.2, the steel plant data extraction was described while in section 5.3, the wildfire data extraction was described. This is followed by a section about the noise data used in section 5.4. The process of labeling and the resulting datasets are described in section 5.5.

5.1. General approach for selecting training data

The project aimed to create a dataset that includes both wildfire and steel plant data. For the machine learning tools to know the difference between plumes and no plumes, the data must consist of data without plumes. In the following sections, no-plume data is sometimes also referred to as noise data. To avoid bias in the machine learning tool, the dataset should be balanced between the wildfire, steel plant and noise data. Therefore, 50% of the data had to include plumes and 50% had not to include plumes. If a 50% split cannot be accomplished the plume detection tool could overestimate one of the classes versus the other. Therefore a 50-50 split would be the most optimal scenario.

Of the plume data, the idea was to include plumes from the steel industry in 50% of the plume data and the other 50% of the plume data to be made available for the wildfire data. An overview of this approach can be found in Figure 5.1. In terms of dataset size, the more information the better, thus the more images the better. A selection in steel plant and wildfire data was made to create the datasets which can be found in section 5.2 and section 5.3 respectively.

During the project, other ideas regarding the labeling of data were considered. However, it became clear that the focus should lie on finding the plumes rather than for example classifying the plumes as wildfire or steel plants. This is because the project would become more complex unnecessarily and would deviate the project from the initial goals. An issue that could arise is for example that there are not enough data points for one group compared to the other groups which would lead to a bias in the detection. The focus is on filtering noise data from plume data to avoid processing noise data. Thus when labeling the steel plant and wildfire plumes both datasets received the same label of 1. The noise data would receive a label of 0. In the following sections, the selection of data for steel plants and wildfire plumes is discussed.

5.2. Steel plant data

An issue surrounding the quantification of steel plant emission is that the location of steel plants is not easily available online. However, the Global Energy Monitor (GEM) keeps track of these locations. They host a Wikipedia-style article archive that can be edited by users. The coordinates of these steel

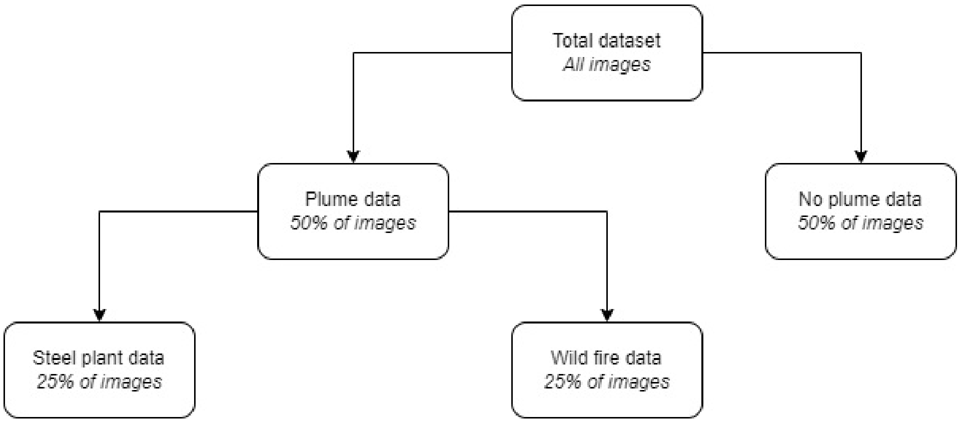

Figure 5.1: Simplified overview of the desired composition of the dataset

plants can be manually validated using available satellite images such as those on Google Maps. A list of steel plant locations created based on GEM entries has 1115 entries[1]. Some of these entries are duplicates or are of nonexistent steel plants. These non-existing steel plants could be planned or their construction was canceled at some point. In this project, these should be filtered as they do not emit CO because they do not exist. An overview of the location of different steel plants based on the available information on the GEM website can be found in Figure 5.2.

Other data taken from GEM includes a rough estimate of the capacity of the plant and whether or not the plant is operational. This data is important as the capacity is an indication of how visible the plumes are in the data. The visibility of the plume can be determined by subtracting the background from the data scene which results in the plume enhancement. The bigger the enhancement, the more visible the plume is. Thus the visibility is a product of the background and the plume enhancement itself. The enhancement of a CO plume depends on the emission source. The more a steel plant emits, the larger the enhancement of the plume will be. The larger the capacity, the more a steel plant could emit. Thus the capacity is loosely connected to the visibility of a CO plume in the data. As described in subsection 2.5.1, the steel plants with high capacity that were operational were prioritized [3].

The data for the steel plants to be used in the machine learning datasets was downloaded by first making a selection for steel plants. To ensure that the steel plant had enough plumes in the dataset, the average divergence was investigated. This average divergence was calculated from the time series of data. The method to produce these images is explained in subsection 2.5.2. It is important to realize that the average divergence over a long period of time should show a peak to indicate an emission source at that location. If this peak is not visible, very few to no plumes from the steel plant are in the dataset. Going through these steel plant datasets is not as interesting because it does not test the plume detection as well. This is because a steel plant is not always operating, therefore for a steel plant that emits a lot of pollution, there will also be no plume granules. Examples of average divergences of steel plants can be seen in Figure 5.3. In Figure 5.3a an example of a peak in the average divergence is visible in the right plot. The plot on the left shows a time series of emissions as calculated on the same data used to calculate the average divergence peak. An example of a no peak in the average divergence peak can be seen in Figure 5.3b. For this steel plant, there could be a few plumes visible in the data but it does not line up with the time series in the right plot. That plot suggests a very active and highly emitting plume which is extremely unlikely. It seems that the no-plume granules in the dataset are processed to large emission estimates. This is because of the issues in APE that detect plumes while there is no plume data. The final example of the average divergence can be seen in Figure 5.3c. In this example, a peak is visible in the data but it is not as clear as in Figure 5.3a.

As stated in subsection 2.5.2, a large amount of data was processed to calculate the average divergence of different steel plant locations. Due to time constraints, it was not possible to go through all of the steel plants. In the end, out of 1115 steel plants, the average divergence of 514 steel plants was calculated. These divergences were then inspected to see if the average had a peak at the emission source. This resulted in a list of 185 steel plants with an average divergence peak. These steel plants

[1] List is created from the entries of `https://www.gem.wiki/Category:Steel_plants` and was retrieved on 25-9-2023

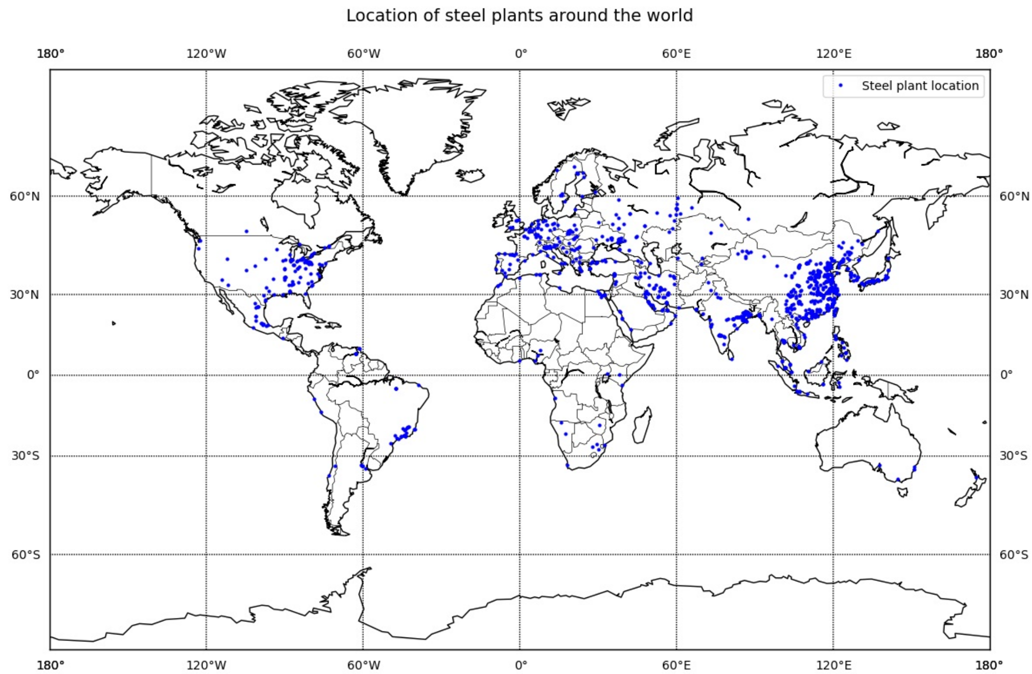

Figure 5.2: Locations of steel plants according to GEM based on the list of steel plants on 25-9-2023.

were primarily considered to create the plume dataset to train the machine learning methods.

5.3. Wildfire data

As denoted in subsection 2.5.1, wildfire plumes can easily be investigated by TROPOMI when the active fire data from VIIRS is considered. Using the VIIRS dataset as input for APE resulted in a large collection of wildfire plumes. This wildfire data was used for the training of the machine learning tools. In particular, the locations found using the active fire data of the years 2018 to 2022 were used. It can be seen in Figure 5.4. This approach resulted in about 8000 images of wildfire plumes that were used in the training data. When comparing the locations of the wildfires in Figure 5.4 with the locations of the steel plants in Figure 5.2, the locations are complementary. In places without wildfires, there are steel plants and vice versa. This excludes areas with deserts and oceans. Combining both datasets effectively would give worldwide coverage.

5.4. No plume data

The no-plume data used in the datasets are in two categories. First is noise data from steel plants or wildfire locations. If no plume was visible when looking at the data from a steel plant or wildfire, it was considered to be noise. The second category of noise data was taken from general orbit data. This data was not taken from an individual industrial source. The idea behind this data was to ensure that various atmospheric effects would be covered as noise data by taking a random selection of locations without plumes.

5.5. Labeling and creation of the datasets

In the early stages of the project different labels it was considered to use different labels for types of plumes. This approach was abandoned for the thesis project, as the different labels for plumes were not considered to be independent of each other. Another factor was that all the different plume labels would be summed together to calculate the number of plumes detected. In this way, the final result did not depend on these individual labels and thus they were removed. The labels are straightforwardly 1 for a granule with a plume and 0 for a granule without a plume.

The process of labeling was relatively tough as there numerous cases that were not straightforwardly

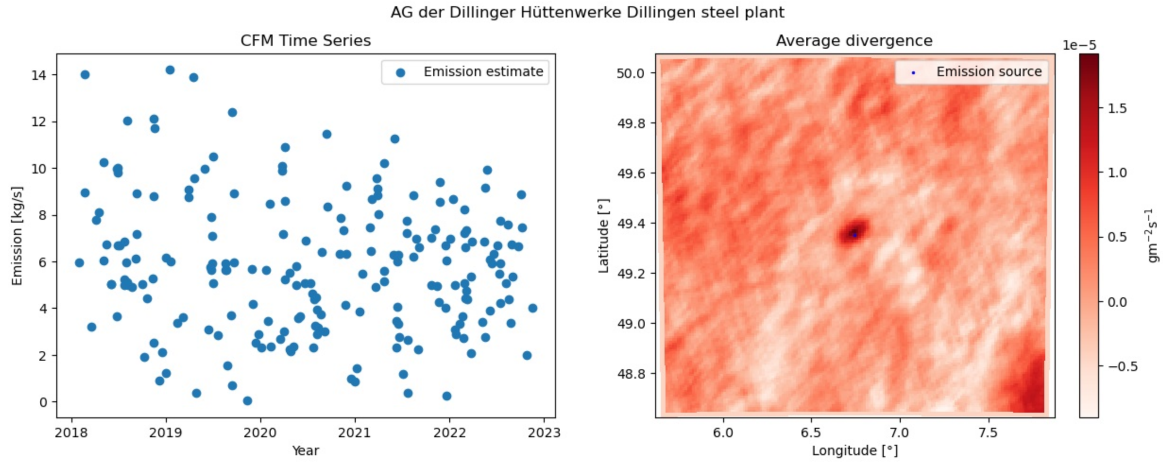

(a) The average divergence peak in the right plot validates the existence of an emission time series on the left plot.

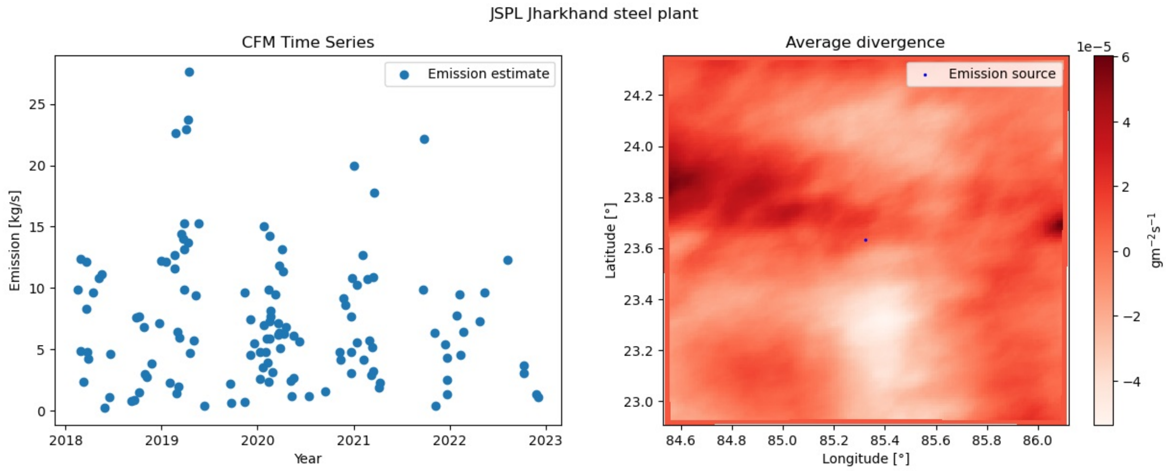

(b) The nonexistence of an average divergence peak in the right plot calls into question the existence of an emission time series on the left plot.

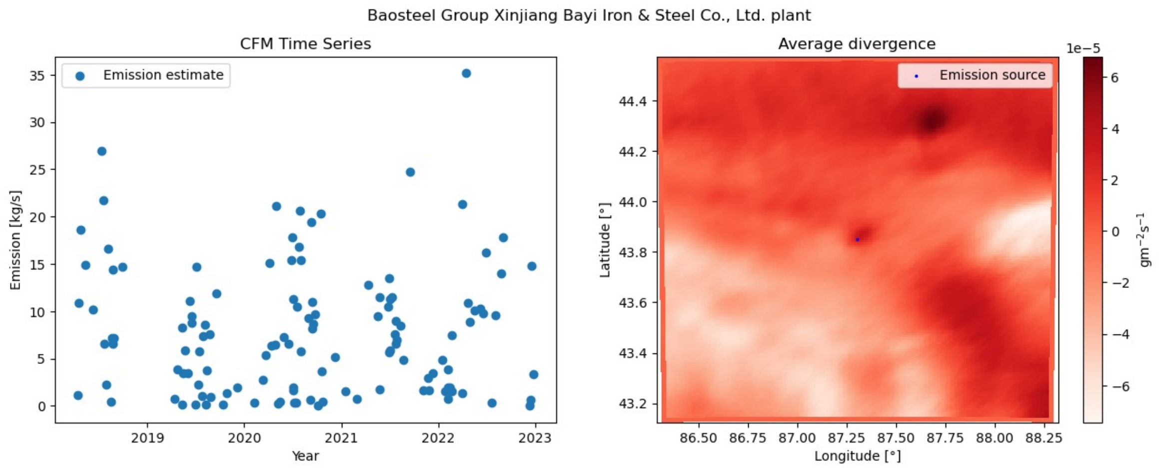

(c) The average divergence peak in the right plot validates the existence of an emission time series on the left plot. However, the peak is less visible so the large values of the emission estimates seem questionable.

Figure 5.3: Overview of different average divergences

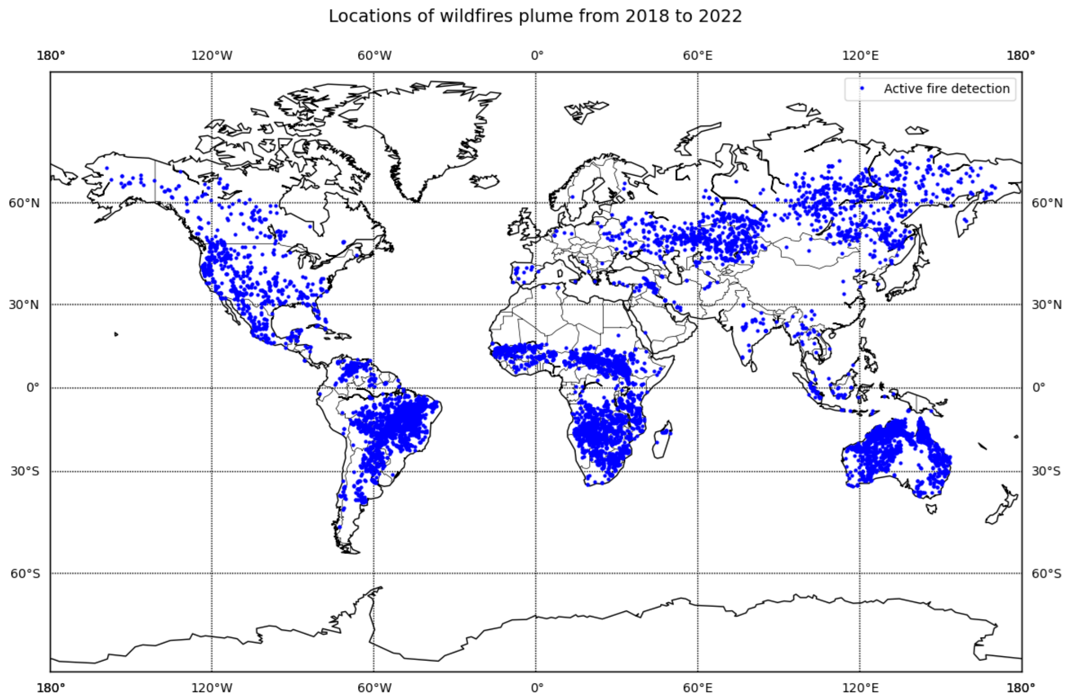

Figure 5.4: Locations of active fire data from VIIRS processed by APE from 2018 to 2022.

noise or plume. Good examples of plumes can be found in Figure 4.4. These are also examples of plumes that experienced very little wind and are therefore pointed straight up or plumes that combine to create a larger enhancement in an irregular form. One can be seen in Figure 5.5. When going through the augmentation algorithm this results in large concentrations at the source but without a plume tail. Other problematic issues came from plumes that were visible in the original data but either disappeared completely or partially after augmentation. This happened due to the plume enhancement concentration not being significantly bigger than the background value. This shows that the augmentation algorithm creates a bias towards large CO concentration plumes in the plume detection tools. Therefore, in the future, a different way should be found to cut plumes out of the background data. This issue is nothing new as it is a version of the background problem described in subsection 2.5.3. The CO concentration background data can vary and this causes issues in determining a good background value.

When going through the initial steel plant sources for plumes to use, the data adaptation algorithm was not finished. Therefore only the concentration data was used to perform the labeling. This resulted

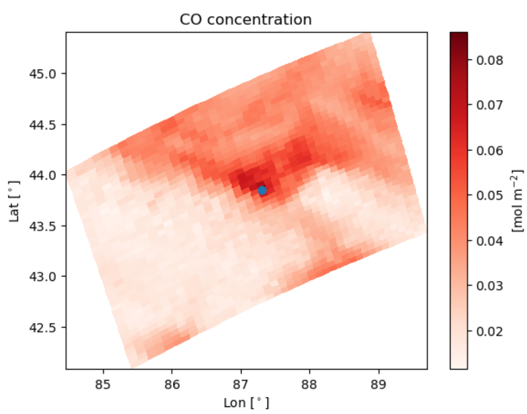

Figure 5.5: Plume with an irregular shape. The blue dot represents the location of a steel plant

in trouble when using the labels from the concentration data on the augmentation data. Thus a second round of labeling happened to see how well the labels held up. In some cases, the labels had to be changed to indicate noise data. Which lowered the number of steel plant data points. This showed the challenge of getting a large and diverse steel plant plume dataset and made the labeling process much longer than it was planned to be. The process of creating the different training datasets is described in the following subsections.

5.5.1. The small dataset

As written in QWP2.1 in chapter 3, the initial aim was to create a dataset of 10.000 images. As written in section 5.1, 25% of this data would have to be steel plant data. The aim was to pick 25 steel plants from the 185 steel plants with an average divergence peak and find 100 plume images. This would amount to 2.500 steel plant plume images. When the labeling process was found to be very time-consuming, a small dataset was created first. This dataset would be used as an initial test. This was needed as before the creation of this dataset, there was some concern that including wildfire data would not result in better accuracy. This was because wildfire plumes are physically different from steel plant plumes. This is because of the different atmospheric conditions and the processes that take place for the two types of plumes to appear.

The dataset includes 2.000 images which is about 300 more images than the internship dataset. About 50% of the images are noise data while the other images are plume data. Half of the plume images are from steel plants while the other images are from wildfires. The noise data was taken from random orbit data. For this dataset, the chosen steel plants can be found in Table 5.1. The wildfire data was limited to 2022 data. Small tests showed that the dataset was capable of detecting plumes. With the inclusion of different data points, it was decided to continue down the road of the project.

5.5.2. The large dataset

To continue in the same direction as the small dataset more data was needed to create the dataset as described in section 5.1. As it became possible to download larger amounts of wildfire data, this data was utilized first. Due to difficulties, at first, it was only possible to slowly download wildfire data. After a fix, this changed dramatically. This resulted in the download of about 8.000 wildfire images. As it took much longer to produce the steel plant data, the wildfire and steel plant data could not be balanced. The dataset also includes 100 granules of 14 steel plants. They can be found in Table 5.1. The original aim was to gather 25 steel plants. Out of the 185 steel plants with an average divergence peak, a selection of 25 steel plants was made. This selection was based on the geographical location of these steel plants. The aim was to have as much worldwide coverage as possible. During the labeling process, it was found that 11 out of the 25 locations had little to no visible plumes. Due to the time-consuming nature of labeling it was decided to proceed with the data of the 14 steel plants.

The large dataset includes about 20.000 images with the vast majority being the newly downloaded wildfire data. Due to the size of the dataset, the quality of the labeling is not as good as for the other datasets. An issue that was found is that the wildfire plumes were processed by APE. That means that if APE detected it as a plume. Since APE has issues detecting plumes, this resulted in a large number of wildfire plumes that are not actually plumes, which caused a lot of relabeling of plumes. This took a lot of time.

5.5.3. The industrial plumes dataset

An issue that arises with the inclusion of wildfire data is the complexity that the wildfire plumes add to the analysis of the plume detection. Since the locations of wildfires are known due to VIIRS, it is not as interesting to focus on finding wildfire plumes as compared to steel plant plumes. During the project, the focus shifted somewhat to investigating emissions of steel plants. Thus a dataset was created to exclude all the wildfire data and only input the steel plant data into a separate dataset. To do this the steel plant data from the large dataset was reused. During the transfer of steel plant data, a mistake was made. This mistake resulted in the exclusion of the three steel plants as compared to the steel plants included in the large data. To avoid confusion, the steel plants included in different datasets can be found in Table 5.1.

Table 5.1: Different steel plants with their coordinates inclusion into the different datasets

Steel plant	Latitude [°]	Longitude [°]	Index	Large dataset	Small dataset	Steel dataset	Internship dataset
AG der Dillinger Hüttenwerke Dillingen steel plant	49.353884	6.746603	18	X			
ArcelorMittal Kryvyi Rih steel plant	47.874411	33.392993	81	X	X		
ArcelorMittal Temirtau steel plant	50.045849	73.040384	95	X			
Bengang Steel Plates Co., Ltd. plant	41.274132	123.722099	129	X	X	X	
Evraz ZSMK steel plant	53.88618	87.257618	257	X		X	
Inner Mongolia BaoTou Steel Union Co.,Ltd. plant	40.647997	109.740898	441	X	X	X	X
Ma'anshan Iron & Steel Co., Ltd. plant	31.698884	118.468194	601	X		X	
Magnitogorsk Iron & Steel Works	53.427593	59.054122	602	X	X	X	
Mechel Chelyabinsk Metallurgical Plant	55.27073	61.436492	613	X		X	
NLMK Lipetsk steel plant	52.557372	39.629574	667	X	X	X	
Shandong Iron and Steel Co., Ltd. Laiwu Branch plant	36.093496	117.837567	789	X		X	
ThyssenKrupp Steel Duisburg steel plant	51.491649	6.733051	931	X		X	
Tonghua Iron & Steel Co., Ltd.	41.779125	126.022349	948	X		X	
Wuhan Iron and Steel Co., Ltd. Qingshan plant	30.616222	114.444975	1009	X		X	

5.5.4. Availability of the dataset

The datasets are available at Zenodo This dataset includes the destriped TROPOMI data as well as the label and the augmentation. By making these datasets available, other researchers are encouraged to use them in their own analysis. A big issue in machine learning is the unavailability of the training data. To aid in solving this issue I have added my datasets to the scientific community [2].

[2]The datasets are available at `https://zenodo.org/records/14604467`

6

A Novel Approach to Estimate Emission without Plume Detection

This chapter explains the development of a new emission estimation method. This method is meant to act as an independent validation tool for the plume detection. In subsection 2.5.2, the current method of calculating the emission of plumes in APE was discussed using either the CFM method or the average divergence method. In this chapter, the process of refining the average divergence method to create a time series of emission estimates is explained in section 6.1. The following section explains the benefits of the newly created method and its use for the plume detection part. It can be found in section 6.2.

6.1. Applying the divergence method for time series analysis

Currently APE has two different emission estimation methods. The average divergence method, as described in subsection 2.5.2, is commonly used to calculate the average emission of a source. The CFM method, also described in subsection 2.5.2, is capable of estimating the emissions of single-plume images. Would it be possible to apply the average divergence method to estimate emissions for single-plume images? As explained in the previously mentioned section on CFM, the CFM method is very sensitive to noise and it would thus be interesting to investigate other alternatives. Since the divergence method is already coded into APE, it served as a starting point for this investigation.

The average divergence method performs three steps. First, the divergences of each granule are calculated. Second, the average of the divergences is taken. Third, the area around the source is integrated to obtain the emission estimate. These steps can also be moved around. It should not matter if you perform step three (integration) before performing step two (averaging). If the steps are performed in this way, a time series of emissions is calculated before the averaging. An overview of the methods can be found in Figure 6.1. The original average divergence method explained in Figure 2.5.2 corresponds to method 1 while the change to this method suggested in this section corresponds to method 2. Performing the first two steps of the second method results in a time series of emission estimates. Thus it is possible to apply the average divergence method to estimate emissions for single-plume images.

The emission estimate procedure for a single granule can be seen in Figure 6.2. Starting from an input granule with a visible plume in Figure 6.2a, the divergence field is calculated. This field can be seen in Figure 6.2b. The field is then divided into the inner circle and the outer circle of data as can be seen in Figure 6.2c and Figure 6.2d. This division is important as a background correction is needed. By correcting for the background, the enhancement due to the source can be found. Without correction, the emission will also include background CO. The data of the outer circle will be used for that purpose. The inner circle is defined as the divergence data around the steel plant location in a radius of 21 km. This distance was chosen because the resolution of TROPOMI data is 5.5 km by 7 km and thus the radius will include at least three TROPOMI pixels of information. This number was picked empirically. This value could be further optimized as for large emitters of CO, it would be better to increase the inner radius to capture more of the emission. Conversely for small polluters of CO, it would be better to make the inner circle smaller to limit the background CO that gets added to its emission. The radius

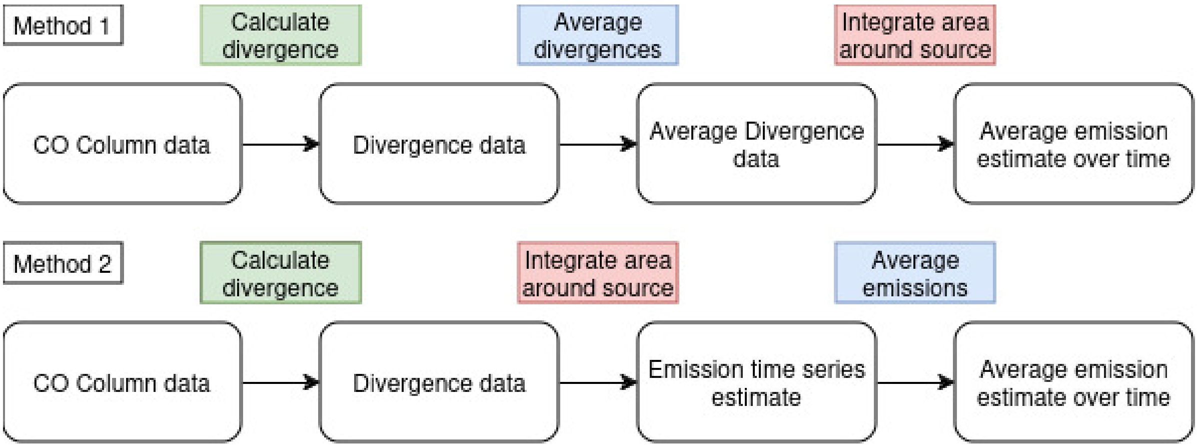

Figure 6.1: Comparing the different methods for average divergence with method 1 corresponding to the original average divergence method described in Figure 2.5.2 and method 2 being the new approach.

here was picked to accommodate a large variety of pollution strengths in sources. When investigating specific emission sources it would be interesting to see the effect of changing the inner radius on the emission estimates.

It has to be noted that there are more pixels included in Figure 6.2c which is due to the interpolation of the data onto a grid with a spatial resolution of 4 km by 4 km. Thus about 6 pixels would be included from the interpolated grid. The interpolation onto the 4 km by 4 km grid is not a necessary step to estimate the emissions. It is done here to be able to check the average divergence. If the average divergence shows a peak at the source, the data has been processed correctly. To get the emission from this collection of pixels, first, the background value should be determined. This is done using the pixel collection in Figure 6.2d. These pixels are taken to be in a radius of 63 km around the steel plant point source without the inclusion of the inner circle. As previously, this number was picked empirically and could be further optimized but this radius should include a large enough area to determine a representable background value. The background value is determined by finding the median value of the values in the background circle. The inner circle values are subtracted by the background value and then multiplied by the resolution of a single pixel. These values are summed and recalculated to kg/s. The resulting emission estimate from this example was calculated to be 16.81 kg/s.

An example of a time series calculated by the new method can be seen in Figure 6.3. This is the time series for the Baosteel Desheng Stainless Steel Co., Ltd. plant in China from late 2019 to mid-2024. The time series itself already shows a certain behavior in the data however, since the individual estimates have biases in the way they are calculated, the averages in the data should also be investigated. Structuring the time series of emission data into a histogram and normalizing it shows a clear structure in the data. As can be seen in Figure 6.4, the data looks to be Gaussian distributed with a slight skew to the positive emission. It is visible however that certain emission estimates make no sense. It is not physically possible for a plume to have a negative emissions as a steel plant cannot emit a negative CO concentration. Therefore the negative values are mistakes that need to be explained. There could be a mistake/inaccuracy in the wind direction data or the CO column data could have too many NaN values or too much low-quality data. These issues would result in inaccurate emission estimations. The large emissions, like the measurement of about 150 kg/s, also seem to be an overestimation.

During the thesis, several steel plants in Ukraine were investigated. The steel plants in question were the Azovstal Iron and Steel Works, Metinvest Ilyich Iron and Steel Works, the ArcelorMittal Kryvyi Rih steel plant and the Metallurgical Plant Kametstal. The first two steel plants are located in the city of Mariupol. During the war in Ukraine, the steel plants sustained heavy damage. Almost since the start of the war, the steel plants have been out of order [46]. As stated in section 2.3, the Sentinel-5P satellite has been in operation since late 2017 and was thus able to capture the change in concentration over Ukraine. The two steel plants in Mariupol are very close to each other and thus it is not possible to separate their emissions. Thus the steel plants are analyzed together. A time series of emissions of both steel plants can be seen in Figure 6.6. The invasion of Russia into Ukraine started in February 2022 and this has been denoted in the figure. As can be seen, there was a stark decrease in emissions

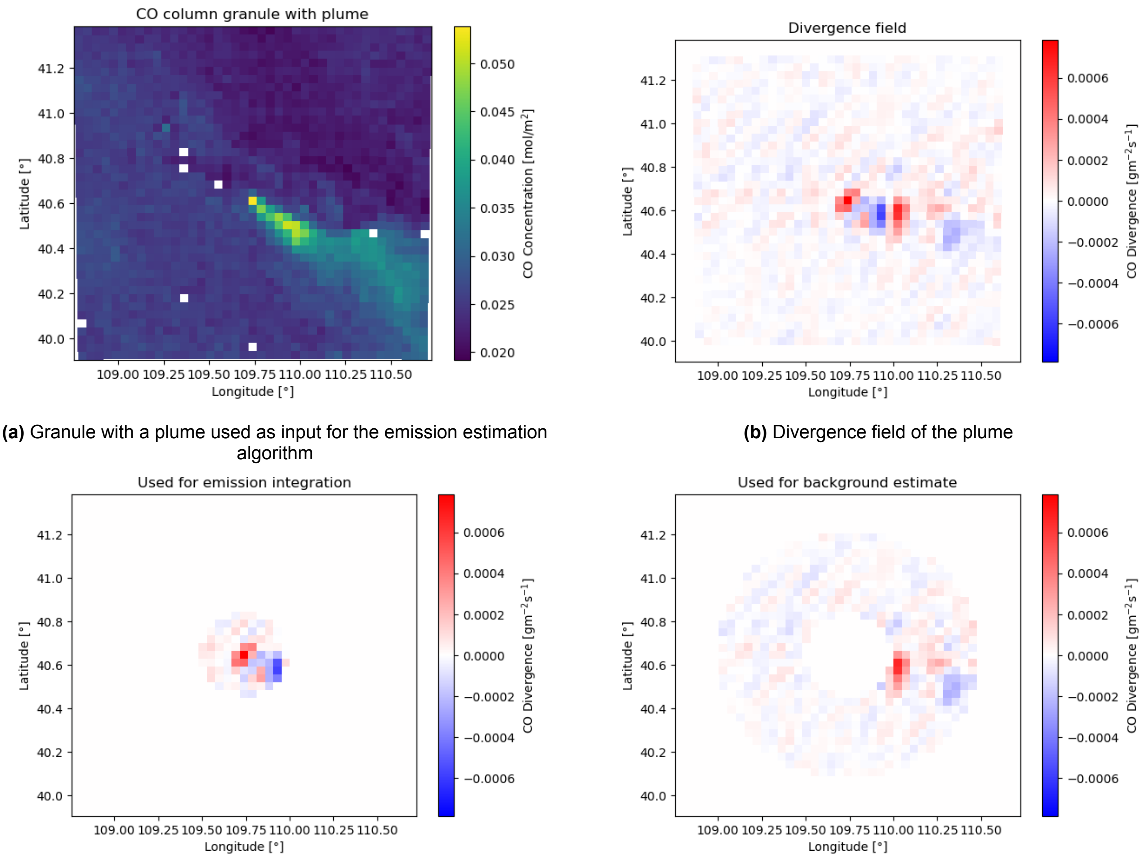

(a) Granule with a plume used as input for the emission estimation algorithm

(b) Divergence field of the plume

(c) Data used for integration in the calculation of the emission estimate

(d) Data used to determine the background value

Figure 6.2: The CO column and the divergence field of a plume granule used in the emission estimation algorithm and the data used to calculate the emission estimate

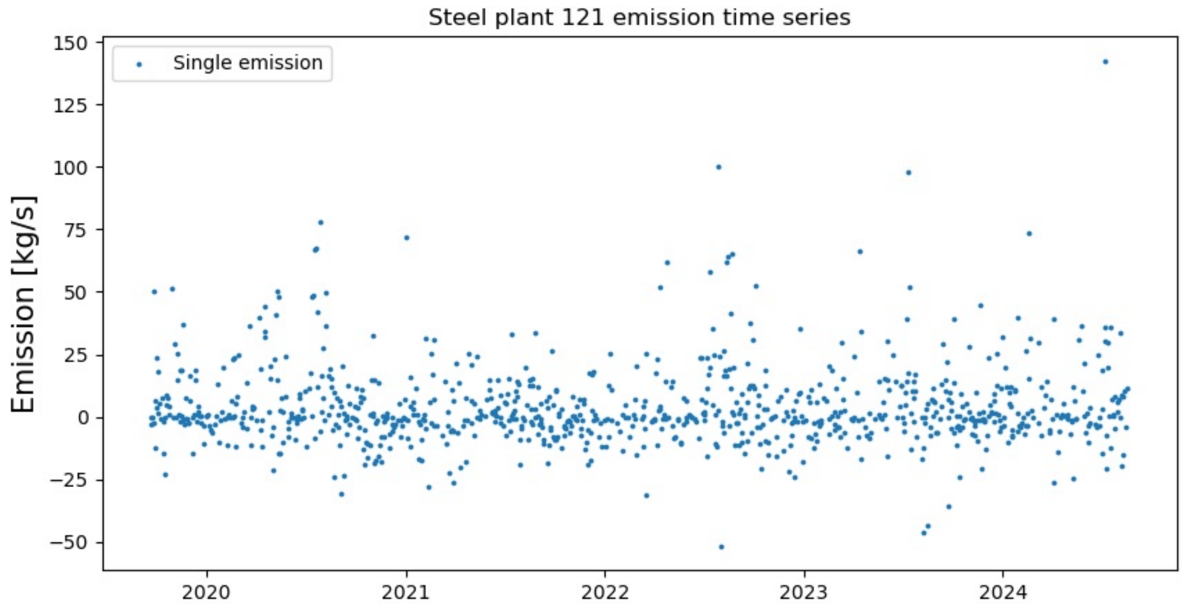

Figure 6.3: Time series of the Baosteel Desheng Stainless Steel Co., Ltd. plant which corresponds to the index of 121 in the steel plant location list.

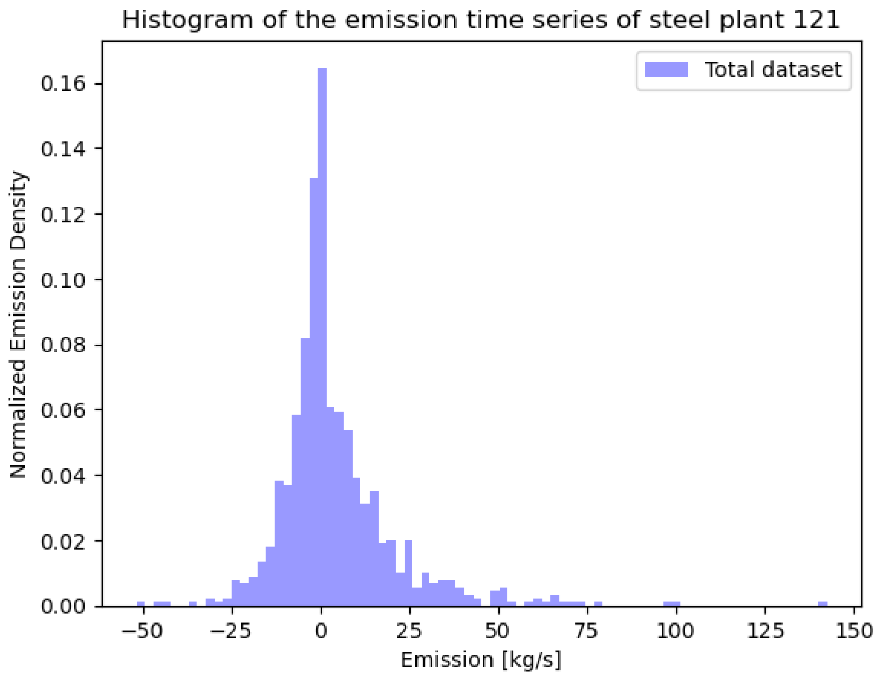

Figure 6.4: Histogram of the time series of the Baosteel Desheng Stainless Steel Co., Ltd. plant which corresponds to the index of 121 in the steel plant location list.

after the invasion. To further investigate this, the emission estimation data has been plotted into a histogram. This resulted in Figure 6.7a and Figure 6.7b for the before-invasion and after-invasion cases, respectively. The contrast between the cases is even more clear in the histograms. The data in the before-invasion case looks like a positively skewed Gaussian distribution while the after-invasion case looks like a Gaussian distribution. Since it is known that the steel plants stopped production quite soon after the start of the invasion, it can be assumed with confidence that the histogram is composed of almost exclusively noise data. This can also be seen in the histograms as there are vastly more emissions around 0 kg/s in the after invasion histogram compared to before the invasion histogram. This gave rise to the idea of modeling the noise data as a Gaussian distribution and the total data as a skewed Gaussian distribution. The point of this is to isolate the plume data. This can be done by simply subtracting the total data distribution function from the noise data. An example of this can be seen in Figure 6.8. This analysis makes it possible to create an estimation of the distribution of both the noise data and the plume data without the use of plume detection.

To fit the total data distribution, the formulas in Equation 6.1 and Equation 6.2 were used. As can be seen in the equations and the histogram, the total fit distribution function is a skewed Gaussian distribution function. In the equation $f_{\text{skewed}}(x)$ is the skewed distribution, N_{skewed} is the amplitude of the function, $f_{\text{norm}}(x)$ is the formula for a normal distribution, F_{norm} is the cumulative normal distribution, α is the skewness factor, μ is the average mean value of the distribution and σ is the standard deviation of the distribution. The noise distribution was fitted using the negative emission estimates. The values of the total distribution that are at the negative emission estimates are mirrored in the positive emission estimate direction. This creates the blue dashed line in Figure 6.8. The plume distribution is then created by subtracting the noise from the total distribution.

$$f_{\text{skewed}}(x) = N_{\text{skewed}} \cdot f_{\text{norm}}(x) \cdot F_{\text{norm}}(\alpha x) \tag{6.1}$$

$$f_{\text{norm}}(x) = e^{-\frac{(x-\mu)^2}{2\sigma^2}} \tag{6.2}$$

In the case of Mariupol, the histograms resulted in Figure 6.5a and Figure 6.5b. The other two steel plants in Ukraine that were mentioned are located slightly further from the front line and have been able to keep producing. Thus the plants do emit CO after the start of the invasion. For the ArcelorMittal Kryvyi

Rih steel plant, the change in emissions can be seen in Figure 6.5c and Figure 6.5d. The skewness in the before-invasion case is much larger than in the after-invasion case. The average emission of the plume distribution also decreases from 7.2 kg/s to 4.4 kg/s. For the Metallurgical Plant Kametstal, the change in emissions can be seen in Figure 6.5e and Figure 6.5f. The skewness in the before-invasion case is much larger than in the after-invasion case. The average emission of the plume distribution also decreases from 6.3 kg/s to 4.7 kg/s.

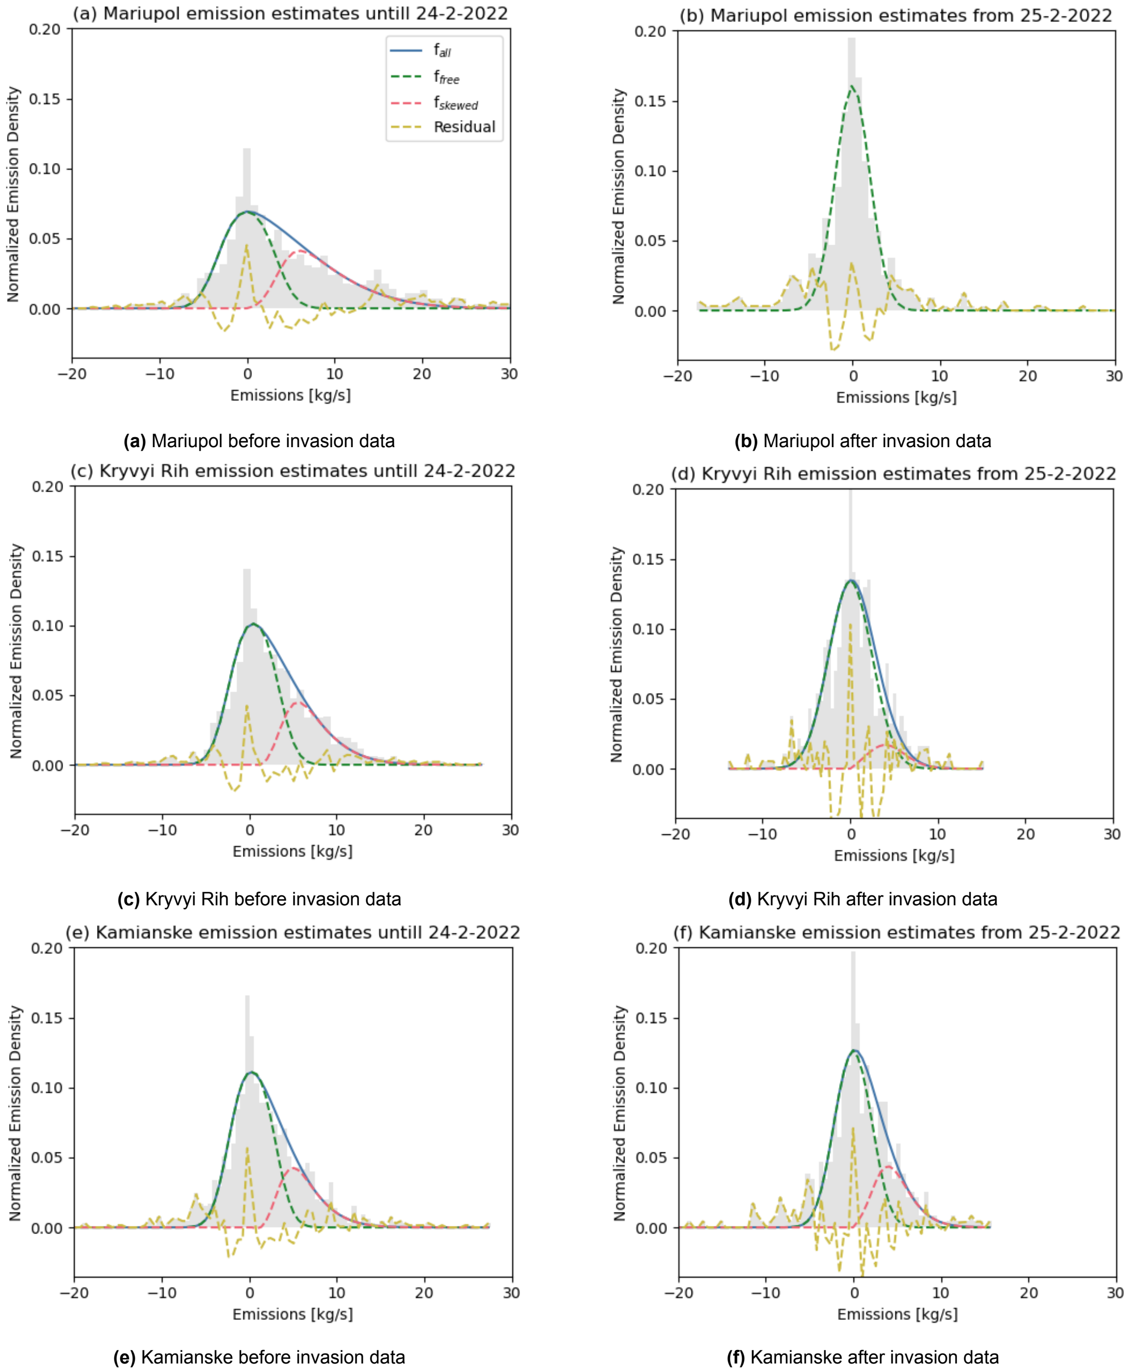

(a) Mariupol before invasion data

(b) Mariupol after invasion data

(c) Kryvyi Rih before invasion data

(d) Kryvyi Rih after invasion data

(e) Kamianske before invasion data

(f) Kamianske after invasion data

Figure 6.5: The data used to calculate the emission estimate

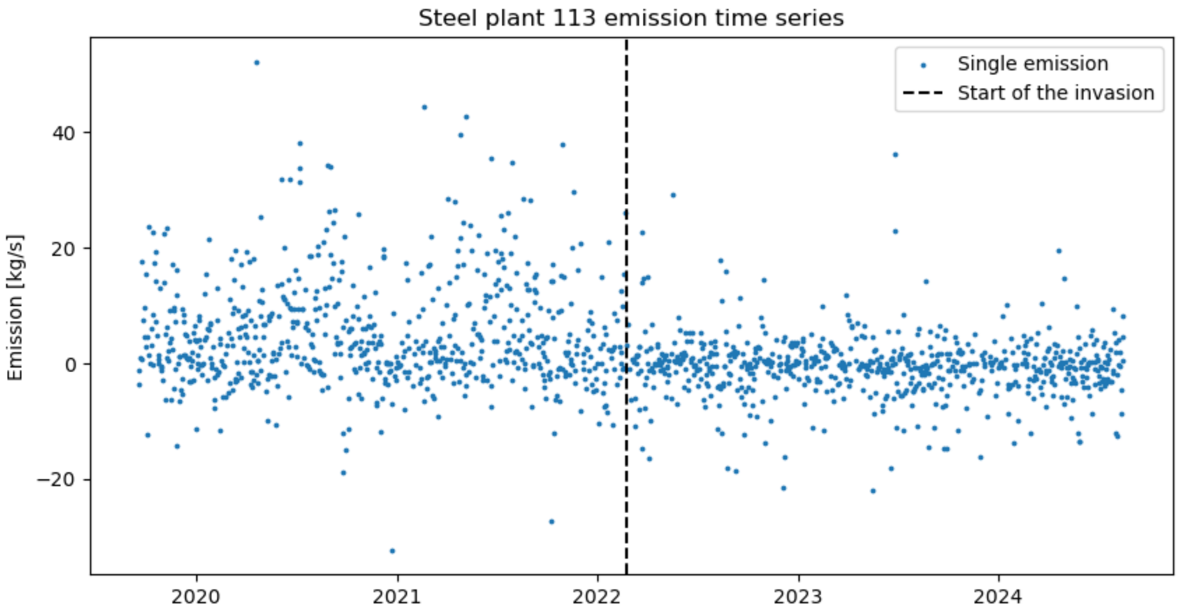

Figure 6.6: The time series of the Metinvest Ilyich Iron & Steel Works and the Azovstal Iron & Steel Works in Mariupol

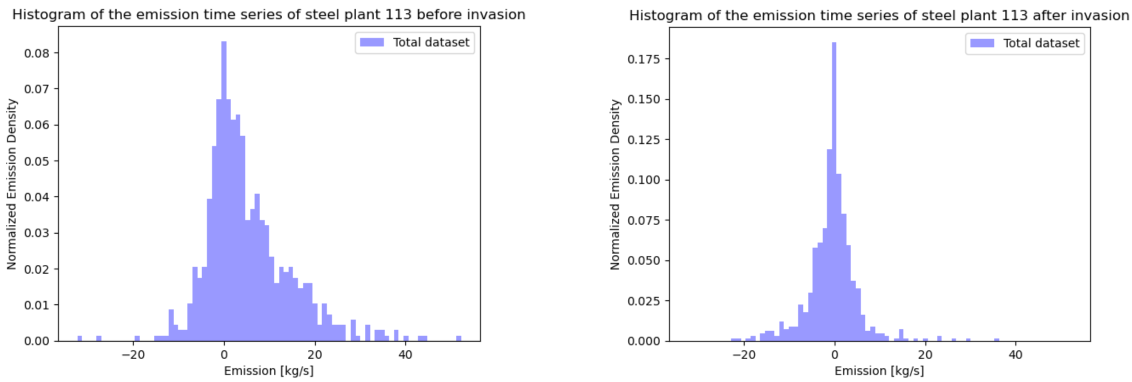

(a) Histogram of the emission time series of two steel plants in Mariupol before the invasion of Ukraine.

(b) Histogram of the emission time series of two steel plants in Mariupol after the invasion of Ukraine

Figure 6.7: Comparison between the emission estimates before the invasion and after the invasion of Ukraine showing the large decrease in emission after the invasion.

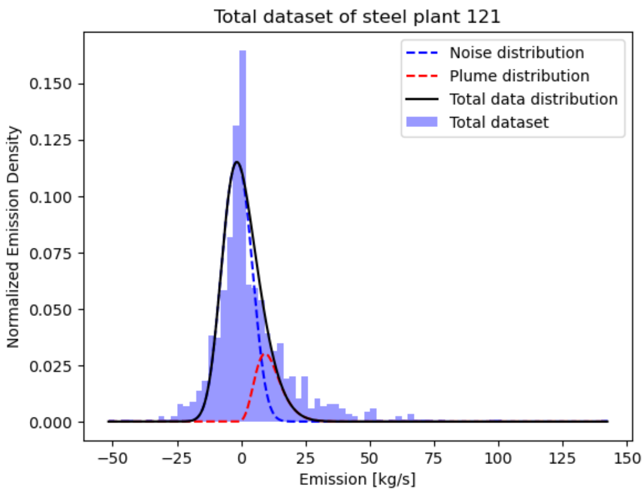

Figure 6.8: Example of the fitted distribution to the Baosteel Desheng Stainless Steel Co., Ltd. plant from Figure 6.4

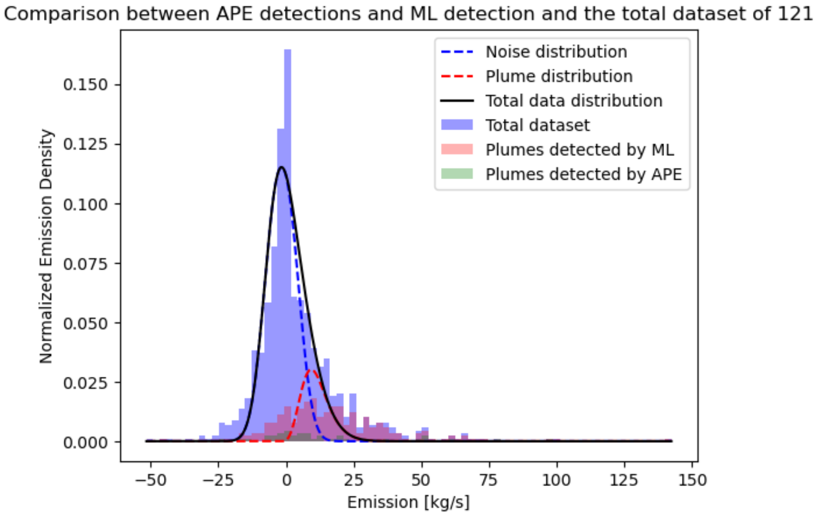

Figure 6.9: Example of how a histogram for steel plant data can be used to validate the accuracy of a machine learning tool

6.2. Validating the plume detection method

The adapted divergence method presented in this chapter is able to create a distribution of both the emission distribution and the measurement noise distribution without the use of a plume detection method. Thus the method for emission estimation as developed in this chapter is independent from plume detection and can function as an independent way to validate the plume detection.

To validate the plume detection tools, first, the required steel plant data should be downloaded. Then machine learning should be performed on the data. APE should be used for the same time period and location. The following step would be to perform the adapted divergence method on the data to get the time series of emission estimates. Using the labels from both APE and the machine learning tool, it is possible to plot the emission estimates of the detected plumes of a dataset into the histograms together with the entire range of emission estimates. An example of this is shown in Figure 6.9. Here it is clear that the machine learning tool in question performs better than APE as it fills up more of the area underneath the plume distribution. The example shows that the ML tool detects far more plumes than APE. Most of these new detections are on the positive emission estimate side. The machine learning tool shows a large improvement in detection capability compared to APE.

The new emission estimation method can also be used to produce an estimate for the average emission distribution. This can be used to compare with APE and the machine learning methods. Doing this would give an independent emission estimation from the plume detection methods. The distribution method is capable of estimating a distribution for the plume data however, it is not known how accurate these distributions are. Comparing these average emission estimate values to validated datasets is therefore important as it would show the accuracy of these distributions.

7

Comparing the machine learning methods with APE

In this chapter, the machine learning approaches for plume detection are compared with the plume detection implemented by APE. This is done by comparing their ability to detect plumes at 180 steel plant locations. If a new approach is shown to outperform the APE plume detection algorithm it should be considered to replace it. First, the compared methods are discussed in section 7.1. Second, the creation of the dataset used to perform the test is discussed in section 7.2. Third, the verification and validation tests of the approaches used in the test are discussed in section 7.3. The results are shown in section 7.4 and discussed in section 7.5. In section 7.6, the recommendations for future work are given.

7.1. The plume detection intercomparison test

The most straightforward way to compare the machine learning tools with APE would be to test the four methods and APE on a large dataset of plume and no plume data. The dataset labels would then be used to determine the accuracy of the methods. This would limit the comparison to the data that was labeled. The process of selecting and labeling data is very time-consuming and thus this would limit the dataset for the test significantly. Therefore another approach was selected. For this approach, a sizable dataset is created and the emission estimations are calculated for the individual granules in the dataset using the divergence method. These can then be used to determine the average emission estimates, the standard deviation of these emission estimates and the number of detections for each location. These values are then used to compare the performance of the different methods. The number of plume detections is further investigated to see how often the different methods agree. Since the test data is not labeled, it is not entirely certain if the detected plumes are actual plumes. Therefore, for this chapter, the number of detections should be seen as the number of granules being flagged as a plume. This data is also analyzed.

In total six methods will be compared in this chapter. The first method is the APE algorithm plume detection. As written in subsection 2.5.1, APE uses the Marker-based watershed transform method [27]. As written before this method has trouble distinguishing the difference between noise data and plume data. This method will be used as the baseline of the plume detection performance. The next four methods are the machine learning methods described in section 4.1. These methods have been implemented for this thesis project and have been trained on the industrial plumes dataset described in subsection 5.5.3. The dataset was augmented using the algorithm described in section 4.2. When implementing the machine learning methods, the training dataset was enlarged by including three rotations of each granule. This was done to avoid bias in the detection of plumes due to rotation. The machine learning methods require augmentation of the TROPOMI data as they were trained on augmented data. The first machine learning approach and the second approach in the comparison test is the SVM. This method has been implemented using the functions in SCIKIT [36]. The third approach in the comparison test is RFC. This method was implemented using Balsa [39]. For this method to work the granule size had to be decreased from 41 by 41 pixels to 15 by 15 pixels. The fourth and fifth methods implemented were the ResNet-26 and ResNet-44 networks. These CNNs were implemented by Peter Sterk, a scientist working at SRON, and Thomas Plewa, a PhD student at the Institute of Environmental Physics of Heidelberg University, respectively, respectively.

The sixth and final method that is included in the intercomparison test is the distribution method. This method is described in chapter 6. For this method, the emission estimates using the divergence method are calculated. The time series of this estimate is then put into a histogram. To determine an average plume emission estimate, an estimate for the noise and plume data distribution is made. The plume distribution is then used to determine the average emission estimate and the standard deviation of the emission estimates. The area of the plume distribution is used to estimate the number of plumes in the data. This is done by multiplying the number of granules with the fraction of the plume distribution data area with the total distribution data area.

With the methods for comparison determined, the focus shifts to the dataset for the intercomparison test. This dataset should include a large variety of sources. The locations used in the comparison test are discussed in the next section.

7.2. Plume detection comparison dataset creation

As written in section 5.2, there are about 185 steel plants that showed a divergence peak. The data from these steel plants were used to perform the intercomparison test. The reason for using these steel plants rather than any of the other steel plants of the total 1115, is the certainty of plume and no plume data in the dataset. Certain steel plants in the larger list are yet to be built or are not in operation. Another reason is time constraints, processing these steel plants is a time-consuming process. The other steel plants could be included in future tests. The TROPOMI data for these steel plants includes the time period of 18 September 2019 to 17 August 2024. The acquired TROPOMI data was then destriped according to the method described in Borsdorff et al. 2024 [21]. In particular, the data for the steel plants were sliced from the TROPOMI overpass data to be 41 by 41 pixel CO concentration data granules with the center pixel at the source. The latitude and longitude as well as the orbit reference time and the delta time data from TROPOMI were also sliced this way. The granule size was chosen to line up with the data processed by APE and the machine learning training data. The ERA5 wind data from ECMWF, as described in Hersbach et al 2020. [29], is included as well to augment the data according to the algorithm described in section 4.2 by using the wind direction. The wind speed from the ERA5 data is used to calculate the divergence of the data. To perform the histogram method, the time series of emission estimates of each steel plant had to be calculated, this can only be done with the divergences of each granule. Therefore, the divergence was calculated for each granule as described in section 6.1. Thus the TROPOMI data was processed separately two times, once to create the augmented version of the data used for the plume detection of the machine learning approaches and a second time to calculate the time series of emission for the histogram analysis.

During the processing of the data, 5 of the 185 steel plants were left out of consideration. This was done because of issues during the processing of the APE plume detection. A figure showing all remaining steel plant locations that are in the dataset can be found in Figure 7.1. As can be seen, most steel plants are located in Asia. The data included is from the period of 18 September 2019 to 17 August 2024. Due to a mistake with the indices of certain files during the processing of the data, the data between November 2017 to early September 2019 was not processed. The data of the remaining 180 steel plants which comprise a timeframe of 6 years, is about 90 GB. This was deemed enough data to continue the comparison test. In its entirety, the total dataset includes about 250.000 granules of 41 by 41 pixel CO concentration. The list of all steel plants included in the analysis can be found in Appendix A. The dataset includes steel plants with various emission strengths.

The dataset used to train the machine learning methods is the industrial plumes dataset which is described in section 5.2. As written in section 5.5, the training dataset was created using data from the same 185 steel plants used to create the comparison test data. The results of the steel plants that were included in both the training and test data will be biased. In both datasets, only 9 steel plants were used twice. The 11 steel plants that were used to create the dataset can be seen in Table 5.1. The 2 steel plants whose data was not included in the comparison test dataset have the index 257 and 667 due to the aforementioned issues with processing the APE results. The bias created by the 9 steel plants is not expected to result in a large bias in the detection accuracy for these steel plants. This is because the data used in the training dataset consist mostly of 2018 data while the dataset used for

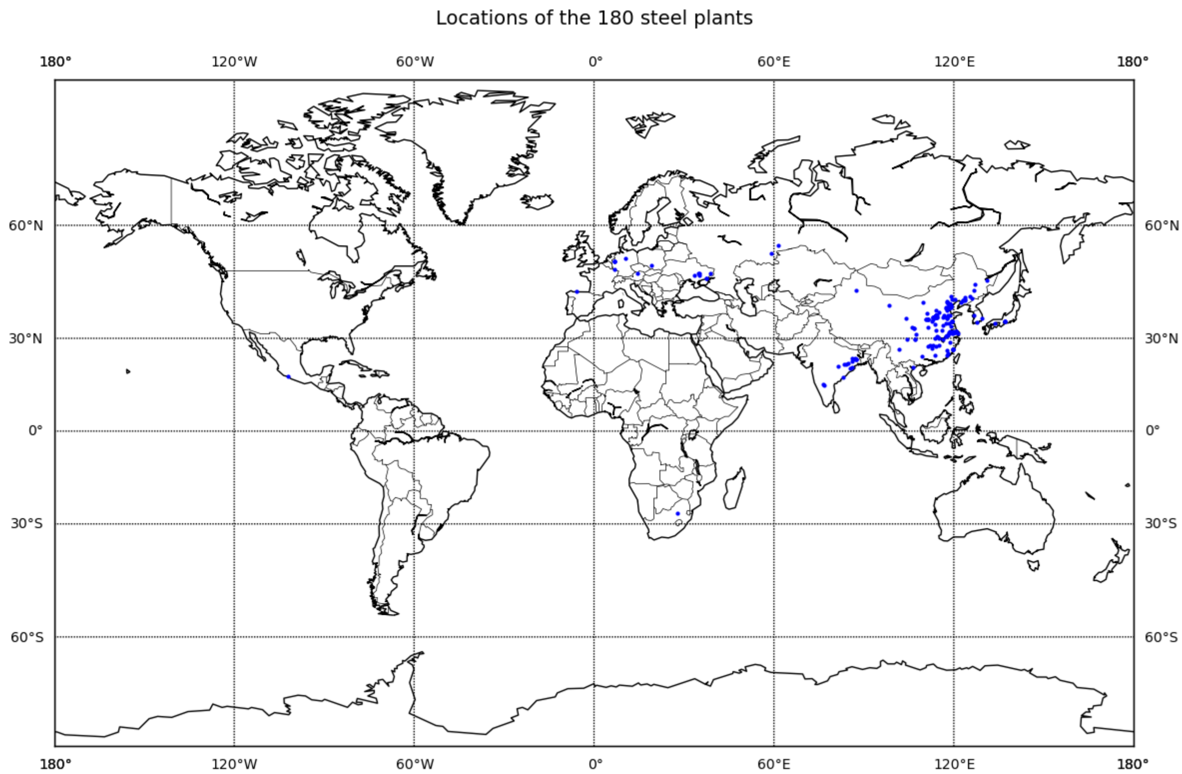

Figure 7.1: Overview of the 180 steel plants that were used in the comparison test

plume detection here starts from September 2019. During the creation of the datasets at SRON, a data migration took place which affected the amount of available data to label. This limited the available data that could be used in the training dataset. The data used per steel plant is at most 100 granules. In the comparison test dataset, there are about 1.400 granules per steel plant. Therefore, the effect of these 9 steel plants is expected to be limited.

7.3. Verification and validation tests

To ensure that the test is performed as intended, the implementation of the methods had to be verified. The first verification test was performed during the training of the machine learning methods. During this phase, the SVM and Balsa machine learning methods are first trained on 70% of the industrial plumes dataset described in subsection 5.5.3 and then tested on the remaining 30% of the data. This resulted in an accuracy of 91.6% and 92.0% for each method respectively. The remaining two machine learning methods are both residual networks. As explained in section 4.1, the ResNet-44 is created following the description in He et al. 2016 [45] while the ResNet-26 has been slightly adapted to work better for smaller images but was otherwise also completely created following the steps outlined in the paper. It can be argued that the two ResNet methods have been verified by the academic community regarding image recognition [45]. The implementation of these methods into Python was performed by Peter Sterk, a scientist working at SRON, and Thomas Plewa, a PhD student at the Institute of Environmental Physics of Heidelberg University, respectively.

The implementation of the divergence method and the augmentation method on the TROPOMI data can be verified by simply looking at them. An example can be found in Figure 7.2. In the figure, an example of a granule of TROPOMI data can be seen in Figure 7.2a with the divergence of the granule in Figure 7.2b and the augmentation of the granule in Figure 7.2c. In the example, the divergence and augmentation were performed correctly. However, verifying every granule would be very time-consuming. The divergence data has been interpolated onto a consistent grid and this makes it possible to calculate the average divergence. Therefore, the divergences can be verified by checking if the location of steel plants is at a divergence peak. This verifies the implementation of the divergence method and also the processing of the TROPOMI data as both need to be done correctly for this to work.

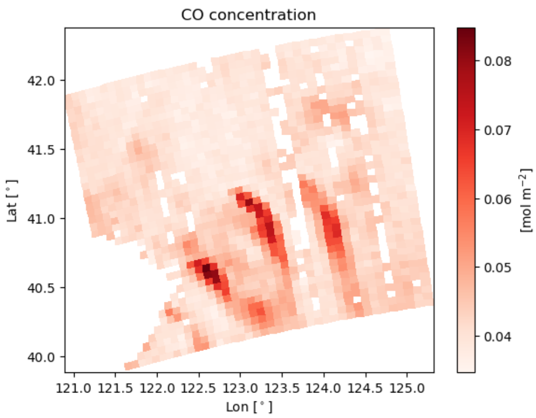

(a) Example of an overpass with a plume

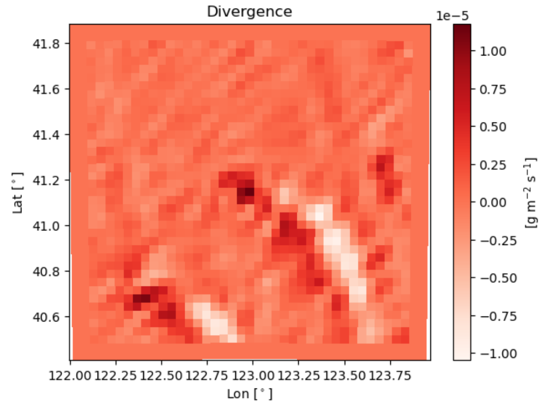

(b) Example of the divergence of the plume from Figure 7.2a

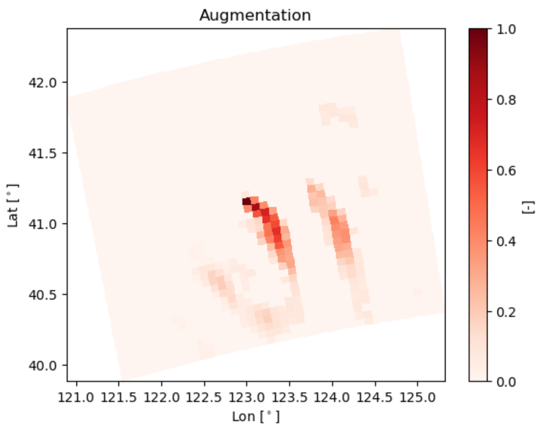

(c) Example of the augmentation of the plume from Figure 7.2a

Figure 7.2: Example of verification of a single overpass

Examples of these can be seen in Figure 7.3. In the figure, two examples of multiple steel plants located on average divergence peaks can be seen, showing that this method was implemented correctly. In the figures, the cyan-coloured dots represent the locations of steel plants. As can be seen in the figure, the locations of the steel plants line up with the peaks of the average divergences. The verification of the augmentation algorithm was already discussed in section 4.2 and shown in Figure 4.4. The figure shows that the algorithm is capable of cutting out the features of the plume from its environment for three different backgrounds.

The implementation of the histogram method needs a longer discussion. When going through the verification, it became clear that out of the 180 steel plants, 11 steel plants were not able to fit the distribution as described in the previous chapter. Out of the 11 steel plants, 8 steel plants have a skewness that is smaller than 0 and 3 steel plants have a skewness that is equal to 0. The eight steel plants that have a negative skewness have the index 122, 375, 386, 473, 779, 784, 814 and 823. The three steel plants with skewness equal to 0 have the indices 736, 829 and 931. The name and coordinates of these steel plants can be found in Appendix A. The skewness is a measure of the bias of the emission estimates towards the positive emissions in the histogram. It is expected that there is a positive skewness in a dataset that includes plumes. This is because calculating the emission estimates of plume granules should result in positive plume emission estimates. If the skewness is smaller than 0, it means that the bias is towards the negative emissions in the dataset. An example of a negative skewness in a histogram is visible in Figure 7.4. As can be seen in the figure, the plume distribution is negative. This is physically impossible and thus this distribution should be disregarded. For 8 out of the 180 steel plants, there was no bias in positive emission estimations which could be due to several reasons. The first reason could be that there are no plumes in the data. This seems unlikely as the reason for picking the 180 steel plants was the fact that they had an average divergence

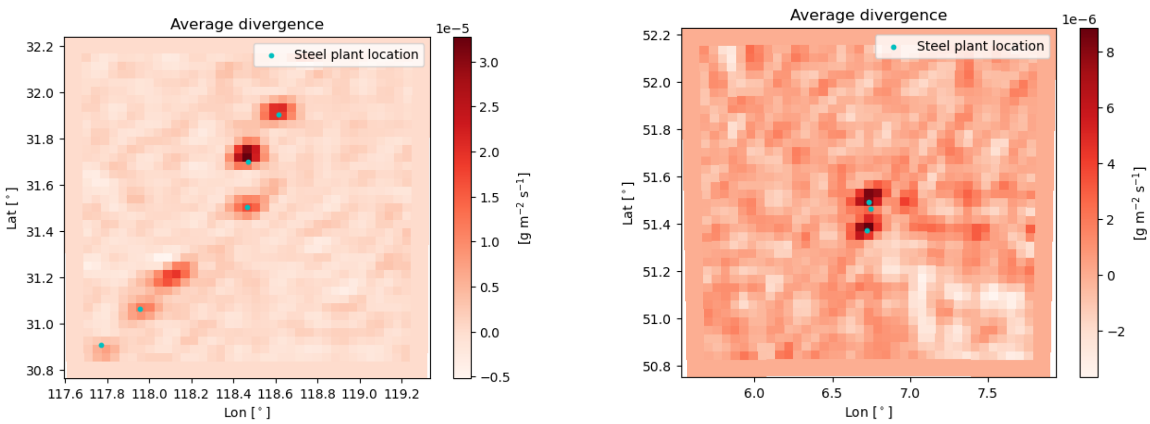

(a) Example of average divergence peaks in China

(b) Example of average divergence peaks in Germany

Figure 7.3: Examples of different steel plant locations with the average divergences

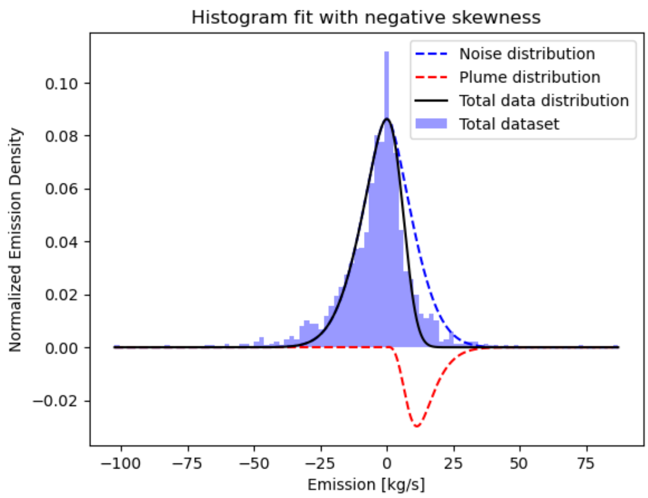

Figure 7.4: Example of a negative skewness factor in a histogram

peak such as in Figure 7.3. The second reason could be that there are plumes in the dataset but for the following reasons are not visible in the augmentation. Either the plume enhancement is not large enough compared to its background so it is not seen in the augmented granule, the plume itself is not large enough or the wind direction is wrong. When looking at the data of one of these steel plants it was hard to find clear plumes. Going through the average divergences of some of these steel plants showed that they were located right next to the sea. This meant that half of the values could not be utilized in the emission estimation process due to the different sensitivity of CO concentration measurements over the sea. This problem was described in subsection 2.5.3. It points to a mistake in processing the data near the sea.

One such steel plant located near the sea is the ArcelorMittal Asturias (Gijón) steel plant located in northern Spain. Looking at the divergence of this steel plant reveals a mistake in the implementation of the method. The average divergence can be seen in Figure 7.5. The line in the average divergence lines up with the coastline which should not be visible. Only the peak at the location of the steel plant should show up. The visibility of the coastline is caused by processing the measurement data above water incorrectly. When interpolating the measurement data onto an equal distant grid to perform the divergence a problem occurs. The function used to perform the interpolation cannot use NaN values. To get around this issue the NaN values in the granule were replaced by the median measurement data point. This should set the values above the ocean to zero. However, the values near the coast are affected when calculating the divergence of the granule. Resulting in the coastline is visible in the average divergence. For this particular steel plant, the data was corrected by taking out the pixels at the

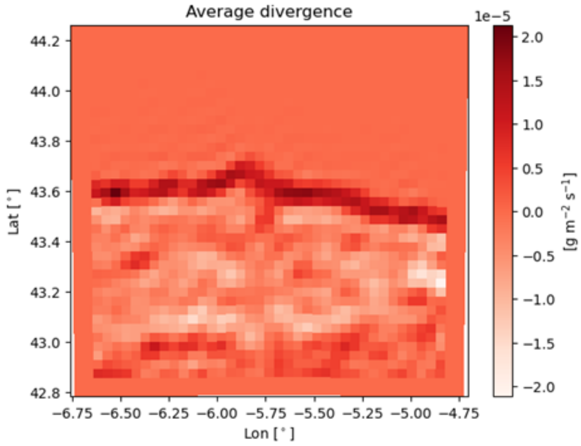

Figure 7.5: An error in processing the data from the ArcelorMittal Asturias (Gijón) steel plant

coastline. For the other steel plants located near large bodies of water, the data can be corrected in a similar way. To avoid issues when calculating the emissions near the ocean, two solutions are possible. The first solution is to use a different interpolation function that is able to process granules with NaN values. The second solution is not to use interpolation and thus to not regrid the data before performing the divergence method. This would also reduce errors in the data caused by the interpolation.

For the three steel plants with a skewness of exactly zero, an error occurred when fitting the distribution functions. In the implementation, the emission estimates are first placed into a histogram. The histogram values are then normalized. Then a distribution on the total normalized emission estimation data is fit. This is done using four parameters which are the skewness α, the amplitude of the skewed distribution N_{skewed}, the mean emission of the distribution μ and the standard deviation of the emission estimates σ. The error occurred when the program took too long to find the optimal parameters for the distributions. This does not mean that these parameters do not exist. If the skewness parameter of the data is not positive, the number of bins in the histogram is increased. This should also have been the case for this situation however due to a bug in the program, it skipped 3 steel plants entirely.

To check if the results of the implemented methods are comparable with reality two emission inventories were used as validation datasets. As described in section 2.1, obtaining comprehensive and reliable emission inventories often presents a significant challenge. Data might not be available or entirely complete. An effort to get a clear overview of the emission inventory is the Emissions Database for Global Atmospheric Research (EDGAR). This dataset was created by the EU Commission's join Research to research the effect of policy on pollution [47]. For this research, the focus is on the CO data from industrial sources. The EDGAR dataset can be subdivided into different emission types or pollution production. This makes the iron and steel sector dataset the most relevant for the comparison test. The specific dataset used is EDGARv8.1 for the CO emissions of the iron and steel industry which can be found on their website [1]. This EDGAR dataset was used as an independent source to validate the results of the average emissions. The APE, machine learning and plume distribution approach try to use a top-down method to estimate the emissions of industrial sources while EDGAR utilizes a bottom-up approach. The dataset is a product of combining statistics about energy usage, industrial activity, transportation and so much more. It can therefore be used to validate the results of the intercomparison test. The EDGAR data has a spatial resolution of 0.1° by 0.1°. This resolution differs from the TROPOMI CO resolution and therefore a procedure is needed to make the data comparable with the other methods. This was done by summing the concentration at the pixels that are in the range of 0.3° in both longitudinal and latitudinal directions with the source at the center. This roughly corresponds to the method described in section 6.1. There the pixels that were in the range of 21 km around the source were integrated. This was not possible for the EDGAR data due to the difference in resolution. The integration was performed for the 180 steel plants. The results of the validation tests are shown in the next section and the discussion is done in section 7.5.

Another dataset that was used to validate the steel plant emission is the European Pollutant Release

[1] Retrieved on 28-1-2025, `https://edgar.jrc.ec.europa.eu/dataset_ap81?utm_source=chatgpt.com`

and Transfer Register (E-PRTR) dataset. This dataset is a collection of emissions of the different industries in the European Union and nearby countries. These industries are forced by law to document the amount of pollutants they produce and release into the air, sea or land. The dataset also includes the emission of CO by steel plants. In this dataset, the emissions per source and location are listed so it makes it easy to check the emissions of individual sources [48]. One downside of using this dataset is that it only accounts for steel plants in the European Union member states and other countries such as the United Kingdom, Serbia, Liechtenstein, Iceland, Norway and Switzerland. In the 180 steel plant list, only 8 steel plants are located in the EU or aforementioned countries. This makes it a very limited validation dataset. Similarly to the validation with the EDGAR dataset, the emission average is calculated by summing the emissions from the sources that are 0.3° around the steel plant locations. The previously mentioned EDGAR data incorporates data from the E-PRTR dataset for specific industries. The iron and steel plant data from E-PRTR data is not incorporated [47].

To ensure that the E-PRTR average emission values are implemented correctly, they are compared with the values from the posterior model as calculated by Leguijt et al. 2025 [48]. To create the posterior model, the researchers used TROPOMI observational data, a chemical transport model and prior emission libraries (E-PRTR and TNO emission inventory). These were used as input for an inverse framework to calculate the posterior emissions. The posterior emissions are an attempt to combine several information sources as well as chemical reactions to increase the accuracy of the emission estimates. Since this method is different from the ones presented in this thesis, it serves as a good comparison to see if the E-PRTR values are implemented well and to see how the average emissions line up with both of these methods. It has to be mentioned that the posterior model used data from 2019 to calculate the emission estimates.

7.4. Intercomparison test results

The results of the test can be split into three parameters, the average emission estimates, the standard deviation of emission estimates and the number of observations. These parameters were calculated for the 180 steel plants. The complete results of the intercomparison test can be found in Appendix B. In this appendix, all the average emission estimates, standard deviations and number of detections can be found per method per steel plant. The results are explored in the following sections.

Plotting the standard deviation of the emission estimates as a function of the average emission results in Figure 7.6. For the plume distribution, there is a seemingly linear relationship. To show this relationship a linear regression was performed on the data of the distribution method. For the machine learning methods and APE, there is no such relationship. A linear regression was also performed for the other methods to compare the relationship between the plotted parameters. The number of detections as a function of the average emission estimate is plotted in Figure 7.7. Once again a linear regression was performed to see the relationship between the plotted parameters. This figure shows that every method has an increasing average emission estimate if the number of detections increases except for APE. For APE the number of detections does not seem to affect the average emission estimate.

For the machine learning methods and APE, plumes are detected which are then used to estimate the emissions. If the methods are working correctly it would be expected that they detect the same plumes. Therefore, it is interesting to see the consistency of plume detections between the different methods. In Table 7.1 the number of shared detections can be found. If method 1 and method 2 are the same then the total detections for that method are written. In this table, it is shown that APE detects significantly fewer plumes than the other methods over the same time period. As a result, the method also has fewer plumes in common with the other methods. To put this into perspective, in Table 7.2 the percentages of commonality are shown. These percentages are calculated by dividing the number of detections by the total detections of the method written in the row. This shows that the machine learning methods seem to agree with each other on the majority of the plumes. The detections of APE seem to agree less with the other methods. Only about 40% of the plumes detected by APE are also found by the machine learning methods while only about 20% of the plumes found by the machine learning methods are found by APE.

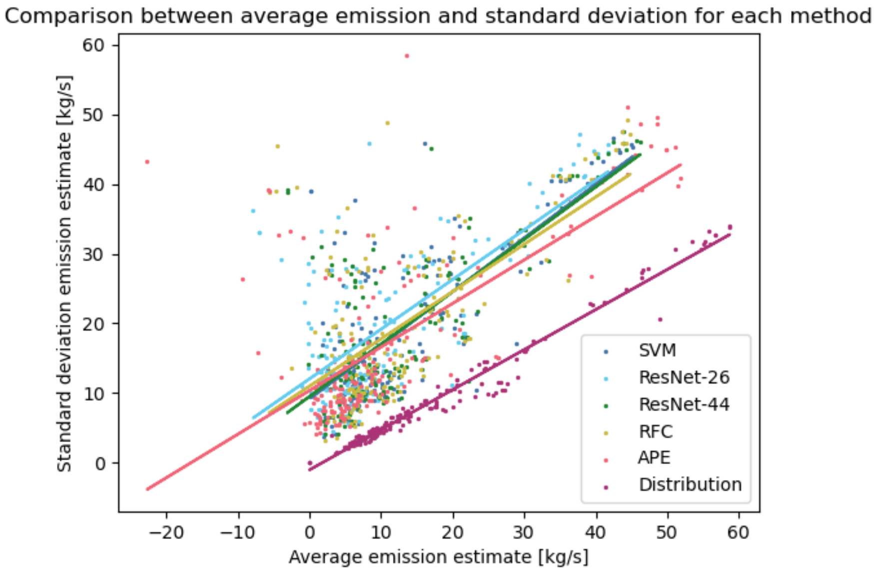

Figure 7.6: Overview of the average emission estimates and the standard deviations of the different methods

Table 7.1: Comparison of the different methods for each other for every steel plant

Method 1/Method 2	SVM	ResNet-26	ResNet-44	RFC	APE
SVM	56912	42415	42708	47146	11475
ResNet-26	42415	68486	36337	38839	12190
ResNet-44	42708	36337	50826	39662	10268
RFC	47146	38839	39662	54457	11148
APE	11475	12190	10268	11148	29013

Table 7.2: Comparison of the different methods for each other for every steel plant in percentage

Method 1/Method 2	SVM	ResNet-26	ResNet-44	RFC	APE
SVM	100	74.5	75.0	82.8	20.2
ResNet-26	61.9	100	53.1	56.7	17.8
ResNet-44	84.0	71.5	100	78.0	20.2
RFC	86.6	71.3	72.8	100	20.5
APE	39.6	42.0	35.4	38.4	100

Table 7.3: Comparison of the granules flagged by the different plume detection methods

Method	Total	Granules found by each method also found by other methods				
		All methods	3 other methods	2 other methods	1 other methods	No other method
SVM	56.912	6.603 (11.6%)	27.363 (48.1%)	14.382 (25.3%)	6.479 (11.4%)	2.085 (3.7%)
ResNet-26	68.486	6.603 (9.6%)	25.983 (37.9%)	8.784 (12.9%)	7.852 (11.5%)	19.264 (28.1%)
ResNet-44	50.826	6.603 (13.0%)	26.283 (51.7%)	9.502 (18.7%)	4.710 (11.5%)	3.728 (7.3%)
RFC	54.457	6.603 (12.1%)	27.116 (49.8%)	12.380 (22.7%)	4.275 (7.9%)	4.083 (7.5%)
APE	29.013	6.603 (22.8 %)	3.403 (11.7%)	1.953 (6.7%)	4.554 (15.7%)	12.500 (43.1%)

When looking further into the plume detection the following statistics were found. The total number of granules is 249.897. All five methods (Machine learning and APE) for plume detection agree that 6.603 granules are plumes and that 144.495 granules are not plumes. This shows that the methods are very confident that there are significantly more noise granules than plume granules. An overview of the granules can be found in Figure 7.8. Here it is shown that 57.8% of the data has not been

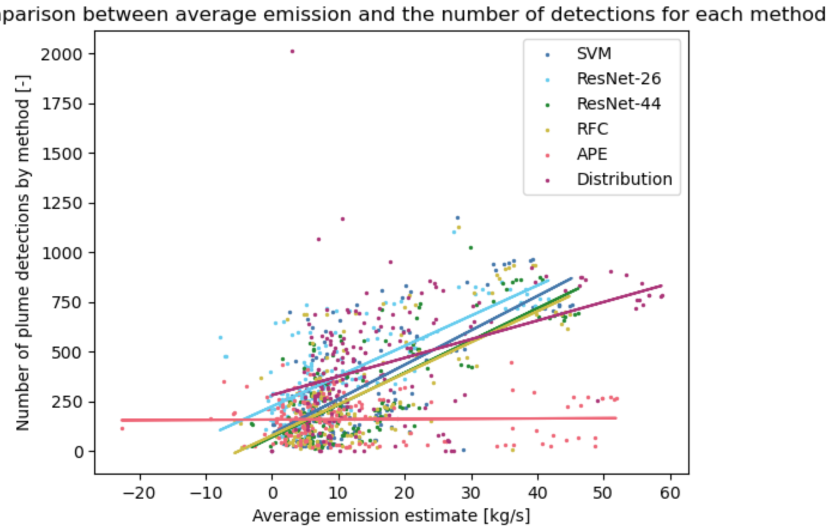

Figure 7.7: Overview of the number of detections as a function of the average emission estimate of the different methods

flagged by any method as containing a plume. A further 16.7% of granules have only been detected by 1 method. Most of these were detected by APE and ResNet-26. The remaining data comprises about 25%. In numbers, four out of five methods agree with each other that 27.537 granules are plumes. Out of these granules, the four machine learning methods agree with each other on 24.134 granules. Only three methods agree with each other on 15.667 granules. Only two methods agree with each other on 13.935 granules. The number of granules that are flagged as plumes by just one method is 41.660. From these granules the SVM method flagged 2.085 granules, the ResNet-26 flagged 19.264 granules, the ResNet-44 method flagged 3.728 granules, the RFC method flagged 4.083 granules and the APE method flagged 12.500 granules. From these statistics the following conclusions can be made, there is a majority of no-plume data in the complete dataset. The APE method and ResNet-26 method flag significantly more unique granules as plumes.

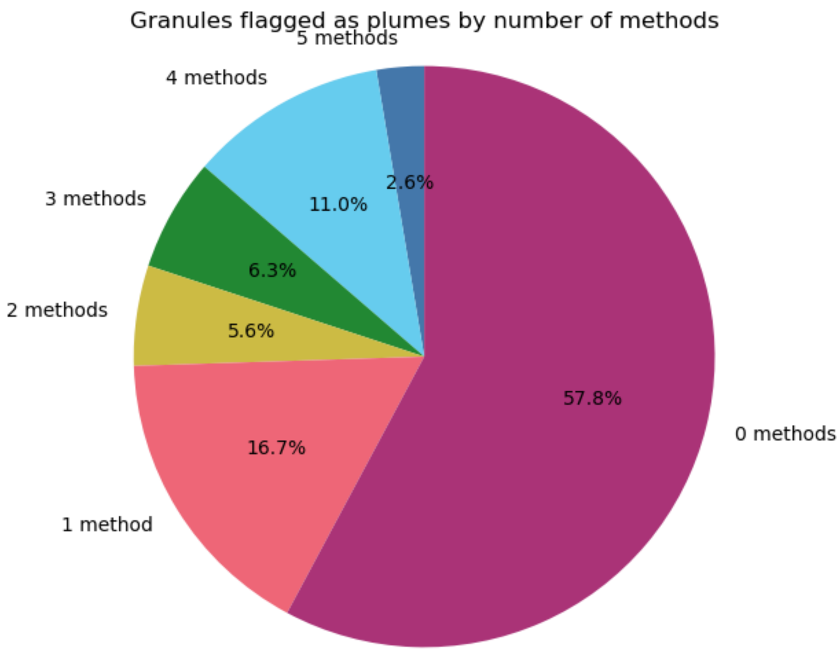

Figure 7.8: Pie chart showing the times individual granules have been flagged by the different plume detection methods.

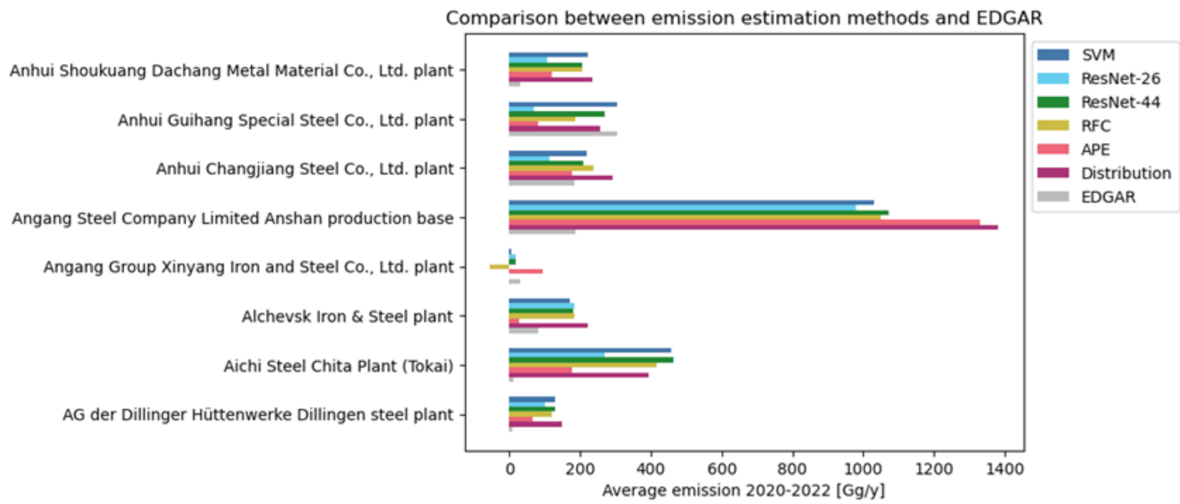

Figure 7.9: Comparison between the average emission estimates of the machine learning approaches, APE, the plume distribution method and the EDGAR data for the first 8 steel plants in the list.

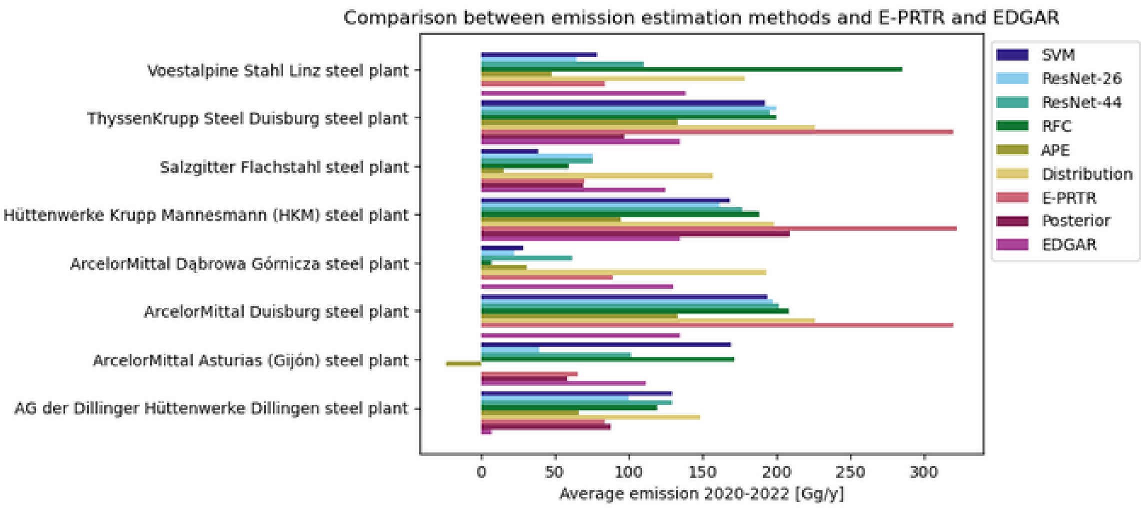

Figure 7.10: Comparison between the average emission estimates of the machine learning approaches, APE, the plume distribution method, EDGAR data, E-PRTR data and results of Leguijt et al. 2025 [48] denoted as Posterior for the 8 steel plants in Europe.

The results of the validation tests are shown in Figure 7.9 and Figure 7.10. For this comparison, the data from the period of 2020 to 2022 was used. This is because the 2019 data was not complete for the machine learning methods, APE and the distribution method and the 2023 data was not complete for E-PRTR and not found for EDGAR. A small overview of the results of the average emission estimates of four steel plants can be found in Figure 7.9. In this figure, the average emission estimate of the four machine learning approaches, APE, the plume distribution and EDGAR are shown side to side. As can be seen, the average emission estimates fluctuate between the different methods in the figure. This is also the case for the other average emission estimates. A comparison of the approaches with the E-PRTR data can be found in Figure 7.10. The figure shows the results for every steel plant as well as the results of Leguijt et al. 2025 [48]. Out of the 8 steel plants that could be validated with E-PRTR only 5 are also included in the paper. These 5 steel plants seem to line up well with the findings in this thesis.

To validate the average emission data between the methods and the validation sets the bias and the standard deviation (Std) were calculated. The results for these can be found in Table 7.4. In the table, the bias is defined as the average difference between the validation and the emission estimation method. A negative bias means that the emission estimation method estimates a higher average emission compared to the validation dataset. The standard deviation is used to show the spread of the average emission estimates. In the table, it is clearly shown that the average emission estimates from the implemented methods line up closer with the E-PRTR and are significantly better than the emission estimates from EDGAR. This is proven by having a smaller bias and standard deviation for almost every method. The E-PRTR data shows a positive bias for each method except for the distribution method. This indicates that most methods underestimate the average emissions of steel plants. The EDGAR validation dataset shows the opposite and shows that the bias is mostly negative which would mean that the machine learning approaches overestimate the average emission of steel plants. The standard deviation for the validation with EDGAR shows a significant difference between the validation set and the average emissions of the other methods. The significant difference between the validation datasets can be explained by noting that the E-PRTR dataset is used to validate the results of 8 steel plants while the EDGAR dataset is used to validate the results of 180 steel plants. If the average emissions of the EDGAR dataset are compared for just the 8 European steel plants the bias and standard deviation are much smaller. The results of this validation test can be seen in Table 7.5. The table shows that the bias is much smaller for the machine-learning approaches compared to APE while the standard deviation is worse for most approaches except for the distribution method and ResNet-44.

Table 7.4: Comparison of the bias and the standard deviation between the average emissions calculated by the different methods and the validation datasets with values in kg/s.

Estimation method	EDGAR Bias	EDGAR Std	E-PRTR Bias	E-PRTR Std
SVM	-1.96	21.61	1.41	2.69
ResNet-26	1.50	21.47	1.96	1.96
ResNet-44	-2.03	21.58	1.20	2.36
RFC	-1.90	21.79	0.46	3.60
APE	-1.25	22.94	3.41	2.40
Distribution	-6.36	22.63	-0.17	3.04

Table 7.5: Comparison of the bias and standard deviation between the average emissions calculated by the different methods and the validation sets for the European steel plants with values in kg/s.

Estimation method	EDGAR (only Europe) Bias	EDGAR (only Europe) Std	E-PRTR Bias	E-PRTR Std
SVM	-0.33	2.42	1.41	2.69
ResNet-26	0.22	2.30	1.96	1.96
ResNet-44	-0.54	1.96	1.20	2.36
RFC	-1.28	2.67	0.46	3.60
APE	1.67	1.99	3.41	2.40
Distribution	-2.37	1.09	-0.17	3.04

7.5. Discussion

The previous section displayed several results showing the issues present in the APE approach to plume detection. The clearest example can be seen in Figure 7.7. Here the number of detections by APE does not increase with the average emission estimation, which does not make sense. When a steel plant has a higher average emission estimate, it would make sense that these steel plants are polluting more. Therefore, their plumes should be more visible for these steel plants. Thus it makes sense that the number of detections increases with the increase in average emission estimate. This shows a flaw in the APE approach compared to the other methods. The plume detections are further explored in Table 7.2. APE is only able to find about 20% of the plumes for the same time period as the other methods and only about 40% of the plumes are found in the datasets of the other approaches which is significantly lower than the other methods. It shows that APE is much less capable of detecting plumes. When comparing APE with the validation datasets in Table 7.5, the bias of APE with the validations set is the second worst with EDGAR and the worst with E-PRTR compared to the other methods. For the standard deviation, APE performs the third best with EDGAR and the second best with the E-PRTR dataset. From the validation results it can be concluded that the other methods perform better than APE.

A deeper look into the EDGAR validation dataset is necessary to explain the large standard deviation in Table 7.4. When looking at the results of the EDGAR validation dataset some steel plants show a significantly high average emission. According to the EDGAR results 9 steel plants have an emission higher than 80 kg/s with 2 higher than 100 kg/s. For context, the other methods do not report any steel plant with a higher emission than 60 kg/s except for the distribution method which has its maximum emission at about 62 kg/s. It seems highly unlikely that steel plants can emit about 40 kg/s more CO without it being found by the other emission estimation methods. Another inconsistency in the results in Table 7.4 is that the negative bias would point towards an overestimation of the emission estimation methods. When comparing the average emissions with the European steel plants in Table 7.5, the standard deviation has the same magnitude as the standard deviation from E-PRTR. This shows that the large standard deviation is due to average emission estimations from steel plants outside Europe. This large change warrants further investigation. Various reasons could explain the large discrepancy between standard deviations. The first reason could be that there is an issue with the emission estimation for the different methods. Perhaps the emissions are incorrectly calculated for steel plants in other regions due to different atmospheric conditions. The second reason is similar to the first but considers the problem to be in the plume detection instead. Perhaps the difference in atmospheric conditions negatively affects the workings of the augmentation algorithm. The third reason considers the problem to be with the EDGAR dataset. Perhaps the information used for other regions is less accurate. Further research is needed to explain the discrepancy. With the large variance in the results in mind, the results of bias and standard deviations of EDGAR will be seen as less significant than the results of the validation with E-PRTR.

When looking at the average emission of the steel plants that are included in the E-PRTR validation in Figure 7.10, it can be seen that three steel plants have a significantly higher emission than the other methods. These are the ArcelorMittal Duisburg steel plant, the Hüttenwerke Krupp Mannesmann (HKM) steel plant and the ThyssenKrupp Steel Duisburg steel plant. For the average emissions of these steel plants, the pollution of the surrounding area was included as well to mirror the 21 km radius that the emission estimation method is using. The three steel plants are located within a 21 km radius of each other. Therefore they affect each other average emission estimations. This also increases the emissions of the E-PRTR dataset. When comparing the emissions of the single steel plant with the data it seems to line up better, however, this would be an unfair comparison as the average emission does include the measurements in the larger 21 km radius.

Certain average emission estimates of the machine-learning approaches seem extremely high. One way to make sense of this is by plotting the locations of these steel plants. This can be seen in Figure 7.11. In Figure 7.11a, the locations of the steel plants are shown on top of the average divergence of the area to show that the steel plant locations are in the right place. What immediately becomes clear is that the high-emitting steel plants, in the figure denoted as $\mu_E > 50$ kg/s, are located extremely close to each other. In Figure 7.11b, the inner radius of the emission calculation method is shown. This radius shows that all the steel plants inside it are used to calculate the emission for the steel plant in the middle of the area. Thus the steel plants that are located closely cannot have an emission estimate that is separate from the others. This explains why the emissions are so high.

When comparing the linear regressions in Figure 7.6 it becomes clear that the four machine learning methods and APE behave very similarly. Compared to the distribution method the other five methods have a much higher standard deviation of emission estimates. This could point to these methods having significantly more false positives, implying that more noise is flagged as if it were plumes. This would end up making the standard deviation larger. The false positives are further explored in Table 7.2. When looking at the percentage of plumes detected by the different methods it becomes clear that the machine learning methods agree on the majority of plumes. One way to get a sense of the number of false positives is by looking at the "confidence" of the different methods. The confidence here is defined as the number of methods that agree that certain granules are plumes. What can be seen is that APE and ResNet-26 flag a significant number of unique granules as containing a plume. These granules would thus have a low confidence as they are only detected by one method. However, when looking at Table 7.4 and Table 7.5, it becomes clear that the ResNet-26 method is the second most accurate

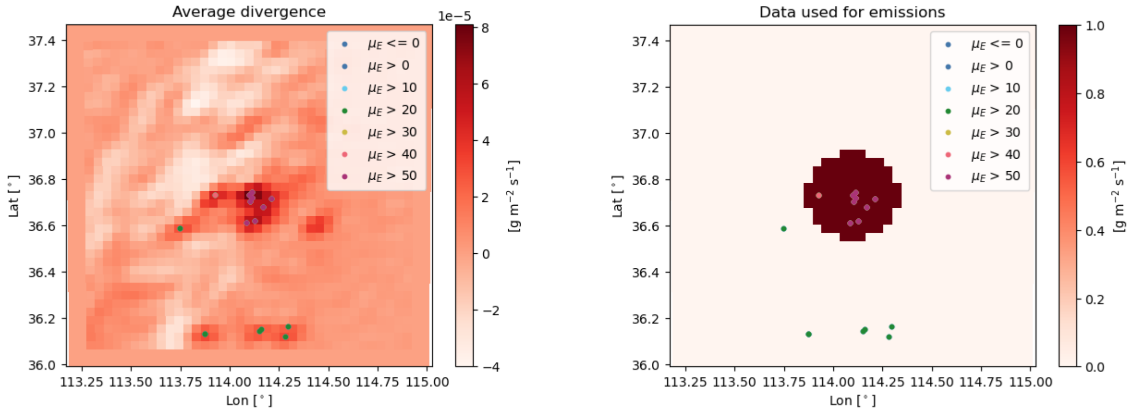

(a) The locations of steel plants and their emission on the average divergence of the area

(b) The radius of the emission estimation method with the location of the steel plants

Figure 7.11: Examples of high emission steel plants. The average emissions are in kg/s

method for both EDGAR and E-PRTR. Therefore it seems that the about 20.000 plumes that were found only by ResNet-26 are not all false positives and probably include a large amount of correctly identified plumes. This would seem to mean that the other machine-learning methods have missed out on a significant amount of plumes. This warrants a further investigation into these unique granules. If these granules are found to contain plumes they should be added to the training dataset. Another look should be taken at the training dataset as it could be that mistakes are present in the labels. The large standard deviation in Figure 7.6 could also be the result of inaccuracies in the emission estimation method which could bias the average emission. When performing the divergence calculation, the CO concentration data is first interpolated onto a grid with a better spatial resolution than the TROPOMI CO data creating interpolation errors.

It is difficult to decide on the most accurate plume detection method. Previously it was discussed that APE should be replaced by one of the machine learning methods. It cannot be replaced by the distribution method as this method does not detect plumes. Out of the machine learning methods, the ResNet approaches outperform APE in terms of bias and standard deviation in both validation tests as can be seen in Table 7.5. This makes both approaches good potential replacements for APE. The plume detection problem of this thesis can be explained as an image recognition exercise and the residual networks were specifically designed to outperform contemporary image recognition approaches as described in He et al. 2016 [45]. Thus it makes sense that they perform well. The SVM and RFC methods have a lower bias compared to APE but a larger standard deviation and thus do not outperform APE. The distribution method has the highest bias and the lowest standard deviation in the validation test with EDGAR. For the validation comparison test with E-PRTR, the bias is the lowest with the second highest standard deviation. These results are confusing. This could be explained by the fact that the average emission values between the validation sets do not line up perfectly.

7.6. Recommendations

During the implementation of the methods and the creation of the datasets, some choices were made that affected the outcome of the test. For example, there are some issues in the normalization of the data in the histogram. These affect the results of the distribution method. This should be fixed in the future Another problem in the implementation occurs when the CO concentration data is interpolated onto a new grid. This causes interpolation errors and also affects steel plants that are located close to large bodies of water as discussed before. If possible the interpolation should be avoided in the future to calculate more accurate emission estimations.

Other recommendations for the implementation include, further refining the training dataset of the machine learning methods to not only include data from 2018 and expand the number of steel plants included. The number of steel plants investigated in Europe should be increased as they can be verified with E-PRTR and there are definitely many more steel plants than the 8 shown in Figure 7.10.

In terms of methods to implement instead of APE, it was already mentioned that the ResNet-44 seemed like a good pick with the SVM approach being a good second choice. Another interesting idea is to investigate the combination of different plume detection methods and see how that would affect emission estimation. The different methods implemented should cover each other's weaknesses and combine their strengths.

It is important to investigate the results of the intercomparison tests for different regions. It could be that the results for Europe are much better than the results for India for example. If the explanation for this could be found it would result in a better understanding of the effects of atmospheric conditions on the plume detection and emission estimation methods. With the E-PRTR dataset, it makes sense to start any investigation of such effects in Europe.

The distribution method should be further explored. As mentioned before the normalization problem should be fixed. Another interesting investigation should be in the division of the different distributions in the noise distribution and the plume distribution. Currently, the implementation is rather arbitrary. This should be further investigated to refine the process. Perhaps other distribution functions fit the data better. It would be interesting to investigate how well the labels line up with the emission estimation. It should be confirmed that all the granules labeled as plumes have high emission estimates.

8. Conclusion and future outlook

In this chapter, the project is concluded by revisiting and answering the research questions. The answers and discussions of these questions are used to conclude the research. These can be found in section 8.1. The research questions that are not answerable are discussed in section 8.2. Finally, an outlook for future development on APE is given in section 8.3

8.1. Answering the research questions

- MQ1 - Can a machine learning approach improve the detection of pollution due to combustion at a global scale?

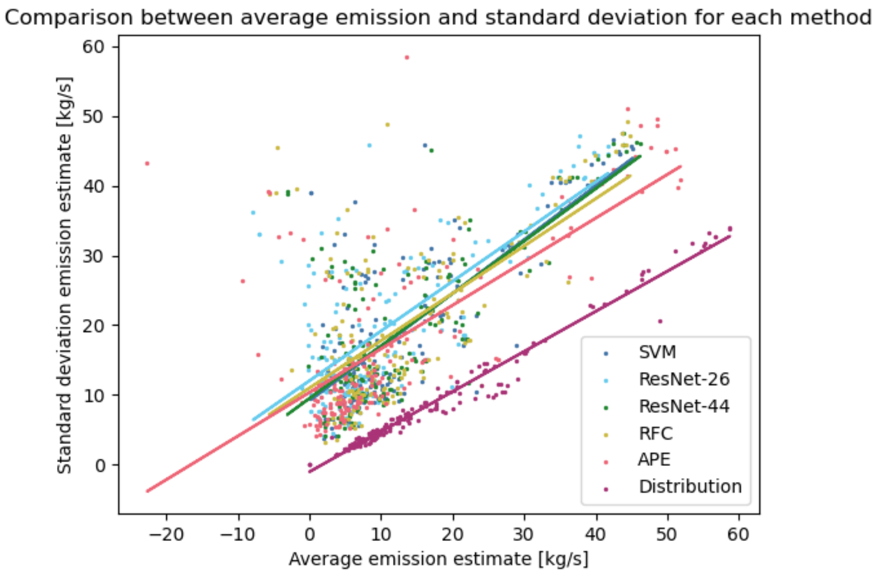

Figure 8.1: Overview The number of detections as a function of the average emission estimate of the different methods

Yes, it has been shown that the machine learning approaches detect more plumes than the currently implemented APE method. However, it has also been shown that they suffer from a not insignificant amount of false positive detections which causes the large standard deviation in Figure 8.1. This shows that the machine learning methods need further refinement. This should be done by improving the training dataset. To check if the machine learning methods outperform APE on a global scale, the intercomparison results were validated using the EDGAR dataset. This showed a large standard deviation for every method which raised suspicions about the results. To confirm that the average emission values for EDGAR were calculated correctly, the average emissions of the 8 steel plants that were also in the E-PRTR dataset were also used to validate the results. As can be seen in Table 8.1 the

Table 8.1: Comparison of the bias and standard deviation between the average emissions calculated by the different methods and the validation sets with values in kg/s.

Estimation method	EDGAR (Global) Bias	Std	EDGAR (only Europe) Bias	Std	E-PRTR Bias	Std
SVM	-1.96	21.61	-0.33	2.42	1.41	2.69
ResNet-26	1.50	21.47	0.22	2.30	1.96	1.96
ResNet-44	-2.03	21.58	-0.54	1.96	1.20	2.36
RFC	-1.90	21.79	-1.28	2.67	0.46	3.60
APE	-1.25	22.94	1.67	1.99	3.41	2.40
Distribution	-6.36	22.63	-2.37	1.09	-0.17	3.04

bias and standard deviations are comparable to E-PRTR's validation test results. This implies that the average emissions of the methods are much better in Europe compared to other areas. This should be further investigated.

- QWP2.1 - Between the machine learning approaches for plume detection as described in subsection 2.6.1, which one detects the most plumes accurately?

The machine learning methods implemented for the project are the SVM, RFC, ResNet-26 and ResNet-44 methods. As can be seen in Table 8.1, the ResNet-26 and ResNet-44 approaches perform the best. These methods have a lower bias and standard deviation for the European steel plants in both validation tests. This shows that out of the methods written in subsection 2.6.1, the deep learning methods detect the plumes most accurately. This makes sense as the deep learning methods implemented here are residual networks that were specifically designed to outperform other image recognition methods as written in He et al. 2016 [45].

- QWP2.2 - Which plume features are the most important for detection?

When implementing machine learning methods, it is important to reduce the noise in the input data. This is to ensure that only useful information is included in the input data. For this purpose, an augmentation algorithm was designed. This algorithm enhances potential plumes in the data by augmenting the data in three ways. The first feature is the wind direction. The plume should be in the direction of the wind. The data out of the wind direction is less relevant and thus removed by the algorithm. The second feature is the distance to the pollution source. The data closer to the source is more relevant and thus the data further away is removed by the algorithm. The third feature is the enhancement of CO in the data compared to the background. If the CO concentration data is twice as big as the background concentration it is considered to be relevant while the rest of the data is removed by the algorithm. This augmentation algorithm was used to create the training data for the machine learning methods which have been shown to outperform the currently implemented APE method. Thus the most important features for detection are the data in the direction of the wind, the data near the source and the enhancement of the measurement data. The augmented data has been shown to outperform the plume detection of APE in Table 8.1.

- QWP3.1 What is the increase in the accuracy of the new plume detection and emission estimation?

To analyze the accuracy of the new methods the bias and standard deviation were calculated with the validation datasets and can be found in Table 8.1. The ResNet-26 and ResNet-44 methods outperform APE in the validation tests. The ResNet-26 detected 136.1% more plumes than APE while ResNet-44 detected 75.2% more plumes. Since not every detection is guaranteed to be a plume these values cannot directly be taken as the accuracy but since the average emissions line up better with the validation values, it can be safe to say that the accuracy is much higher. To find out the exact accuracy increase the entire intercomparison dataset must be labeled. Due to the time constraints this was done.

- QWP3.2 What phenomena are visible in the data?

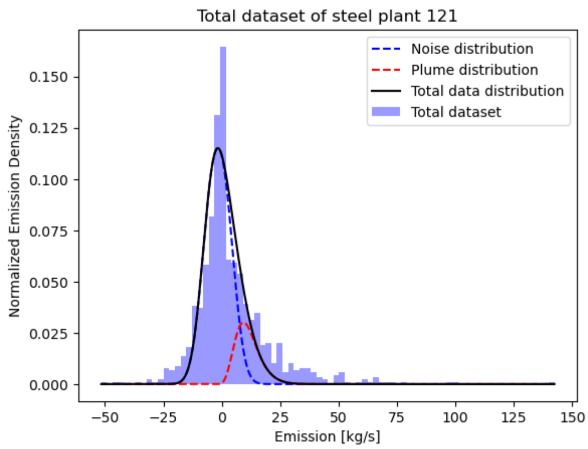

Figure 8.2: Example of the fitted distribution to the Baosteel Desheng Stainless Steel Co., Ltd. plant from Figure 6.4

There are quite some phenomena visible in the data. One important phenomenon found in the data during the project was that if you calculate the emissions and then put them into a histogram it is possible to estimate the characteristics of the plume data without using plume detection. The skewness of the data can be seen quite clearly in most histograms. An example of a histogram can be seen in Figure 8.2. Here the total dataset is fitted by a skewed normal distribution. The noise data is fitted as a normal distribution around the peak. The plume data is determined by subtracting the noise distribution from the skewed normal distribution. This new method to determine the average emission lines up relatively well with the E-PRTR dataset. As seen in Table 8.1, the distribution method has a lower RMSE than APE. Further refinement of the method is needed to make it more accurate.

8.2. Unanswered research questions

- MQ2 - Does combining different trace gases improve the quantification of emissions?
- QWP1.1 What is the effect of adding NO_2 and NO_x data for emission quantification?

When reading the research questions and justification for them in chapter 3 one would expect a large part of the thesis to focus on including NO_x into the emission estimation process of APE. However, in the chapters that follow, this research focus is not mentioned. This will naturally result in questions from the reader. It might be assumed that no developments took place. This is not the case.

In Figure 3.1, the first two tasks were getting access to the chemical model from Kuhlman et al., 2021 [44] and making it possible for APE to download NO_2 column TROPOMI data. Both of these tasks were completed. An example of APE downloading NO_2data can be seen in Figure 8.3. This figure gives an example of a CO plume and an NO_2 plume in Figure 8.3a and Figure 8.3b respectively. Overlaying the NO_2 plume over the CO plume results in Figure 8.3c. The figure shows that the two trace gases both appear at steel plants.

The Kuhlmann model focuses on four large emitters. To simulate the relationship between NO_2 and NO the model relies on background concentration values that were simulated using MicroHH. Thus to utilize this model for other locations, these values had to be calculated individually using MicroHH. This was deemed too time-consuming and unfeasible for this project. Other avenues to tackle the problem were investigated but due to time constraints, the development and research for this topic must be left for a different project.

8.3. Future outlook

With the current project coming to an end the focus must be put on the next steps of research. The project was able to show that machine-learning methods are better for plume detections but further refinement and implementation are needed. This should be done by expanding the number of steel plants that are used in the training dataset and reducing the number of false positives. A critical look must also be cast on the augmentation algorithm to see if it does not distort the plume feature enhancement. One issue in the current augmentation algorithm is the inclusion of the ERA5 wind data which

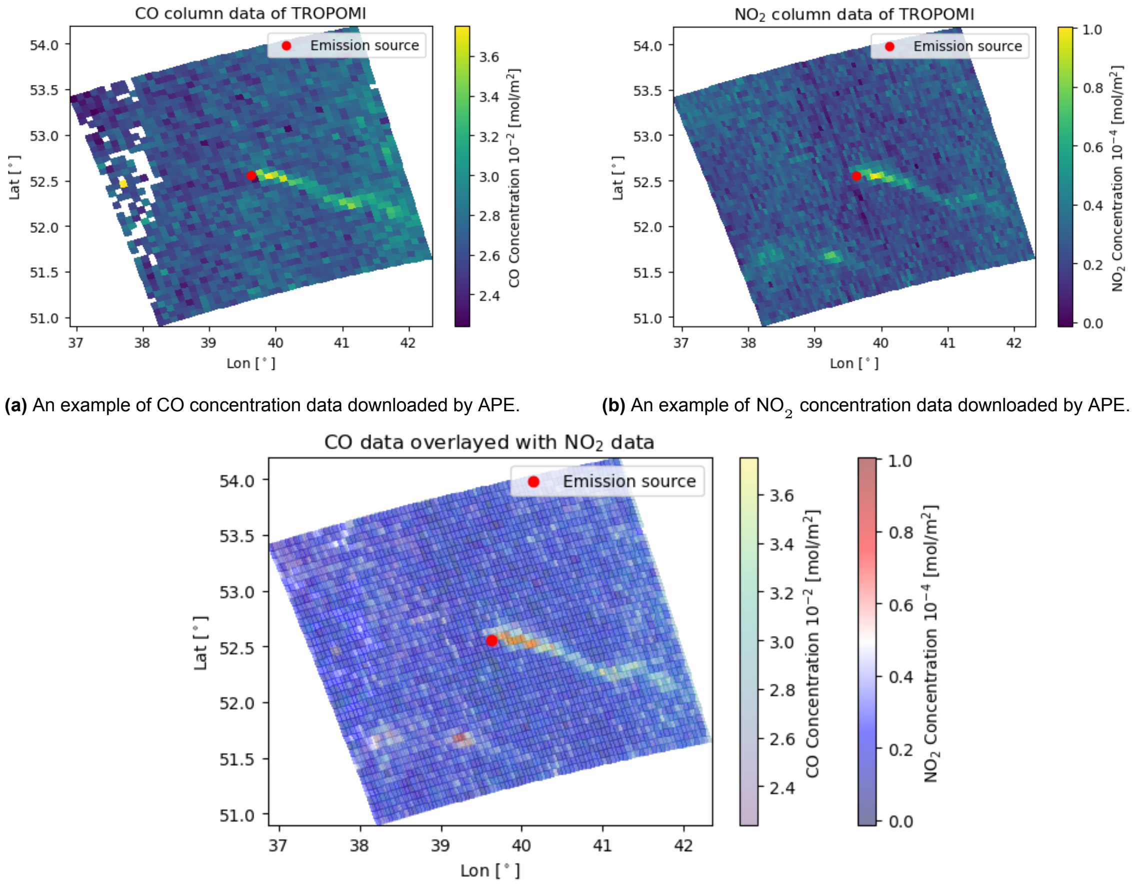

(a) An example of CO concentration data downloaded by APE. **(b)** An example of NO_2 concentration data downloaded by APE.

(c) Combination of both column data to show that they line up with each other.

Figure 8.3: Example of APE's ability to download CO and NO_2 data. The plumes were measured on 2-7-2018 over Lipetsk.

has a worse spatial resolution compared with the TROPOMI CO measurement data. To use the ERA5 data on the same grid as the CO data, it has to be interpolated which introduces interpolation errors. It can happen that due to uncertainties in the data or interpolation errors, plumes get obscured by the augmentation algorithm.

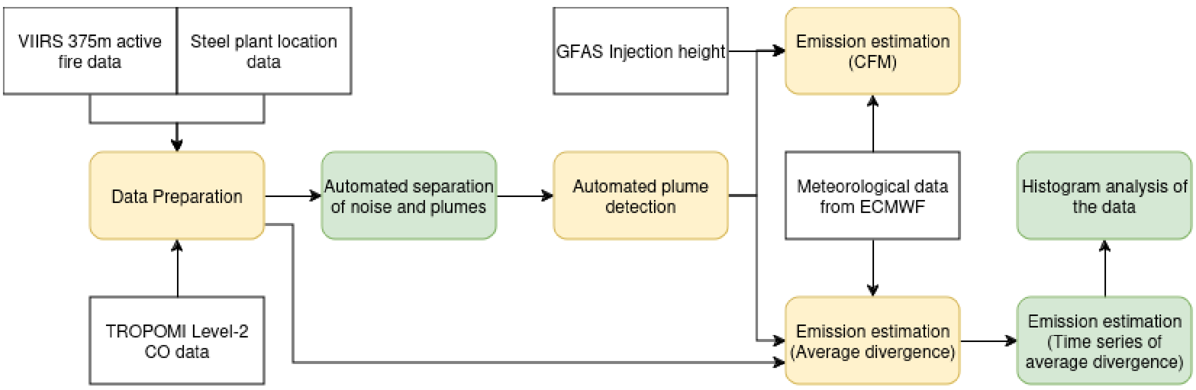

Figure 8.4: Overview of the new functionalities if they were integrated into APE. The yellow blocks denote the already existing and implemented functions of APE and the green blocks show the implementations that were investigated in this project.

The work that is left for the future is implementing the new methods into APE. These include the new emission estimation method to calculate the time series as well as the histogram method and ultimately the machine learning method. However, that last one is dependent on the outcome of other research.

If all steps are completed, APE will be able to function as outlined in Figure 8.4. In the figure, the green blocks show the methods explored in this thesis.

The final goal for APE is to be able to find plumes from industrial or wildfire sources without the need for a predetermined location while processing TROPOMI data in real time. Currently APE is not able to do the first point well and the second point at all. The new plume detection using machine learning should make it possible to input granules from every location on Earth. After the implementation of the improved plume detections, the development of APE should focus on the real-time processing of the data. This process should include creating a routine to get rid of granules that include only noise data like granules with only ocean measurements or granules that consist mostly of measurements with low quality values. An overview of this version of APE can be found in Figure 8.5. The blue block represents a new implementation that should be researched. This version of APE is smaller than the one presented in Figure 8.4 as this future version of APE is only focused on real-time evaluation. A second version of APE would be used to perform further analysis of emissions from certain locations over specified periods of time.

With the completion of these implementations, the focus could be placed on the inclusion of other trace gases into APE such as NO_x. These trace gases could be used in several ways. They could be used to refine the plume masking, improve plume detection, or improve the emission estimation. All three of these implementations should be researched.

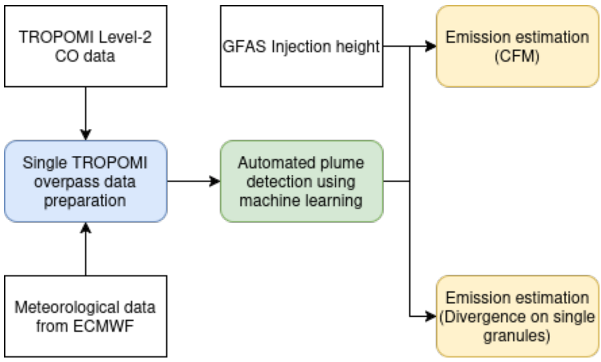

Figure 8.5: Potential future version of APE that is used for near real-time processing. The yellow blocks denote the already existing and implemented functions of APE, the green blocks show the implementations that were investigated in this project and the blue block indicates an area for further development.

References

[1] ESA. *Sentinel-5P*. Jan. 2024. URL: https://sentinels.copernicus.eu/web/sentinel/missions/sentinel-5p.

[2] SRON. *TROPOMI*. Jan. 2024. URL: https://earth.sron.nl/project/tropomi/.

[3] Arthur Bronstring. "Automated Assessment of the TROPOMI CO Data Product for Plume Detection and Emission Estimation". Report made for the course AE5051 Internship. 2024.

[4] James Hansen et al. "Climate change and trace gases". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* (2007).

[5] Carl-Friedrich Schleussner et al. "Science and policy characteristics of the Paris Agreement temperature goal". In: *Nature Climate Change* (2016).

[6] IPCC 2023. "Climate Change 2023: Synthesis Report." In: *IPCC* (2023). Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)], pp. 35–115.

[7] Copernicus. *Global Climate Highlights 2024*. Jan. 2025. URL: https://climate.copernicus.eu/global-climate-highlights-2024.

[8] André Berger and Ch Tricot. "The greenhouse effect". In: *Surveys in geophysics* (1992).

[9] Glenn Althor, James EM Watson, and Richard A Fuller. "Global mismatch between greenhouse gas emissions and the burden of climate change". In: *Scientific reports* (2016).

[10] José L Domingo et al. "Health risks for the population living near petrochemical industrial complexes. 1. Cancer risks: a review of the scientific literature". In: *Environmental research* (2020).

[11] Jose M Baldasano, Eduardo Valera, and Pedro Jiménez. "Air quality data from large cities". In: *Science of the Total Environment* (2003).

[12] Jonathan M Harris and Brian Roach. *Environmental and natural resource economics: A contemporary approach*. Fifth. Routledge, 2017.

[13] Jack Davison et al. "Verification of a national emission inventory and influence of on-road vehicle manufacturer-level emissions". In: *Environmental Science & Technology* (2021).

[14] Nickolay A Krotkov et al. "Aura OMI observations of regional SO 2 and NO 2 pollution changes from 2005 to 2015". In: *Atmospheric Chemistry and Physics* (2016).

[15] Manfred Gottwald and Heinrich Bovensmann. *SCIAMACHY-Exploring the changing Earth's Atmosphere*. Springer Science & Business Media, 2010.

[16] Haiyan Luo et al. "Greenhouse gases monitoring instrument on GaoFen-5 Satellite-II: optical design and evaluation". In: *Remote Sensing* (2023).

[17] Carl Malings et al. *Remote Sensing of Trace Gases, NASA Air Quality Remote Sensing Training for EPA, March 21-23, 2023*. 2023. URL: https://appliedsciences.nasa.gov/sites/default/files/2023-03/D1P7_TraceGases_Final.pdf.

[18] CEOS Atmospheric Composition Virtual Constellation, the CEOS Working Group on Calibration, and Validation. *Geostationary Satellite Constellation for Observing Global Air Quality: Geophysical Validation Needs*. 2019. URL: https://ceos.org/observations/documents/GEO_AQ_Constellation_Geophysical_Validation_Needs_1.1_2Oct2019.pdf.

[19] Michael Berger et al. "ESA's sentinel missions in support of Earth system science". In: *Remote Sensing of Environment* (2012).

[20] J Pepijn Veefkind et al. "TROPOMI on the ESA Sentinel-5 Precursor: A GMES mission for global observations of the atmospheric composition for climate, air quality and ozone layer applications". In: *Remote sensing of environment* (2012).

[21] Tobias Borsdorff et al. “Random Forest Classifier for Cloud Clearing of the Operational TROPOMI XCH4 Product”. In: *Remote Sensing* (2024).

[22] Manu Goudar et al. “Plume detection and emission estimate for biomass burning plumes from TROPOMI carbon monoxide observations using APE v1. 1”. In: *Geoscientific Model Development* (2023).

[23] Arjuman R Reshi, Subbarao Pichuka, and Akshar Tripathi. “Applications of Sentinel-5P TROPOMI Satellite Sensor: A Review”. In: *IEEE Sensors Journal* (2024).

[24] Daniel A Potts et al. “Diagnosing air quality changes in the UK during the COVID-19 lockdown using TROPOMI and GEOS-Chem”. In: *Environmental Research Letters* (2021).

[25] Nenghan Wan et al. “Estimation of biomass burning emission of NO 2 and CO from 2019–2020 Australia fires based on satellite observations”. In: *Atmospheric Chemistry and Physics* (2023).

[26] A Apituley et al. *Sentinel-5 precursor/TROPOMI level 2 Product user manual Carbon Monoxide*. Tech. rep. ESA, 2022.

[27] Manu Goudar et al. “Plume detection and emission estimate for biomass burning plumes from TROPOMI carbon monoxide observations using APE v1.1”. In: *Geoscientific Model Development* (Aug. 2023).

[28] Jochen Landgraf et al. “The TANGO mission: A satellite tandem to measure major sources of anthropogenic greenhouse gas emissions”. In: *EGU General Assembly Conference Abstracts*. 2020, p. 19643.

[29] Hans Hersbach et al. “The ERA5 global reanalysis”. In: *Quarterly Journal of the Royal Meteorological Society* 146.730 (2020), pp. 1999–2049. DOI: `https://doi.org/10.1002/qj.3803`.

[30] Steffen Beirle et al. “Pinpointing nitrogen oxide emissions from space”. In: *Science advances* (2019).

[31] Erik Koene et al. *Documentation of plume detection and quantification methods*. Tech. rep. Tech. Rep., Empa. CoCO2: Prototype System for a Copernicus CO2 Service …, 2021.

[32] J Landgraf et al. *S5P Mission Performance Centre Carbon Monoxide [L2_CO_] Readme*. Tech. rep. ESA, 2023.

[33] David J Lary et al. “Machine learning in geosciences and remote sensing”. In: *Geoscience Frontiers* (2016).

[34] Aaron E Maxwell, Timothy A Warner, and Fang Fang. “Implementation of machine-learning classification in remote sensing: An applied review”. In: *International journal of remote sensing* (2018).

[35] María C Alonso, José A Malpica, and A Martínez de Agirre. “Consequences of the Hughes phenomenon on some classification techniques”. In: *ASPRS 2011 Annual Conference*. 2011.

[36] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* (2011).

[37] Yuan Tian et al. “Quantifying CO emission rates of industrial point sources from Tropospheric Monitoring Instrument observations”. In: *Environmental Research Letters* (2022).

[38] Sarmad Dashti Latif et al. “Prediction of atmospheric carbon monoxide concentration utilizing different machine learning algorithms: A case study in Kuala Lumpur, Malaysia”. In: *Environmental Technology & Innovation* (2023).

[39] Jigsaw B.V. “Balsa: A Fast C++ Random Forest Classifier”. Read me file written for the SRON Eurth Group. 2024.

[40] Leo Breiman. “Random Forests—Random Features. Statistics Department”. In: *University of California, Berkeley, Technical Report* (1999).

[41] Mahesh Pal. “Random forest classifier for remote sensing classification”. In: *International journal of remote sensing* (2005).

[42] Douglas P Finch, Paul I Palmer, and Tianran Zhang. “Automated detection of atmospheric NO 2 plumes from satellite data: a tool to help infer anthropogenic combustion emissions”. In: *Atmospheric Measurement Techniques* (2022).

[43] Bo Zheng et al. “Observing carbon dioxide emissions over China’s cities and industrial areas with the Orbiting Carbon Observatory-2”. In: *Atmospheric Chemistry and Physics* (2020).

[44] Gerrit Kuhlmann et al. “Quantifying CO2 emissions of power plants with CO2 and NO2 imaging satellites”. In: *Frontiers in Remote Sensing* (2021).

[45] Kaiming He et al. “Deep residual learning for image recognition”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 770–778.

[46] Schwirtz M. “Last Stand at Azovstal: Inside the Siege That Shaped the Ukraine War”. In: *New York Times* (July 2022). https://www.nytimes.com/2022/07/24/world/europe/ukraine-war-mariupol-azovstal.html (Accessed: 5 May 2024).

[47] M. Crippa et al. “Insights into the spatial distribution of global, national, and subnational greenhouse gas emissions in the Emissions Database for Global Atmospheric Research (EDGAR v8.0)”. In: *Earth System Science Data* (2024).

[48] G. Leguijt et al. “Comparing space-based to reported carbon monoxide emission estimates for Europe’s iron and steel plants”. In: *Atmospheric Chemistry and Physics* (2025).

A

Steel plants in the comparison test result overview

This appendix displays the locations of the 180 steel plants used in the comparison test. The names and locations can be found in Table A.1. The name, latitude and longitude were taken from GEM. The latitude and longitude are given in degrees.

Table A.1: The 180 steel plants that were used in the comparison test

Steel plant name	Index	Latitude	Longitude
AG der Dillinger Hüttenwerke Dillingen steel plant	18	49.3539	6.7466
Aichi Steel Chita Plant (Tokai)	21	35.0446	136.901
Alchevsk Iron & Steel plant	29	48.4781	38.7708
Angang Group Xinyang Iron and Steel Co., Ltd. plant	38	32.4887	114.038
Angang Steel Company Limited Anshan production base	40	41.1483	122.983
Anhui Changjiang Steel Co., Ltd. plant	42	31.5023	118.465
Anhui Guihang Special Steel Co., Ltd. plant	43	30.5311	117.251
Anhui Shoukuang Dachang Metal Material Co., Ltd. plant	48	32.3292	115.955
Anshan Baode Iron & Steel Co., Ltd. plant	51	41.1736	122.946
Ansteel Group Chaoyang Steel & Iron Co., Ltd. plant	53	41.5289	120.37
Anyang Huixin Special Steel Co., Ltd. plant	54	36.1524	114.158
Anyang Iron & Steel Co., Ltd. plant	55	36.1221	114.283
Anyang Xinpu Steel Co., Ltd. plant	56	36.1632	114.296
ArcelorMittal Asturias (Gijón) steel plant	64	43.5251	-5.73193
ArcelorMittal Duisburg steel plant	69	51.463	6.74461
ArcelorMittal Dąbrowa Górnicza steel plant	71	50.3424	19.285
ArcelorMittal Kryvyi Rih steel plant	81	47.8744	33.393
ArcelorMittal Lázaro Cárdenas steel plant	82	17.9307	-102.201
ArcelorMittal Vanderbijlpark Steel Works	97	-26.6574	27.8223
Azovstal Iron & Steel Works	113	47.1004	37.5959
Baoshan Iron and Steel Co., Ltd. Baoshan headquarters	119	31.4162	121.44
Baosteel Desheng Stainless Steel Co., Ltd. plant	121	26.4875	119.654
Baosteel Group Xinjiang Bayi Iron & Steel Co., Ltd. plant	122	43.8521	87.3046
Bengang Steel Plates Co., Ltd. plant	129	41.2741	123.722
Benxi-Beiying Iron & Steel (Group) Co., Ltd. plant	131	41.2248	123.608
Cangzhou China Railway Equipment Manufacture Material Co., Ltd. plant	145	38.2892	117.814
Changzhi Xingbao Steel Co., Ltd. plant	162	36.3324	113.155
Changzhou Eastern Special Steel Co., Ltd. plant	163	31.6026	119.724

Continued on the next page...

Steel plant	Index	Latitude	Longitude
Chengde Iron and Steel Group Co., Ltd. plant	166	40.9477	117.72
Chengde Jianlong Special Steel Co., Ltd. plant	167	40.515	117.596
Chengyu Vanadium & Titanium Technology Co., Ltd. steel plant	171	29.7321	104.504
Chifeng Yuanlian Steel Co., Ltd. plant	172	42.2781	119.015
Chizhou Guichi Guihang Metal Products Co	177	30.5328	117.252
Chongqing Iron & Steel Co., Ltd. plant	180	29.7903	107.04
Chongqing Yonghang Steel Group Co., Ltd. plant	181	29.8362	106.988
Chubu Steel Plate Nagoya plant	183	35.1188	136.868
Daye Xinye Special Steel Co., Ltd. plant	218	30.0675	114.934
Dniprovskiy Metallurgical Plant	229	48.528	34.6409
ESL Steel Ltd plant	251	23.6329	86.2937
Fujian Luoyuan Minguang Iron and Steel Co., Ltd. plant	279	26.4838	119.671
Fujian Quanzhou Minguang Iron and Steel Co., Ltd. plant	281	25.2444	118.03
Fujian Sanbao Iron and Steel Co., Ltd. plant	282	24.6358	117.604
Fujian Tsingtuo Nickel Industry Co., Ltd. plant	283	26.7689	119.765
Fujian Yixin Steel Co., Ltd. plant	285	26.4913	119.636
Fushun Special Steel Co., Ltd. plant	288	41.8388	123.807
Gansu Jiu Steel Group Hongxing Iron and Steel Co., Ltd. plant	290	39.8121	98.2956
Hebei Anfeng Iron & Steel Co., Ltd.	370	39.6585	118.895
Hebei Donghai Special Steel Co., Ltd. plant	373	39.6106	118.475
Hebei Huaxi Special Steel Co., Ltd. plant	375	39.2608	119.006
Hebei Jinxi Iron & Steel Group Co., Ltd. plant	377	40.2098	118.221
Hebei Longfengshan Casting Co., Ltd.plant	378	36.7051	114.105
Hebei Puyang Iron and Steel Co., Ltd. plant	380	36.7337	113.926
Hebei Rongxin Steel Co., Ltd. plant	381	39.8807	118.559
Hebei Taihang Iron and Steel Group Co., Ltd. plant	382	36.6137	114.082
Hebei Tianzhu Iron and Steel Group Special Steel Co., Ltd. plant	386	39.27	119.014
Hebei Xin Wu’an Steel Group Wen’an Iron and Steel Co., Ltd. plant	389	36.6794	114.169
Hebei Xinda Iron and Steel Co., Ltd. plant	392	39.8984	118.579
Hebei Xinghua Iron and Steel Co., Ltd. plant	393	36.7449	114.112
Hebei Xinjin Iron and Steel Co., Ltd. plant	394	36.7156	114.209
Hebei Xinxing Ductile Iron Pipes Co., Ltd. plant	395	36.6215	114.128
Hebei Zongheng Group Fengnan Iron & Steel Co., Ltd. plant	398	39.2226	118.094
Heilongjiang Jianlong Iron and Steel Co., Ltd. plant	399	46.6023	131.072
Hejin Huaxinyuan Steel & Iron Co., Ltd. plant	400	35.6055	110.652
Henan Fengbao Special Steel Co., Ltd. plant	403	36.1317	113.872
Henan Jiyuan Iron & Steel (Group) Co., Ltd. plant	404	35.0864	112.554
Henan Xinjinhui Stainless Steel Industry Co., Ltd. plant	405	34.2635	113.854
Hengyang Valin Steel Tube Co., Ltd. plant	407	26.8655	112.577
Hoa Phat Hai Duong Steel plant	415	21.0154	106.536
Hubei Jinshenglan Metallurgical Technology Co., Ltd. plant	421	29.9167	113.823
Hubei Wucheng Iron and Steel Group Co., Ltd. plant	424	30.3412	114.858
Hunan Valin Lianyuan Iron and Steel Co., Ltd. plant	426	27.7478	111.967
Hunan Valin Xiangtan Iron and Steel Co., Ltd. plant	427	27.8141	112.897
Hyundai Steel Dangjin steel plant	433	36.9863	126.697
Hyundai Steel Pohang steel plant	435	36.005	129.385
Hüttenwerke Krupp Mannesmann (HKM) steel plant	436	51.3713	6.72331
Ilyich Iron & Steel Works	437	47.1425	37.5864
Inner Mongolia BaoTou Steel Union Co.,Ltd. plant	441	40.648	109.741
Interpipe Nyzhnyodniprovskyi Tube Rolling Plant	445	48.4921	35.0927
JFE West Japan Works (Kurashiki) steel plant	472	34.5012	133.722
Jianglong Acheng Iron & Steel Co., Ltd.	473	45.562	126.968
Jiangsu Binxin Steel Group Co., Ltd.	474	35.0956	119.28

Continued on the next page...

Steel plant	Index	Latitude	Longitude
Jiangsu Changqiang Iron & Steel Co., Ltd. plant	475	31.9521	120.19
Jiangsu Hongtai Iron and Steel Co., Ltd. plant	477	32.1384	119.555
Jiangsu Shagang Group Huaigang Special Steel Co., Ltd. plant	478	33.566	118.988
Jiangsu Shente Steel Co. Ltd. plant	479	31.4711	119.486
Jiangsu Xugang Iron and Steel Group Co., Ltd.	481	34.577	117.34
Jiangsu Yonggang Group Co., Ltd. steel plant	482	31.8552	120.731
Jiangyin Xingcheng Special Steel Works Co., Ltd. plant	486	31.9502	120.328
Jilin Jianlong Steel Co., Ltd. plant	490	44.012	126.537
Jilin Xinda Iron and Steel Co., Ltd. plant	494	42.3249	125.247
Jinan Iron & Steel Group Co., Ltd. plant	495	36.719	114.112
Jincheng Fusheng Iron & Steel Co., Ltd. plant	496	35.663	112.882
Jinding Heavy Industry Co., Ltd. steel plant	501	36.719	114.112
Jingye Iron and Steel Co., Ltd. plant	502	38.3914	114.14
Jiugang Group Yuzhong Iron & Steel Co., Ltd. plant	506	36.0336	104.036
Jiujiang Pinggang Iron and Steel Co., Ltd. plant	507	29.7784	116.275
JSPL Chhattisgarh steel plant	509	21.9227	83.3472
JSW Odisha steel plant	515	21.764	84.0225
JSW Steel Vijayanagar steel plant	523	15.1804	76.6631
Kalyani Steels Hospet plant	527	15.3391	76.2517
Laiwu Iron and Steel Group Yinshan Section Steel Co., Ltd. plant	561	36.1061	117.82
Lengshuijiang Steel Co., Ltd. plant	570	27.6904	111.437
Lingyuan Iron & Steel Co., Ltd. plant	581	41.271	119.414
Liuzhou Iron & Steel Co., Ltd.(Liuzhou Base□plant	592	24.3853	109.38
Ma’anshan Iron & Steel Co., Ltd. plant	601	31.6989	118.468
Magnitogorsk Iron & Steel Works	602	53.4276	59.0541
Mechel Chelyabinsk Metallurgical Plant	613	55.2707	61.4365
Mesco Steel Kalinganagar plant	617	20.9763	86.0444
Minyuan Iron and Steel Group Co., Ltd.	626	34.0134	116.423
Nanjing Iron & Steel Co., Ltd. plant	638	32.2008	118.746
Nanyang Hanye Special Steel Co., Ltd. plant	640	33.2693	111.535
Nippon East Japan Works (Nagoya) steel plant	653	35.0276	136.871
Panzhihua Steel Group Panzhihua Steel Vanadium Co., Ltd. plant	713	26.5653	101.674
Pingxiang Pinggang Anyuan Iron & Steel Co., Ltd. plant	716	27.6456	113.902
POSCO Gwangyang steel plant	722	34.9201	127.749
POSCO Zhangjiagang Stainless Steel Co., Ltd. plant	725	31.9802	120.571
Puyang Linzhou Iron and Steel Co., Ltd. plant	728	36.1333	113.873
Qian’an Jiujiang Wire Co., Ltd. steel plant	731	39.9522	118.559
Qinhuangdao Hongxing Iron and Steel Co., Ltd. plant	735	39.7007	118.871
Quzhou Yuanli Metal Products Co., Ltd. steel plant	736	28.8934	118.867
Rizhao Steel Holding Group Co., Ltd. plant	754	35.1654	119.367
Rizhao Steel Yingkou Medium Plate Co., Ltd. steel plant	755	40.6711	122.395
SAIL Alloy Steel Plant	762	23.5229	87.2759
SAIL Bhilai steel plant	763	21.1852	81.3942
SAIL Bokaro steel plant	764	23.6717	86.1069
SAIL Durgapur steel plant	765	23.5484	87.2452
SAIL IISCO steel plant	766	23.6732	86.9262
SAIL Rourkela steel plant	768	22.2108	84.869
Salzgitter Flachstahl steel plant	769	52.1618	10.4094
Sansteel Minguang Co., Ltd. Fujian plant	771	26.2611	117.615
SGIS Songshan Co.,Ltd.	778	24.7082	113.636
Shaanxi Hanzhong Iron & Steel Group Co., Ltd.	779	33.1324	106.658
Shaanxi Longmen Steel Co., Ltd. plant	782	35.6185	110.576
Shaanxi Lueyang Iron and Steel Co., Ltd. plant	783	33.3237	106.167

Continued on the next page...

Steel plant	Index	Latitude	Longitude
Shaanxi Steel Group Hanzhong Iron and Steel Co., Ltd. plant	784	33.1332	106.672
Shagang Group Anyang Yongxing Special Steel Co., Ltd. plant	785	36.1436	114.15
Shandong Chuanyang Group Co., Ltd. steel plant	786	36.8948	117.888
Shandong Fulun Iron and Steel Co., Ltd. plant	787	36.3084	117.54
Shandong Iron and Steel Co., Ltd. Laiwu Branch plant	789	36.0935	117.838
Shandong IRON&STEEL Group Rizhao Co., Ltd. plant	792	35.189	119.384
Shandong Laigang Yongfeng Steel Co., Ltd. plant	793	36.8135	116.758
Shandong Shouguang Juneng Special Steel Co., Ltd. plant	797	36.9391	118.786
Shandong Taishan Steel Group Co., Ltd. plant	798	36.2214	117.628
Shanghai Meishan Iron and Steel Co., Ltd. plant	799	31.903	118.617
Shanxi Changxin Industrial Co., Ltd. steel plant	800	36.3266	113.117
Shanxi Gaoyi Steel Co., Ltd. plant	801	35.5868	111.258
Shanxi Tongcai Industry and Trade Co., Ltd. steel plant	810	35.6943	111.429
Shanxi Zhongsheng Iron and Steel Co., Ltd. plant	813	35.7278	111.33
Shanxi Zhongyang Iron and Steel Co., Ltd. plant	814	37.3708	111.161
Shiheng Special Steel Group Co., Ltd. plant	815	36.2292	116.534
Shougang Changzhi Iron and Steel Co., Ltd. plant	821	36.3564	113.061
Shougang Jingtang Iron & Steel United Co., Ltd. plant	823	38.9537	118.503
Shougang Qian'an Iron and Steel Co., Ltd. plant	824	39.9767	118.559
Shyam Steel Durgapur plant	826	23.5084	87.2823
Sichuan Dazhou Iron & Steel Group Co., Ltd. plant	829	31.1895	107.453
Tangshan Donghai Iron and Steel Group Co., Ltd. plant	882	39.7484	118.625
Tangshan Donghua Iron & Steel Enterprise Group Co., Ltd. plant	883	39.4684	118.257
Tangshan Ganglu Iron and Steel Co., Ltd. plant	886	40.1839	118.074
Tangshan Songting Iron & Steel Co., Ltd. plant	894	39.9389	118.577
Tangshan Yanshan Iron and Steel Co., Ltd. plant	899	39.9255	118.673
Tata Steel BSL Dhenkanal plant	904	20.7961	85.2604
Tata Steel Jamshedpur steel plant	906	22.7886	86.1996
ThyssenKrupp Steel Duisburg steel plant	931	51.4916	6.73305
Tianjin Iron & Steel Group Co., Ltd. plant	932	39.0311	117.499
Tianjin Iron Works Co., Ltd. plant	933	36.5897	113.746
Tianjin Rockcheck Steel Group Co., Ltd. plant	937	38.9706	117.496
Tonghua Iron & Steel Co., Ltd.	948	41.7791	126.022
Tongling Fuxin Iron and Steel Co., Ltd. plant	949	30.9076	117.773
Tongling Xuanli Special Steel Co., Ltd. plant	950	31.0628	117.956
Vizag Steel plant	987	17.6128	83.1919
Voestalpine Stahl Linz steel plant	991	48.274	14.3343
Wu'an Yuhua Iron and Steel Co., Ltd. plant	1005	36.7318	114.098
Wuhan Iron and Steel Co., Ltd. Qingshan plant	1009	30.6162	114.445
Xinji Aosen Iron & Steel Co., Ltd. plant	1023	37.7524	115.185
Xinyu Steel Group Co., Ltd. plant	1031	27.7869	114.922
Zaporizhstal steel plant	1059	47.8684	35.1618
Zenith Steel Group Co., Ltd. plant	1061	31.7077	120.08
Zhangjiagang Hongchang Steel Co., Ltd.	1063	31.9831	120.639
Zhangjiagang Rongsheng Special Steel Co., Ltd.	1064	31.9855	120.645
Zhong Xin Iron and Steel Group Co., Ltd. plant	1070	34.3686	118.316

B

Emission estimates per steel plant

This appendix displays the results of performing the plume detection on the 180 steel plants of the comparison test which can be found in Appendix A. The steel plants are listed by index number rather than name to ensure that the tables fit on the page. The names that belong to the index numbers can be found in Appendix A. The results for the APE, SVM and ResNet-26 methods are shown in Table B.1 while the results for ResNet-44, RFC and the distribution method can be found in Table B.2. In the tables, the parameters μ, σ and # denote the average emission estimates in kg/s, the standard deviations of the emission estimates in kg/s and the number of detections, respectively.

Table B.1: Overview of the results of performing the APE, SVM and ResNet-26 methods on the 180 steel plants

Steel plant	APE			SVM			ResNet-26		
index	μ	σ	#	μ	σ	#	μ	σ	#
18	1.52	5.61	234	3.57	4.17	126	3.23	4.54	129
21	6.72	8.77	28	12.43	9.09	127	7.54	12.66	386
29	1.37	6.41	201	5.45	5.85	129	5.00	8.05	74
38	3.00	7.53	153	5.03	13.70	57	1.68	11.28	122
40	36.22	27.04	305	29.41	27.51	848	28.02	29.25	828
42	4.54	10.58	124	6.18	12.60	229	3.22	13.02	264
43	2.87	8.26	150	7.05	12.13	81	2.37	10.48	172
48	3.56	5.53	168	5.70	8.94	81	3.85	12.85	126
51	39.43	26.82	293	32.41	28.82	837	30.55	29.80	823
53	3.39	10.96	318	7.69	27.19	33	3.53	21.97	68
54	16.14	21.78	258	17.91	26.96	204	13.20	27.77	369
55	13.03	19.07	246	16.26	24.15	191	11.95	24.57	360
56	10.91	16.39	248	16.00	22.83	198	10.60	23.45	345
64	1.33	13.31	43	11.28	11.98	90	6.58	10.92	334
69	4.82	6.39	246	6.87	5.71	341	6.89	5.97	302
71	0.86	5.66	162	2.24	6.62	75	1.09	6.64	92
81	4.11	5.59	294	6.20	5.41	220	5.69	5.48	236
82	19.71	12.64	46	23.25	10.51	235	19.92	11.10	363
97	5.53	5.30	233	5.47	6.45	301	4.16	8.73	343
113	0.06	5.81	36	11.75	10.68	223	8.62	10.24	339
119	11.74	18.19	26	12.41	17.63	96	10.89	15.62	162
121	14.03	26.14	36	18.31	21.36	241	10.51	20.26	407
122	-7.25	15.81	351	-3.25	19.28	88	-3.13	15.43	391
129	19.23	28.43	316	19.71	27.20	602	15.25	28.29	620
131	20.06	32.28	317	19.51	29.45	622	15.80	29.65	564

Continued on the next page...

Steel plant index	APE μ	σ	#	SVM μ	σ	#	ResNet-26 μ	σ	#
145	11.33	17.18	20	18.04	29.26	118	10.95	23.55	213
162	6.02	14.66	112	7.32	19.30	43	2.85	17.64	125
163	0.96	10.65	68	5.93	19.85	67	1.32	15.19	132
166	14.24	15.06	351	20.71	23.55	96	15.52	21.31	157
167	5.18	13.67	46	5.95	10.97	17	-3.90	29.14	16
171	4.31	6.47	57	7.53	9.05	28	1.34	13.26	74
172	5.38	8.10	412	6.14	16.90	33	5.20	11.24	84
177	2.92	8.21	151	6.90	12.17	82	2.32	10.80	173
180	4.46	6.94	92	4.18	6.67	117	1.49	6.60	292
181	2.15	7.50	91	4.21	6.71	89	0.20	7.66	290
183	4.35	8.76	35	10.77	11.50	95	5.40	11.55	390
218	8.74	12.63	156	11.19	13.12	223	8.28	12.78	346
229	2.31	5.88	308	5.20	6.60	185	4.31	7.18	154
251	10.08	12.58	151	8.23	12.35	486	7.01	12.90	519
279	15.14	26.95	34	18.53	21.63	261	11.72	20.24	399
281	1.95	8.06	59	6.98	7.49	12	1.87	7.52	131
282	2.46	6.59	77	5.23	8.26	110	1.47	7.70	257
283	4.28	9.87	32	7.12	17.62	198	2.28	15.05	443
285	13.15	25.50	36	17.19	21.30	230	8.99	19.96	403
288	5.78	17.68	386	7.14	27.25	262	-0.21	25.61	401
290	9.10	13.67	349	15.39	17.08	54	9.52	14.09	599
370	36.42	33.95	29	20.86	33.92	802	18.32	35.04	742
373	23.41	28.13	33	16.61	28.95	806	14.14	29.72	696
375	-4.29	32.72	24	0.63	42.15	281	-5.87	40.42	405
377	2.10	26.51	257	6.34	37.69	287	4.05	35.22	233
378	51.03	45.22	271	45.20	45.33	733	41.69	45.58	739
380	13.47	58.47	219	16.06	45.83	608	8.28	45.84	716
381	48.63	49.61	66	39.35	40.10	962	38.39	41.14	788
382	44.46	41.37	223	43.68	47.58	691	37.80	47.23	737
386	-2.66	33.22	23	-1.03	42.98	280	-5.79	41.30	407
389	51.50	39.70	260	44.89	45.69	732	41.65	44.56	765
392	46.23	48.65	67	38.87	40.02	961	38.21	41.11	812
393	45.45	44.21	269	39.31	43.14	725	35.10	42.17	719
394	51.87	40.91	268	43.21	44.92	728	40.35	44.35	765
395	46.44	39.17	232	42.38	46.53	703	37.53	45.67	750
398	11.20	14.95	20	4.67	25.43	282	-0.25	28.56	278
399	5.34	7.89	226	8.71	11.60	100	5.94	14.33	72
400	10.02	27.51	214	5.30	22.55	665	3.86	23.02	599
403	-5.85	39.24	190	0.24	39.00	135	-7.08	33.12	475
404	4.14	26.47	200	13.04	29.64	68	-0.64	23.11	251
405	1.32	11.48	177	4.08	17.48	33	1.92	17.70	76
407	1.18	6.71	132	8.49	18.73	19	1.60	10.56	117
415	2.15	4.93	31	9.36	17.90	65	4.17	18.28	135
421	3.11	7.31	103	4.94	9.98	25	1.67	9.70	106
424	2.83	10.43	146	4.82	12.47	239	2.52	12.83	261
426	5.61	9.10	104	9.04	10.48	133	5.28	10.43	236
427	8.25	8.88	146	9.01	9.01	168	6.01	10.81	263
433	10.12	9.89	33	10.18	11.44	200	5.93	12.45	430
435	5.99	11.76	19	9.12	17.33	96	3.42	15.12	197
436	3.41	7.23	252	5.78	5.99	341	5.46	5.63	254
437	1.89	7.49	50	10.67	9.59	224	8.02	9.62	334
441	36.11	32.91	446	27.99	31.28	1174	27.33	31.18	1102

Continued on the next page...

Steel plant index	APE μ	APE σ	APE #	SVM μ	SVM σ	SVM #	ResNet-26 μ	ResNet-26 σ	ResNet-26 #
445	-0.65	5.67	214	3.22	6.56	57	0.22	7.60	90
472	8.75	12.18	16	7.67	10.35	340	5.68	12.31	485
473	-3.90	12.24	328	-2.63	23.51	140	-8.00	23.71	162
474	48.63	48.59	24	29.07	33.01	510	24.88	32.11	604
475	5.09	12.23	56	4.98	13.04	267	2.30	12.97	266
477	5.00	12.60	102	5.56	14.64	81	NaN	NaN	0
478	2.84	8.77	79	4.87	11.65	56	1.80	12.00	112
479	-1.08	9.02	82	2.48	13.36	34	-1.04	11.07	107
481	7.20	11.08	148	9.18	11.65	154	5.47	12.88	203
482	20.21	18.40	33	23.29	21.31	461	21.40	21.18	403
486	4.93	13.81	49	7.18	14.66	371	4.95	13.88	360
490	1.58	9.63	435	3.35	16.41	211	0.26	19.32	214
494	6.34	11.40	333	15.67	16.08	57	8.96	12.14	133
495	49.76	44.90	268	43.27	44.51	733	39.65	44.67	732
496	8.76	14.47	74	28.84	17.20	8	6.83	20.03	31
501	49.76	44.90	268	43.27	44.51	733	39.65	44.67	732
502	7.66	14.10	289	11.22	18.34	369	8.09	17.19	528
506	1.06	13.47	261	7.49	14.52	287	1.71	13.10	652
507	4.61	9.36	122	7.88	10.50	146	4.38	10.83	264
509	4.49	7.56	173	7.57	9.65	189	7.18	10.39	171
515	5.09	7.96	171	6.45	9.89	272	4.52	9.56	278
523	6.42	8.17	244	6.70	7.73	470	5.92	7.67	413
527	0.70	7.05	165	4.61	8.70	209	2.24	8.38	285
561	13.51	18.13	237	18.40	20.72	298	15.91	19.35	344
570	3.88	8.80	79	7.07	11.45	49	1.79	8.72	209
581	9.78	11.97	286	13.25	18.91	45	10.34	13.88	110
592	7.89	8.21	149	10.17	10.76	260	7.74	11.99	312
601	10.89	10.29	176	9.04	13.26	347	8.28	13.98	356
602	10.74	8.86	297	10.25	8.89	656	9.64	9.02	704
613	4.69	6.27	250	5.66	7.11	199	5.07	6.68	280
617	1.43	4.25	42	2.46	12.15	77	1.41	10.77	128
626	4.32	7.91	163	5.34	7.79	66	2.97	8.55	117
638	9.61	11.68	183	7.67	14.59	322	5.96	14.91	295
640	3.92	9.73	159	10.34	13.03	59	2.05	10.71	225
653	7.03	8.87	30	13.10	10.30	136	7.87	12.61	397
713	15.79	14.73	157	12.80	15.42	734	11.72	15.83	703
716	8.67	11.30	145	11.72	12.63	167	7.12	11.98	261
722	16.10	15.00	27	16.45	15.04	509	14.43	15.94	559
725	25.98	15.28	49	21.47	17.12	548	20.71	17.36	499
728	-5.58	38.91	192	0.17	39.03	135	-7.01	33.04	477
731	39.91	40.60	103	34.95	40.48	943	32.57	38.80	798
735	35.22	38.38	31	20.22	35.01	832	19.35	34.82	751
736	0.61	8.65	148	3.38	7.03	148	-0.11	9.20	282
754	41.96	41.26	26	27.72	31.59	506	23.14	31.96	580
755	20.97	19.14	88	20.64	25.56	595	17.86	25.81	590
762	6.98	9.76	177	6.55	10.57	331	5.42	10.10	418
763	6.13	8.44	210	6.07	7.87	618	5.70	8.48	595
764	8.96	10.33	170	8.11	10.14	472	7.20	10.18	494
765	6.76	9.13	177	6.06	10.87	330	5.31	10.47	425
766	3.61	10.88	147	2.94	12.95	278	1.45	12.99	292
768	4.18	8.89	190	7.21	9.99	266	5.01	10.23	299
769	0.85	4.21	215	2.06	3.72	70	1.92	4.19	69

Continued on the next page...

Steel plant index	APE μ	σ	#	SVM μ	σ	#	ResNet-26 μ	σ	#
771	4.81	7.47	156	6.78	8.25	182	4.11	8.08	306
778	4.04	8.97	158	6.33	9.52	219	3.35	10.62	251
779	1.12	9.12	110	2.89	9.81	213	-0.50	9.03	388
782	10.75	33.74	213	8.12	26.67	604	3.03	27.62	583
783	1.05	6.90	88	4.54	6.35	26	0.14	10.15	151
784	1.05	9.17	112	2.36	9.85	219	-1.04	9.07	386
785	16.81	22.39	259	17.35	29.01	208	13.48	27.86	373
786	12.50	28.81	229	13.08	26.67	444	9.77	24.62	500
787	8.92	23.77	195	20.23	19.30	216	11.20	20.58	336
789	11.89	18.31	226	15.99	20.58	292	14.23	19.43	333
792	33.89	32.48	25	26.91	30.79	503	22.68	30.18	569
793	1.75	11.22	185	3.64	19.52	51	-0.18	17.14	157
797	3.39	14.52	108	2.79	27.11	159	-0.27	20.11	214
798	18.30	19.22	215	21.05	21.17	289	18.51	19.42	335
799	6.53	11.53	174	6.17	13.21	346	5.25	13.54	321
800	8.07	13.32	115	9.31	13.82	36	2.42	16.67	121
801	23.90	26.86	249	19.62	28.97	771	17.15	28.57	748
810	-0.86	32.34	272	2.62	27.66	697	0.50	27.46	646
813	8.16	32.68	270	9.11	27.68	698	5.13	27.90	733
814	-9.31	26.42	164	7.16	18.64	121	-2.39	20.27	240
815	3.23	11.79	212	1.74	20.08	38	0.34	15.53	96
821	8.90	15.89	127	9.16	20.85	28	2.05	15.46	123
823	3.52	11.48	22	6.57	23.67	107	-4.98	25.51	172
824	36.20	37.88	112	33.75	40.56	911	30.62	40.13	791
826	6.79	9.67	175	6.37	10.23	327	5.03	10.01	431
829	1.04	5.02	99	4.95	6.38	36	0.63	6.53	169
882	40.73	40.30	36	33.28	37.13	943	31.73	36.86	815
883	14.56	36.51	22	7.84	26.29	547	5.95	28.25	440
886	6.87	20.76	286	4.94	33.05	269	1.70	33.18	207
894	42.43	42.28	80	36.36	40.66	959	35.84	40.66	814
899	44.41	51.03	52	35.68	40.39	949	34.60	41.15	849
904	4.88	8.46	115	6.12	9.33	460	4.79	10.29	449
906	9.28	9.81	198	10.39	9.48	240	8.88	9.89	359
931	4.93	6.36	244	7.00	5.72	334	7.02	5.89	294
932	8.13	22.53	49	15.94	28.13	168	7.01	25.93	240
933	-22.67	43.28	116	3.47	31.65	452	-7.86	36.20	571
937	8.06	17.37	46	16.90	30.94	175	6.52	27.89	197
948	9.28	12.88	338	16.34	15.05	136	11.87	16.23	207
949	0.68	9.41	153	1.83	16.89	111	-0.13	13.63	195
950	4.19	10.69	165	6.29	13.80	154	2.82	13.08	251
987	9.66	7.23	29	15.68	11.97	182	12.91	11.38	325
991	1.52	6.01	223	3.16	5.33	25	1.22	7.59	231
1005	47.68	45.43	273	41.02	43.31	719	36.89	42.87	720
1009	13.97	11.78	173	12.48	12.38	458	10.38	13.16	467
1023	5.66	14.92	232	9.10	31.66	84	5.01	25.28	135
1031	8.66	8.99	141	9.55	9.44	240	7.78	10.35	272
1059	3.89	5.53	282	4.71	6.08	187	4.41	5.96	174
1061	13.12	11.07	98	9.01	14.87	217	5.63	15.50	224
1063	26.31	15.14	48	21.80	18.09	573	21.62	18.63	516
1064	26.28	15.06	48	21.76	18.02	576	21.49	18.62	517
1070	4.97	9.04	89	8.52	11.28	42	4.86	10.29	117

Table B.2: Overview of the results of performing the ResNet-44, RFC and Distribution methods on the 180 steel plants

Steel plant	ResNet-44			RFC			Distribution		
index	μ	σ	#	μ	σ	#	μ	σ	#
18	3.90	4.23	105	3.71	3.68	71	5.12	2.23	307
21	12.77	9.42	102	11.78	9.00	79	11.52	6.54	697
29	4.97	5.29	94	5.73	5.66	90	6.95	2.89	290
38	4.04	12.22	48	5.76	12.24	32	6.76	3.05	240
40	30.74	27.95	735	30.36	27.12	796	39.58	22.74	848
42	5.95	12.26	197	6.57	12.63	173	9.77	4.94	275
43	5.10	11.87	63	5.62	10.56	65	8.43	3.95	315
48	5.53	9.11	75	5.65	9.40	49	7.53	3.32	312
51	33.38	28.94	754	33.08	29.10	793	41.22	23.01	877
53	5.43	22.86	45	6.83	27.44	31	12.61	4.76	73
54	19.29	26.98	226	18.25	26.87	230	24.36	13.37	646
55	16.09	24.34	206	17.32	23.82	215	19.77	10.48	599
56	15.50	22.63	207	15.88	21.70	210	20.25	10.71	520
64	10.84	9.33	81	11.38	12.20	77	11.10	5.45	515
69	6.10	12.62	279	7.28	5.65	267	6.91	3.92	1067
71	2.06	6.07	70	2.55	5.91	40	6.15	2.39	86
81	6.12	5.37	168	6.35	5.33	176	6.06	2.89	481
82	22.16	11.83	215	23.99	11.43	152	17.81	10.96	950
97	5.53	7.30	235	5.44	6.57	246	6.52	3.78	690
113	11.34	10.89	188	12.36	10.95	178	6.47	2.73	268
119	14.57	15.30	101	13.57	14.64	89	15.60	8.38	600
121	17.07	22.16	220	18.16	21.98	230	11.26	5.31	225
122	1.88	21.57	77	-0.93	23.25	71	NaN	NaN	0
129	20.85	26.36	508	19.53	26.33	587	29.96	16.08	625
131	21.27	29.53	539	20.10	30.25	550	30.18	16.08	730
145	16.23	28.33	110	18.57	28.99	109	12.53	5.95	360
162	5.11	16.09	100	10.50	15.88	22	17.72	8.37	308
163	4.51	17.44	61	4.57	20.66	57	8.62	3.56	78
166	21.37	20.04	124	20.77	23.59	89	14.28	8.04	794
167	0.41	29.63	44	-4.50	45.42	21	11.17	5.35	332
171	10.39	21.10	31	15.88	26.87	16	5.10	2.21	147
172	6.81	14.90	48	4.76	17.18	36	8.26	4.60	714
177	5.63	13.05	65	5.37	10.51	67	8.42	4.02	318
180	2.69	6.47	95	3.50	6.06	94	3.77	1.62	204
181	4.00	6.70	75	3.31	6.71	73	3.61	1.45	108
183	10.94	12.48	93	13.15	12.97	49	9.58	4.92	414
218	12.46	13.30	161	12.31	13.53	188	9.25	4.53	498
229	5.01	6.32	131	4.90	6.51	121	5.54	2.82	641
251	8.21	12.95	337	8.19	12.82	402	13.68	7.04	423
279	16.51	22.49	244	19.18	21.67	241	11.28	5.28	248
281	6.44	8.12	12	7.06	9.35	14	5.57	2.48	154
282	4.22	8.47	82	4.35	9.25	93	7.10	3.04	196
283	7.49	17.20	159	6.00	17.88	180	8.77	3.73	265
285	16.22	22.03	212	16.83	21.77	223	10.47	4.83	170
288	6.00	28.30	218	7.27	25.26	299	16.28	8.43	379
290	11.98	13.63	82	11.99	15.57	41	14.04	6.75	492
370	22.17	33.71	725	21.63	34.71	737	31.18	17.49	745
373	16.51	29.11	713	16.69	29.40	744	28.97	14.92	592
375	4.69	33.50	269	1.97	39.44	248	NaN	NaN	0
377	4.83	29.20	418	5.25	36.59	331	18.98	7.43	164

Continued on the next page...

Steel plant index	ResNet-44 μ	ResNet-44 σ	ResNet-44 #	RFC μ	RFC σ	RFC #	Distribution μ	Distribution σ	Distribution #
378	45.69	46.29	688	44.18	45.87	700	58.70	33.74	777
380	17.06	45.14	557	10.86	48.90	600	48.99	20.63	252
381	40.74	40.98	860	39.74	40.60	932	53.46	31.68	889
382	43.92	47.54	656	44.48	49.12	614	55.95	31.48	734
386	3.73	35.28	270	1.40	39.43	250	NaN	NaN	0
389	46.21	46.04	694	44.86	47.24	690	55.73	33.26	844
392	40.30	41.06	873	39.31	40.68	936	51.09	30.57	905
393	40.28	43.85	677	38.51	43.70	687	54.80	30.70	717
394	44.17	44.98	677	43.72	45.81	688	58.74	33.92	781
395	42.49	46.03	684	43.71	47.44	637	55.28	31.68	760
398	3.16	26.90	237	2.78	27.13	241	19.96	7.97	127
399	7.41	12.07	86	8.73	11.91	53	9.89	4.74	461
400	6.28	22.40	538	4.77	23.10	688	21.03	9.03	325
403	-3.00	38.73	136	-5.64	38.77	143	26.42	11.55	55
404	6.33	25.20	73	7.82	29.35	82	18.51	7.66	27
405	4.25	11.09	28	5.48	15.21	23	7.69	2.93	51
407	7.50	15.89	17	5.77	18.56	16	7.08	2.70	92
415	8.36	17.47	68	8.87	13.63	52	8.59	3.56	216
421	6.10	11.72	31	6.43	9.72	24	5.30	2.22	243
424	4.85	12.65	191	4.89	12.50	183	8.62	3.90	288
426	8.94	11.48	110	8.89	12.28	116	7.46	3.70	459
427	9.78	9.33	141	8.34	8.93	153	10.54	5.42	470
433	9.42	12.98	158	10.43	10.95	109	11.39	5.76	497
435	10.66	13.31	79	13.19	13.54	70	8.65	3.61	149
436	5.94	6.07	276	6.31	6.07	237	3.01	3.92	2013
437	10.25	9.93	181	11.53	9.85	187	6.36	2.68	280
441	29.92	31.24	1022	28.10	31.77	1128	39.18	22.82	921
445	2.33	6.17	53	4.08	5.44	32	5.29	2.09	125
472	7.67	10.03	253	8.73	10.20	257	9.08	4.64	488
473	-0.86	21.40	115	0.07	16.75	97	NaN	NaN	0
474	30.14	33.44	420	30.83	34.37	431	20.62	11.89	697
475	4.35	11.99	250	4.35	13.28	202	13.04	5.39	203
477	5.28	14.58	71	4.17	15.19	73	10.09	5.07	300
478	5.67	12.11	65	3.02	15.40	45	7.03	3.54	408
479	-0.23	14.11	27	-2.55	13.52	29	8.52	3.24	34
481	8.80	12.16	124	9.66	11.92	130	10.59	4.69	364
482	23.45	21.76	417	24.34	22.94	343	22.53	13.50	798
486	6.73	13.74	337	7.07	15.43	281	13.80	6.28	377
490	2.49	18.70	167	3.57	19.75	135	7.36	3.86	326
494	14.83	14.96	56	15.49	13.38	51	10.74	5.74	565
495	43.87	45.79	681	42.80	44.85	697	56.74	32.72	781
496	7.48	23.71	24	36.17	26.23	8	17.65	8.00	181
501	43.87	45.79	681	42.80	44.85	697	56.74	32.72	781
502	11.35	18.02	376	10.16	18.63	357	17.63	8.87	536
506	9.30	13.92	235	9.37	14.19	242	9.14	3.87	250
507	7.74	9.09	126	8.58	10.08	107	8.84	4.18	358
509	7.22	10.05	164	7.21	9.36	165	6.30	3.22	388
515	6.14	10.37	221	6.67	10.41	229	7.67	3.54	392
523	6.77	7.81	395	7.11	7.56	368	9.42	5.14	517
527	3.93	9.17	201	3.49	9.59	196	8.38	3.12	72
561	18.35	20.91	241	19.30	21.25	247	20.53	11.28	607
570	6.13	11.57	46	6.85	10.88	49	6.77	2.89	276

Continued on the next page...

Steel plant index	ResNet-44			RFC			Distribution		
	μ	σ	#	μ	σ	#	μ	σ	#
581	13.46	13.87	47	13.23	17.46	44	13.33	7.29	601
592	11.16	10.88	209	10.26	10.82	205	12.23	6.00	463
601	8.66	13.34	299	8.51	13.24	305	13.81	7.52	647
602	10.33	8.51	554	10.82	9.03	505	10.61	5.85	1167
613	5.60	7.13	165	5.81	7.25	128	7.52	3.32	577
617	2.80	10.56	92	2.13	12.24	63	4.80	2.38	396
626	5.55	6.90	55	3.92	9.47	44	7.15	3.32	334
638	8.27	14.15	265	7.75	15.49	279	13.21	6.60	521
640	7.29	12.36	78	7.97	13.81	45	8.74	3.66	68
653	12.28	10.25	112	12.57	10.26	84	11.31	6.47	714
713	13.86	16.04	557	12.66	15.54	695	10.56	5.40	541
716	12.88	14.04	114	11.67	12.95	141	6.92	3.45	410
722	16.31	14.94	404	16.66	15.06	446	18.22	9.97	687
725	21.22	18.01	489	21.73	17.41	478	23.68	14.15	854
728	-3.07	39.18	140	-4.76	39.08	142	26.67	11.57	47
731	36.31	41.02	851	35.07	41.07	915	46.37	26.56	833
735	22.39	35.17	742	20.73	35.39	782	32.90	17.91	653
736	2.92	7.94	107	4.06	7.55	107	NaN	NaN	0
754	29.47	33.35	430	29.28	31.87	469	18.25	9.71	622
755	20.46	26.50	479	21.17	25.15	549	25.73	14.25	805
762	6.95	9.87	284	6.21	10.10	292	8.87	4.82	556
763	6.24	8.10	511	6.12	7.92	544	8.26	4.43	567
764	7.71	9.97	341	7.64	10.09	393	12.27	6.69	526
765	6.71	10.67	282	5.99	10.65	280	8.94	4.94	584
766	2.67	11.95	250	2.49	13.77	232	7.94	3.85	252
768	7.01	9.98	202	7.20	9.92	213	7.85	3.40	266
769	2.10	4.20	58	2.12	3.30	36	4.96	2.05	273
771	6.63	8.47	154	6.76	8.17	160	6.51	2.75	251
778	5.89	9.69	165	7.20	10.06	149	9.03	3.78	221
779	2.55	9.55	168	1.67	10.03	217	NaN	NaN	0
782	9.22	25.79	484	5.17	27.50	688	27.63	11.45	254
783	1.91	7.18	42	3.13	7.54	30	5.38	2.19	110
784	1.86	9.55	180	1.05	10.06	219	NaN	NaN	0
785	19.53	27.25	220	18.00	27.91	225	24.61	13.64	665
786	13.71	27.28	357	13.79	28.26	387	16.11	9.00	719
787	17.53	20.16	209	17.06	22.08	182	23.96	11.33	383
789	16.73	20.20	255	18.09	20.51	238	19.23	10.45	552
792	28.49	31.08	433	28.18	31.10	465	20.88	10.28	521
793	4.21	18.53	50	3.17	19.61	56	9.36	3.96	149
797	3.11	28.47	141	2.89	29.49	127	15.20	6.36	210
798	20.47	20.54	231	21.26	22.03	233	20.52	11.93	771
799	5.80	13.96	289	6.29	13.95	296	12.25	6.10	450
800	6.06	16.94	104	11.69	14.26	21	16.90	7.78	360
801	21.72	28.79	668	20.01	29.30	722	30.42	16.36	650
810	5.70	27.17	544	2.26	27.67	666	23.03	9.45	220
813	10.30	28.07	576	8.19	28.13	686	29.02	12.34	285
814	-3.64	19.11	307	2.00	17.94	210	NaN	NaN	0
815	5.47	15.10	47	4.21	18.61	35	9.72	4.72	379
821	7.45	17.27	88	14.64	19.57	23	16.73	7.77	384
823	4.35	21.95	103	8.65	23.99	78	NaN	NaN	0
824	34.80	41.22	844	33.92	41.30	887	44.18	24.82	808
826	7.44	10.19	292	6.38	10.25	287	8.81	4.75	534

Continued on the next page...

Steel plant index	ResNet-44			RFC			Distribution		
	μ	σ	#	μ	σ	#	μ	σ	#
829	4.52	5.19	36	5.77	6.22	20	NaN	NaN	0
882	35.29	36.55	816	34.24	37.23	867	47.23	27.79	864
883	6.30	25.98	506	8.19	27.60	471	25.69	9.76	85
886	6.37	26.28	368	3.80	33.21	288	17.46	8.03	253
894	37.46	41.16	863	36.37	41.24	934	46.46	27.27	874
899	36.73	41.21	860	35.70	40.70	924	46.62	27.64	872
904	7.31	9.62	345	6.63	9.72	362	8.39	4.01	346
906	10.26	10.35	222	10.35	8.98	186	8.70	4.74	599
931	6.01	14.17	275	7.17	5.95	268	NaN	NaN	0
932	16.22	26.80	153	14.71	28.05	167	16.03	8.87	670
933	1.65	31.02	462	-1.72	39.52	370	24.98	9.91	78
937	16.28	29.22	175	15.54	30.32	167	16.32	8.96	625
948	15.82	14.90	121	17.91	15.23	113	12.06	6.03	405
949	1.68	15.40	92	0.68	18.72	74	9.02	3.63	190
950	5.55	14.25	138	6.42	15.53	105	11.18	4.88	258
987	15.72	13.48	155	15.90	12.41	148	13.04	7.79	694
991	2.21	5.37	33	4.27	6.89	13	5.24	2.29	307
1005	41.79	44.31	668	39.91	44.36	690	54.99	31.45	760
1009	12.70	12.46	399	12.74	12.46	402	12.58	6.94	707
1023	8.81	30.74	70	8.94	32.61	78	10.66	5.41	445
1031	9.97	9.88	183	9.78	9.59	202	9.71	4.42	351
1059	4.28	6.13	156	4.82	5.95	139	5.91	3.06	574
1061	8.98	13.77	202	8.34	15.16	195	13.94	6.82	495
1063	22.11	18.00	510	22.38	17.88	495	24.73	14.52	828
1064	21.99	17.90	513	22.23	17.89	498	24.43	14.45	843
1070	9.17	11.23	36	7.96	10.83	35	9.11	4.08	313

C

Time series of European steel plants

This appendix displays the time series of emission estimates of the different methods for the 8 steel plants in the European Union from Table A.1. The different steel plants are listed in order of their steel plant index number from the table.

#18: AG der Dillinger Hüttenwerke Dillingen steel plant

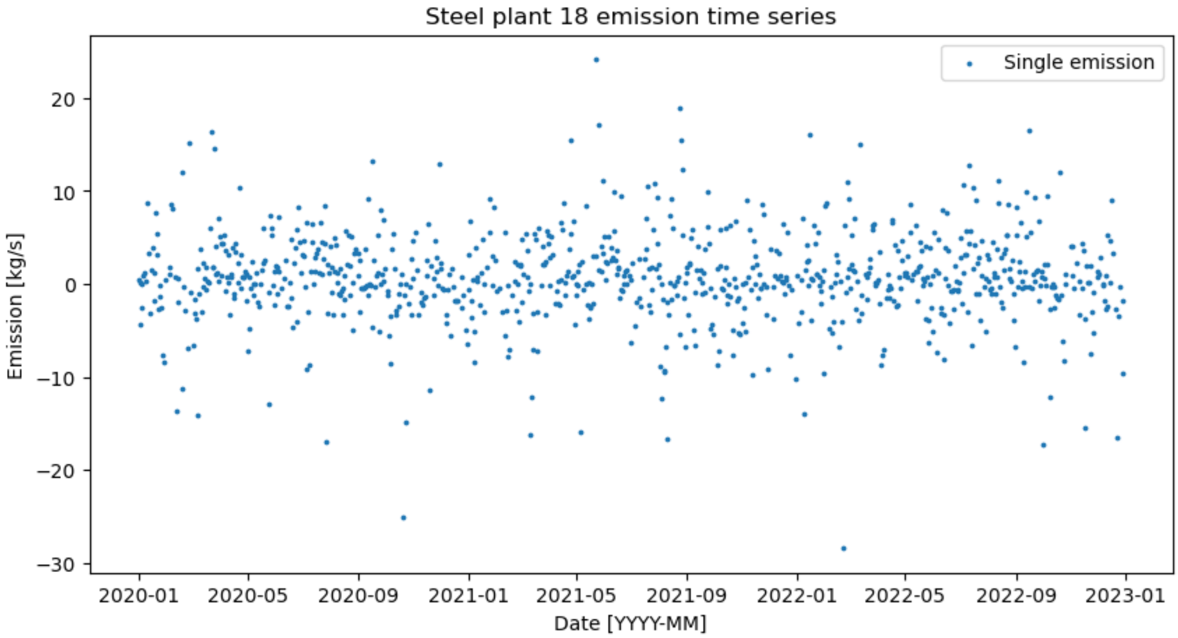

Figure C.1: Time series of the emission estimates of the AG der Dillinger Hüttenwerke Dillingen steel plant

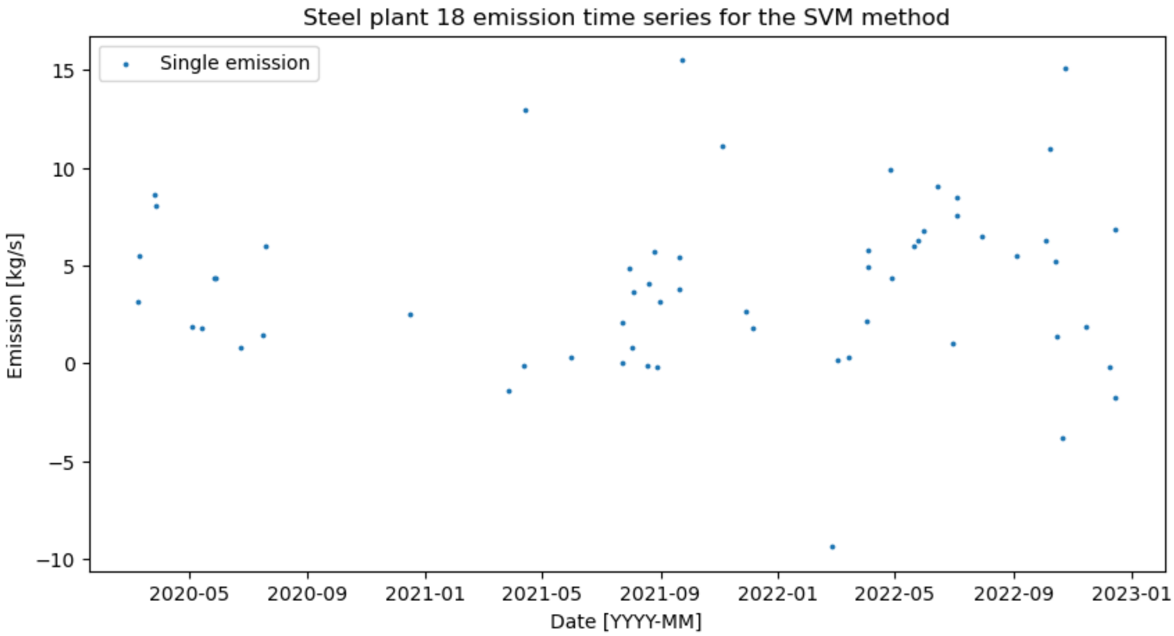

Figure C.2: Time series of the emission estimates of the AG der Dillinger Hüttenwerke Dillingen steel plant for the SVM method

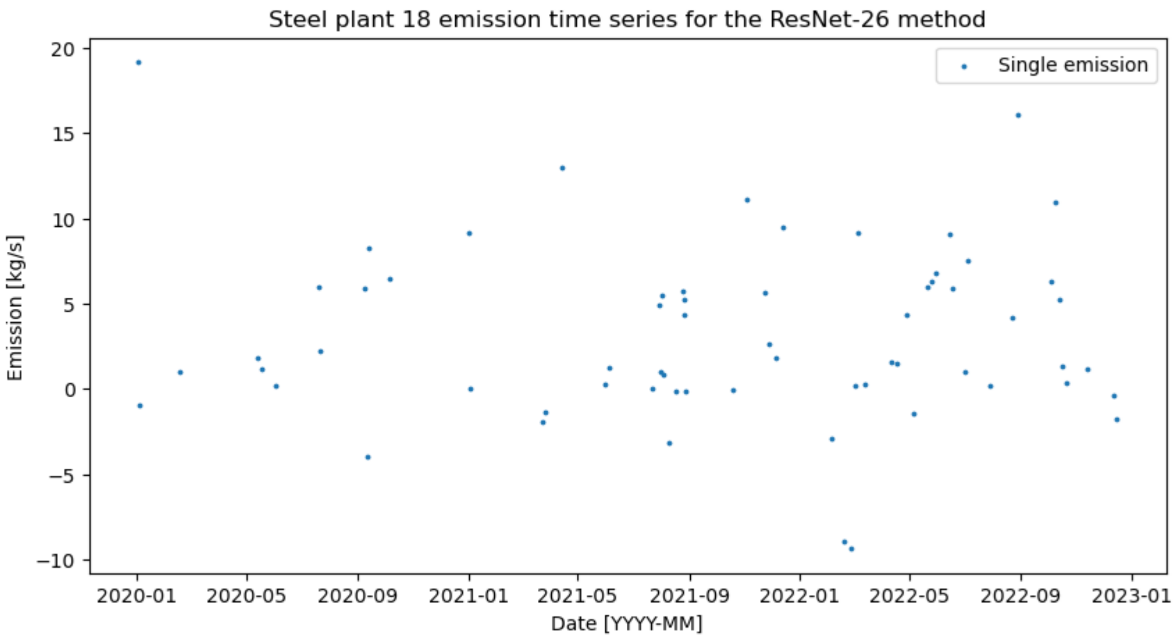

Figure C.3: Time series of the emission estimates of the AG der Dillinger Hüttenwerke Dillingen steel plant for the ResNet-26 method

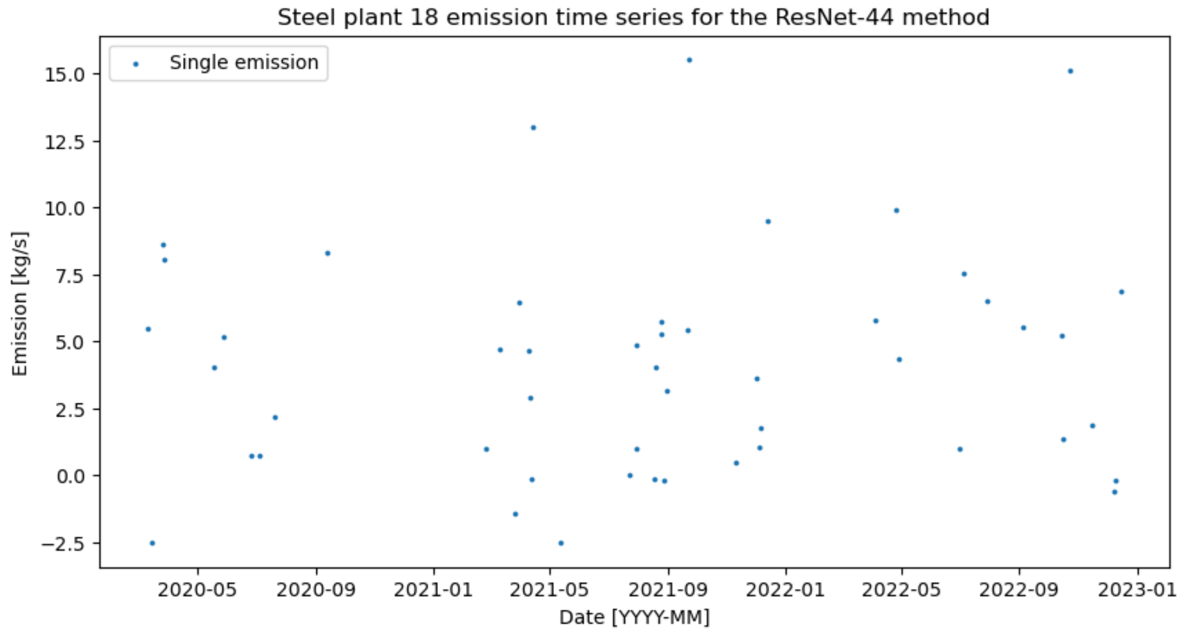

Figure C.4: Time series of the emission estimates of the AG der Dillinger Hüttenwerke Dillingen steel plant for the ResNet-44 method

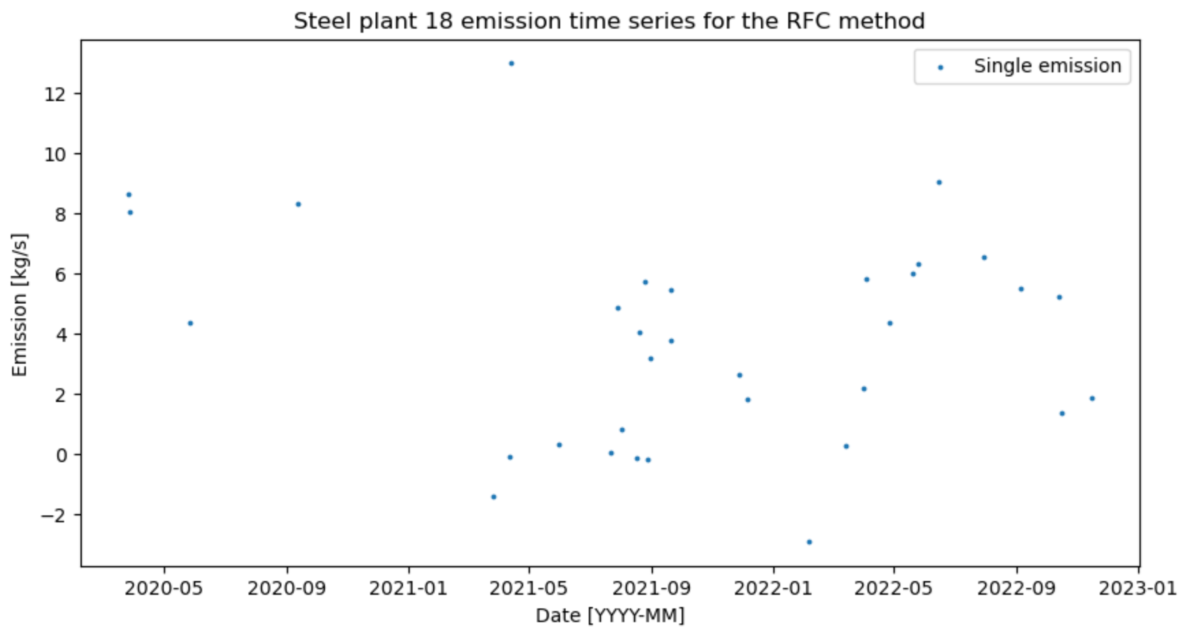

Figure C.5: Time series of the emission estimates of the AG der Dillinger Hüttenwerke Dillingen steel plant for the RFC method

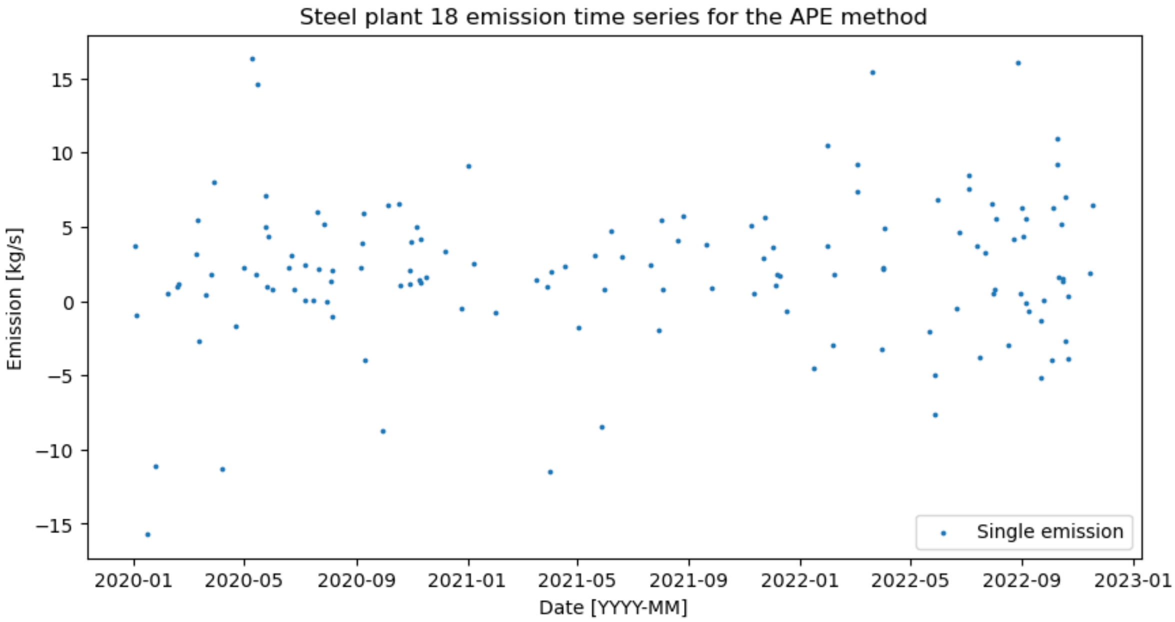

Figure C.6: Time series of the emission estimates of the AG der Dillinger Hüttenwerke Dillingen steel plant for the APE method

#64: ArcelorMittal Asturias (Gijón) steel plant

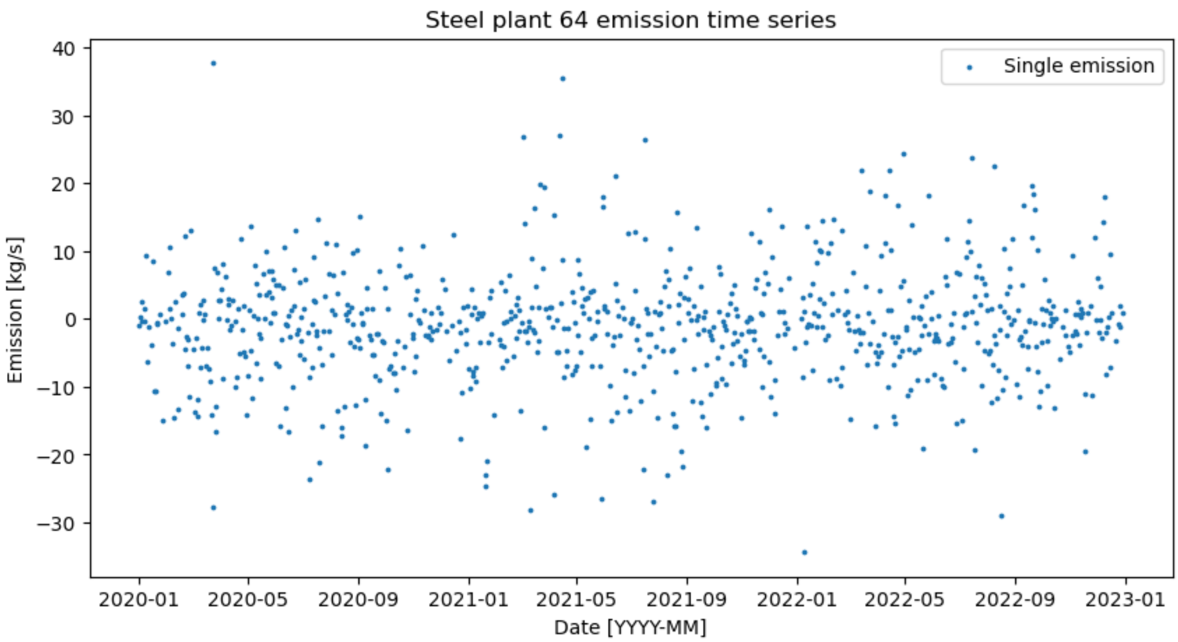

Figure C.7: Time series of the emission estimates of the ArcelorMittal Asturias (Gijón) steel plant

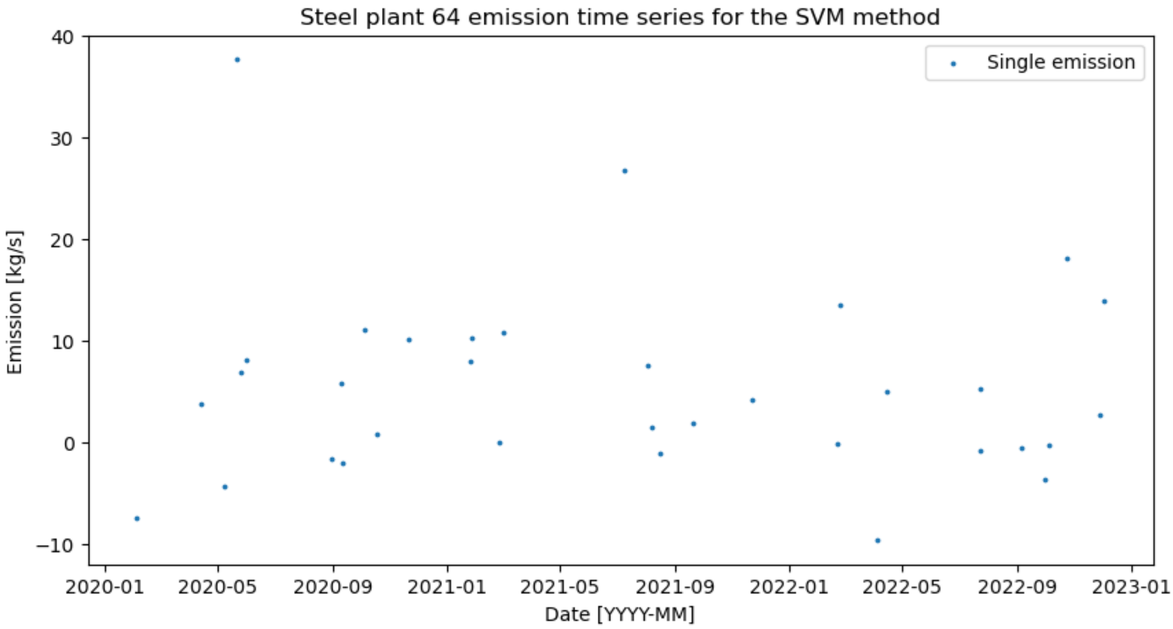

Figure C.8: Time series of the emission estimates of the ArcelorMittal Asturias (Gijón) steel plant for the SVM method

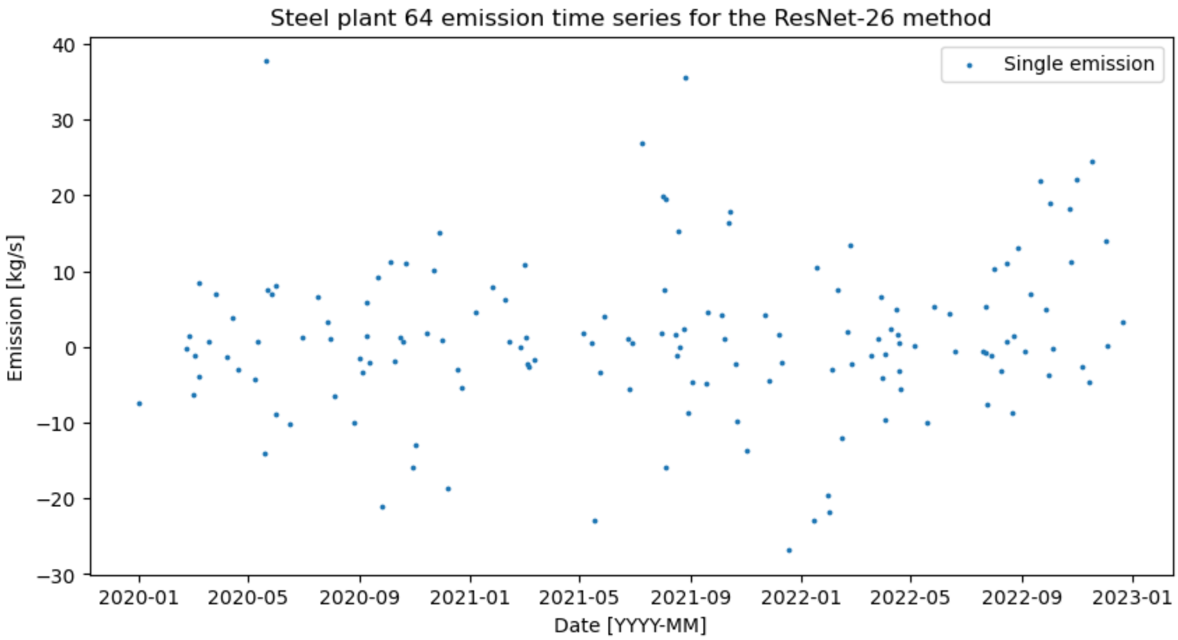

Figure C.9: Time series of the emission estimates of the ArcelorMittal Asturias (Gijón) steel plant for the ResNet-26 method

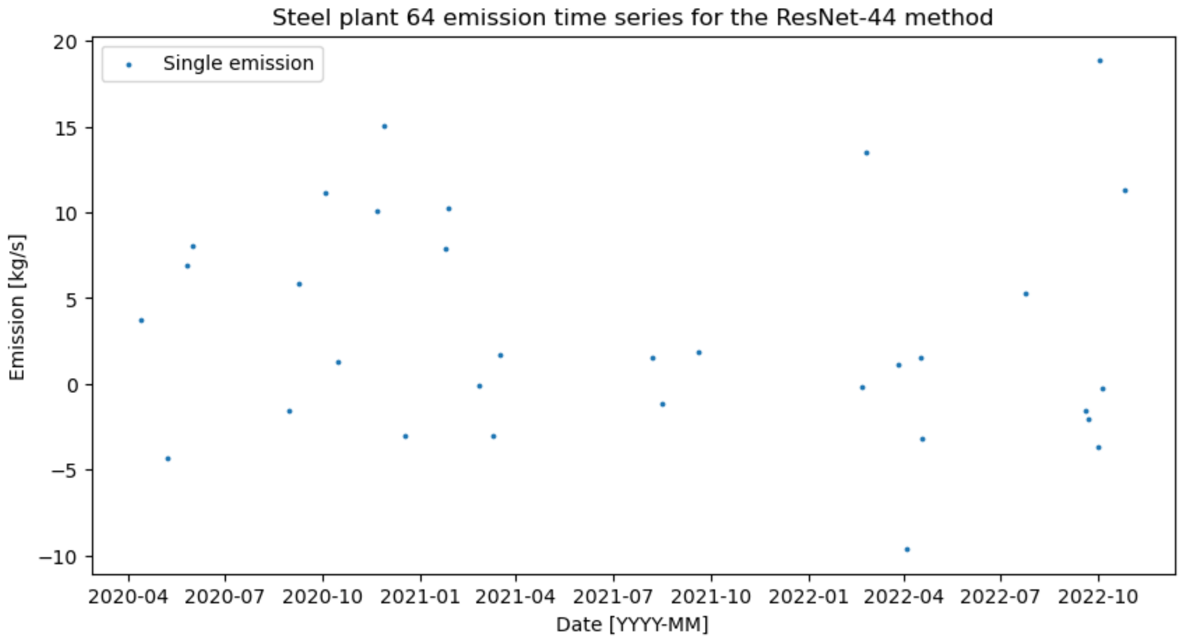

Figure C.10: Time series of the emission estimates of the ArcelorMittal Asturias (Gijón) steel plant for the ResNet-44 method

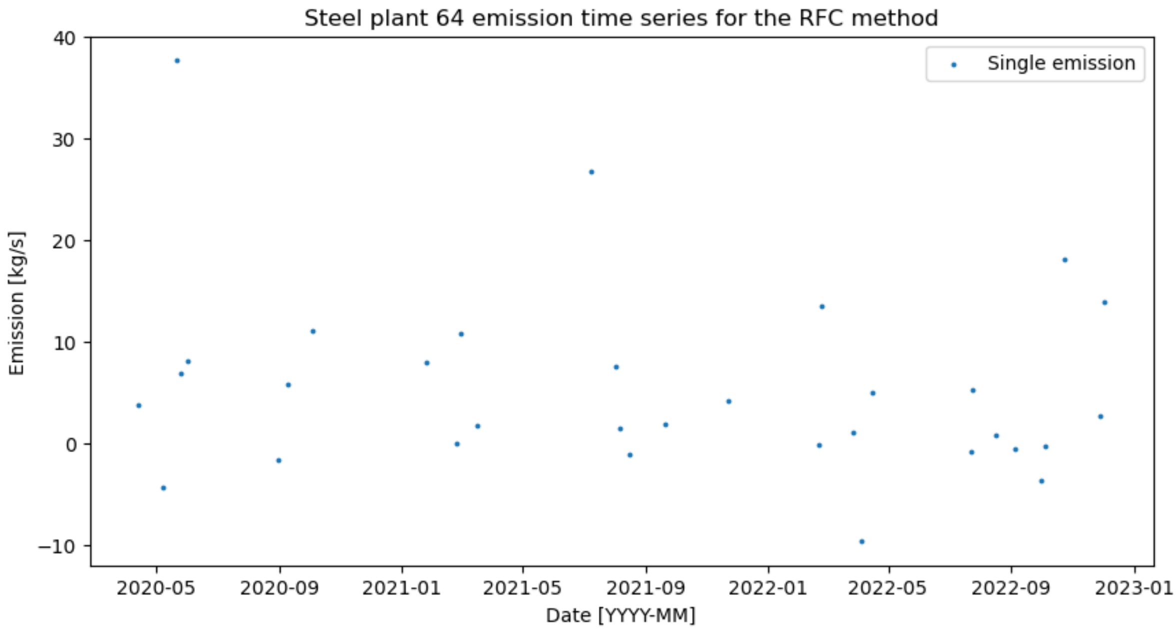

Figure C.11: Time series of the emission estimates of the ArcelorMittal Asturias (Gijón) steel plant for the RFC method

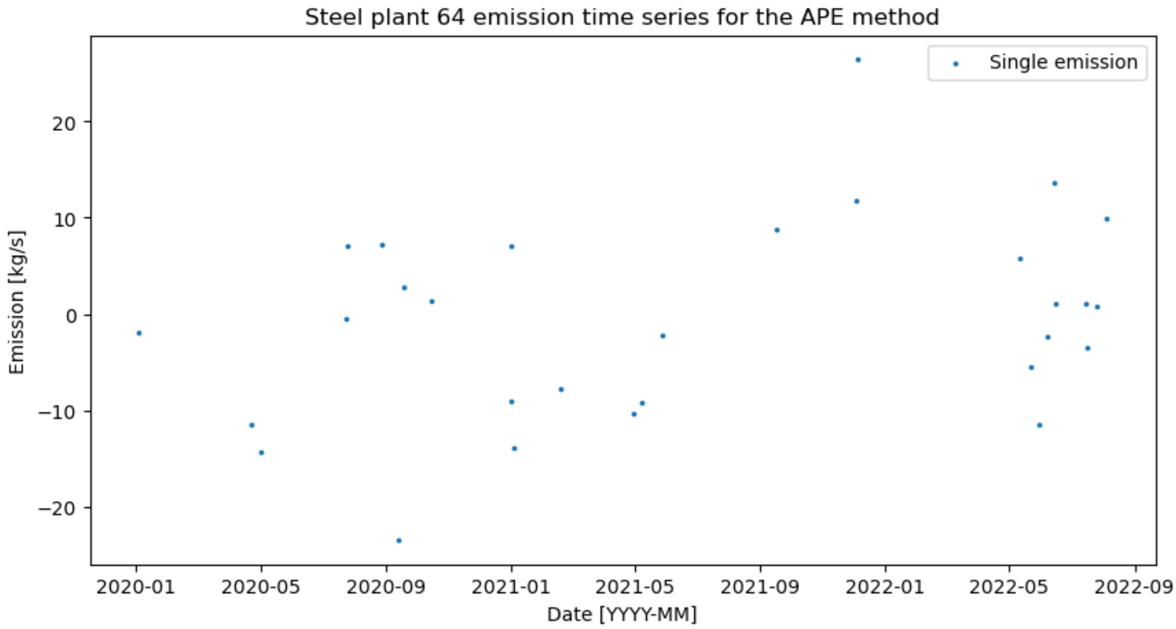

Figure C.12: Time series of the emission estimates of the ArcelorMittal Asturias (Gijón) steel plant for the APE method

#69: ArcelorMittal Duisburg steel plant

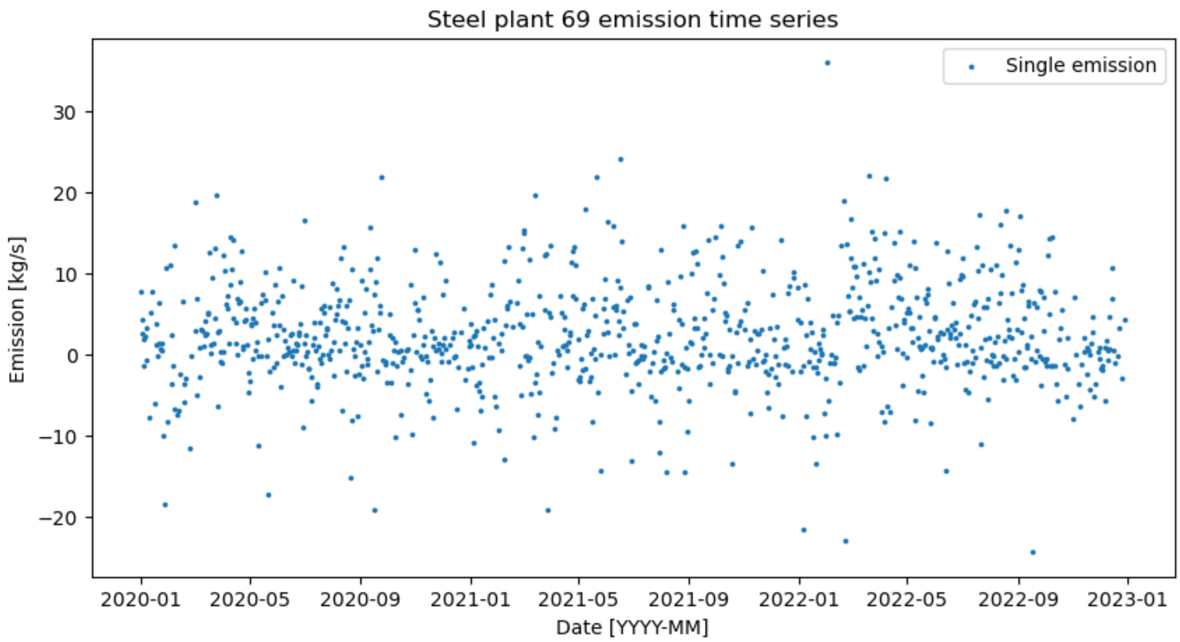

Figure C.13: Time series of the emission estimates of the ArcelorMittal Duisburg steel plant

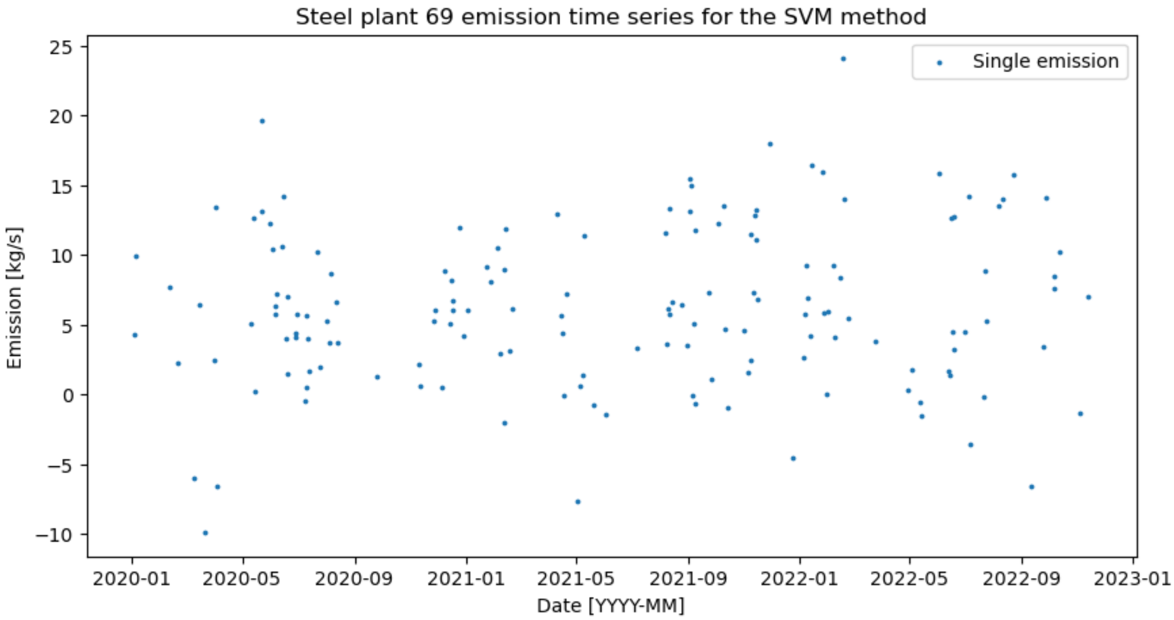

Figure C.14: Time series of the emission estimates of the ArcelorMittal Duisburg steel plant for the SVM method

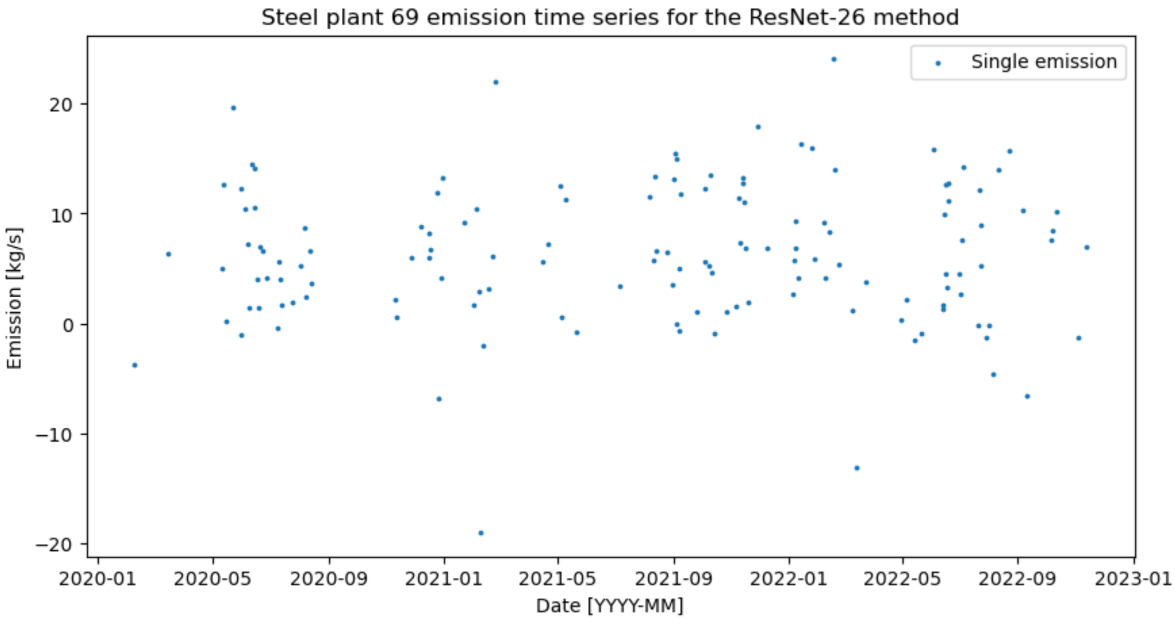

Figure C.15: Time series of the emission estimates of the ArcelorMittal Duisburg steel plant for the ResNet-26 method

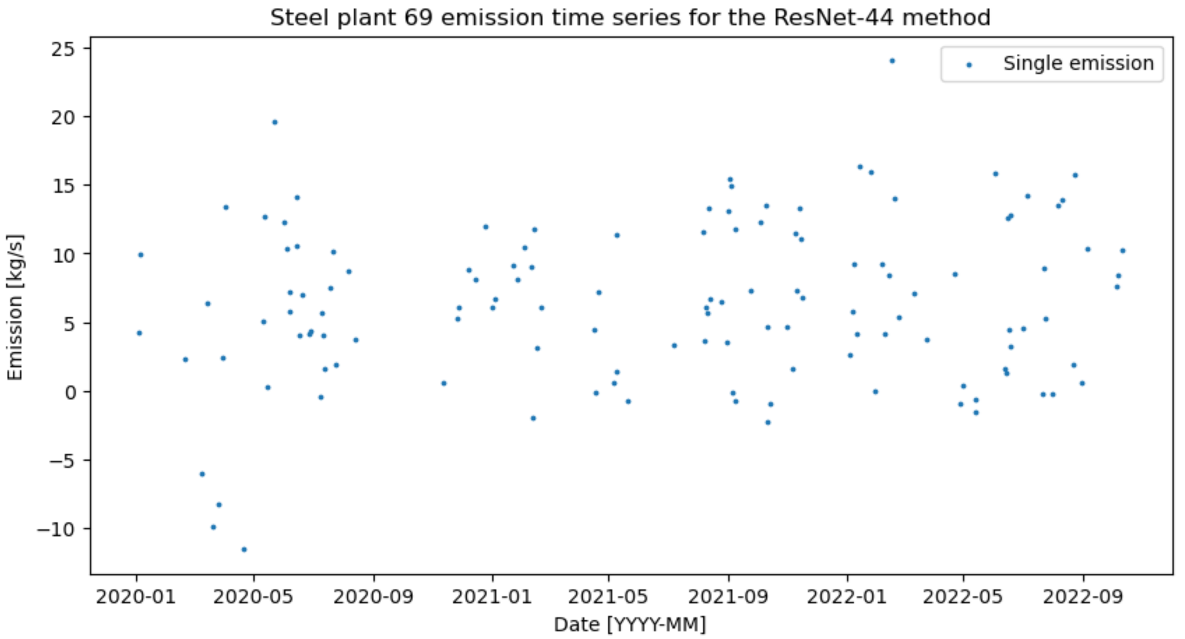

Figure C.16: Time series of the emission estimates of the ArcelorMittal Duisburg steel plant for the ResNet-44 method

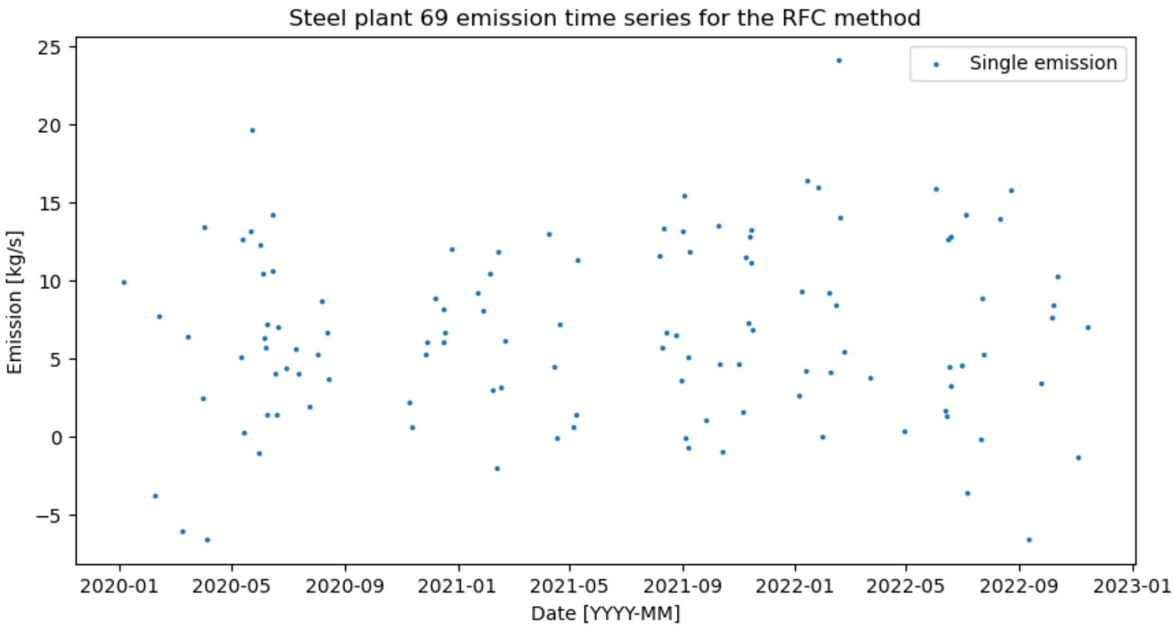

Figure C.17: Time series of the emission estimates of the ArcelorMittal Duisburg steel plant for the RFC method

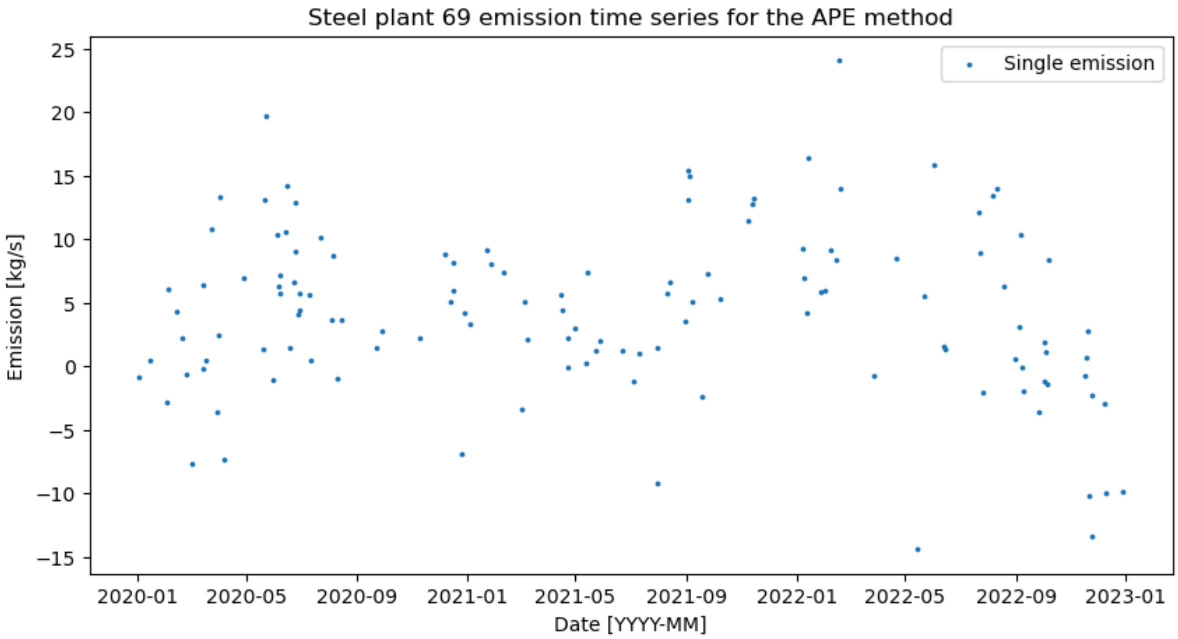

Figure C.18: Time series of the emission estimates of the ArcelorMittal Duisburg steel plant for the APE method

#71: ArcelorMittal Dąbrowa Górnicza steel plant

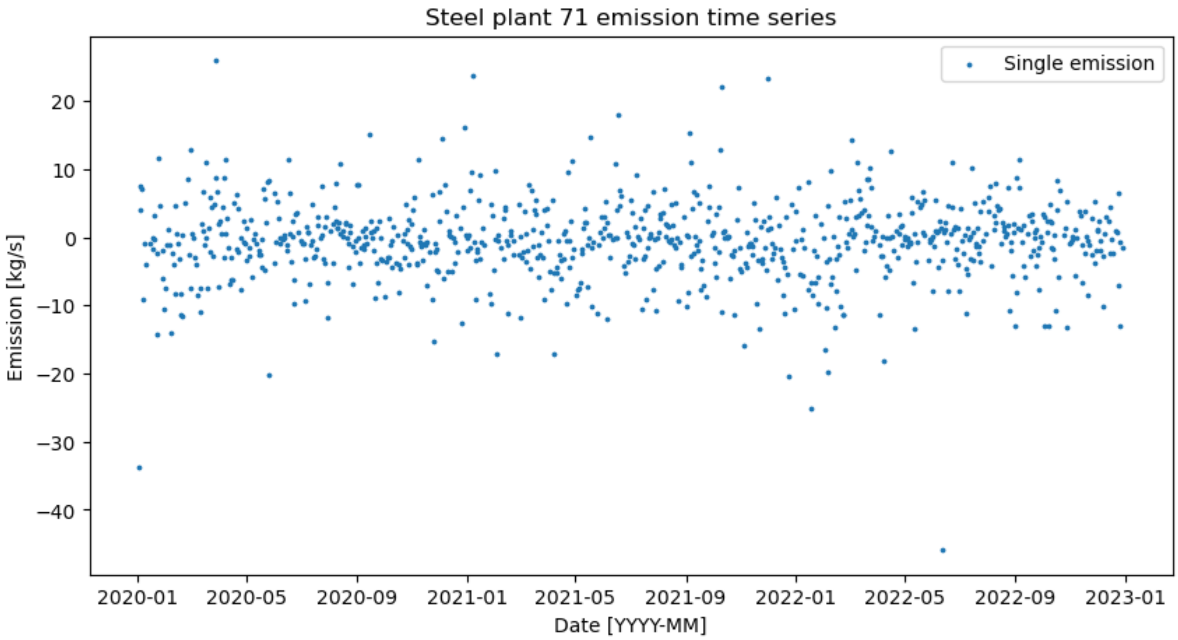

Figure C.19: Time series of the emission estimates of the ArcelorMittal Dąbrowa Górnicza steel plant

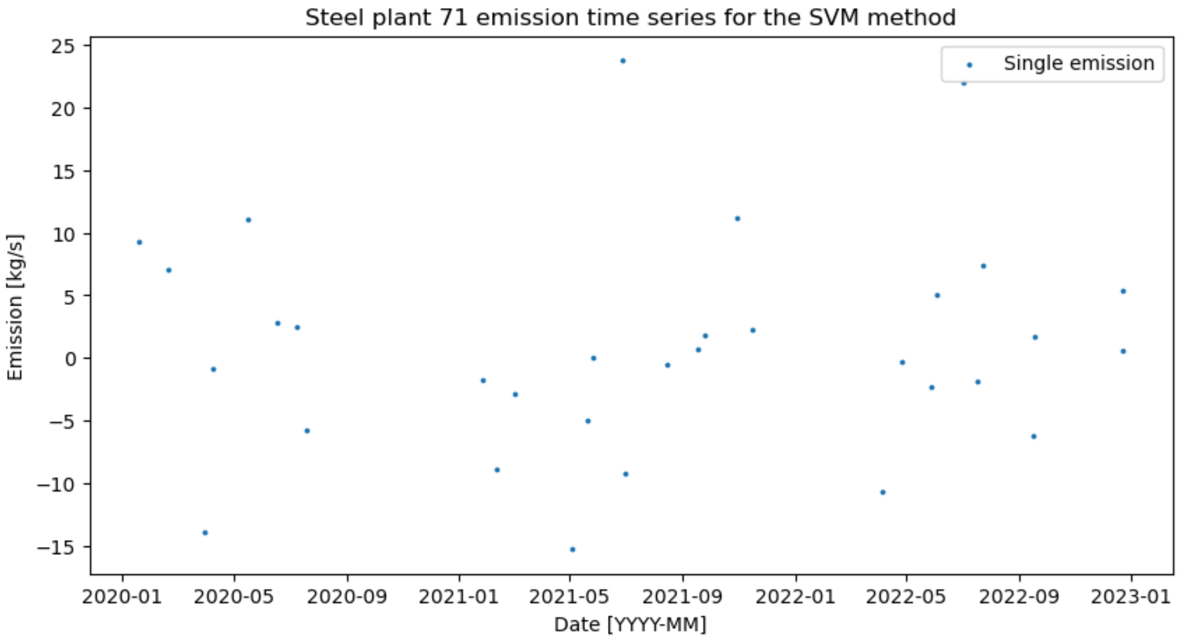

Figure C.20: Time series of the emission estimates of the ArcelorMittal Dąbrowa Górnicza steel plant for the SVM method

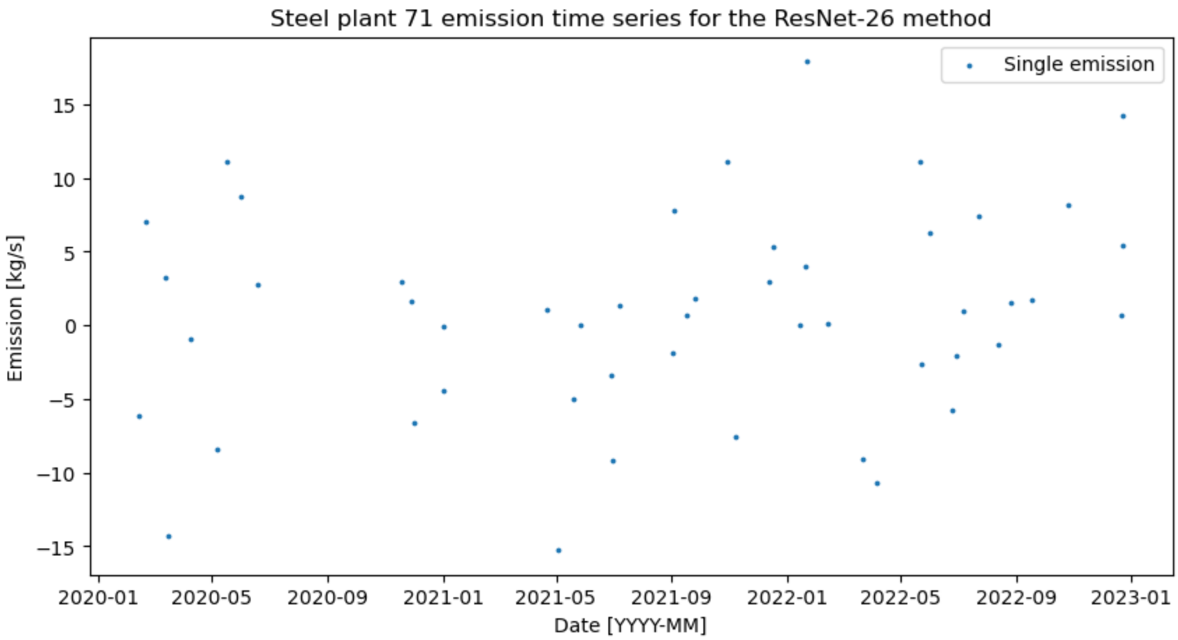

Figure C.21: Time series of the emission estimates of the ArcelorMittal Dąbrowa Górnicza steel plant for the ResNet-26 method

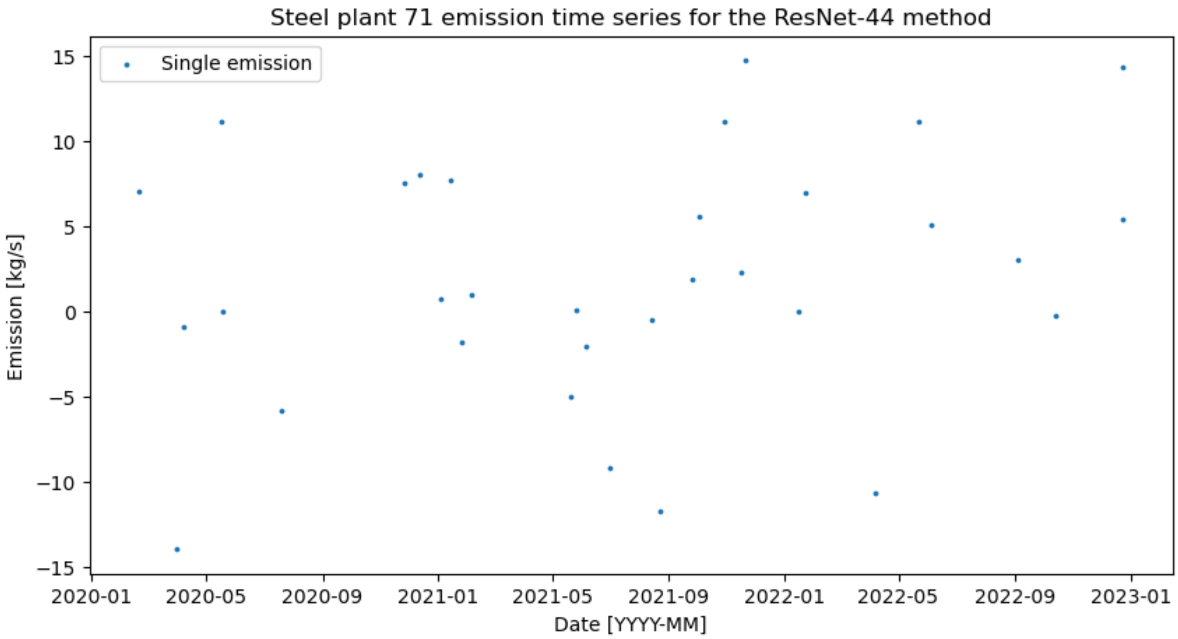

Figure C.22: Time series of the emission estimates of the ArcelorMittal Dąbrowa Górnicza steel plant for the ResNet-44 method

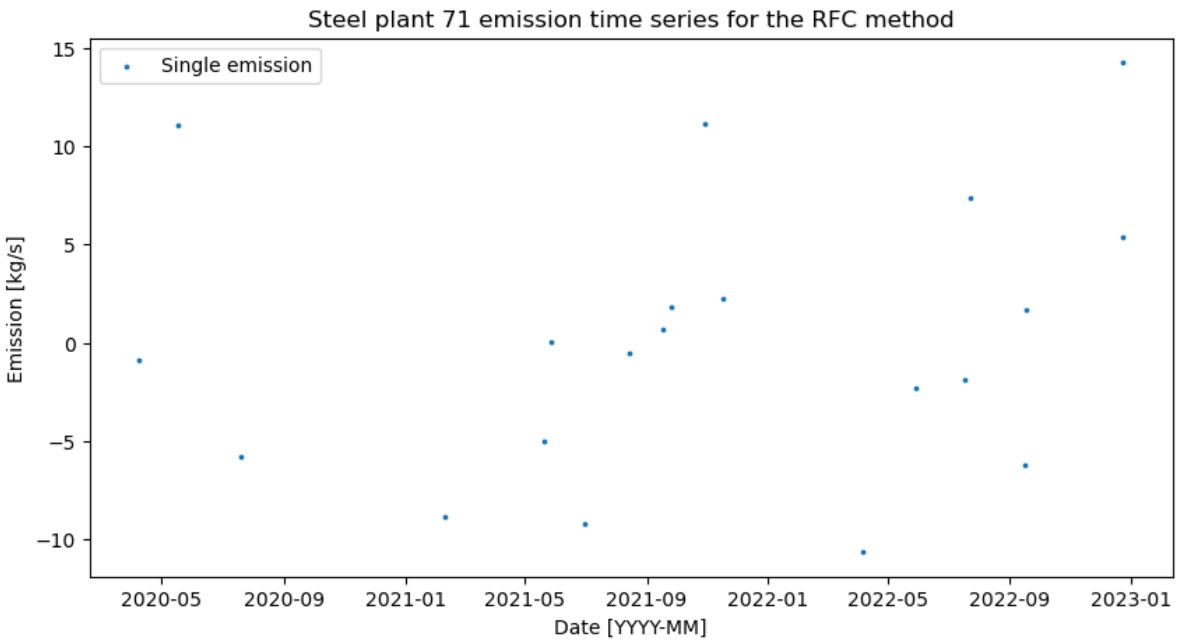

Figure C.23: Time series of the emission estimates of the ArcelorMittal Dąbrowa Górnicza steel plant for the RFC method

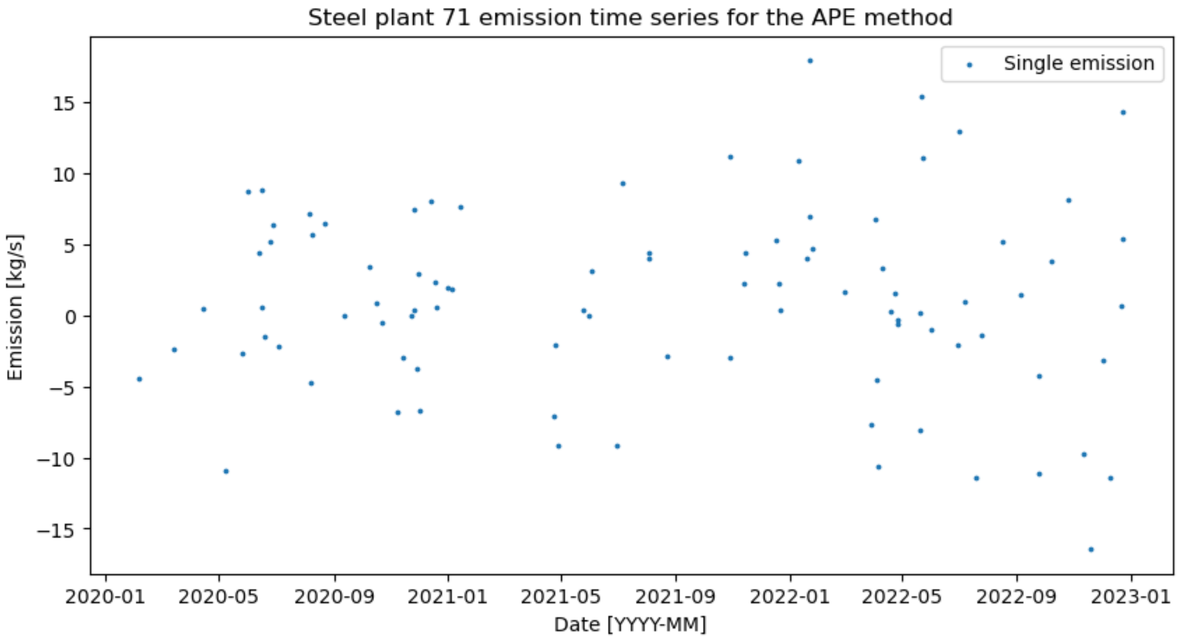

Figure C.24: Time series of the emission estimates of the ArcelorMittal Dąbrowa Górnicza steel plant for the APE method

#436: Hüttenwerke Krupp Mannesmann (HKM) steel plant

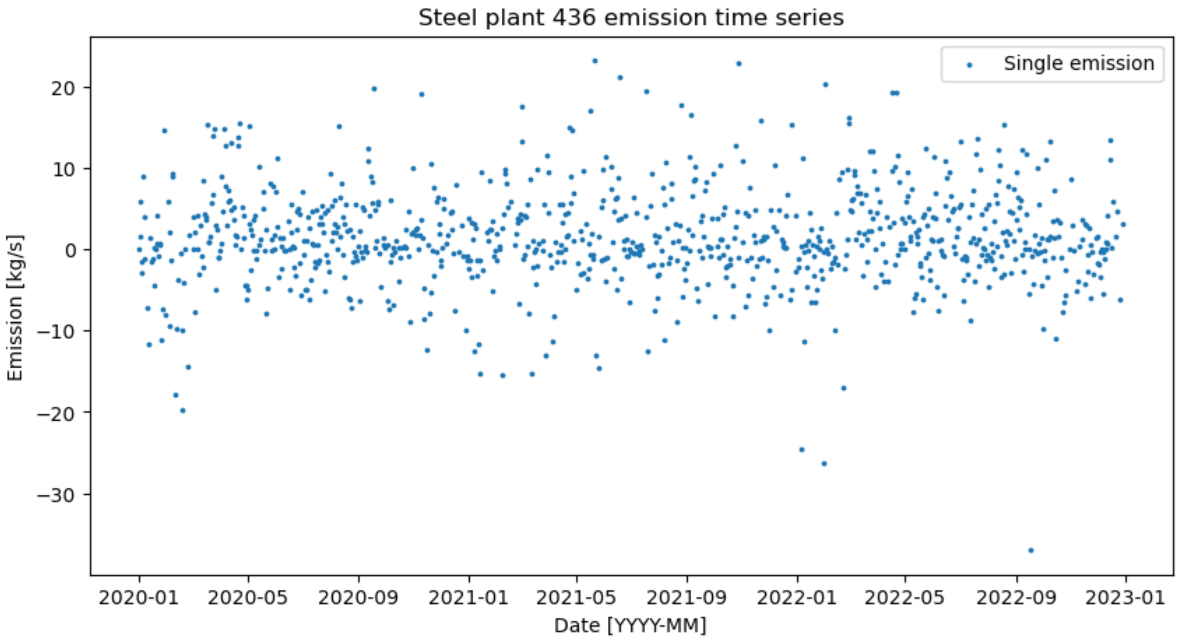

Figure C.25: Time series of the emission estimates of the Hüttenwerke Krupp Mannesmann (HKM) steel plant

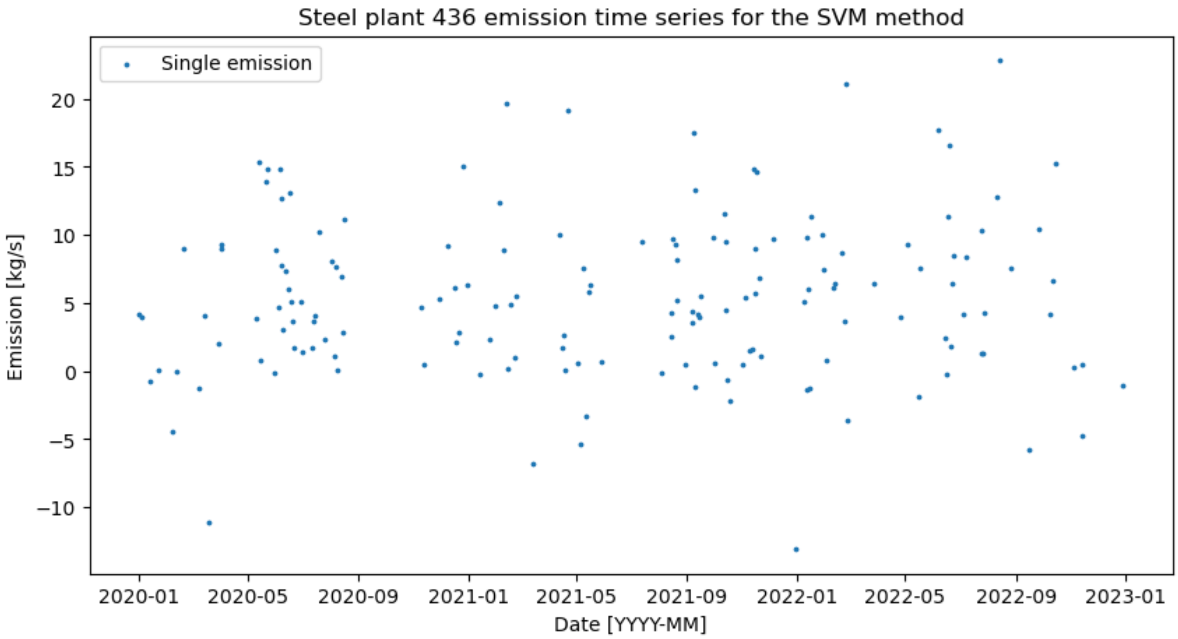

Figure C.26: Time series of the emission estimates of the Hüttenwerke Krupp Mannesmann (HKM) steel plant for the SVM method

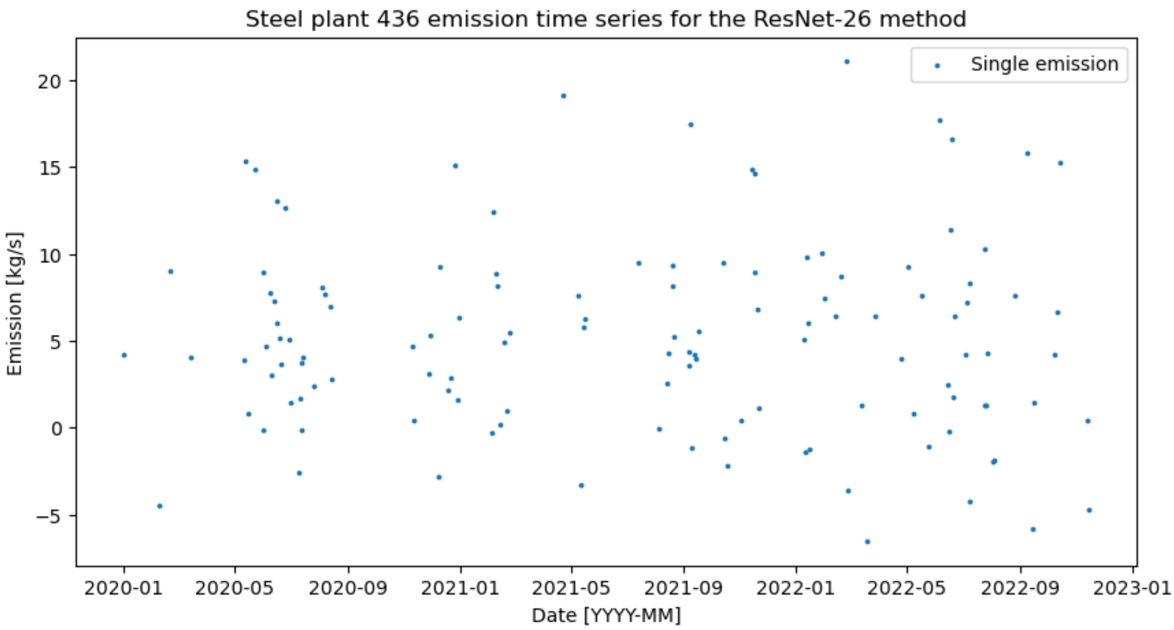

Figure C.27: Time series of the emission estimates of the Hüttenwerke Krupp Mannesmann (HKM) steel plant for the ResNet-26 method

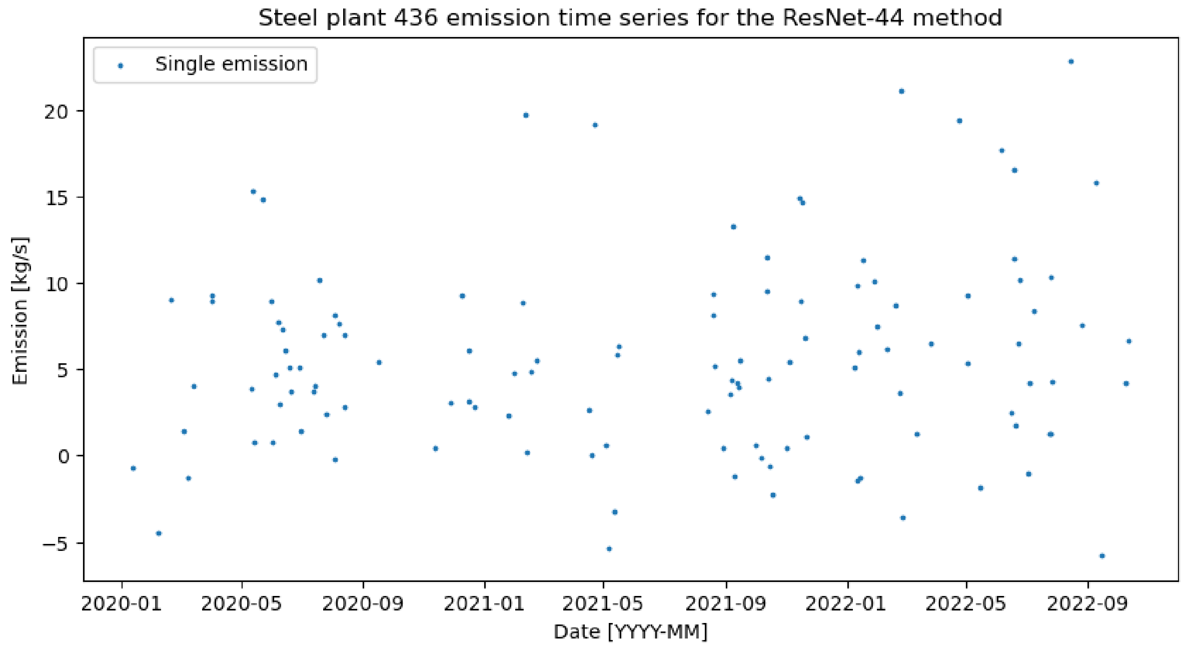

Figure C.28: Time series of the emission estimates of the Hüttenwerke Krupp Mannesmann (HKM) steel plant for the ResNet-44 method

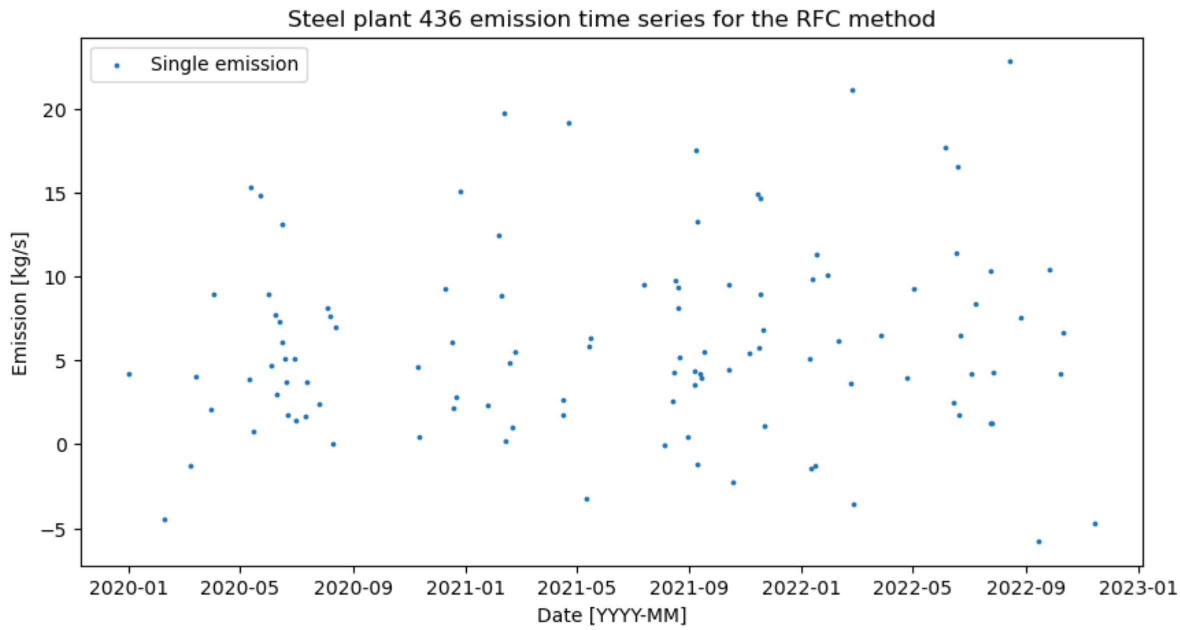

Figure C.29: Time series of the emission estimates of the Hüttenwerke Krupp Mannesmann (HKM) steel plant for the RFC method

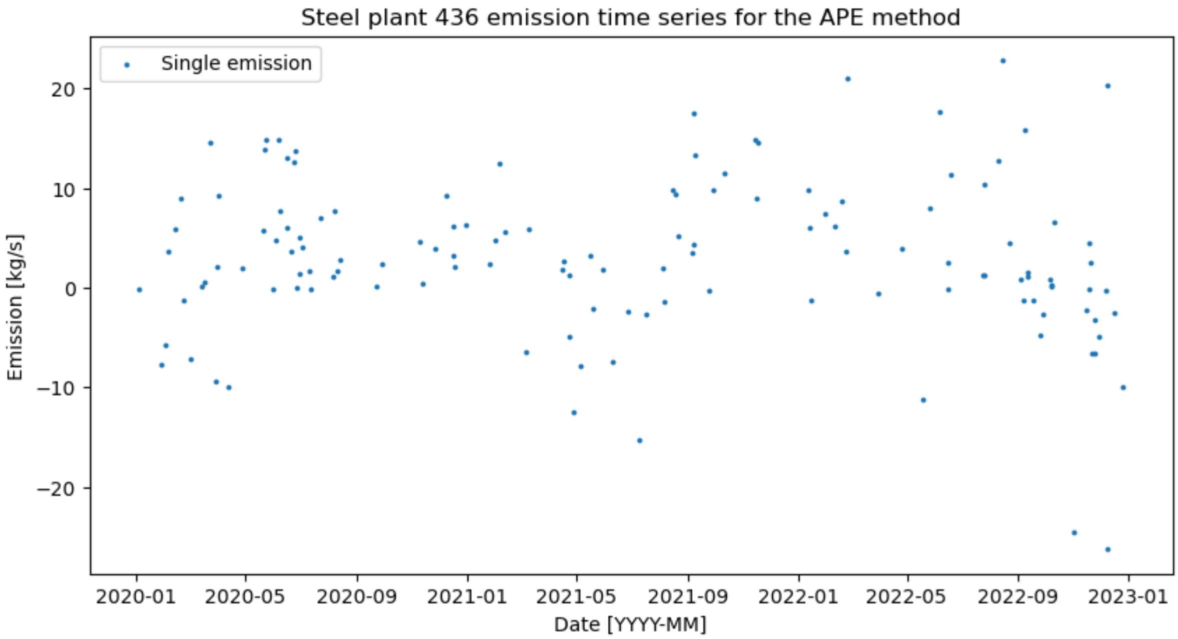

Figure C.30: Time series of the emission estimates of the Hüttenwerke Krupp Mannesmann (HKM) steel plant for the APE method

#769: Salzgitter Flachstahl steel plant

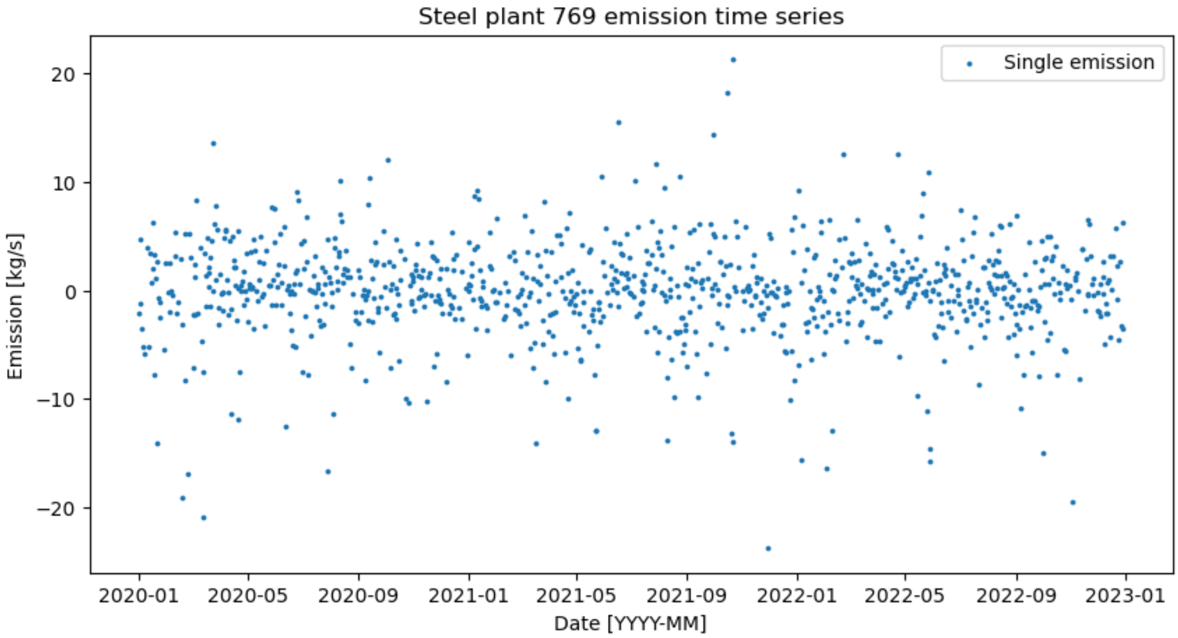

Figure C.31: Time series of the emission estimates of the Salzgitter Flachstahl steel plant

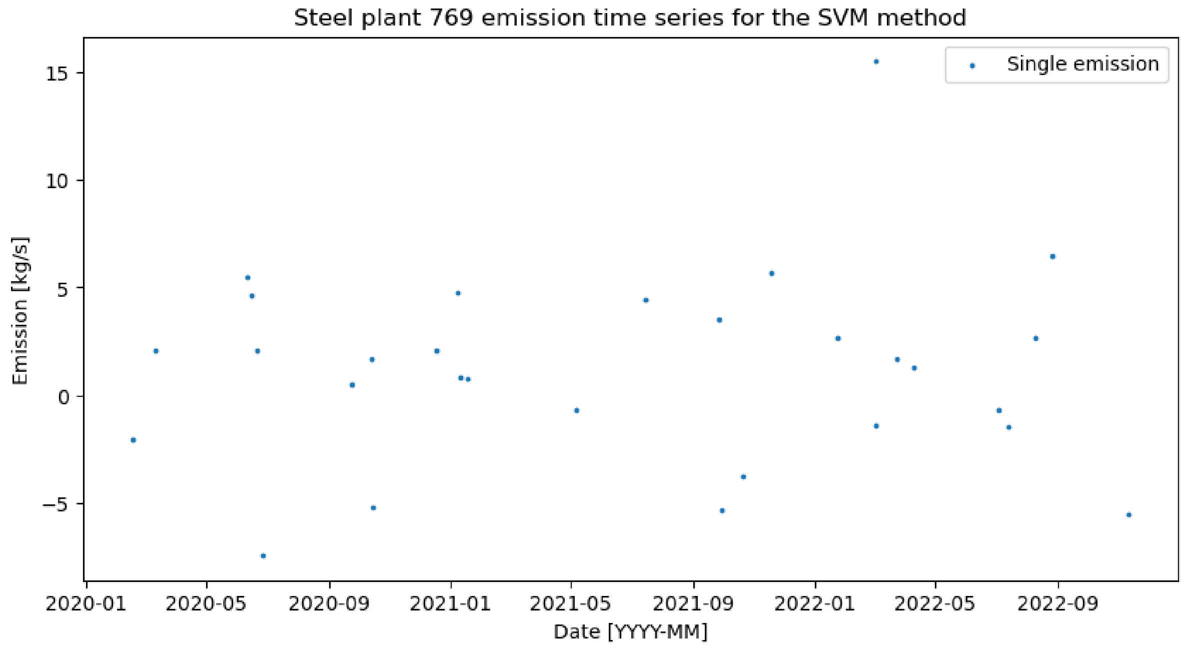

Figure C.32: Time series of the emission estimates of the Salzgitter Flachstahl steel plant for the SVM method

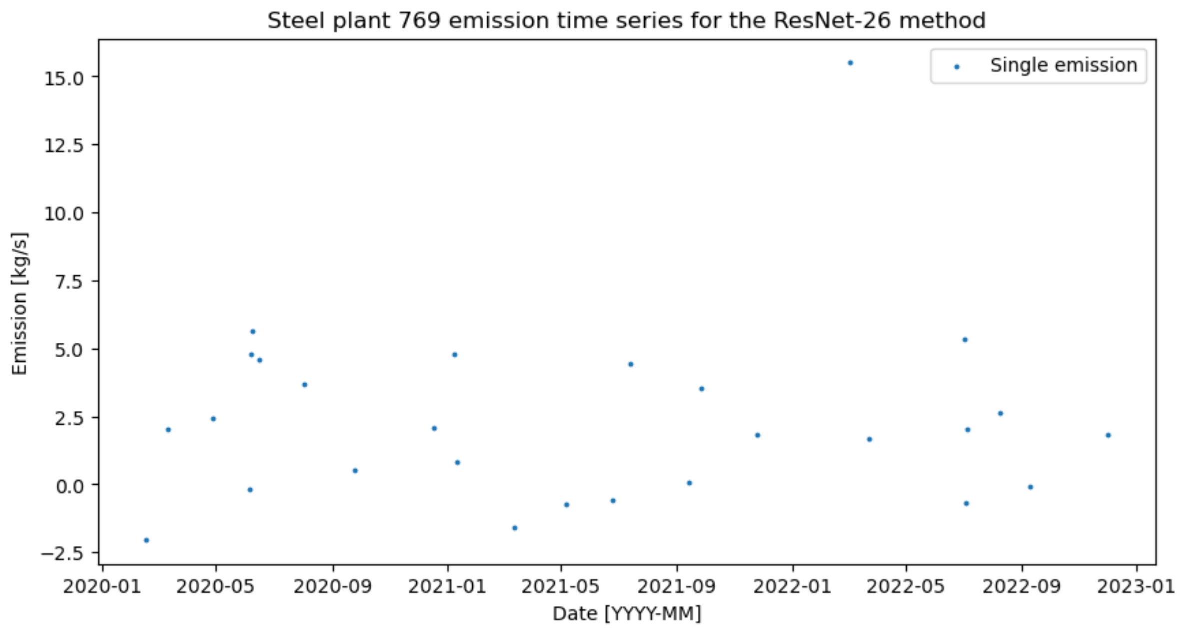

Figure C.33: Time series of the emission estimates of the Salzgitter Flachstahl steel plant for the ResNet-26 method

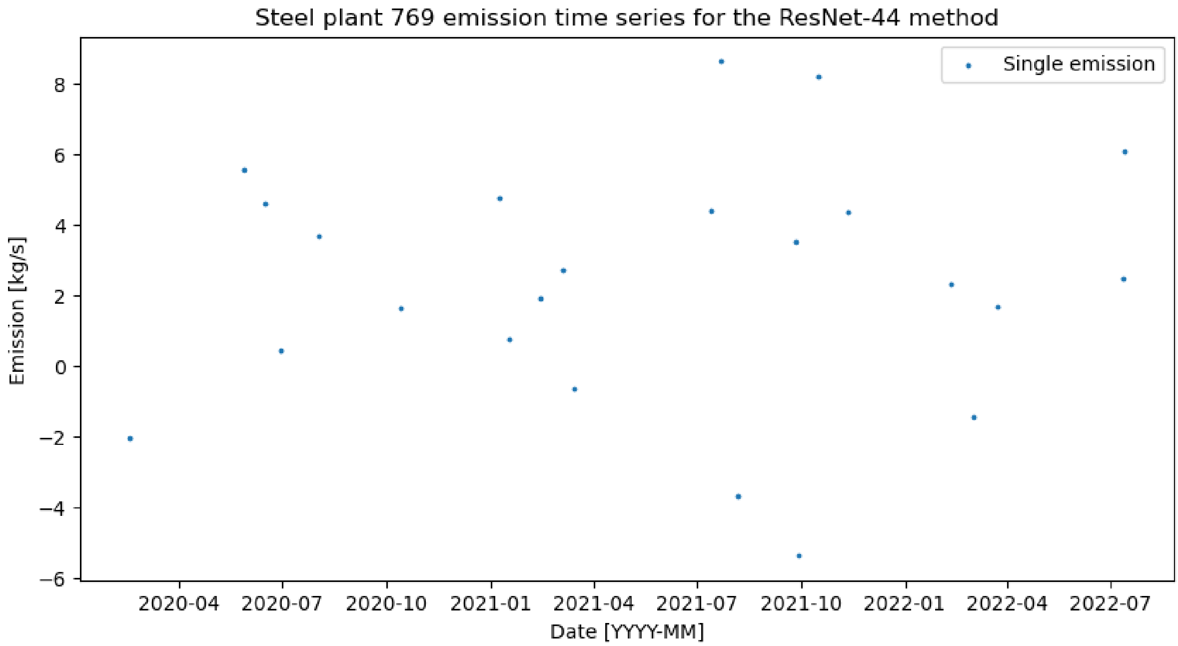

Figure C.34: Time series of the emission estimates of the Salzgitter Flachstahl steel plant for the ResNet-44 method

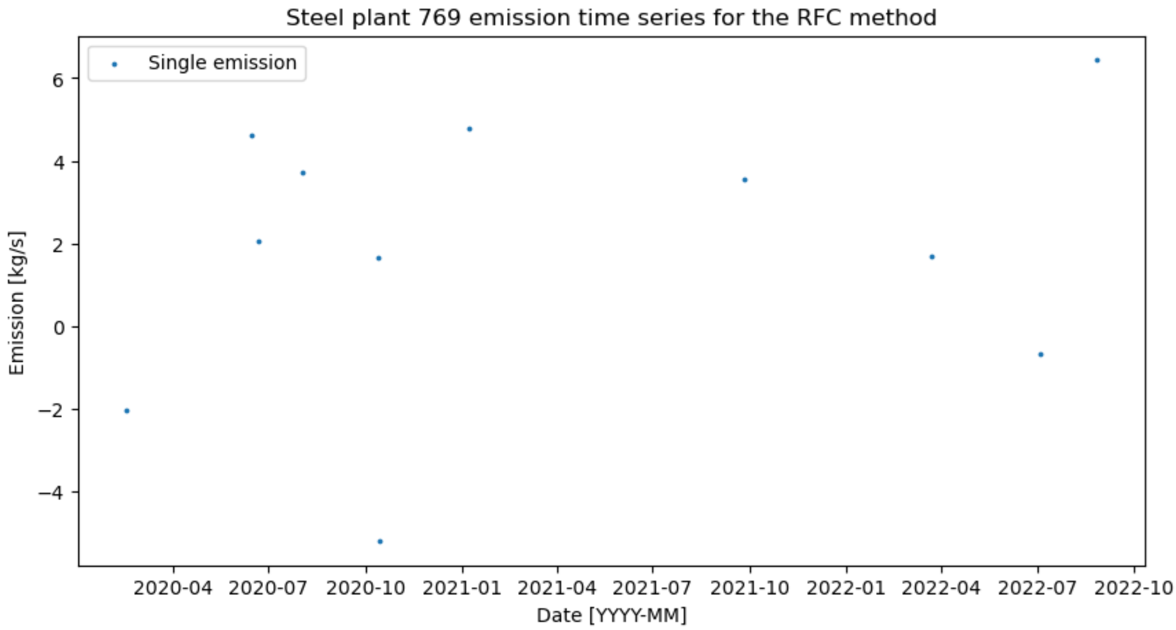

Figure C.35: Time series of the emission estimates of the Salzgitter Flachstahl steel plant for the RFC method

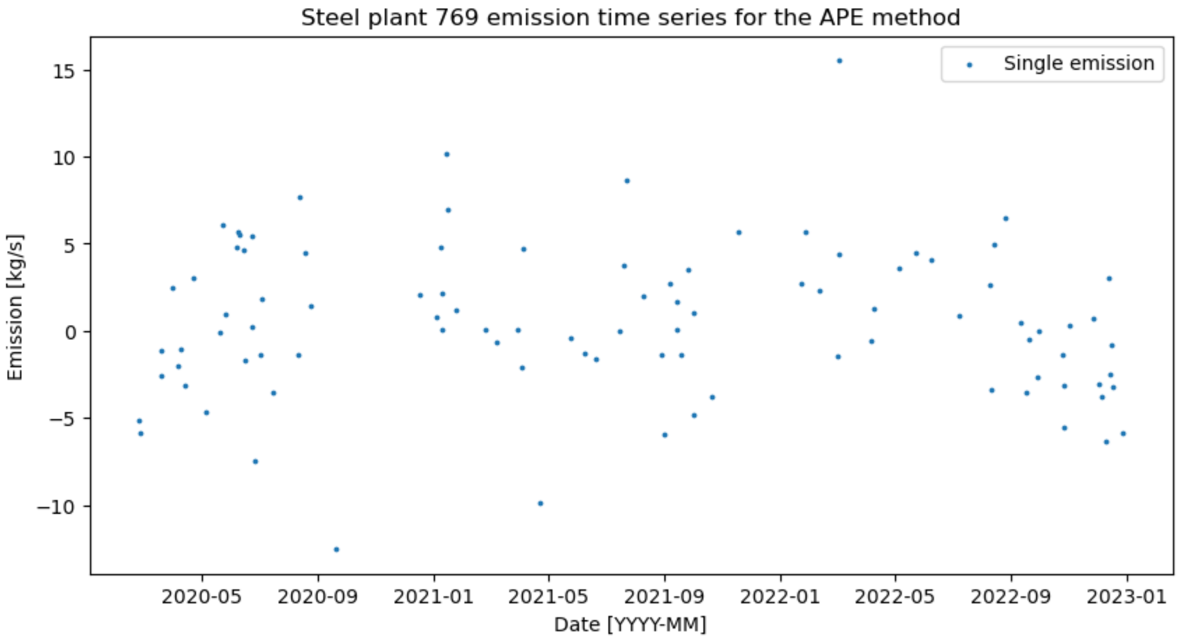

Figure C.36: Time series of the emission estimates of the Salzgitter Flachstahl steel plant for the APE method

#931: ThyssenKrupp Steel Duisburg steel plant

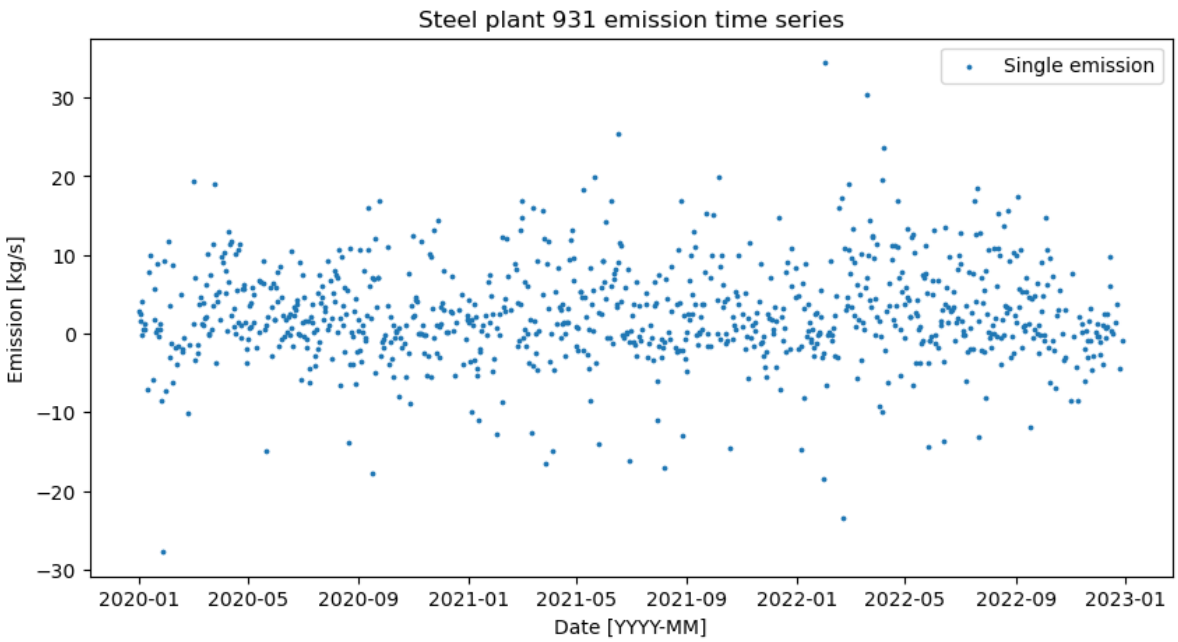

Figure C.37: Time series of the emission estimates of the ThyssenKrupp Steel Duisburg steel plant

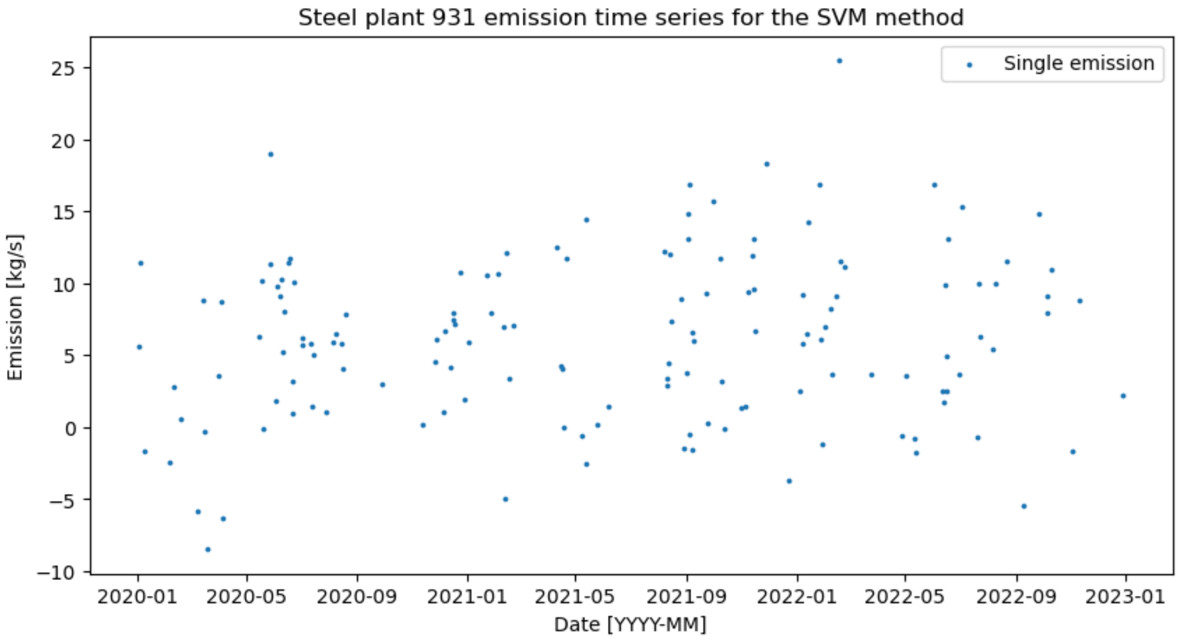

Figure C.38: Time series of the emission estimates of the ThyssenKrupp Steel Duisburg steel plant for the SVM method

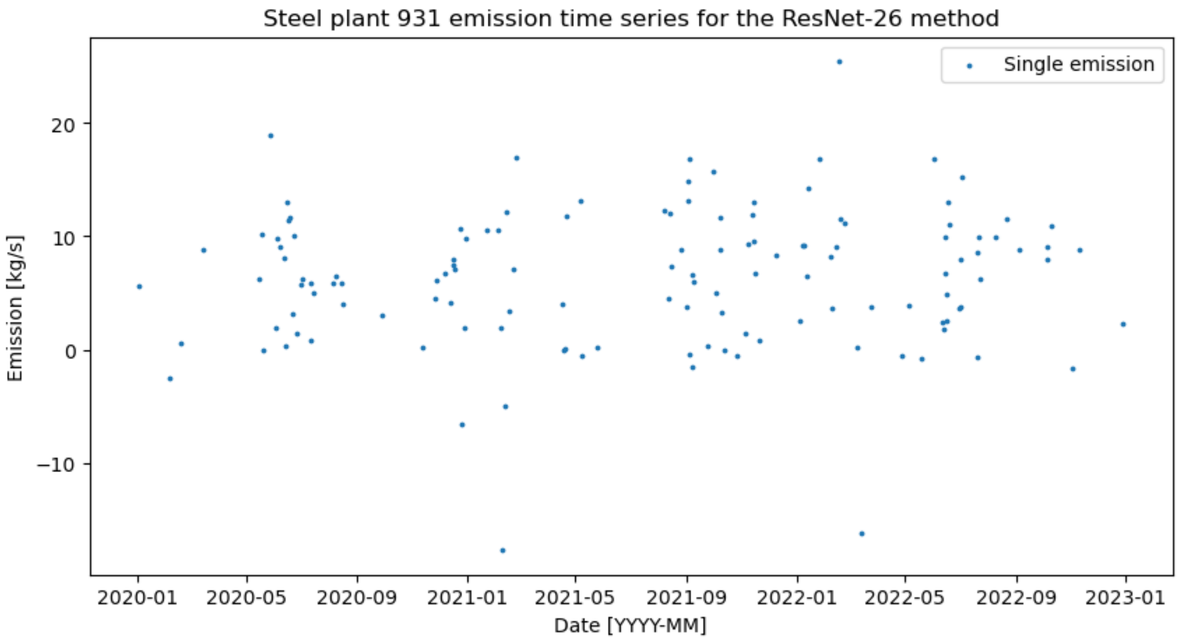

Figure C.39: Time series of the emission estimates of the ThyssenKrupp Steel Duisburg steel plant for the ResNet-26 method

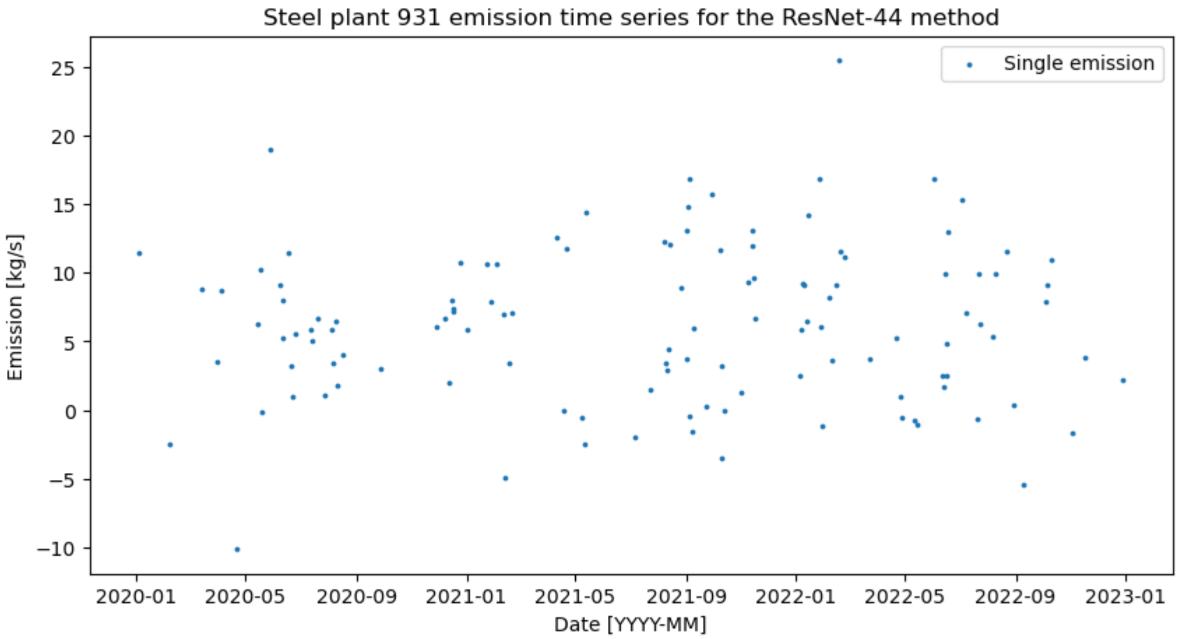

Figure C.40: Time series of the emission estimates of the ThyssenKrupp Steel Duisburg steel plant for the ResNet-44 method

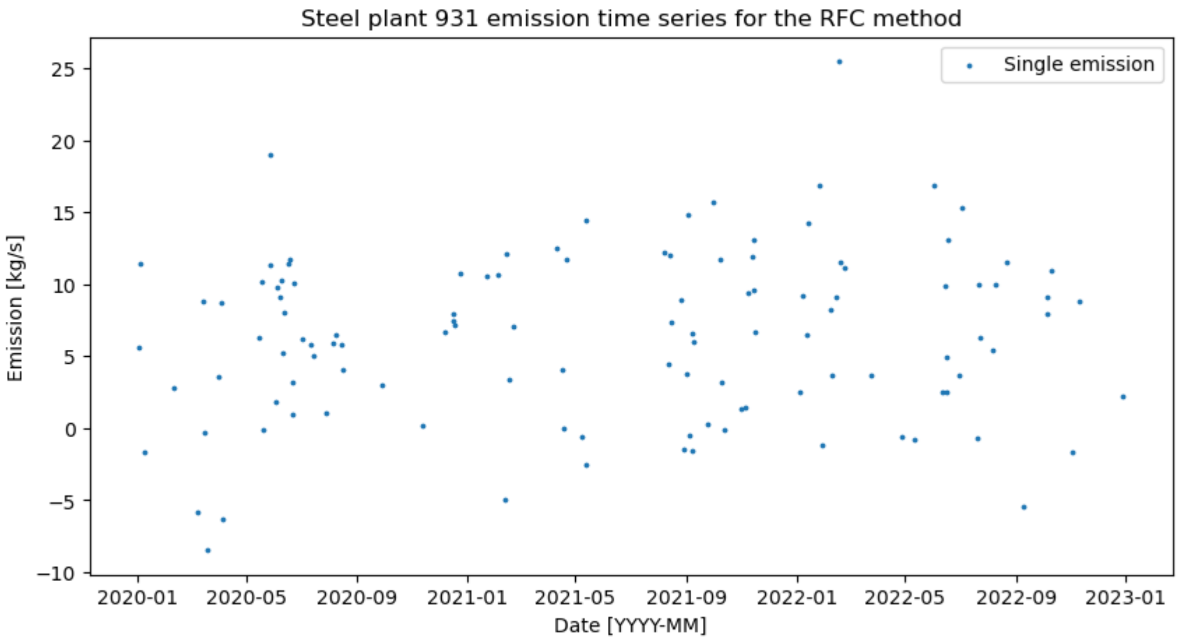

Figure C.41: Time series of the emission estimates of the ThyssenKrupp Steel Duisburg steel plant for the RFC method

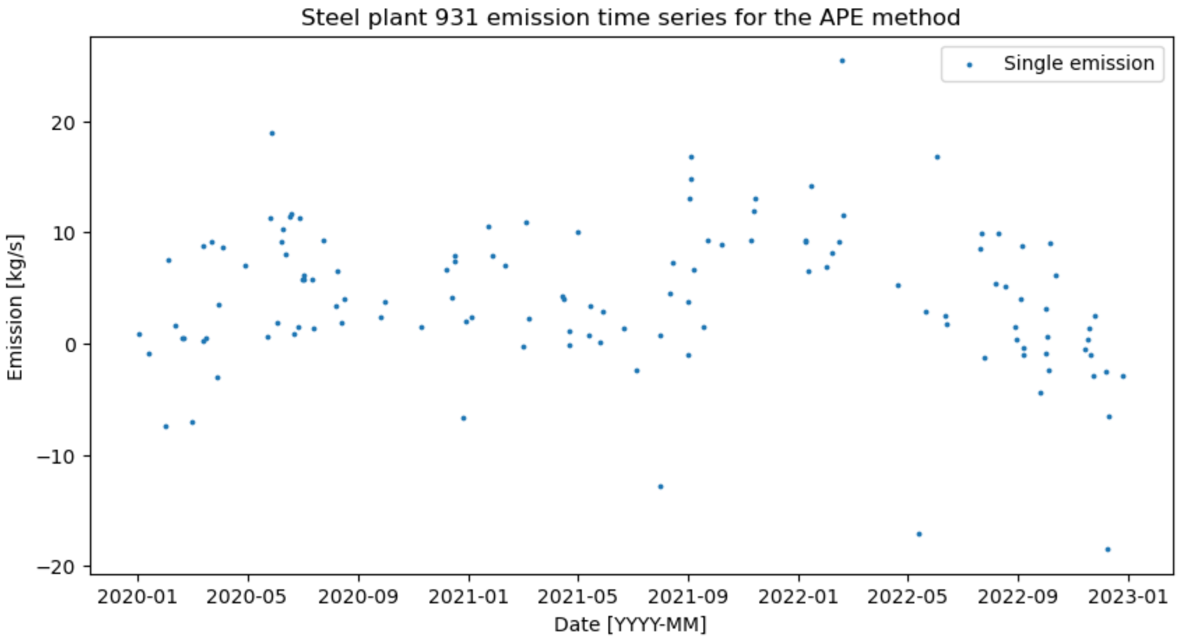

Figure C.42: Time series of the emission estimates of the ThyssenKrupp Steel Duisburg steel plant for the APE method

#991: Voestalpine Stahl Linz steel plant

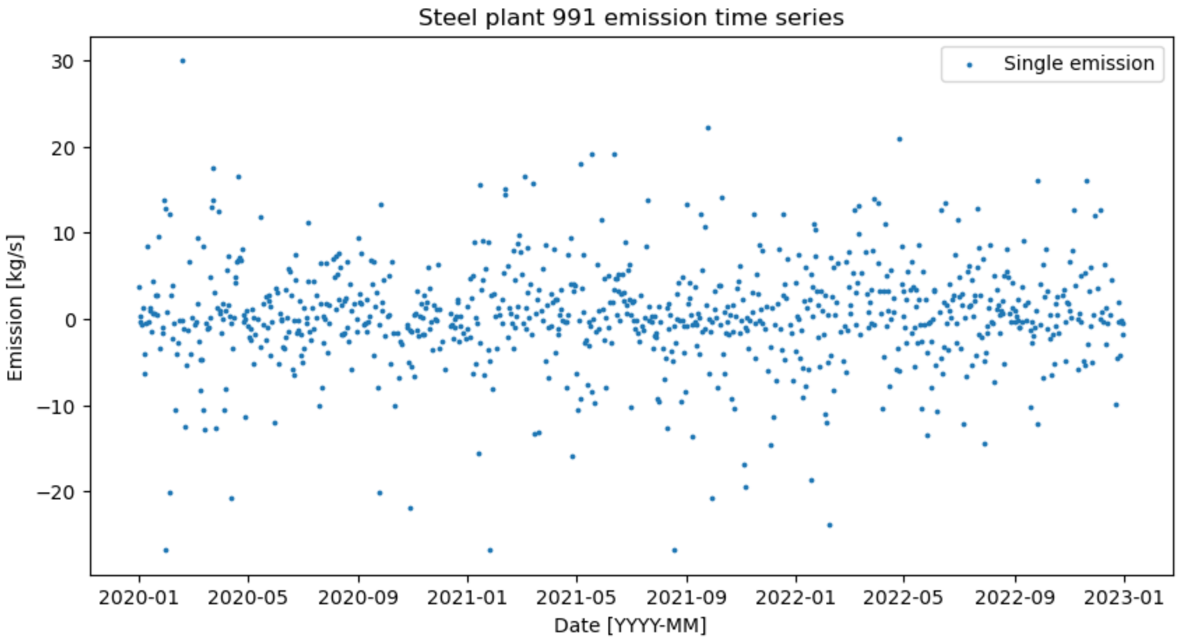

Figure C.43: Time series of the emission estimates of the Voestalpine Stahl Linz steel plant

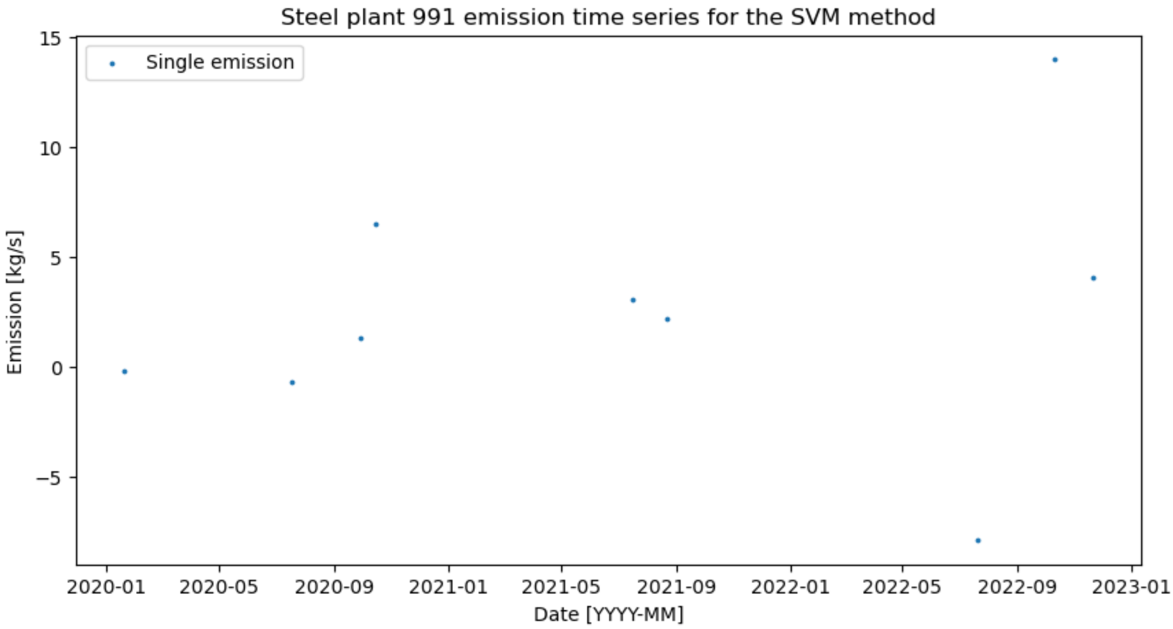

Figure C.44: Time series of the emission estimates of the Voestalpine Stahl Linz steel plant for the SVM method

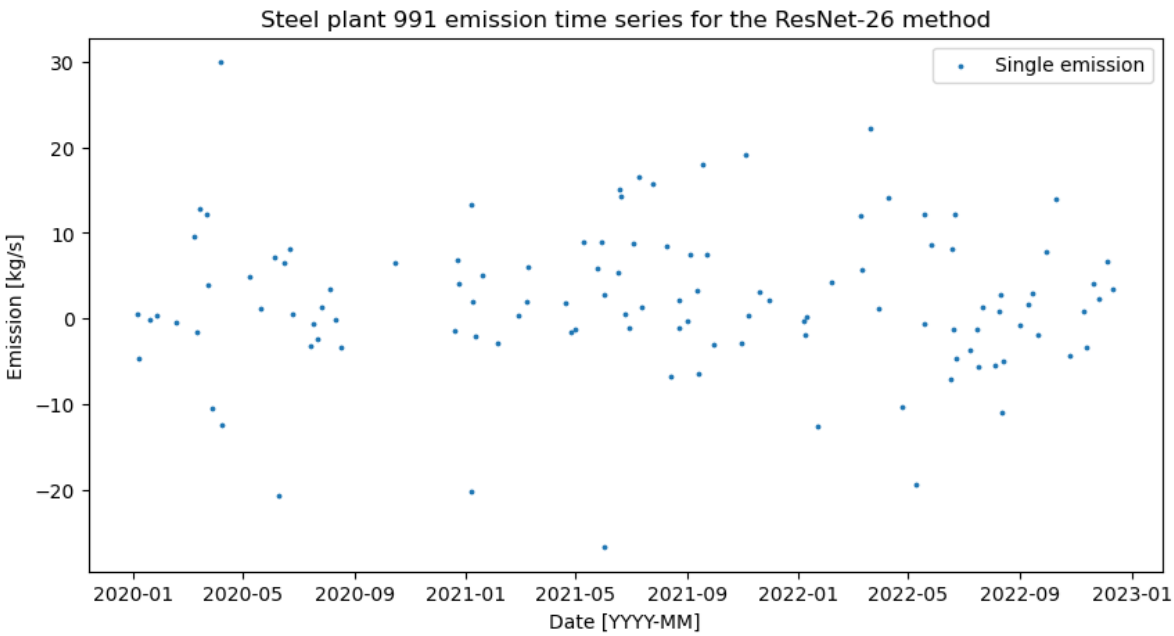

Figure C.45: Time series of the emission estimates of the Voestalpine Stahl Linz steel plant for the ResNet-26 method

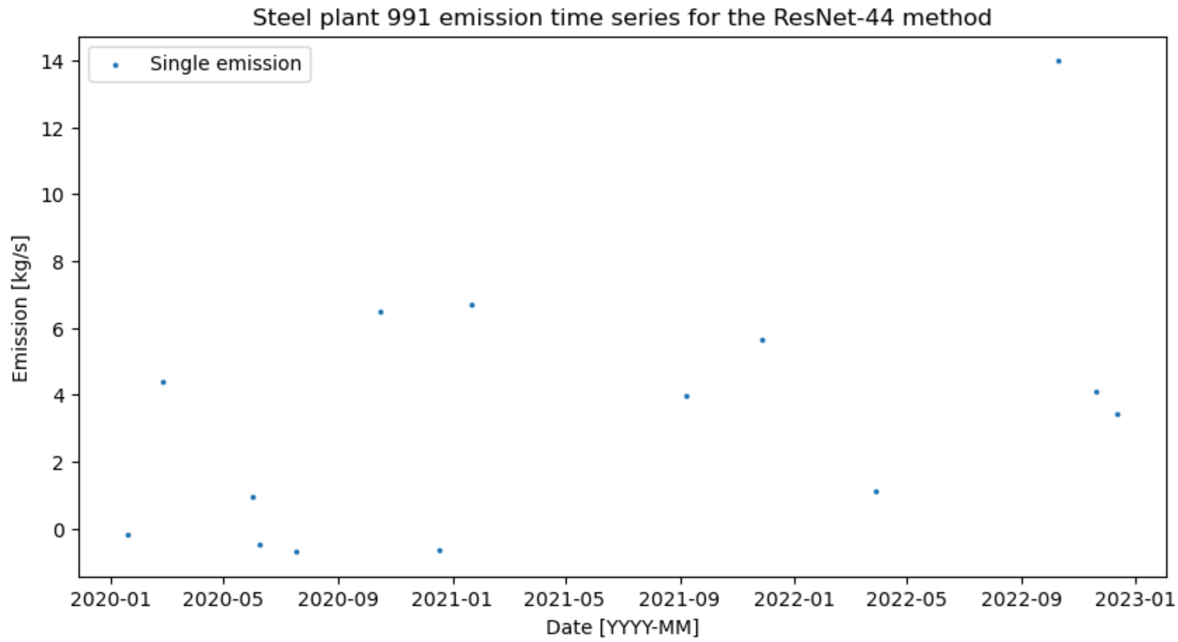

Figure C.46: Time series of the emission estimates of the Voestalpine Stahl Linz steel plant for the ResNet-44 method

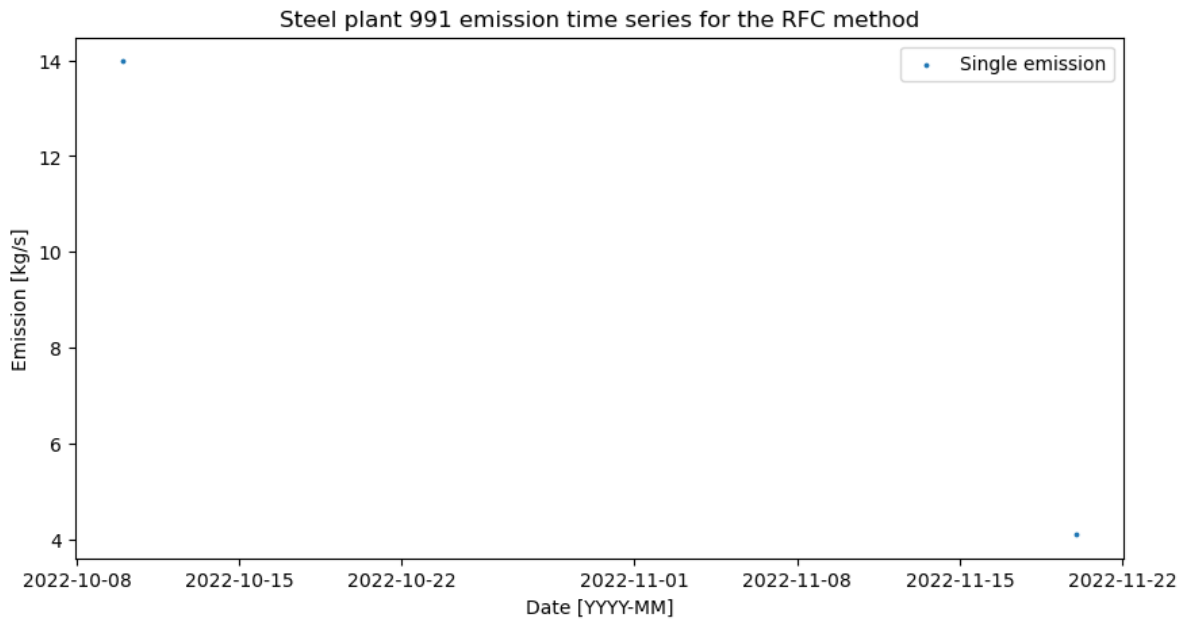

Figure C.47: Time series of the emission estimates of the Voestalpine Stahl Linz steel plant for the RFC method

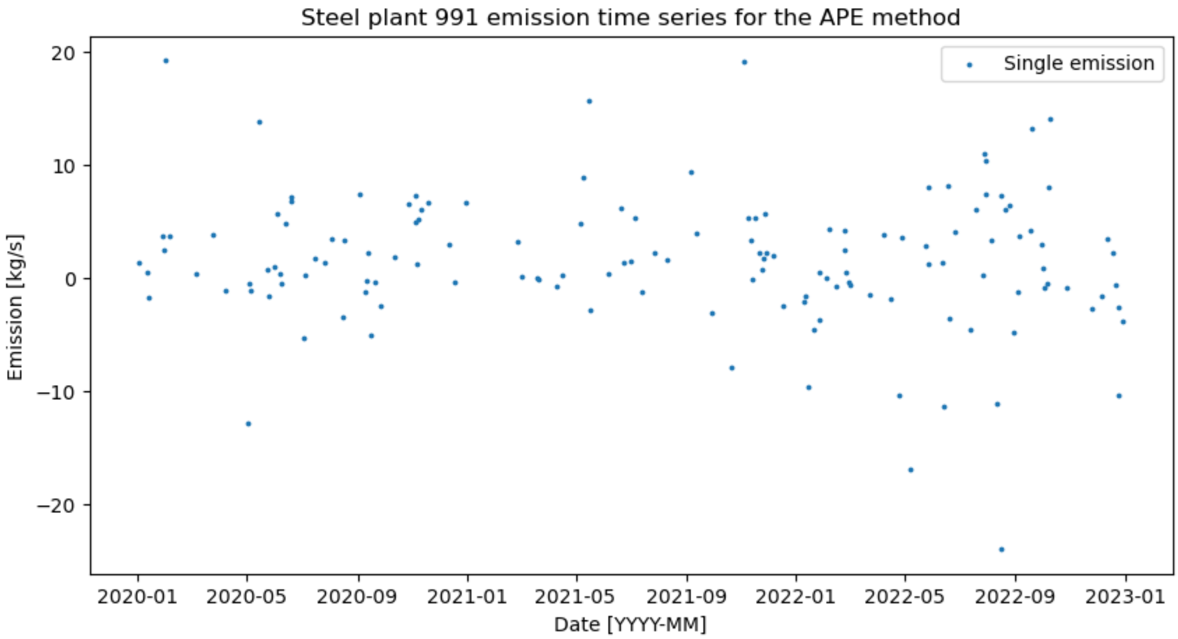

Figure C.48: Time series of the emission estimates of the Voestalpine Stahl Linz steel plant for the APE method