# D4.5 – Report on SILVANUS advanced detection capabilities

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No 101037247

| | |
|---|---|
| **Project Acronym** | SILVANUS |
| **Grant Agreement number** | 101037247 (H2020-LC-GD-2020-3) |
| **Project Full Title** | Integrated Technological and Information Platform for Wildfire Management |
| **Funding Scheme** | IA – Innovation action |

**DELIVERABLE INFORMATION**

| | |
|---|---|
| **Deliverable Number:** | D4.5 |
| **Deliverable Name:** | Report on SILVANUS advanced detection capabilities: Publication of a consolidated review of all the technologies released within SILVANUS and subsequent improvements being carried out following the demonstration activities as outlined in WP9. |
| **Dissemination level:** | PU |
| **Type of Document:** | DEM |
| **Contractual date of delivery:** | 31/09/2024 (M36) |
| **Date of submission:** | XXX |
| **Deliverable Leader:** | CMCC Foundation |
| **Status:** | Final |
| **Version number:** | V0.6 |
| **WP Leader/ Task Leader:** | CMCC Foundation |
| **Keywords:** | Earth Observations; Weather/Climate Forecasting; Internet of Things; Social Media Sensing; UGV; UAVs |
| **Abstract:** | This document reports the final implementation and demonstration of the systems (interfaces, methods, and tools) to collect and aggregate data from heterogenous sources (Earth Observations, Weather and Climate Modes, in-situ devices, Social Media Sensing, UGVs and UAVs) and subsequent pre-processing capabilities to detect wildfires. |

| | |
|---|---|
| **Lead Author(s):** | Gabriele Accarino; Shahbaz Alvi |
| **Reviewers:** | • ATOS: Jose Ramon Martinez<br>• INTRA: Nelly Leligou<br>• EXAI: Ciro Caterino, Simone Martin Marotta |

| | |
|---|---|
| | • KEMEA: George Sakkas; John Tsaloukidis; Nikos Sokas |

**Disclaimer**

| Document History | | | |
|---|---|---|---|
| **Version** | **Date** | **Contributor(s)** | **Description** |
| V0.1 | 10/07/2024 | Gabriele Accarino (CMCC Foundation) | Assigned contributors; Table of Contents and basic structure. |
| V0.2 | 11/09/2024 | All contributors | Initial draft of the document. |
| V0.3 | 13/09/2024 | Gabriele Accarino and Shahbaz Alvi (CMCC Foundation) | All the contributing partners have provided their inputs, and CMCC has additionally integrated them, making the document ready for internal review. |
| V0.4 | 16/09/2024 | All contributors | Harmonization of the manuscript. |
| V0.5 | 19/09/2024 | Gabriele Accarino and Shahbaz Alvi (CMCC Foundation) | Ready for internal review |
| V0.6 | 26/09/2024 | Gabriele Accarino and Shahbaz Alvi (CMCC Foundation) | Ready for submission |

**List of Contributors**

| Partner | Author(s) |
|---|---|
| CMCC | Gabriele Accarino, Shahbaz Alvi, Giusy Fedele, Leo Loprieno, Pasquale Schiano, Giuseppe Giugliano, Carmen Alvarez-Castro |
| CTL | Konstantinos Avgerinakis, Maria I. Maslioukova, Vangelis Mathioudis |
| UTH | Panagiotis Oikonomou, Georgios Boulougaris, Kostas Kolomvatos |
| FINC | Antonio Santovito, Alessandro Annese, Marcello Scipioni |
| DELL | Mustafa Al-Bado, Merry Globin, Michal Sworzeniowski |
| UISAV | Zoltán Balogh, Emil Gatial, Ján Zelenka, Tomáš Kasanický, Stefan Dlugolinský |
| HB | Sándor Darányi, Johan Eklund, Thomas van Erven |
| TRT | Yann Semet, Grégoire Boussu, Bruno Marcon, Mathieu Marchand |

| ACRONYM | Description |
|---|---|
| 10-daily | Every 10 days |
| ADAM | Adaptive Moment Estimation |
| AI | Artificial Intelligence |
| AQA | Air Quality Assessment |
| BERT | Bidirectional Encoder Representations from Transformers |
| BoW | Bag of Words |
| BR | Born Classifier |
| C3S | Copernicus Climate Change Service |
| CLC | Corine Land Cover |
| CMCC | Euro-Mediterranean Center on Climate Change |
| CSV | Comma Separated Value |
| cURL | URL client |
| DBSCAN | Density-Based Spatial Clustering of Applications with Noise |
| DDS | Data Delivery System |
| DEM | Digital Elevation Model |
| DIP | Data Ingestion Pipeline |
| DTM | Dynamic Topic Modelling |
| EAQI | European Air Quality Index |
| ECMWF | European Centre for Medium-Range Weather Forecasts |
| EMDCs | Edge Micro-Data Centers |
| EO(s) | Earth Observation(s) |
| ERA5 | Fifth generation of ECMWF reanalysis |
| FCCs | Forward Command Centers |
| FDI | Fire Danger Index |
| FIFO | First in, First out. |
| FWI | Fire Weather Index |
| GDAL | Geospatial Data Abstraction Library |
| GeoTiff | Geo Tag Image File Format |
| GFS | Global Forecasting System |
| GPS | Global Positioning System |
| GSM | Global System for Mobile Communications |

| | |
|---|---|
| HaZZ | Firefighters and Rescue Units in Slovakia |
| HRAPIE | Health risks of air pollution in Europe |
| HTTP | Hypertext Transfer Protocol |
| I2C | Inter-Integrated Circuit |
| IFS | Integrated Forecasting System |
| IoT | Internet of Things |
| IP | Internet Protocol |
| JJA | June July August |
| JSON | JavaScript Object Notation |
| KDE | Kernel Density Estimation |
| KPI(s) | Key Performance Indicator(s) |
| LDA | Latent Dirichlet Allocation |
| LLM(s) | Large Language Model(s) |
| LST | Land Surface Temperature |
| LSTM | Long Short-Term Memory |
| LTE | Long-Term Evolution |
| MAE | Mean Absolute Error |
| MBE | Mean Bias Error |
| MCU | Micro-Controller Unit |
| ML | Machine Learning |
| MODIS | Moderate Resolution Imaging Spectroradiometer |
| MQTT | Message Queue Telemetry Transport |
| NCAR | National Center for Atmospheric Research |
| NCEP | National Centers for Environmental Prediction |
| NCEP-GFS | National Centers for Environmental Prediction -Global Forecast System |
| NDVI | Normalized Difference Vegetation Index |
| NE | Named Entity |
| NER | Named Entity Recognition |
| NetCDF | Network Common Data Form |
| NIR | Near infrared |
| NLP | Natural Language Processing |
| NWP | Numerical Weather Prediction |

| | |
|---|---|
| PDF (s) | Probability Density Function(s) |
| RDF | Resource Description Framework |
| RK | Runge-Kutta |
| RMSE | Road Mean Squared Error |
| SAL | Storage Abstraction Layer |
| SCC | Supercomputing Center |
| SemKB | Semantic Knowledge Base |
| SMI | Soil Moisture Index |
| SON | September October November |
| TF-IDF | Term Frequency-Inverse Document Frequency |
| TM | Topic Modelling |
| TOML | Tom's Obvious Minimum Language |
| UART | Universal Asynchronous Receiver-Transmitter |
| UAVs | Unmanned Aerial Vehicles |
| UGVs | Unmanned Ground Vehicles |
| UI | User Interface |
| UP | User Product |
| USB | Universal Serial Bus |
| UX | User Experience |
| VHR | Very High Resolution |
| VPN | Virtual Private Network |
| WGS84 | World Geodetic System 1984 |
| WP | Work Package |
| WRF | Weather Research and Forecasting |

**List of beneficiaries**

| No | Partner Name | Short name | Country |
|---|---|---|---|
| 1 | UNIVERSITA TELEMATICA PEGASO | PEGASO | Italy |
| 2 | ZANASI ALESSANDRO SRL | Z&P | Italy |
| 3 | INTRASOFT INTERNATIONAL SA | INTRA | Luxembourg |
| 4 | THALES | TRT | France |
| 5 | FINCONS SPA | FINC | Italy |
| 6 | ATOS IT SOLUTIONS AND SERVICES IBERIA SL | ATOS IT | Spain |
| 6.1 | ATOS SPAIN SA | ATOS SA | Spain |
| 7 | EMC INFORMATION SYSTEMS INTERNATIONAL | DELL | Ireland |
| 8 | SOFTWARE IMAGINATION & VISION SRL | SIMAVI | Romania |
| 9 | CNET CENTRE FOR NEW ENERGY TECHNOLOGIES SA | EDP | Portugal |
| 10 | ADP VALOR SERVICOS AMBIENTAIS SA | ADP | Portugal |
| 11 | TERRAPRIMA - SERVICOS AMBIENTAIS SOCIEDADE UNIPESSOAL LDA | TP | Portugal |
| 12 | 3MON, s. r. o. | 3MON | Slovakia |
| 13 | CATALINK LIMITED | CTL | Cyprus |
| 14 | SYNTHESIS CENTER FOR RESEARCH AND EDUCATION LIMITED | SYNC | Cyprus |
| 15 | EXPERT SYSTEM SPA | EAI | Italy |
| 16 | ITTI SP ZOO | ITTI | Poland |
| 17 | Venaka Treleaf GbR | VTG | Germany |
| 18 | MASSIVE DYNAMIC SWEDEN AB | MDS | Sweden |
| 19 | FONDAZIONE CENTRO EURO-MEDITERRANEOSUI CAMBIAMENTI CLIMATICI | CMCC F | Italy |
| 20 | EXUS SOFTWARE MONOPROSOPI ETAIRIA PERIORISMENIS EVTHINIS | EXUS | Greece |
| 22 | Micro Digital d.o.o. | MD | Croatia |
| 23 | POLITECHNIKA WARSZAWSKA | WUT | Poland |
| 24 | HOEGSKOLAN I BORAS | HB | Sweden |
| 25 | GEOPONIKO PANEPISTIMION ATHINON | AUA | Greece |
| 26 | ETHNIKO KENTRO EREVNAS KAI TECHNOLOGIKIS ANAPTYXIS | CERTH | Greece |
| 27 | PANEPISTIMIO THESSALIAS | UTH | Greece |
| 28 | ASSOCIACAO DO INSTITUTO SUPERIOR TECNICO PARA A INVESTIGACAO E DESENVOLVIMENTO | IST | Portugal |
| 29 | VELEUCILISTE VELIKA GORICA | UASVG | Croatia |
| 30 | USTAV INFORMATIKY, SLOVENSKA AKADEMIA VIED | UISAV | Slovakia |

| No | Partner Name | Short name | Country |
|---|---|---|---|
| 31 | POMPIERS DE L'URGENCE INTERNATIONALE | PUI | France |
| 32 | THE MAIN SCHOOL OF FIRE SERVICE | SGSP | Poland |
| 33 | ASSET - Agenzia regionale Strategica per lo Sviluppo Ecosostenibile del Territorio | ASSET | Italy |
| 34 | LETS ITALIA srls | LETS | Italy |
| 35 | Parco Naturale Regionale di Tepilora | PNRT | Italy |
| 36 | FUNDATIA PENTRU SMURD | SMURD | Romania |
| 37 | Romanian Forestry Association - ASFOR | ASFOR | Romania |
| 38 | KENTRO MELETON ASFALEIAS | KEMEA | Greece |
| 39 | ELLINIKI OMADA DIASOSIS SOMATEIO | HRT | Greece |
| 40 | ARISTOTELIO PANEPISTIMIO THESSALONIKIS | AHEPA | Greece |
| 41 | Ospedale Israelitico | OIR | Italy |
| 42 | PERIFEREIA STEREAS ELLADAS | PSTE | Greece |
| 43 | HASICSKY ZACHRANNY SBOR MORAVSKOSLEZSKEHO KRAJE | FRB MSR | Czechia |
| 44 | Hrvatska vatrogasna zajednica | HVZ | Croatia |
| 45 | TECHNICKA UNIVERZITA VO ZVOLENE | TUZVO | Slovakia |
| 46 | Obcianske zdruzenie Plamen Badin | PLAMEN | Slovakia |
| 47 | Yayasan AMIKOM Yogyakarta | AMIKOM | Indonesia |
| 48 | COMMONWEALTH SCIENTIFIC AND INDUSTRIAL RESEARCH ORGANISATION | CSIRO | Australia |
| 50 | FUNDACAO COORDENACAO DE PROJETOS PESQUISAS E ESTUDOS TECNOLOGICOS COPPETEC | COPPETEC | Brazil |
| 51 | RINICOM LIMITED | RINICOM | United Kingdom |

# Index of figures

Figure 1. TOML configuration snapshot for consuming NDVI data from DDS. ....5
Figure 2. List of 44 CLC classes [7]. ....7
Figure 3. Land surface temperature (a) daytime, (b) nighttime, (top) at 1km resolution (bottom) at downscaled to 2.2 km resolution. The missing values in the picture correspond to sea or areas covered by clouds. ....9
Figure 4. Soil Moisture Index (left) at 5km resolution (right) at downscaled to 2.2 km resolution. ....10
Figure 5. (a) The original CLC map consumed from DDS over the pilot region of Apulia. (b)-(k) shows the 10 groups to which the original 44 classes are mapped. ....12
Figure 6. Digital elevation model data over the Apulia region. (top left) the original DEM at 10 m of spatial resolution, (top right) the remapped to 2.2 km and (bottom) the slope computed from the remapped DEM. ....13
Figure 7. Road and waterway distance over the Apulia region. Panel (a) shows the original 100 m resolution of road (top) and waterway (bottom) distance, whereas panel (b) shows the same data remapped to the 2.2 km target resolution. ....14
Figure 8. PDFs of daily (a) 2m Mean Temperature, (b) 2m Minimum Temperature, (c) 2m Maximum Temperature, (d) 10m Wind Speed and (e) Accumulated Precipitation for E-OBS (black), VHR (blue), WRF@CMCC forced by GFS analysis (red) and WRF@CMCC forced by IFS analysis (orange). These analyses are performed over the period April-May 2019. ....16
Figure 9. WRF_2km@CMCC weather forecasts operational scheme. ....16
Figure 10. Forecast simulation domains and topography (in meters above sea level) of the Tepilora, Sardinia (left), and Cova da Beira, Portugal (right), pilot sites. The red dots in the Sardinia domain (left) denote the locations of the station observations used for the model evaluation. ....19
Figure 11. Probability seasonal FWI for CMCC model in an example of hindcast, JJA of 2016. There is a colour scale for each category of the FWI with light or dark colours in base of the lower probability or higher probability of being in that category. Each point has also information about the probability to be below, normal or above normal conditions to know how normal or extreme is a forecast/hindcast. ....21
Figure 12. PDFs of the 2-meter air temperature (in °C) averaged over all stations (solid line) and over all single model grid points closest to each station (dashed line). ....22
Figure 13. PDFs of the daily accumulated precipitation (in mm/day) over all available stations within the Tepilora Park (solid line) and over all single model grid points closest to each station (dashed line). ....22
Figure 14. PDFs of the daily mean wind speed (in m/s) of all station observations (solid line) as well as of all single model grid cells closest to each station (dashed line). ....23
Figure 15. PDFs of the daily mean relative humidity (in %) of all available in-situ observations (solid line) as well as of all single model grid points located the closest to each station (dashed line). ....24
Figure 16. MAE, MBE and RMSE of daily 2m air temperature, daily accumulated precipitation, daily mean wind speed, and daily mean relative humidity of the WRF_2km@CMCC model against observations (where available). ....24
Figure 17. Probability density functions (PDFs) of daily two-meter air temperature (in °C) calculated over all WRF_2km@CMCC land grid points (dashed line) against E-OBS gridded data (solid line) over the same domain. The investigated period is April-October 2020. To allow for a more consistent evaluation, the model grid was remapped onto the E-OBS grid. ....25
Figure 18. Cumulative frequency of precipitation (in mm/d) as simulated by the WRF_2km@CMCC model (dashed line) and the equivalent domain of E-OBS gridded observations. For the model values, the grid was remapped onto that of E-OBS and only land cells were considered. The period investigated is April-October 2020. ....26

Figure 19. PDFs of daily mean wind speed (in m/s) calculated over values simulated by the WRF_2km@CMCC model (dashed line) and values of E-OBS gridded observations (solid line). For the PDFs, the model grid was re-gridded onto the E-OBS grid and only land values were considered. The values were calculated over the period April-October 2020. .....26
Figure 20. PDFs of near-surface relative humidity (in %) as recorded by WRF_2km@CMCC (dashed line) and E-OBS gridded observations (solid line). .....27
Figure 21. Mean absolute error (MAE), mean bias error (MBE), and root mean squared error (RMSE) of daily 2m air temperature, daily accumulated precipitation, daily mean wind speed, and daily mean relative humidity of the WRF_2km@CMCC model against E-OBS gridded *observations.* .....28
Figure 22. Correlations (ACC) between JJA CMCC seasonal FWI (start date of May) and ERA5 FWI. Red values correspond to those regions where the seasonal forecast system can reproduce the interannual variability of the FWI in JJA. .....28
Figure 23. WRF_2km@CMCC forecast for the 28/07/2024 at 00:00:00 UTC+2 (initialized 27/07/2024 at 00:00:00 UTC+2) of the temperature at 2m (left) and relative humidity at 2m (right) over the numerical domain of the Apulia pilot. .....30
Figure 24. WRF_2km@CMCC forecast for the 28/07/2024 at 00:00:00 UTC+2 (initialized 27/07/2024 at 00:00:00 UTC+2) of the temperature at 2m (left) and relative humidity at 2m (right) over the numerical domain of the Tepilora pilot. .....30
Figure 25. WRF_2km@CMCC forecast for the 28/07/2024 at 00:00:00 UTC+2 (initialized 27/07/2024 at 00:00:00 UTC+2) of the temperature at 2m (left) and relative humidity at 2m (right) over the numerical domain of the Portugal pilot. .....31
Figure 26. (Left) New case for our IoT, (Right) Previous IoT case. .....33
Figure 27. Example of SILVANUS dashboard for Smoke/fire detection (IoT) layer. .....33
Figure 28. Process of computing drone flight routes at the FCC. .....36
Figure 29. UP4a IoT device connected to the MESH in the Sky node. .....40
Figure 30. UP4a enhanced with smoke detection algorithm. .....41
Figure 31. UP4a mounted on CSIRO's UGV. .....42
Figure 32. Visualizing UP4a layer on dashboard. .....43
Figure 33. UP9b and MESH in the Sky integration. .....44
Figure 34. Installation of UP9b on an UGV. .....44
Figure 35. Architectural scheme of the East-West API first demonstration. .....45
Figure 36. Social media sensing Framework. .....46
Figure 37. Crawling Architecture. .....48
Figure 38. Results for K-Means method. .....49
Figure 39. Results for Isolation Forest method. .....50
Figure 40. Results for LSTM method. .....51
Figure 41. Results for Moving Average. .....52
Figure 42. Example of image-based fire/smoke detection API call, using cURL command. .....60
Figure 43. Example response of the image-based fire/smoke detection API. .....61
Figure 44. Integration flow of the Fire Event Detection Module within the SILVANUS platform .....64
Figure 45. RDF Types and Object Properties. .....65
Figure 46. Visualizing Interdependencies within the Knowledge Base. .....66
Figure 47. Visualizing the social sensing ontological network. .....66
Figure 48. Visualizing social media in the form of graph database. .....67
Figure 49. (left) Semantic intensity contour map of 4K tweets with 117 tweets, in red, indexed by high sentiment value phrases (bigrams) for comparison. Elevation corresponds to the heated nature of messaging. (right) Social importance contour map of 4K tweets. Elevation corresponds to the perceived importance (e.g., number of retweets) of messages. .....68
Figure 50. Workflow of the 'TweetAnalyzer for Information Fusion' web service module. .....69

Figure 51. PageRank-based “landscape” of tweets, expansion phase No 4 (at 4 K): the sentiment intensity of marked up tweets does not overlap with the attractor structure. ....71
Figure 52. SILVANUS dashboard. ....72
Figure 53. Visualization example of a fire event in SILVANUS dashboard. ....73
Figure 54. Multiple UGVs demonstrating online map merging at the Australia pilot. ....75
Figure 55. The CSIRO UGV platform autonomously navigating in complex, natural terrain in the presence of (artificial) smoke, during a simulated mid-fire scouting operation. ....75
Figure 56. Semantic classification showing coloured point cloud (left) and predicted class labels (right). Data were collected during demonstrations at the Australia pilot. ....76
Figure 57. Data recorded by the robot during an autonomous exploration mission at the Australia pilot. Records were automatically uploaded to the SILVANUS platform and can be retrieved and plotted on an aerial map view. Information includes location, time, and *forest analytics results*. ....77
Figure 58. Place of trajectory optimization module within the integrated Fire Management approach. ....78
Figure 59. Monitored area for Czech pilot, green polygon represents the monitored area, and red polygon represents no-fly zone. ....79
Figure 60. Results of drone path generation for two drones in Czech pilot. ....79
Figure 61. Results of drone path generation for five drones in Romania pilot. Dashboard visualization and visualization in Google earth pro. ....79
Figure 62. Results of drone path generation for five drones in Romania pilot with non-fly zone. Individual path or polygon. ....80
Figure 63. UAV drone flight data collection preparation, execution and aggregation at the Czech pilot site. ....81

# Index of tables

Table 1. List of ancillary variables consumed from the DDS for the FDI forecast. ..... 6
Table 2. List of ancillary variables consumed from the DDS for the FDI forecast. ..... 7
Table 3. Grouping the native 44 CLC classes and the 10 groups. ..... 10
Table 4. List of weather variables used for daily FDI forecast. The weather variables are stored and consumed from the DDS. ..... 17
Table 5. Process of computing drone flight routes at the FCC. ..... 39
Table 6. Comparison of Anomaly Detection Methods. ..... 53
Table 7. Pilots and tweets collected for each pilot. ..... 54
Table 8. NE classes recognized by the new enhanced NER model of the Facebook Post Analyser. ..... 56
Table 9. Distribution of entity classes within the dataset. ..... 57
Table 10. Precision, Recall and F1 statistics for individual classes and models. ..... 57
Table 11. Overall performance of individual models. ..... 57
Table 12. Number of relevant and irrelevant tweets per language. ..... 59
Table 13. Performance metrics for all models. ..... 59
Table 14. Fire event detection evaluation. ..... 63

# Table of Contents


***Executive Summary*** ..... *1*

**1 INTRODUCTION** ..... **2**

**2 DATA COLLECTION, AGGREGATION AND PRE-PROCESSING OF EO DATASETS** ..... **4**

2.1 TOOL DESCRIPTION ..... 4
*2.1.1 The CMCC DDS* ..... *4*
2.2 INNOVATIONS AND UPDATES ..... 4
*2.2.1 Ingestion of EO and ancillary data in the DDS* ..... *4*
*2.2.2 Consuming data from DDS* ..... *5*
*2.2.3 Overview and description of EO datasets* ..... *5*
2.2.3.1 Normalized Difference Vegetation Index ..... 6
2.2.3.2 Land Surface Temperature ..... 6
2.2.3.3 Soil Moisture Index ..... 6
2.2.3.4 Corine Land Cover classes ..... 6
*2.2.4 Ancillary variables* ..... *7*
2.2.4.1 Digital Elevation Model (DEM) and Slope ..... 7
2.2.4.2 Road distance ..... 8
2.2.4.3 Waterway distance ..... 8
2.2.4.4 World population ..... 8
2.3 SCIENTIFIC RESULTS AND DRAWBACKS ..... 8
*2.3.1 Weaknesses of the data gathering pipeline* ..... *8*
2.4 DEMONSTRATION REPORT ..... 8
*2.4.1 Land surface temperature (LST)* ..... *9*
*2.4.2 Soil Moisture Index* ..... *9*
*2.4.3 Corine Land Cover* ..... *10*
*2.4.4 Digital Elevation Model* ..... *12*
*2.4.5 WorldPop data* ..... *13*

**3 WEATHER/CLIMATE DATA SERVICES FOR FOREST FIRE THREAT RISK ASSESSMENT** ..... **15**

3.1 TOOLS DESCRIPTION ..... 15
*3.1.1 Numerical models for weather nowcasting* ..... *15*
*3.1.2 Numerical model for seasonal forecasts* ..... *16*
*3.1.3 Ingestion of weather variables in the DDS* ..... *17*
*3.1.4 Processing of weather variables* ..... *17*
3.1.4.1 Disintegration of cumulative data ..... 18
3.1.4.2 Conversion of units ..... 18
3.1.4.3 Computing the wind speed ..... 18
3.1.4.4 Weather-derived variables ..... 18
3.1.4.5 Computing summary variables ..... 18
3.2 INNOVATIONS AND UPDATES ..... 18
*3.2.1 Deployment and evaluation of numerical weather models in other pilot sites* ..... *18*
*3.2.2 Reusability of the data processing pipeline* ..... *20*
*3.2.3 Probabilistic FWI based on seasonal forecast* ..... *20*
3.3 SCIENTIFIC RESULTS AND DRAWBACKS ..... 21
*3.3.1 Evaluation of WRF_2km@CMCC performances over the Tepilora domain* ..... *21*
*3.3.2 Evaluation of WRF_2km@CMCC performances over the Portugal domain* ..... *24*
*3.3.3 Evaluation of CMCC seasonal FWI at European domain* ..... *28*
3.4 DEMONSTRATION REPORT ..... 28
*3.4.1 Apulia-Italy Pilot* ..... *29*
*3.4.2 Tepilora-Italy Pilot* ..... *30*
*3.4.3 Cova de Beira-Portugal Pilot* ..... *31*

**4 DATA COLLECTION, AGGREGATION, AND PRE-PROCESSING OF IN-SITU DEVICES ..... 32**

4.1 TOOL DESCRIPTION ..... 32
4.2 INNOVATIONS AND UPDATES ..... 32
*4.2.1 Fire detection from IoT devices ..... 32*
*4.2.2 Fire detection from edge devices ..... 34*
*4.2.3 IoT for air quality assessment ..... 34*
*4.2.4 Forward Command Centers ..... 35*
4.2.4.1 FCC East-West Communication ..... 35
4.3 SCIENTIFIC RESULTS AND DRAWBACKS ..... 36
*4.3.1 Fire detection from IoT devices ..... 36*
*4.3.2 Fire detection from edge devices ..... 37*
*4.3.3 IoT for air quality assessment ..... 37*
*4.3.4 Forward Command Centers ..... 38*
4.4 DEMONSTRATION REPORT ..... 39
*4.4.1 Fire detection from IoT devices ..... 39*
**4.4.1.1** Croatia Pilot - 2023 ..... 39
**4.4.1.2** France Pilot - 2023 ..... 40
**4.4.1.3** Italy Tabletop - 2023 ..... 41
**4.4.1.4** Greece Tabletop - 2023 ..... 41
**4.4.1.5** Australia Pilot - 2023 ..... 41
**4.4.1.6** Czech Pilot - 2024 ..... 42
*4.4.2 Fire detection from edge devices ..... 43*
*4.4.3 IoT for air quality assessment ..... 43*
*4.4.4 Forward Command Centers ..... 44*

**5 SOCIAL MEDIA SENSING AND CONCEPT EXTRACTION ..... 46**

5.1 TOOL DESCRIPTION ..... 46
5.2 INNOVATIONS AND UPDATES ..... 47
*5.2.1 Social Media Crawling ..... 47*
5.2.1.1 Implementation of the X Crawler ..... 47
5.2.1.2 Monitoring and anomaly detection ..... 48
5.2.1.3 X Crawling ..... 53
*5.2.2 Social Media Analysis Toolkit ..... 54*
5.2.2.1 Textual analysis of social media content ..... 54
5.2.2.2 Visual analysis of social media content ..... 60
5.2.2.3 Fire Events Detection ..... 62
5.2.2.4 Integration with SAL ..... 64
5.2.2.5 Fire Events stored to knowledge database ..... 64
5.3 SCIENTIFIC RESULTS AND DRAWBACKS ..... 67
*5.3.1 Social media Crawling ..... 67*
*5.3.2 Textual Analysis ..... 68*
5.4 DEMONSTRATION REPORT ..... 71

**6 UGV MONITORING FOR WILDFIRE BEHAVIOUR ..... 74**

6.1 TOOL DESCRIPTION ..... 74
6.2 INNOVATIONS AND UPDATES ..... 74
6.3 SCIENTIFIC RESULTS AND DRAWBACKS ..... 77
6.4 DEMONSTRATION REPORT ..... 77

**7 UAV DEPLOYMENT FOR REMOTE SENSING AND IDENTIFICATION OF WILDFIRE ..... 78**

7.1 TOOL DESCRIPTION ..... 78
7.2 INNOVATIONS AND UPDATES ..... 78
7.3 SCIENTIFIC RESULTS AND DRAWBACKS ..... 80
7.4 DEMONSTRATION REPORT ..... 80

**8 CONCLUSION ..... 82**

**9 REFERENCES ..... 83**

**Executive Summary**

This document outlines the activities conducted within Work Package (WP) 4, which focuses on utilizing Earth Observations (EOs) techniques for wildfire preparedness and prevention. Additionally, it highlights the technology developed under the SILVANUS project to combat forest fires in the eventuality they occur. Each task is detailed with descriptions of the technology developed, scientific outcomes, and results achieved by the respective partners. The sections pertaining to the tasks within WP4 are organized into four sub-sections: Tool Description, Innovation and Updates, Scientific Results and Drawbacks, and Demonstration Report. The Tool Description section provides a comprehensive overview of the tools and technologies developed for each task. In the Innovation and Updates section, we delve deeper into the scientific background and specifics of the technologies employed and developed to achieve each task. This section also highlights both the current state-of-the-art advancements in the technology and innovations that exceed existing standards. The Scientific Results and Drawbacks section summarizes the scientific findings from previous tasks while addressing any known limitations associated with each. Since scientific innovation is central to all activities within SILVANUS, it is crucial to highlight the current state-of-the-art for each task, citing relevant articles and resources, as well as detailing accomplishments that extend beyond this standard. This section also ensures coherence by discussing any shortcomings of the solutions developed. The Demonstration Report concludes with an overview of demonstrations of the technologies and tools developed during each task, including pilot exercises conducted in previous reporting periods. The conclusion of this deliverable summarizes the activities and achievements presented in D4.5. It also clarifies how the various tasks within WP4 interconnect to form a cohesive, integrated approach to forest fire prevention, detection, and response coordination.

# 1 Introduction

Wildfires pose a significant threat to both ecosystems and human settlements, with the frequency and intensity of these destructive events increasing in recent years. This surge in wildfires has led to substantial economic losses and widespread social disruption. Driven by factors such as intensified human activities and the escalating impacts of climate change, the risk of wildfires is expected to grow in the coming years. Fortunately, as the threat has increased, so do the advancements in technology, offering new possibilities to address and mitigate these threats.

This document outlines the innovative technologies developed within the SILVANUS framework, which aims to redefine the state-of-the-art in wildfire detection and countermeasures. In particular, D4.5 aims at providing the most updated overview of the capabilities and tools developed in SILVANUS up to M36.

Historically, fire indices, such as the Canadian Fire Weather Index (FWI), were widely used in fire management as key tools to estimate the risk of wildfire occurrence. They help in understanding fire behaviour and are critical for decision-making in both wildfire detection and response. However, the FWI only relies on weather variables, such as temperature, humidity, wind speed, and precipitation, which does not fully account for the complex stochastic nature of wildfires. For instance, it is well understood that, in addition to weather, there is a strong connection among wildfires occurrence and anthropogenic activities, local geography, as well as vegetation. For this reason, in the context of the SILVANUS framework, a cutting-edge data-driven Fire Danger Index (FDI) has been developed and deployed in WP5 to account for the underlying complex, potentially non-linear, relationship among variables under play.

To support the collection, pre-processing and consumption of a large set of variables to serve as input to the data-driven FDI, WP4 provides capabilities and tools to gather variables from heterogeneous data sources (e.g., EOs and weather forecasts), harmonize and store them. In particular, Section 2 provides an overview of EO repositories, as well as the transformations performed to aggregate and process the selected variables to be used as input predictors for the data-driven approach. In Section 3 the technical details and setup of operational numerical weather models over different pilot sites are described.

The Euro-Mediterranean Center on Climate Change's (CMCC) Data Delivery System (DDS), deployed in the CMCC facilities, is employed to gather the EOs and weather forecast data and process them to serve the data-driven FDI. The technical details of the processing functions supported, as well as aggregation and harmonization are described throughout Section 2 and Section 3. Despite the DDS serves the wider weather and climate research community, yet it is fully integrated with the SILVANUS Storage Abstraction Layer (SAL) to address the specific goals of the project, especially related to the Machine Learning (ML)-based FDI estimation that makes up the phase A of the project (prevention and preparedness).

In addition, WP4 provides enhanced tools and capabilities that cover both phases A and B (fire detection and response coordination). Specifically, Section 4 highlights the Internet of Things (IoT) devices developed for two key User Products, UP4a and UP9b. The former focuses on detecting fire and smoke, employing ML, whereas the latter is designed to assess air quality, similarly relying on ML to interpret environmental data and provide real-time insights on hazardous conditions. Additionally, Edge Micro-Data Centers (EMDCs), UP4b, play a crucial role in gathering and processing images and videos to identify fire events. On the response side, Forward Command Centers (FCCs), UP10, are used to aid commanders with key operational services.

In addition to the previously mentioned tools, the SILVANUS project leverages the widespread use of social media platforms such as X (formerly Twitter) and Facebook to enhance wildfire monitoring and response capabilities. Section 5 introduces a powerful Social Media Sensing framework, a robust system specifically designed to collect, analyze, and process social media posts related to fire incidents. This framework plays a crucial role in supporting stakeholders, such as emergency responders, fire management teams, and government authorities, by providing timely and relevant information during a fire event. By continuously monitoring social media platforms, it gathers real-time data from posts, photos, and user reports that mention or show fire-related activities.

The SILVANUS project not only relies on advanced detection tools and in-situ monitoring devices, but also integrates unmanned vehicles, which play a pivotal role in field operations for real-time fire detection and surveillance. Both Unmanned Ground Vehicles (UGVs) and Unmanned Aerial Vehicles (UAVs), as detailed in Section 6 and Section 7, respectively, are critical components of this system, providing unique capabilities for wildfire management. Improvements in UGVs focused on the multi-robot navigation system with automatic map merging and place recognition. Additionally, the system added mobile manipulation, allowing a to collect ground-level data. Regarding UAVs, a toolbox for optimizing flight plans has been developed. It allows generating optimal sweeping trajectories based on several sensing parameters, such as camera shooting angle, as well as flight parameters like altitude and turning radius.

Beyond describing all the technical details of the tools and the improvements performed in the last reporting period, D4.5 also describes the scientific results and drawbacks to point out potentiality and limitation of the proposed solutions, as well as future improvements. Moreover, for each developed tool the D4.5 provides a dedicated section to demonstrate the outcomes of each tool in different pilot sites, as further outlined in WP9.

# 2 Data collection, aggregation and pre-processing of EO datasets

## 2.1 Tool description

The EO tool developed in SILVANUS enables the collection, aggregation, and pre-processing of remotely sensed data, serving as key predictors for the data-driven FDI developed in WP5. By processing various EO variables, actionable intelligence is generated to monitor and assess the dynamic natural and human-made environment.

Within the SILVANUS project, EO plays a pivotal role in the prevention, mitigation, and management of wildfire risks. This data-driven approach relies on EO observations to evaluate wildfire occurrences, offering critical insights into fire-prone areas.

This section describes the tools developed to support the collection, aggregation, and pre-processing of the EOs. It also details how these variables are sourced from repositories, stored, and utilized to fulfill the project's objectives, particularly those outlined in T5.1 of WP5. Furthermore, an overview of the EOs that support SILVANUS applications is also provided.

The following sections describes the datasets, processing steps and tools used to develop functionalities aimed at supporting the advanced detection capabilities, such as the estimation of FDI within SILVANUS.

### *2.1.1 The CMCC DDS*

For the purpose of data-driven FDI (WP5), the DDS [1] serves as the primary tools for the ingestion and consummation of the EOs from the parent repositories. The DDS disseminates scientific dataset produced by CMCC Research Divisions and other external data providers which are of interest to activities carried out by the CMCC research divisions. DDS relies on a cloud-native microservices-based architecture that has been deployed — using Kubernetes. The core of the DDS is geokube [2]; it is a Python developed with the aim to improve the User Experience (UX) as well as for the analysis and visualisation of climate and earth science data. It provides high-level abstractions in terms of both Data Model (inspired by Climate Forecast and Unidata Common Data Models) and Application Programming Interface.

The DDS is deployed and managed on the CMCC facilities, and it has been integrated with the SAL for the purpose of the SILVANUS project.

## 2.2 Innovations and updates

### *2.2.1 Ingestion of EO and ancillary data in the DDS*

Variables coming from EO repositories and additional ancillary datasets serve as predictors for estimating the FDI (WP5). As can be seen in Table 1 and Table 2, most of the variables are already ingested into the DDS (see the Parent repository column) and ready to be gathered and processed, whereas others are not. This is the case of land surface temperature (daytime and nighttime) and soil moisture index. For these variables the DDS uses a Prefect [3] pipeline to automatically synchronize every day with these repositories and ingest the newest data available on the data holder's repository. Each Prefect Task uses appropriate APIs to interact with the specific service, including authentication, search for new data, select the geographical domain of interest, formulate the request and download data in Network Common Data Form (NetCDF) format. Once data is downloaded, the DDS performs indexing and caching of metadata so that it becomes accessible to users. Contrary, Worldpop indicators, such as world population and road and waterway distance are static and were downloaded once and ingested manually into the DDS.

### *2.2.2 Consuming data from DDS*

The pipeline for consuming data (i.e. gathering and processing), previously ingested in the DDS, follows a specific template. Data supported by DDS are listed in Table 1 and Table 2. The pipeline is highly configurable as it uses Tom's Obvious Minimum Language (TOML) configuration files to define the variables to be gathered and how they should be processed. This flexibility allows the same variable to be processed differently, specifying the appropriate processing function(s) to be applied as needed. By default, and in the context of supporting the estimation of the data-driven FDI in WP5, the pipeline retrieves 10 days in the past starting from the day in which the FDI should be estimated. Indeed, the temporal correlation between the dynamic fire predictors is also an important ingredient along with the spatial correlation around a given pixel. Regarding the spatial dimension, for each point in space (pixel/grid point) we collect the 25 x 25 points around it to account for spatial characteristics of the predictors. In Figure 1 we show screenshot from the Tom's Obvious Minimum Language (TOML) file for consuming NDVI data from DDS.

```
[data.ndvi]
    [data.ndvi.gathering]
        longname="Normalized difference vegetation index"
        source="dds"
        dataset="ndvi-silvanus"
        frequency="puglia"
        variable=["ndvi"]
        format="netcdf"
        var_name=["ndvi"]
        type="dynamic"
    [data.ndvi.processing]
        functions=["geokube_resample","fill_time_dimension","interp"]
        kwargs=["{'operator':'mean','frequency':'1D','closed':'right','to_lonlat':False}","{'dates':'dates'}","None"]
```

**Figure 1.** TOML configuration snapshot for consuming NDVI data from DDS.

In the *gathering* section (Figure 1) all the parameters related to retrieving a variable from the DDS are defined. In the *gathering* section a variable can also be configured to be retrieved from a locally stored file.

The *processing* section instead defines the set of functions that must be performed on each variable so that it can serve as an input to the ML model used in WP5 for estimating the FDI. These functions are listed, and the transformations are performed in a "first in, first out" (FIFO) manner. The transformations are implemented as Python functions, and the arguments (i.e., kwargs) for each Python function are provided in a list as string representations of dictionaries corresponding to each transformation at the same index.

### *2.2.3 Overview and description of EO datasets*

Among different EOs repositories that have been identified in T4.1, we considered a set of 5 EO-based variables that are likely to be related to wildfire occurrence. These variables are selected following the work of [4] and are used as predictors to provide a fully data-driven FDI, which is part of WP5 activities. The EO variables from which these predictors are derived, and their parent repository, are summarized in Table 1.

**Table 1.** List of ancillary variables consumed from the DDS for the FDI forecast.

| No. | Variables | Parent repository |
|---|---|---|
| 1 | Normalized difference vegetation index (NDVI) | DDS |
| 2 | Land surface temperature in daytime | NASA MODIS [5] |
| 3 | Land surface temperature in nighttime | NASA MODIS |
| 4 | Soil moisture index | European Drought Observatory [6] |
| 5 | Corine Land Cover classes | DDS |

The datasets mentioned in Table 1 are ingested from their parent repository, harmonized and saved in the DDS repository. The repository, furthermore, is automatically synced as the parent repository is updated. From the DDS, these variables are consumed, processed and utilized for the computation of the FDI.

*2.2.3.1 Normalized Difference Vegetation Index*

Normalized Difference Vegetation Index (NDVI) is an index which quantifies the health and diversity of vegetation. It is a widely used index for assessing the state of vegetation as it is highly correlated with the true state of vegetation. The theoretical basis of the NDVI measurement relies on the fact that live green plants appear bright in the near infrared (NIR) region (strongly reflecting this wavelength) of the electromagnetic spectrum. However, chlorophyll in plants strongly absorbs visible light for photosynthesis. The NDVI is therefore computed according to the following equation:

$$NDVI = \frac{NIR - Red}{NIR + Red}$$

NDVI is constrained between -1 and 1.

*2.2.3.2 Land Surface Temperature*

Land Surface Temperature (LST) is the surface temperature of the Earth sourced by the solar radiation and measured by the radiation remitted by the Earth. LST contrasts with the temperature at 2 meters which represents the air temperature at 2-meter height from the surface of the Earth.

*2.2.3.3 Soil Moisture Index*

The Soil Moisture Index (SMI) is the proportion of the difference between the soil moisture present and the permanent wilting point of the vegetation to the field capacity and the residual soil moisture. For the case of fire danger prediction, it shows the state of drought as areas which have had drought are more prone to catch fires. The values of the soil moisture index ranges from 0 (maximum drought) to 1 (minimum drought).

*2.2.3.4 Corine Land Cover classes*

In 1985 the European Community initiated a programme called Corine under which the Corine Land Cover (CLC) project has the objective of mapping the use of European land to assist policy making among many other objectives. The essence of the CLC programme is to provide up-to-date information on land cover throughout Europe in a regular cycle and to show changes occurring between successive cycles (2018 being the latest cycle). The land cover is a gridded dataset at 100 m resolution and is divided into 44 classes. These classes are shown in Figure 2.


**Figure 2.** List of 44 CLC classes **[7]**.

### *2.2.4 Ancillary variables*

In addition to the EO variables listed in Table 1, other ancillary variables are collected, pre-processed and ingested in the DDS. Table 2 list these variables with the corresponding parent repository.

**Table 2.** List of ancillary variables consumed from the DDS for the FDI forecast.

| No. | Variables | Parent repository |
|---|---|---|
| 1 | Digital Elevation Model | DDS |
| 2 | Road distance | Worldpop [8] |
| 3 | Waterway distance | Worldpop [8] |
| 4 | World population | Worldpop [8] |

#### *2.2.4.1 Digital Elevation Model (DEM) and Slope*

The digital elevation represents the state of the bare ground and its topography. From the DEM output the surface slope can be computed using the Geospatial Data Abstraction Library (GDAL) library. The slope is an important parameter in the FDI computation, as a steep slope is more prone to wildfires propagation than a shallower one.

#### *2.2.4.2 Road distance*

Road distance dataset include the distance to major roads in the given region. The dataset is available in Geo Tag Image File Format (Geotiff) format at a resolution of 3 arc seconds (approximately 100m at the equator). The projection is Geographic Coordinate System 84 (WGS84). The values of the raster are the distance (in kilometres) from the cell centre to the nearest feature.

#### *2.2.4.3 Waterway distance*

Waterway distance dataset include the distance to major waterways in the given region. The dataset is available in Geotiff format at a resolution of 3 arc seconds (approximately 100m at the equator). The projection is WGS84, and the values of the raster are the distance (in kilometres) from the cell centre to the nearest feature.

#### *2.2.4.4 World population*

World population is the estimated total number of people per grid-cell. The dataset is available to download at a resolution of 30 arc seconds (approximately 1 km at the equator). The projection is WGS84, and the units are number of people per pixel.

## 2.3 Scientific results and drawbacks

### *2.3.1 Weaknesses of the data gathering pipeline*

The repositories which host the variables listed in Table 1 and Table 2 operate autonomously. Therefore, the availability of data for a given region cannot always be ensured, which is a potential drawback for applications which use these variables, such as the data-driven based FDI method over the traditional methods. This is especially important for dynamic variables which change within a year and year-by-year. The pipeline is equipped with strategies to partly circumvent this problem. In case a variable is not available in the DDS datastore, the pipeline can be instructed, defining the variable "contingency" for a variable in the TOML file, to either,

- Retrieve the data for the latest date: In this case, the data for the latest available period in the datastore is used. This option works well for variables which do not change significantly over a short period of time, e.g. NDVI.
- Retrieve the data for the same period in the previous year: In this case, the data from the previous year in the same period is used. In this case we assume that the statistical characteristics of the variables remain consistent across years. This option works well for variables which have cyclic variation, such as the LST and soil moisture index.

Another limitation comes from the fact that not all satellite data are useful for wildfire nowcasting due to missing of significant variables like wind speed and air temperature which are crucial for predicting the speed of propagation. To address these gaps, it is necessary to integrate satellite data with the first few hours of forecast data from meteorological models, which provide the missing information and improve the accuracy of wildfire nowcasting estimates.

## 2.4 Demonstration report

In this section, we demonstrate how EO variables are processed to serve as input for the data-driven approach developed in T5.1. Indeed, these variables, together with those related to weather forecasts (described in Section 3), are the key components of UP2 which aims at delivering a fully data-driven model for estimating the FDI in different pilot sites, which act as demonstrators in WP9. We also show the results of the processing that can be applied on these variables in the pipeline that has been developed for T5.1. It is worthwhile to mention that the pipeline for preprocessing is customizable and therefore other kind of transformations can be applied to the variables.

### *2.4.1 Land surface temperature (LST)*

The variables LST is already described in Section 2.2.3.2. From the NASA MODIS [5] data repository we ingest two kinds of LST variables, i.e. LST daytime and LST nighttime. NASA provides these variables at the resolution of 1 km which must be remapped to the spatial resolution of weather variables of 2.2 km. To remap the variables, we use the NetCDF function *interp_like* [9] with the method argument set to *nearest*. The original LST variable and the remapped version are shown in Figure 3.

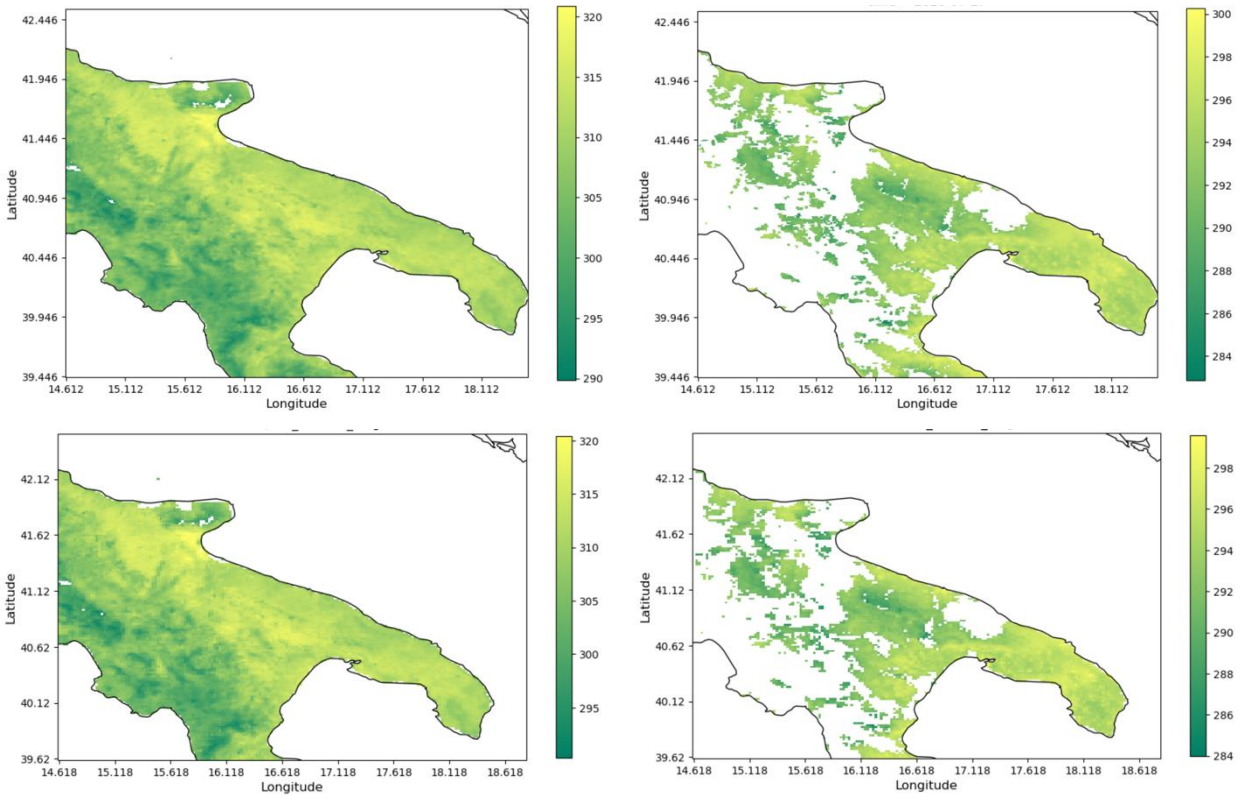


(a) (b)

**Figure 3.** Land surface temperature (a) daytime, (b) nighttime, (top) at 1km resolution (bottom) at downscaled to 2.2 km resolution. The missing values in the picture correspond to sea or areas covered by clouds.

### *2.4.2 Soil Moisture Index*

The SMI is ingested into the DDS from the European Drought Observatory [6]. The frequency at which the SMI is updated is 10-daily, whereas the spatial resolution is of 5 km. In the Figure 4, the remapped SMI over the Apulia region is shown at 2.2 km resolution.

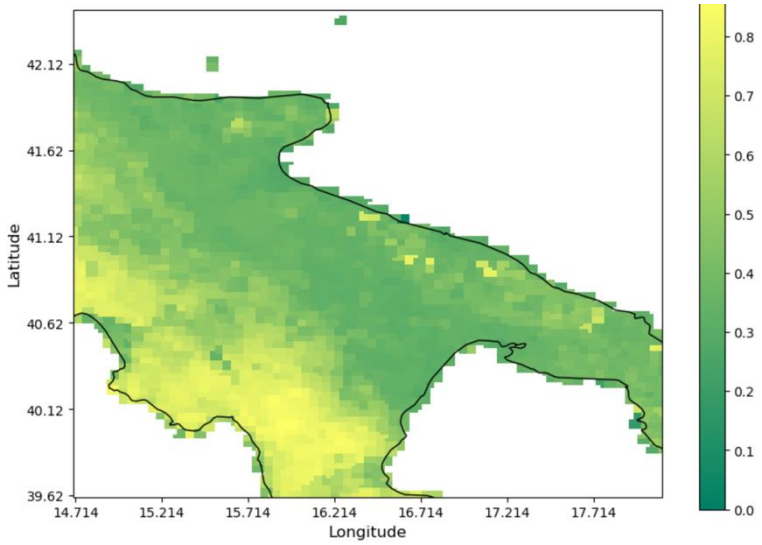


**Figure 4.** Soil Moisture Index (left) at 5km resolution (right) at downscaled to 2.2 km resolution.

### *2.4.3 Corine Land Cover*

CLC has special use when it comes to the problem of prevention and assessment of the fire danger. The variable has already been described in detail in the Section 2.2.3.4. CLC is processed in a particular way allowing the extraction of summarized information from the original CLC classes in a way which serves the estimation of the FDI.

The spatial resolution of the CLC dataset is 100 m, while the resolution of the weather variables is 2.2 km, meaning that there are 484, potentially different, CLC classes (each pixel is associated to one class).

We start by grouping the 44 CLC classes (1 – 44) into 10 groups (0 – 9). It is important to note that the category for each CLC class is described in Figure 2. The 10 groups in which the 44 CLC classes a grouped is show in Table 3.

**Table 3.** Grouping the native 44 CLC classes and the 10 groups.

| Mapped groups | Native CLC Classes |
|---|---|
| 0 | 1,3,4,5,6,7,8,9,10,11 |
| 1 | 2 |
| 2 | 12,13,14 |
| 3 | 15,16,17 |
| 4 | 18 |
| 5 | 19,20,21,22 |
| 6 | 23,24,25 |
| 7 | 26,27,28,29 |
| 8 | 30,31,32,33,34 |
| 9 | 35,36,37,38,39,40,41,42,43,44 |

Then, for each pixel on the 2.2 km resolution map, we compute the percentage of each one of the 10 classes in the group, therefore for each pixel the sum of the ratio over the 10 groups is 1. The original CLC map over the Apulia region of Italy (location of one of the SILVANUS pilot sites), and the 10 groups to which the CLC variable is remapped is shown in Figure 5.

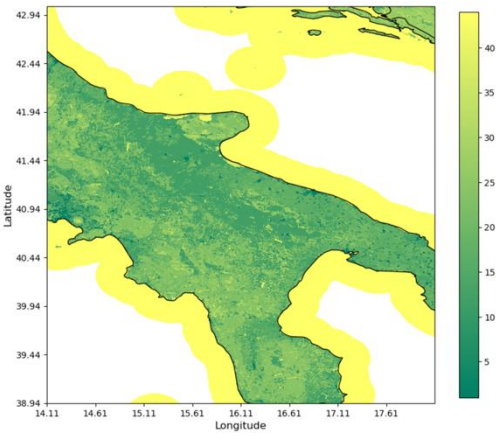


(a)

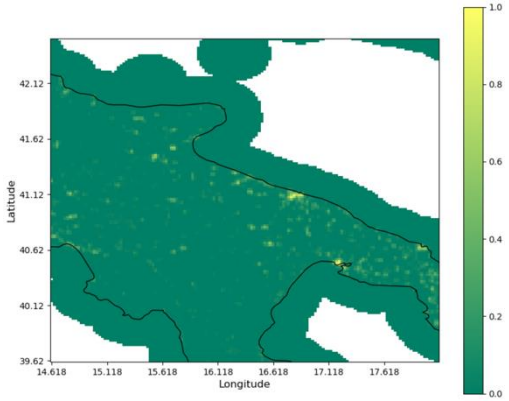


(b)

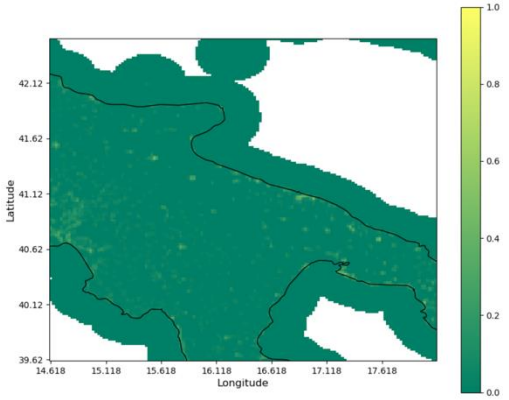


(c)

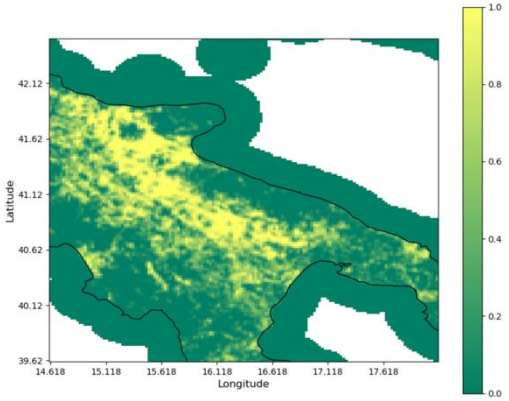


(d)

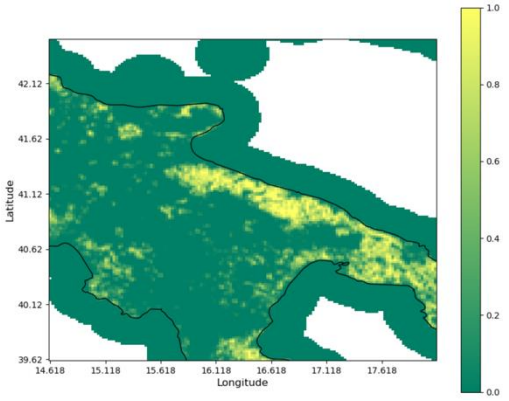


(e)

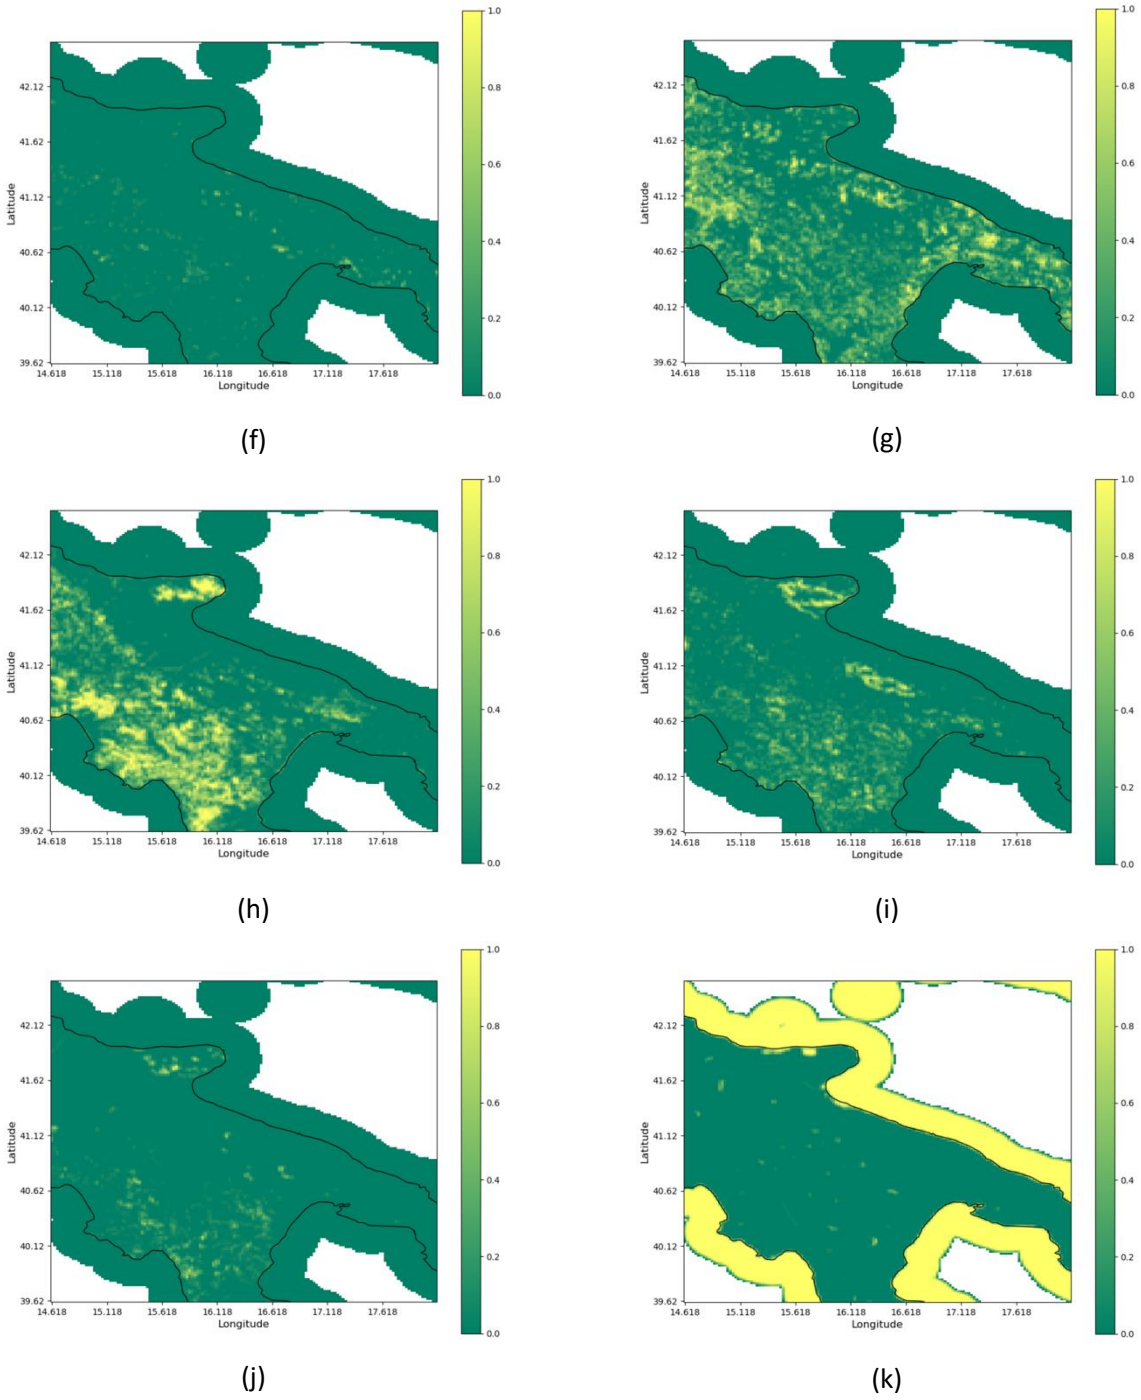


(f) (g)

(h) (i)

(j) (k)

**Figure 5.** (a) The original CLC map consumed from DDS over the pilot region of Apulia. (b)-(k) shows the 10 groups to which the original 44 classes are mapped.

### *2.4.4 Digital Elevation Model*

The Digital Elevation Model (DEM) is originally consumed from the DDS at 10 m resolution which is further remapped to 2.2 km of spatial resolution. Both the maps, before and after the resolution are shown in Figure 6 (upper left and right, respectively). The Digital Elevation is also used to calculate the Slope of the topography which serves as an important variable in concerning the spread of fire. The slope computed from the remapped DEM is also shown in Figure 6 (bottom).

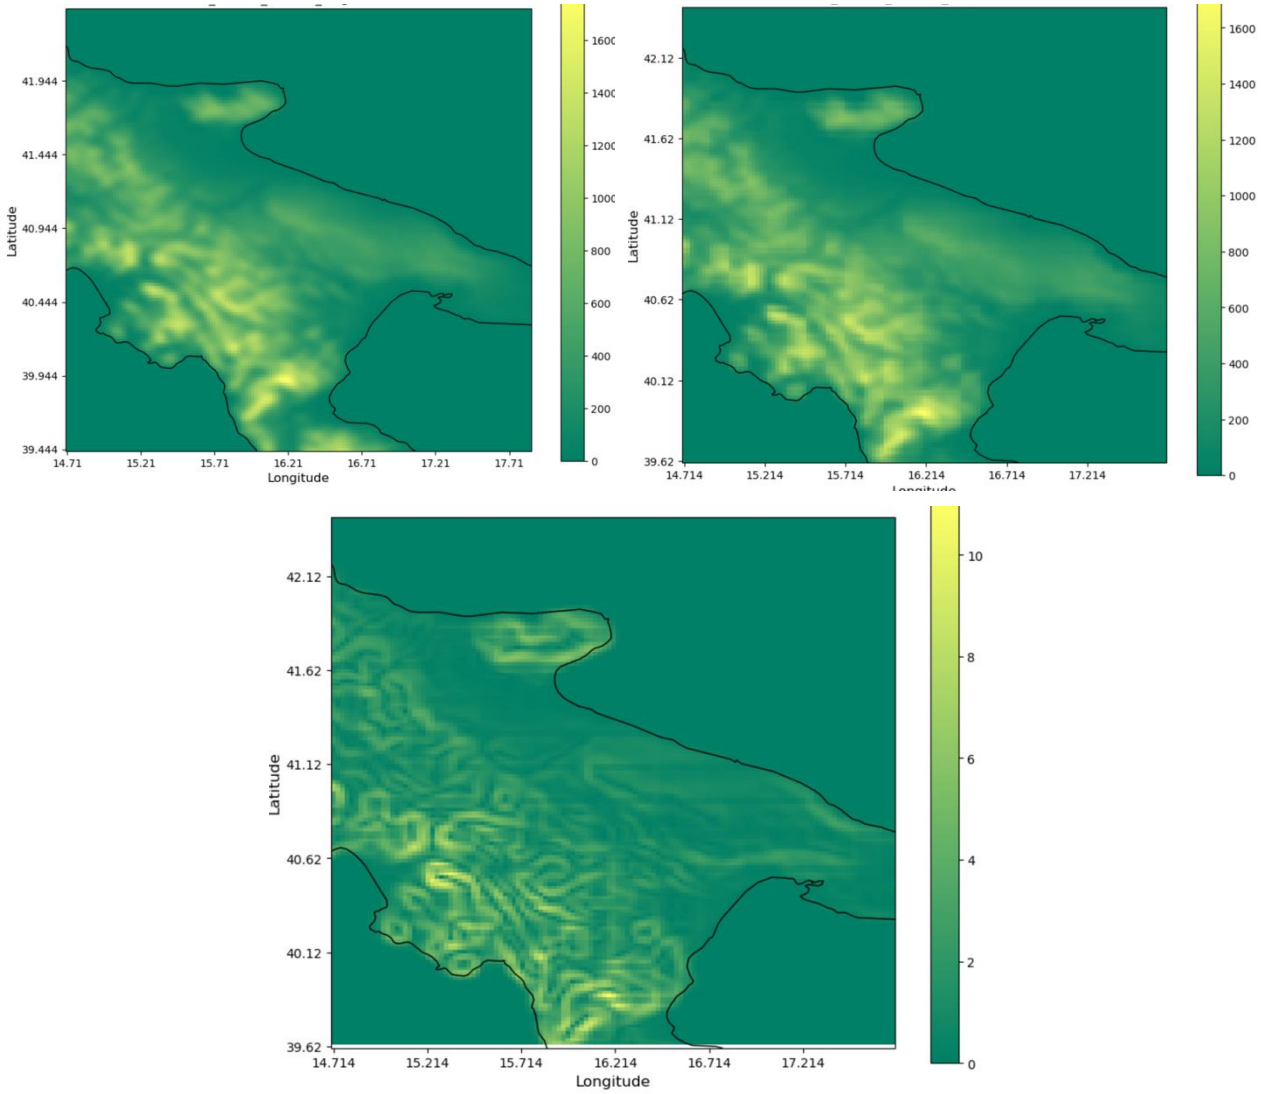


**Figure 6.** Digital elevation model data over the Apulia region. (top left) the original DEM at 10 m of spatial resolution, (top right) the remapped to 2.2 km and (bottom) the slope computed from the remapped DEM.

### *2.4.5 WorldPop data*

As explained in Section 2.2.4.4, the data from WorldPop.org include variables concerning population density, road distance and waterway distance. These variables are remapped using the same technique described above for other variables. For completeness and demonstration, we show map of variables road distance and waterway distance in Figure 7

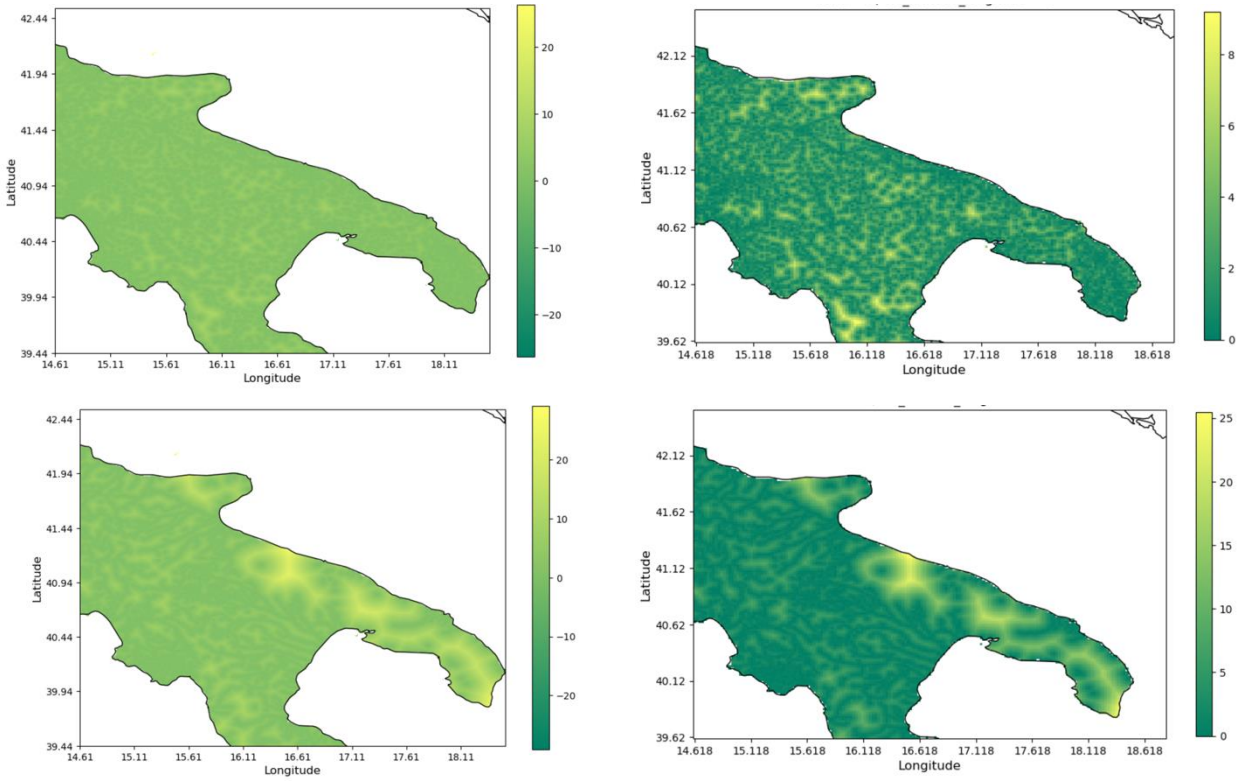


(a) (b)

**Figure 7.** Road and waterway distance over the Apulia region. Panel (a) shows the original 100 m resolution of road (top) and waterway (bottom) distance, whereas panel (b) shows the same data remapped to the 2.2 km target resolution.

# 3 Weather/Climate data services for forest fire threat risk assessment

## 3.1 Tools description

### *3.1.1 Numerical models for weather nowcasting*

The short-term weather forecasts developed for estimating the threat risk of wildfires have been generated using the numerical Weather Research and Forecasting (WRF) model [10] version 4.2.1. This model was chosen over other Numerical Weather Prediction (NWP) models because it has previously demonstrated promising results, in other projects focusing on the Apulia region in Italy, one of the pilot sites for this project.

The WRF model was developed in the late 1990s through a collaboration among various universities and research centers, including the National Center for Atmospheric Research (NCAR) in Boulder, Colorado. WRF is a non-hydrostatic model that describes fully compressible atmospheric flow. Its flexible atmospheric circulation system allows it to operate across various spatial scales. The WRF configuration developed by CMCC at about 2km resolution (hereafter WRF_2km@CMCC) has been run over the pilot domains centred on the Apulia region which includes the Gargano promontory, over the north-eastern corner of Sardinia comprising the Tepilora Regional Natural Park, and finally on the Cova da Beira case study subregion, part of the Portuguese mainland. More details about the configuration of the forecast runs are listed in Section 3.4. In the initial part of the project, which focused solely on the Apulia region, different forcings have been used to run the historical simulations to evaluate the model configuration performances: the Integrated Forecasting System (IFS) of the European Centre for Medium-Range Weather Forecasts (ECMWF) and the Global Forecasting System (GFS) of the National Centers for Environmental Prediction (NCEP). Results reported in Figure 8 demonstrated consistent performances over the investigated period (April-May 2019) between the simulations' outputs driven by IFS and GFS analyses and two reference datasets: the gridded observational data E-OBS [11] at ~11 km resolution and the Fifth generation ECMWF reanalysis (ERA5) downscaling at ~2.2 km over Italy [12] [13]. Specifically, the skills of the numerical models were statistically assessed by computing the Probability Density Functions (PDFs) of the following variables: daily mean, minimum and maximum temperature, wind speed and precipitation. In general, the WRF_2km@CMCC can forecast nearly 200 variables at an hourly temporal frequency.

Since comparable performances are achieved by both IFS- and GFS-driven simulations, GFS forecasts have been utilized for operational weather forecasting due to data availability Figure 9.

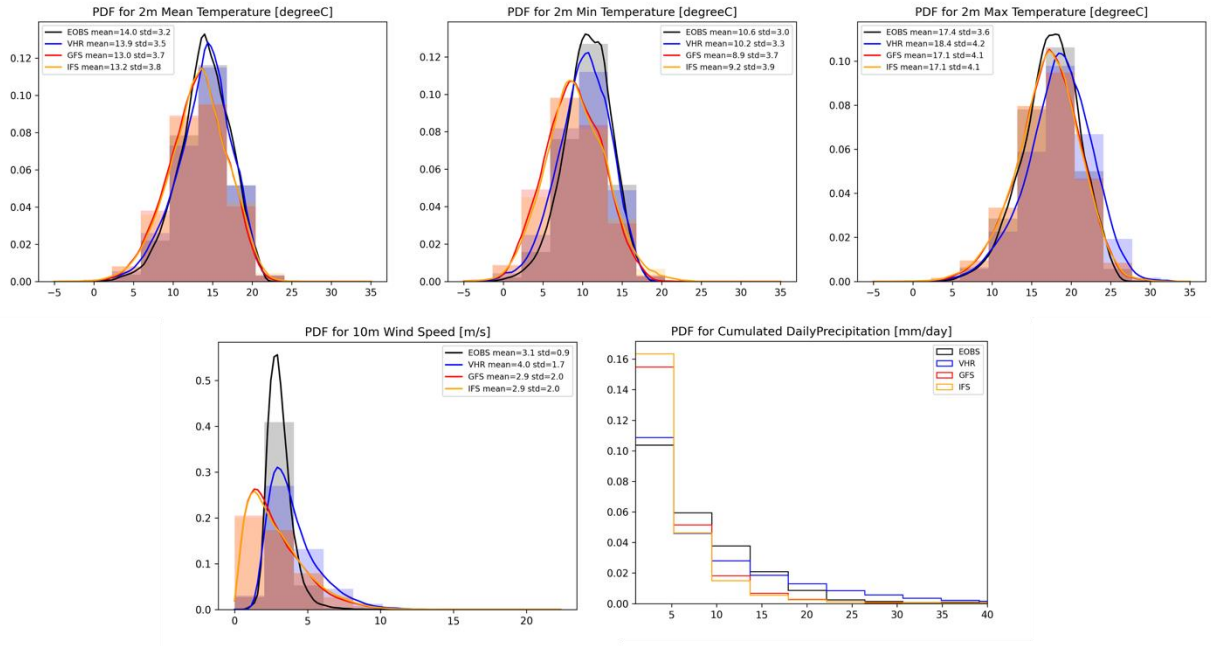


**Figure 8.** PDFs of daily (a) 2m Mean Temperature, (b) 2m Minimum Temperature, (c) 2m Maximum Temperature, (d) 10m Wind Speed and (e) Accumulated Precipitation for E-OBS (black), VHR (blue), WRF@CMCC forced by GFS analysis (red) and WRF@CMCC forced by IFS analysis (orange). These analyses are performed over the period April-May 2019.

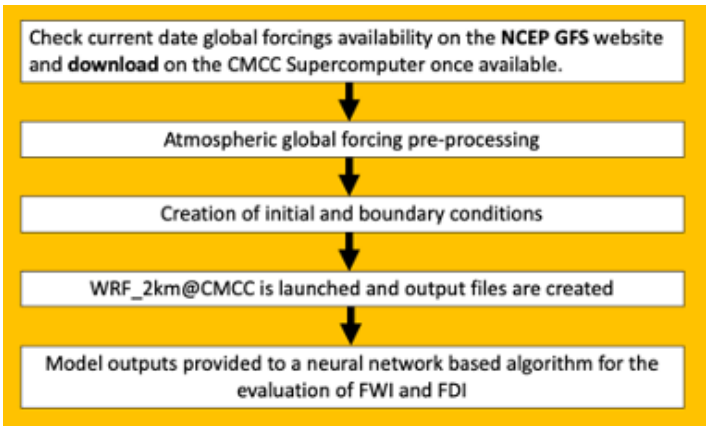


**Figure 9.** WRF_2km@CMCC weather forecasts operational scheme**.**

### *3.1.2 Numerical model for seasonal forecasts*

For the seasonal time scales, probabilistic forecasts of FDIs are provided using ensemble seasonal forecasts produced by the Copernicus Climate Change Service (C3S) operational multi-model prediction system, for the model CMCC sps3.5 (Developed at CMCC, [14]). The System is based on a coupled Ocean-Atmosphere Global Climate Model, complemented by several additional modules. The System is operated monthly in Ensemble seasonal mode (6-month predictions) and is completed by a database of monthly ensemble hindcasts covering the period 1993-2016 which can be used to evaluate the performance of the System and to apply bias removal techniques from operational forecasts. There are 50 members in operational and preoperational prediction mode and 40 members in hindcast mode. The seasonal forecast models are

produced and provided globally but to cover all the pilot cases we have compute the seasonal FWI at the European domain.

The variables of the seasonal forecast model CMCC sps3.5 are available at a horizontal resolution of 1° x 1°, at daily/sub-daily time resolution. The variables used to compute the FWI at seasonal time scales are:

- Temperature at 2 meters level
- Dew point Temperature at 2 meters
- Total precipitation
- 10m u-component of wind
- 10m v-component of wind

The variables used to compute the method of the subsampling of members at seasonal times scales are:

- Sea surface temperature
- Sea level pressure

### *3.1.3 Ingestion of weather variables in the DDS*

Weather is one of the most important components on which the occurrence of wildfire depends. Indeed, weather has a significant role in the spread of fire, once a spark has started.

Among the huge number of variables simulated by the WRF_2km@CMCC, which was described in the previous section, only a subset of those were ingested in the DDS. The list of the ingested variables is provided in Table 4, and they natively come at an hourly temporal resolution and in NetCDF format.

**Table 4.** List of weather variables used for daily FDI forecast. The weather variables are stored and consumed from the DDS.

| No. | Variables |
|---|---|
| 1 | 2-meter temperature |
| 2 | U component of wind at 10-meter |
| 3 | V component of wind at 10-meter |
| 4 | Accumulated total grid scale precipitation |
| 5 | Surface pressure |
| 6 | Water vapor mixing ratios at 2-meter |

The WRF_2k@CMCC numerical model has been setup on the CMCC Supercomputing facilities, where also the DDS platform is deployed. Numerical model forecasts of weather variables, reported in Table 4, are stored in different directories, one for each pilot site, which are automatically synchronized with the DDS for ingestion. Once new data are produced, the DDS extract metadata, cache it and update the catalogue making it visible to be consumed.

### *3.1.4 Processing of weather variables*

In the context of SILVANUS, weather variables are typically used to estimate the FDI by means of a data-driven model. Nevertheless, the ML model requires these variables to be engineered in a certain way before making the inference (i.e., estimating the FDI). Indeed, the daily probability of fire occurrence in each area is sensitive to summary variables than the instantaneous data (i.e., hourly forecasts), such as the daily maximum or minimum values in given area/pixel. However, the variables retrieved from the DDS are not

homogenous on temporal (some observations are available hourly, others are observed on daily or 10-daily etc.) and spatial scales (different resolution). This is because the DDS can deal with different variables at the same time, each one eventually coming from different domains, such as EO, weather, ancillary data. While this provides a high flexibility in many applications, on the other side, all these variables should undergo processing to ensure that –at the end– they are consistent to be consumed in the desired application. Similarly to EO variables (see Section 2.2.2), the processing steps of weather variables can be defined through TOML configuration files.

The procedures for processing weather variables are defined in the following sections.

#### *3.1.4.1 Disintegration of cumulative data*

Forecasts for precipitation consists of hourly cumulative sum. Therefore, it must be disintegrated to obtain the hourly precipitation from which the summary variable (i.e., minimum or maximum) can be computed. The routine however is not developed specifically for precipitation data, and it can handle another cumulative variable as well.

#### *3.1.4.2 Conversion of units*

Often the retrieved variable has a different unit than the data on which the ML-model is trained. This conversion must be handled for each variable separately. We use xclim [15] to handle unit conversion.

#### *3.1.4.3 Computing the wind speed*

Wind speed is computed from the U and V components of the wind velocity using the *uas_vas_2_sfcwind* function of xclim.

#### *3.1.4.4 Weather-derived variables*

Some variables are not directly retrieved from the DDS but instead they are derived from other weather variables. The 2-meter dew point temperature is one of the weather variables that we use in WP5 for estimating the FDI. We compute the dew point temperature through surface pressure and water vapour mixing a 2m.

#### *3.1.4.5 Computing summary variables*

The two kinds of summary variables have been used in the model is the daily *minimum* and the daily *maximum*. In each case, the maximum or minimum of a given variable in each pixel is computed from hourly forecasts belonging to that day. We compute:

- Daily maximum temperature
- Daily maximum surface pressure
- Daily maximum total precipitation
- Daily maximum wind speed
- Daily minimum relative humidity

For computing the summary variables, we use the python library geokube [2].

## 3.2 Innovations and updates

### *3.2.1 Deployment and evaluation of numerical weather models in other pilot sites*

In addition to the operational simulations available over the Apulia Region started on April 2023 and still ongoing, in the second stage of the project, two new pilot sites were introduced: the Tepilora Regional Natural Park (Figure 10, left) located in Sardinia, and the Cova da Beira subregion (Figure 10, right) in central-eastern Portugal. The high-resolution operational weather forecasts were generated in the same manner as for the Apulia pilot region. The Tepilora domain consists of 64 staggered grid points in north-

south direction and 60 staggered grid points in west-east direction. The Portuguese domain is slightly larger and spans 180 staggered points in north-south direction and 137 staggered grid cells in west-east direction.

Two different run types have been performed (hindcast and operational forecasts), which were forced by National Centers for Environmental Prediction-Global Forecast System (NCEP-GFS) data and centred over the respective pilot site. The NCEP-GFS forecasts have been preferred to ECMWF-IFS forcings, due to data availability as mentioned in Section 3.1. Operational forecasts over Tepilora and Cova da Beira regions started progressively on 21/05/2024 and 26/07/2024 respectively and are currently ongoing.

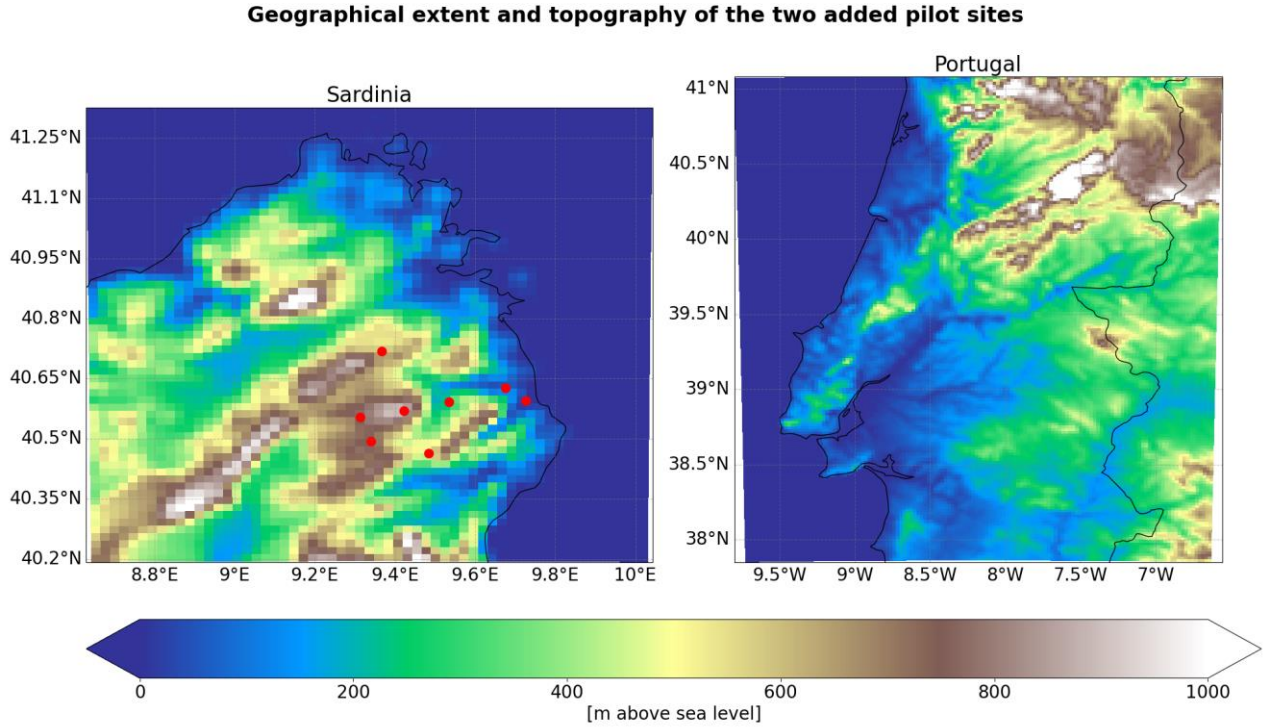


**Figure 10.** Forecast simulation domains and topography (in meters above sea level) of the Tepilora, Sardinia (left), and Cova da Beira, Portugal (right), pilot sites. The red dots in the Sardinia domain (left) denote the locations of the station observations used for the model evaluation.

The hindcast simulations were driven by hourly GFS forecasts provided by the NCEP. GFS consists in a global model with a horizontal resolution of roughly 28 km (for forecasts in the range of less than one week). The hindcasts were run over the 6-month period 01/04/2024-31/10/2024. The operational forecast runs were forced by the same data, which could be downloaded just a few hours after the forecast initialization, allowing for a more instantaneous generation of the simulations. Each forecast simulation was run at a range of 72 hours.

The performance of the WRF_2km@CMCC model in reproducing the atmospheric weather variability over the Sardinian domain and over the considered period has been preliminarily assessed using in-situ observations of eight stations located in the Tepilora park, provided to us by its administration. The locations of the single stations are represented by the red dots in Figure 10. The observations feature daily values for various temperature, precipitation, humidity, and wind variables. For the evaluation of the Cova da Beira domain, on the other hand, E-OBS [16] gridded observations were consulted.

The model's skill was investigated inspecting the most relevant meteorological variables, which are essential for the derivation of wildfire weather indicators: daily values of the mean 2-meter temperature, the accumulated precipitation, the mean wind speed, and the mean relative humidity. These variables were

assessed through PDFs and different statistical metrics mentioned below. To validate the WRF_2km@CMCC model against the in-situ observations, the grid cell closest to real location of the respective station was considered; for the validation against gridded observations, the E-OBS grid was remapped onto the WRF_2km@CMCC model grid to allow for a better comparison.

The statistical metrics used for the validation are the ones commonly used in NWP model assessment. These include the Mean Absolute Error (MAE), the Mean Bias Error (MBE), and the Root Mean Squared Error (RMSE). The equations for each of these metrics are reported below.

$$MAE = \frac{\sum_{i=1}^{n}|y_i - x_i|}{n}$$

$$MBE = \frac{1}{n}\sum_{i=1}^{n}(y_i - x_i)$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - x_i)^2}$$

with $n$ representing the number of data points, $y_i$ being the actual target values for data point $i$, and $x_i$ the model value.

### 3.2.2 *Reusability of the data processing pipeline*

As stated in sections 3.1.3 and 3.1.4, the pipeline for collecting and processing data in the DDS (Section 2.2.2) is highly configurable and modular. For this reason, it can be deployed even outside the scope of ML application within SILVANUS. The pipeline is configured using standard TOML files which can be configured according to the need. Variables which need to be consumed from the DDS datastore are configured in a TOML file where information related to each variable in stored. Functions for processing weather variables, such as those listed in Section 3.1.4, can be defined and listed in a FIFO order in the TOML file. The simplicity of the pipeline allows it to be reused by other partners who might want to consume variables from DDS and then applied a series of transformation to the data.

### *3.2.3 Probabilistic FWI based on seasonal forecast*

Using the variables outlined in Section 3.1.2, we computed the FWI [17] at a seasonal scale following the method described in [18] for the CMCC sps3.5 model. We have also downscaled the horizontal resolution in the CMCC seasonal forecast model using the functions CST_Calibration and CST_BiasCorrection developed in the Rpackage CSTools ( [19], [20]). We have use ERA5 at 0.25º of horizontal resolution for the downscaling.

To improve the skill of the forecast it is applied some recently developed sub-sampling techniques based on the representation of atmospheric patterns in the model with respect to the ERA5 reanalysis. The sub–sampling approach it is further strengthened and complemented by the design and development of novel additional techniques of subsampling based upon physics-informed ML/AI methods (dynamical systems theory, [21], [22] and [23], using the function Proxiesattractor within CSTools Rpackage). In that way we diagnose the forecast trajectory of each ensemble member and, again by using ML techniques trained on the CMCC hindcasts, algorithms are produced to eliminate from the ensemble those trajectories which are heading toward the “wrong” sectors of the real-world climate attractor.

After the subsampling of members, the probability of the forecast as in Figure 11 is performed. In the figure it is shown the different categories of the FWI in colours and the probability in the ensemble for each category. This is performed in the hindcast period (1993-2016) but also in the forecast period (i.e. case of June July August (JJA) of 2024).

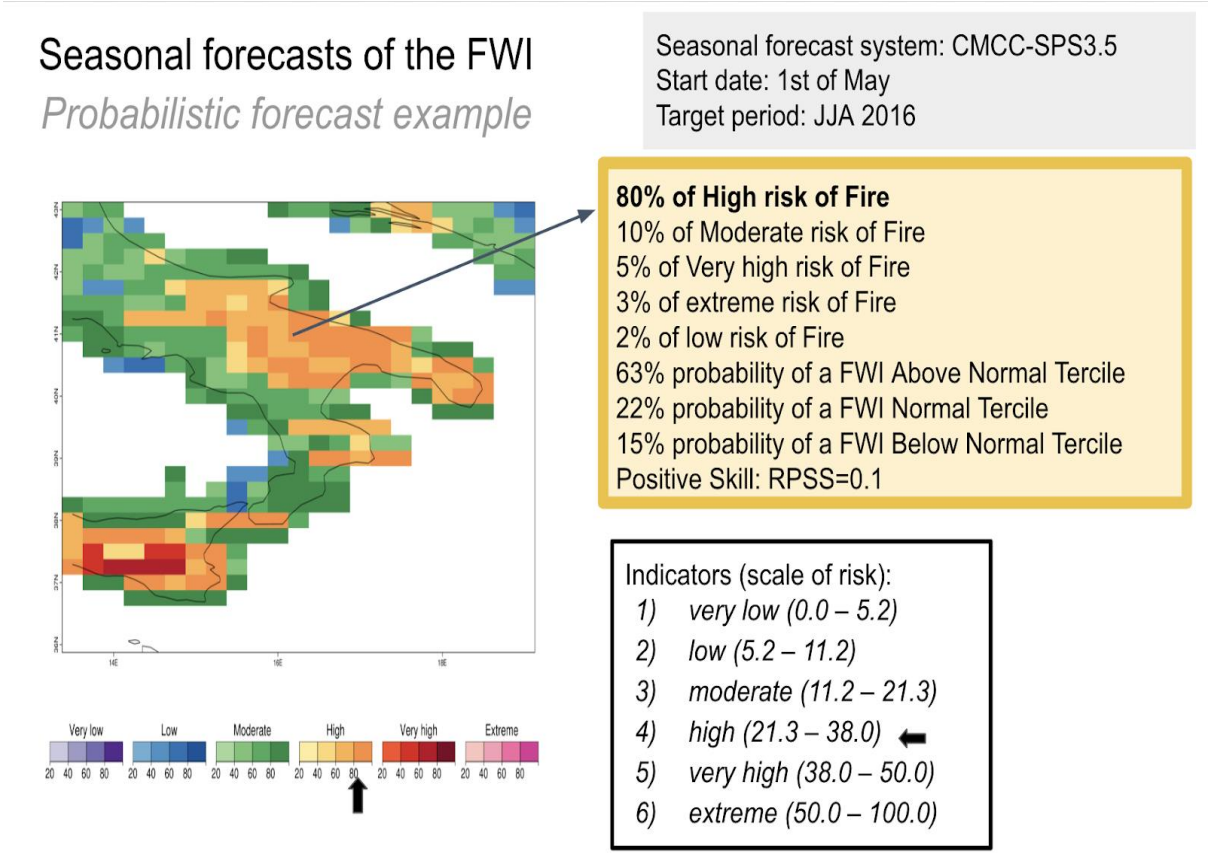


**Figure 11.** Probability seasonal FWI for CMCC model in an example of hindcast, JJA of 2016. There is a colour scale for each category of the FWI with light or dark colours in base of the lower probability or higher probability of being in that category. Each point has also information about the probability to be below, normal or above normal conditions to know how normal or extreme is a forecast/hindcast.

## 3.3 Scientific results and drawbacks

### *3.3.1 Evaluation of WRF_2km@CMCC performances over the Tepilora domain*

Figure 12 displays the PDFs of the daily mean near-surface temperature for the in-situ observations (solid line) and the WRF_2km@CMCC model (dashed line). The solid line describes the distribution considering the values of all stations, while the dashed line includes all the values related to the grid cells closest to each of the single stations. Overall, the modelled function shows a good agreement with the observed one, especially for temperatures below 8°C. The stations as well as the WRF_2km@CMCC model record temperatures in a range of 3°C and 32°C. Both observed and modelled distributions identify two temperature peaks, yet at slightly shifted values; the stations register a local maximum at around 15°C and a somewhat higher probability density of temperatures around 22°C. The two peaks of the WRF_2km@CMCC model, on the other hand, lie at temperatures around 19°C and 25°C, both values showing very similar probability density values. Furthermore, the model underestimates the probability of moderate temperatures between 10°C and 17°C and overestimates the occurrence of mean daily

temperatures warmer than 25°C. Despite these small differences, the model well represents the range of value and its skewness toward higher values.

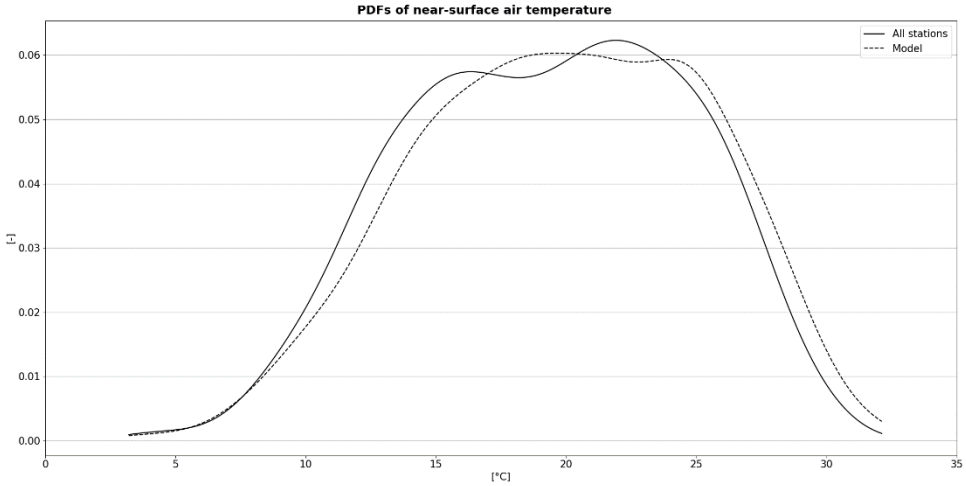


**Figure 12.** PDFs of the 2-meter air temperature (in °C) averaged over all stations (solid line) and over all single model grid points closest to each station (dashed line).

Figure 13 shows the PDFs of daily accumulated precipitation of the WRF_2km@CMCC model (dashed line) compared to the distribution of the in-situ observations (solid line). Again, all station observations have been included, while in the model, the grid point closest to each station was considered. In the case of precipitation, the model distinctly overestimates the non-precipitation events with respect to observations, whereas the probability density of light-rainfall events well reproduces the reference precipitation events recorded by the weather stations.

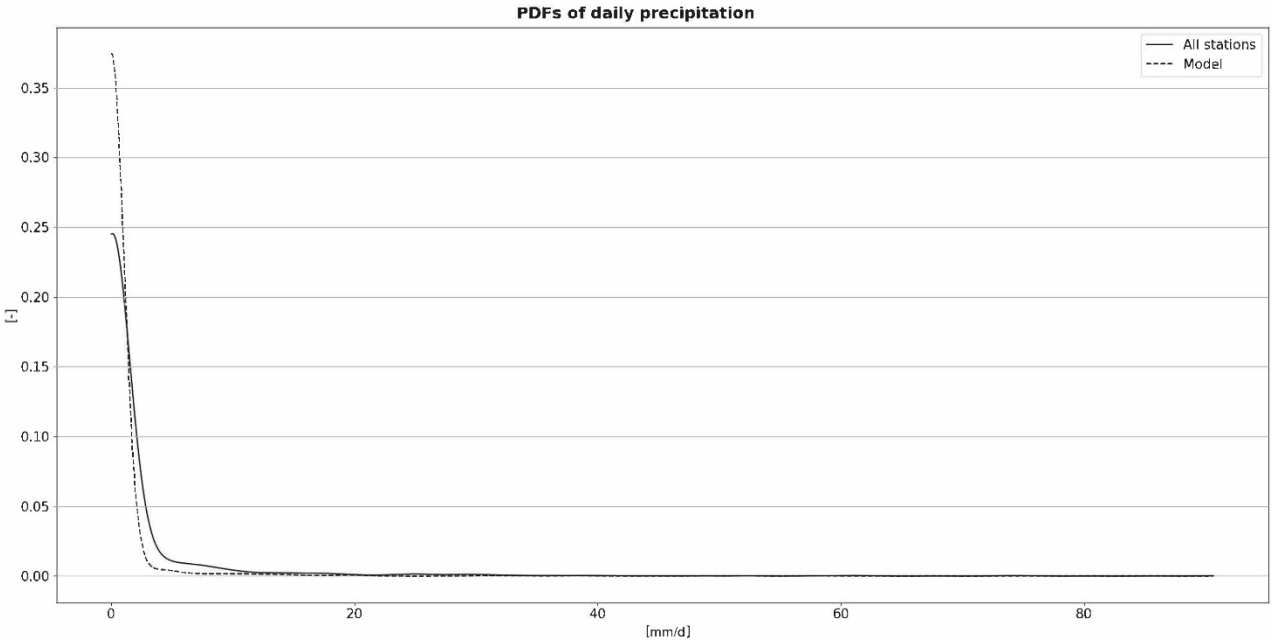


**Figure 13.** PDFs of the daily accumulated precipitation (in mm/day) over all available stations within the Tepilora Park (solid line) and over all single model grid points closest to each station (dashed line).

Figure 14 shows the PDFs of the mean daily wind speed as recorded by the weather stations (solid line) and as simulated by the WRF_2km@CMCC model (dashed line). The solid line represents the distribution over all in-situ observations, while the dashed line contains the values from the grid cell located the closest to each measurement station. The shapes of the two distributions are roughly of similar appearance, with

slightly shifted probabilities. For mean wind speeds of 12 m/s and higher, the observed and modelled PDFs agree with one another. Observations as well as the model recognize a maximum at wind speeds of around 2.4 m/s, whereupon WRF_2km@CMCC slightly underestimates these peak wind speeds with respect to stations. Moreover, WRF_2km@CMCC lightly overestimates wind speeds between 3 m/s and 12 m/s. Although the wind is strongly dependent on location, the model results look anyway consistent with the observations.

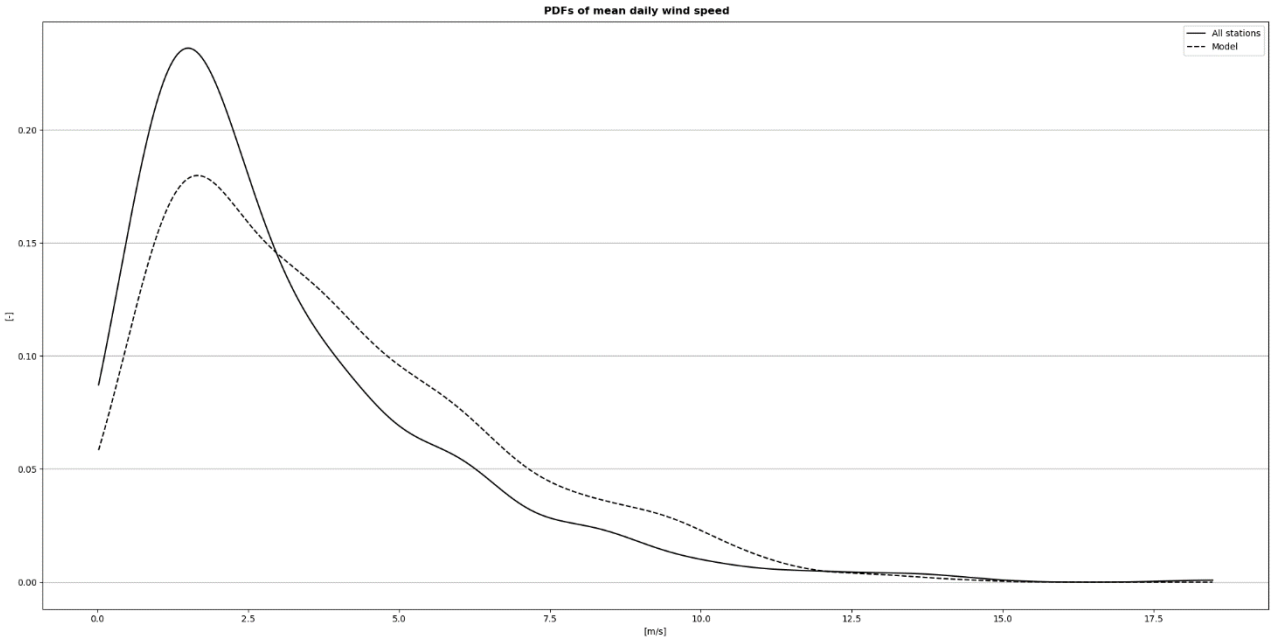


**Figure 14.** PDFs of the daily mean wind speed (in m/s) of all station observations (solid line) as well as of all single model grid cells closest to each station (dashed line).

Figure 15 displays the PDFs of the daily mean relative humidity of the in-situ observations (solid line) and the WRF_2km@CMCC model (dashed line). As in the previous examples, the PDFs include the values registered at the available stations and those simulated at the grid point that lies closest to each station, respectively. Both curves have a shape very similar to a normal distribution and agree on the probability density of low levels of relative humidity of roughly 32% and lower. The model, however, overestimates the observable maximum; while the relative humidity recorded at the stations lies principally around levels of 65%, the model primarily simulates values around 55% at the same locations. On the other hand, the model underestimates humid situations with levels of 60% and above.

Further insight into the model behaviour over the Sardinian domain can be gained by looking at the MAE, the MBE and the RMSE. Figure 16 reports the values of these metrics for the investigated variables at each observation station. It shall be noted, however, that not all stations record the variables used for the model evaluation. The overall picture reveals only small model errors with respect to observations. In line with the similar PDFs for the two-meter air temperature (Figure 17), there is a relatively low bias in this variable with respect to in-situ observations. The error is highest with respect to the station at Siniscola, Nuoro (40.57 °N, 9.7 °E) with a mean bias of 0.9°C. This station, however, lies close to the sea and might thus be strongly influenced by marine dynamics, which cannot be well represented by the model since no atmosphere-ocean coupling is here implemented. The overall error for the accumulated precipitation is of similar magnitude. This variable represents the only one that is registered by every station. Figure 13 reveals how there is a light underestimation of rainfall by WRF_2km@CMCC, in line with the overestimation of non-precipitation events visible in Figure 16. The error has a moderate magnitude, lying between -0.6 and -1.4 mm/day across all stations. As for the daily mean wind speed and the daily mean relative humidity, here

too the errors remain relatively small. These variables, however, are measured only at three of the seven available stations. The MBE for wind is in the range -0.3 to 1.3 m/s, while relative humidity is being underestimated by WRF_2km@CMCC with respect to the three stations, showing an MBE between -2.8% and -9.4%. This is consistent with the negative MBE of accumulated precipitation when evaluating against all seven stations.

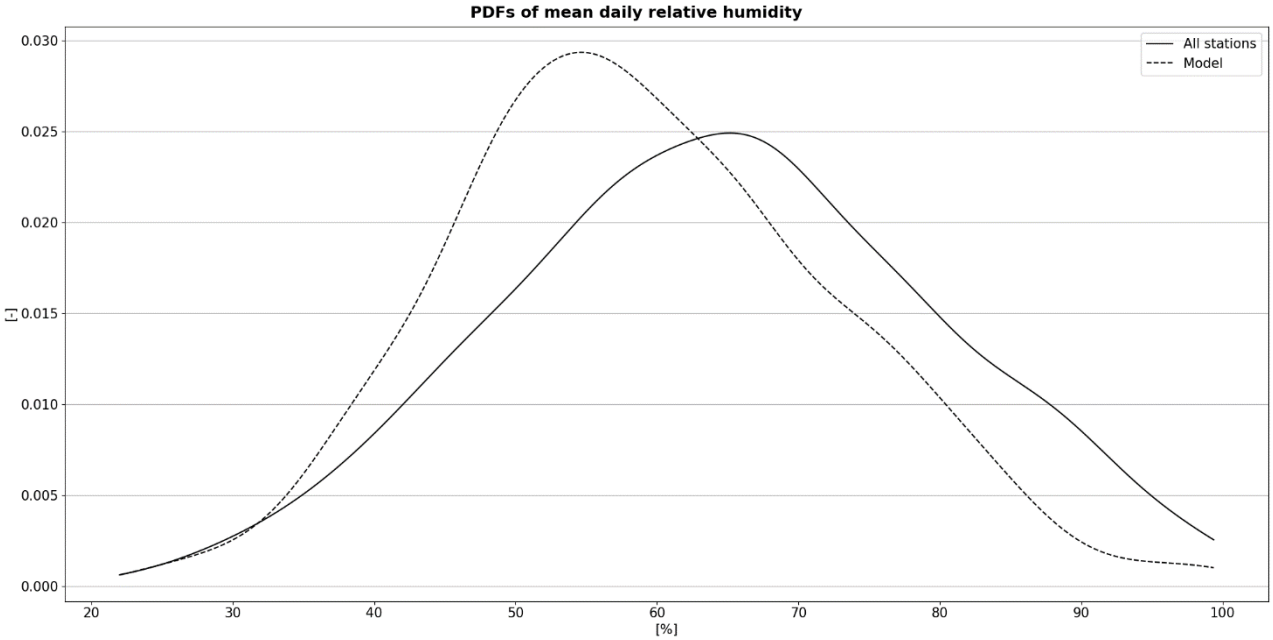


**Figure 15**. PDFs of the daily mean relative humidity (in %) of all available in-situ observations (solid line) as well as of all single model grid points located the closest to each station (dashed line).

| ***2m-temperature*** | ALA' DEI SARDI RU | BITTI RU | DIGA SOS CANALES | LODE' RU | LULA RF | ONANI' MAMONE | SINISCOLA RU | TORPE' RU |
|---|---|---|---|---|---|---|---|---|
| MAE | 0.6 | 0.7 | - | 0.7 | - | - | 1.3 | 0.8 |
| MBE | 0.3 | -0.3 | - | 0.3 | - | - | 0.9 | 0.5 |
| RMSE | 0.8 | 0.9 | - | 1 | - | - | 1.5 | 1 |
| | | | | | | | | |
| ***Daily accum. precipitation*** | | | | | | | | |
| MAE | 1.7 | 1.4 | 1.1 | 1 | 1.2 | 1.6 | 1 | 1.4 |
| MBE | -1.3 | -0.9 | -0.6 | -0.8 | -0.8 | -1.1 | -0.9 | -1.4 |
| RMSE | 5.5 | 4.4 | 3.7 | 3.8 | 5 | 5.3 | 4.5 | 6.2 |
| | | | | | | | | |
| ***Daily mean wind*** | | | | | | | | |
| MAE | 1.2 | 0.7 | - | - | - | - | 1.4 | - |
| MBE | 1.1 | -0.3 | - | - | - | - | 1.3 | - |
| RMSE | 1.5 | 1 | - | - | - | - | 1.8 | - |
| | | | | | | | | |
| ***Daily mean rel. humid.*** | | | | | | | | |
| MAE | 4.7 | 5.8 | - | - | - | - | 10.3 | - |
| MBE | -2.8 | -4.4 | - | - | - | - | -9.4 | - |
| RMSE | 6.4 | 7.4 | - | - | - | - | 12 | - |

**Figure 16.** MAE, MBE and RMSE of daily 2m air temperature, daily accumulated precipitation, daily mean wind speed, and daily mean relative humidity of the WRF_2km@CMCC model against observations (where available).

### *3.3.2 Evaluation of WRF_2km@CMCC performances over the Portugal domain*

Figure 17 displays the probability density functions (PDFs) of the daily mean near-surface air temperature for the WRF_2km@CMCC model (dashed line) and E-OBS gridded observations (solid line). For the calculations, the model grid was remapped onto the E-OBS grid and ocean grid points were masked. The two PDFs have quite similar shape; in both cases, local maxima are identified, albeit at slightly different

values. E-OBS values show a higher probability of occurrence of temperatures around 15 °C, 21 °C, and 24 °C. The model, on the other hand, slightly overestimates the first of the three maxima (as it does in general with temperatures below the 17 °C mark). The second maximum, in order of increasing temperature, has a slightly higher probability density of temperatures around 20 °C. Values between 21 °C and 30 °C are being underestimated by the model. Furthermore, it does not simulate the local maximum registered by the gridded observations at around 24 °C. Very elevated temperatures of > 30 °C occur slightly more in the model than in E-OBS.

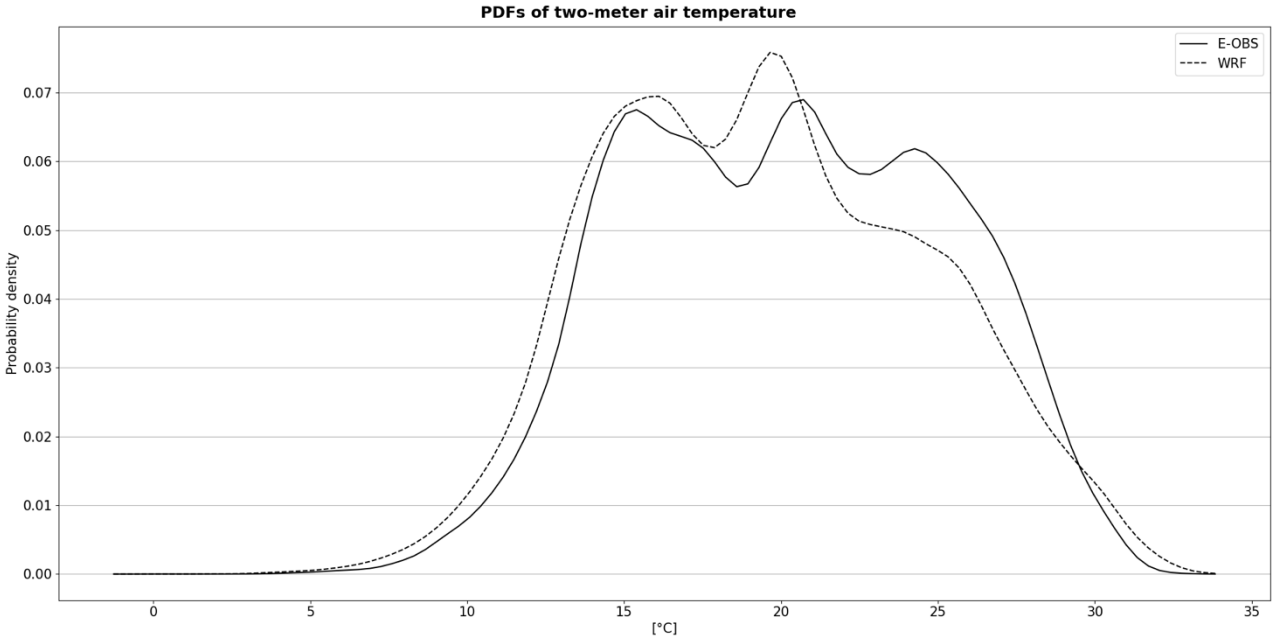


**Figure 17.** Probability density functions (PDFs) of daily two-meter air temperature (in °C) calculated over all WRF_2km@CMCC land grid points (dashed line) against E-OBS gridded data (solid line) over the same domain. The investigated period is April-October 2020. To allow for a more consistent evaluation, the model grid was remapped onto the E-OBS grid.

Figure 18 shows the cumulative frequency function of precipitation of the WRF_2km@CMCC model (dashed line) and E-OBS data. To account for the relatively low probability of partly very elevated rainfall estimates, this kind of distribution was preferred to PDFs as done for the remaining variables. Again, the model data was remapped onto the grid of the observations and only land grid points were considered. The two distributions are in relative agreement with one another until precipitation amounts of roughly 50 mm/day, after which the frequencies drastically diverge. Observations record values of around 60 mm/day; the model, on the other hand, simulates values of more than 100 mm/day. At this point, it needs to be highlighted that E-OBS data is homogenized spatially and highly dependent on the distribution of in-situ observations within the geographical domain considered. Values and statistics of potentially local phenomena such as precipitation might thus not be fully representative, which is an aspect to be considered in the analysis.

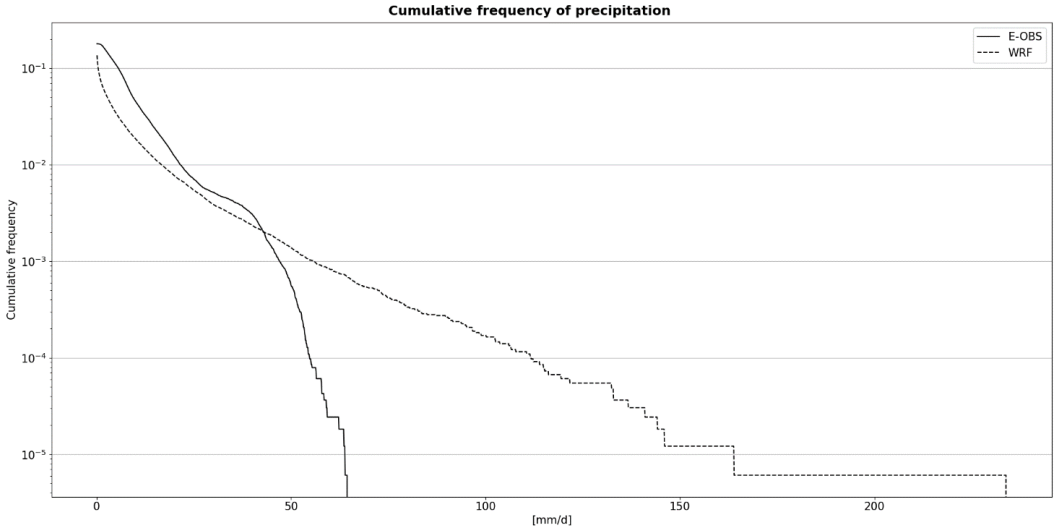


**Figure 18.** Cumulative frequency of precipitation (in mm/d) as simulated by the WRF_2km@CMCC model (dashed line) and the equivalent domain of E-OBS gridded observations. For the model values, the grid was remapped onto that of E-OBS and only land cells were considered. The period investigated is April-October 2020.

Figure 19 displays the PDFs from values of the mean daily wind speed as estimated by the WRF_2km@CMCC model (dashed line) and as recorded by E-OBS gridded observations (solid line). Since E-OBS gridded observations only include values over land territories, the model grid was remapped onto the observations' grid and values over oceanic grid cells masked out. This kind of distribution is relatively similar to what can be observed in Figure 8 in the evaluation of the Tepilora domain. In the E-OBS data, there is a peak probability density at wind speeds around 2 m/s. The distribution of the values simulated by WRF_2km@CMCC register a similar peak, albeit at a lower density and slightly higher wind speeds; here, the maximum lies at around 3 m/s. Furthermore, the respective ranges of the largest fraction of wind speed values present in the datasets slightly differ across the two data sets. Most of the values recorded by the gridded observations lie between 0 m/s and 4 m/s; the values simulated by WRF_2km@CMCC, however, often reach speeds of up to 6 m/s.

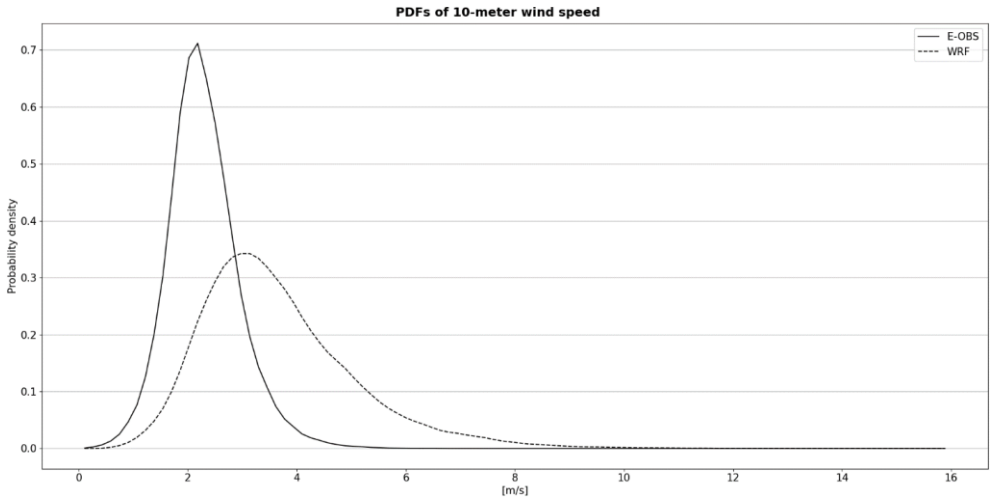


**Figure 19.** PDFs of daily mean wind speed (in m/s) calculated over values simulated by the WRF_2km@CMCC model (dashed line) and values of E-OBS gridded observations (solid line). For the PDFs, the model grid was re-gridded onto the E-OBS grid and only land values were considered. The values were calculated over the period April-October 2020.

Figure 20 visualizes the PDFs of near-surface relative humidity calculated from WRF_2km@CMCC model (dashed line) and observed (solid line) values. As before, simulated values were remapped onto the observations' grid and ocean grid points were masked out to be in line with E-OBS values. The overall structure of both curves is in coarse agreement. Both functions agree on an increasing probability density towards medium-to-high humidities of around 70%. Low humidities < 50% are slightly overestimated by WRF_2km@CMCC, as are values around the 70% mark. Medium and elevated humidities around 60% and > 80%, respectively, are being underestimated compared to E-OBS.

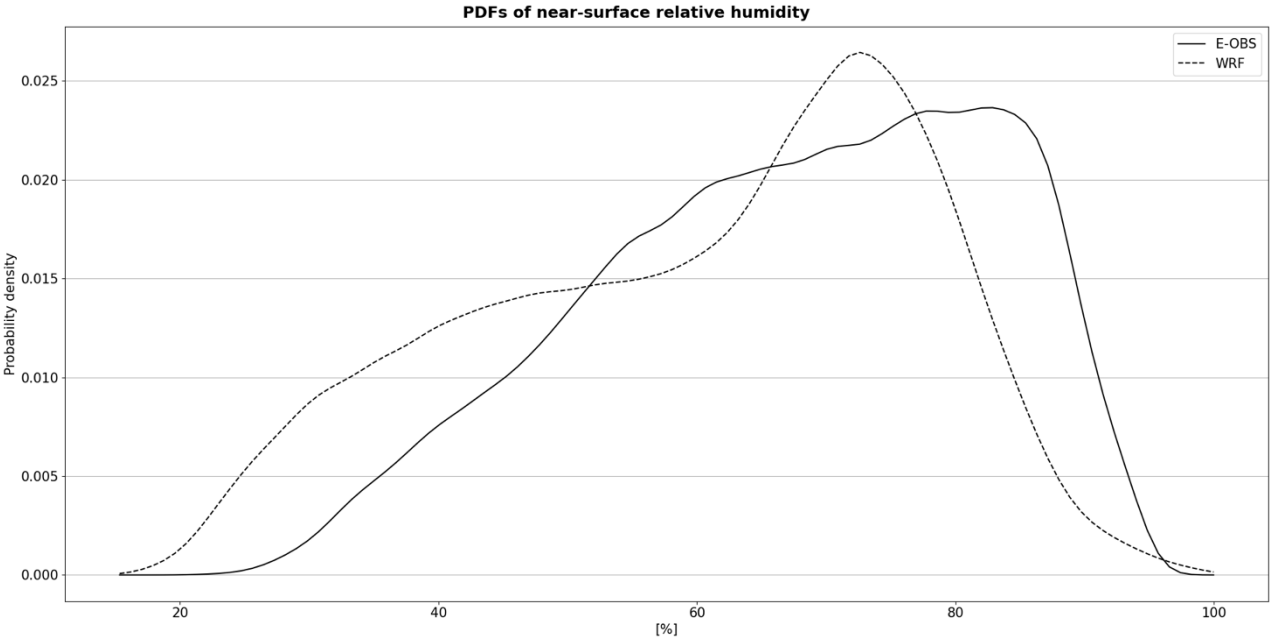


**Figure 20.** PDFs of near-surface relative humidity (in %) as recorded by WRF_2km@CMCC (dashed line) and E-OBS gridded observations (solid line).

Further insight into the behaviour of the WRF_2km@CMCC model over the extended Cova da Beira domain can again be obtained by investigating the mean absolute error (MAE), the mean bias error (MBE), and the root mean squared error (RSME). Figure 21 displays the values of these metrics for the evaluated variables with respect to the E-OBS dataset. As with the Tepilora domain, the large picture reveals rather small model errors with respect to gridded observations. In line with the rather similar PDF for the two-meter air temperature (Figure 12), there is a relatively low MBE in this variable with respect to E-OBS data; on average, the model reports a temperature 0.7 K cooler than what is recorded by observations. It needs to be kept in mind, however, that although oceanic grid points of WRF_2km@CMCC have been masked out, coastal grid points might still be influenced by the proximity to the sea. Furthermore, Figure 21 clarifies that results are similar for the daily accumulated precipitation, although the model has proven to simulate some single values to be much higher than what can be observed just from E-OBS (Figure 18). With values of 1.2 mm/day, -0.8 mm/day, and 4 mm/day for MAE, MBE, and RSME, respectively, the model error for this variable is of moderate magnitude. As for the daily mean wind speed and the daily mean relative humidity, here too the errors remain of low magnitude. The mean bias for wind is around 1.4 m/s, whereas relative humidity is, on average, being underestimated by the model, lying at -7.3%. Again, this is consistent with the slightly negative MBE of daily accumulated precipitation when evaluating against gridded observations.

| | *2m-temperature* | *Daily accum. precipitation* | *Daily mean wind* | *Daily mean rel. humid.* |
|---|---|---|---|---|
| **MAE** | 1.3 | 1.2 | 1.4 | 8.8 |
| **MBE** | -0.7 | -0.8 | 1.4 | -7.3 |
| **RMSE** | 1.6 | 4 | 1.8 | 10.9 |

**Figure 21.** Mean absolute error (MAE), mean bias error (MBE), and root mean squared error (RMSE) of daily 2m air temperature, daily accumulated precipitation, daily mean wind speed, and daily mean relative humidity of the WRF_2km@CMCC model against E-OBS gridded *observations.*

In general, despite the slight differences between the model and the reference data, WRF_2km@CMCC achieve good performances and therefore may be used for the Fire Danger and Fire Weather Risk Indicators estimations.

Finally, it is worth to mention that the goal resolution of few hundred meters has not been achieved through statistical downscaling techniques over the pilot sites starting from the output of WRF model, since the weather stations on the ground did not cover longer historical period (at least 30 years) and the distance among them did not fit the desired resolutions.

### *3.3.3 Evaluation of CMCC seasonal FWI at European domain*

For each start date of the forecast that we have explored (with focus on May for the study of JJA season and August for the study of September October November (SON) season) we have evaluated the ability of the model in reproducing the FWI as in the reanalysis. We perform the ACC between the FWI computed with the CMCC model and the FWI of ERA5 in JJA and, as is shown in the Figure 22, the red colours indicates where the seasonal forecast system can reproduce the internal variability of ERA5. Southwestern Europe, the Northern of Africa and Turkey has capability of reproduced it while the southeastern of Europe has more difficulties

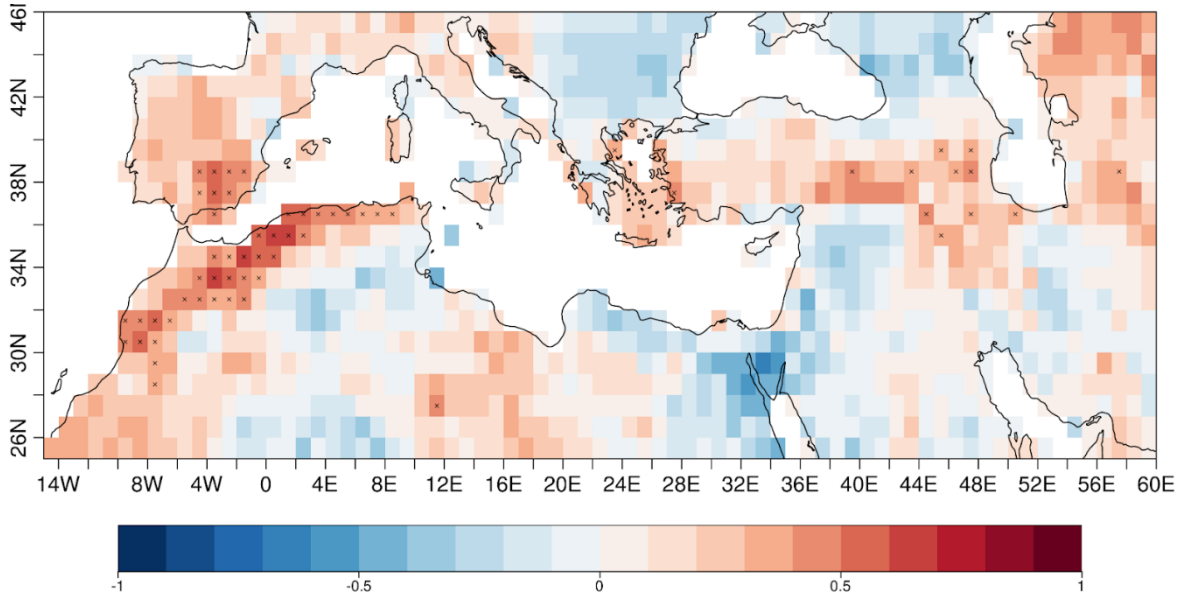


**Figure 22.** Correlations (ACC) between JJA CMCC seasonal FWI (start date of May) and ERA5 FWI. Red values correspond to those regions where the seasonal forecast system can reproduce the interannual variability of the FWI in JJA.

## 3.4 Demonstration report

The use of technologies developed in this task, related to weather forecast and estimation of FWI have been developed based on the interaction with the stakeholders involved in the pilots collected during several meetings. Specifically, below is the detail about the activities performed for the pilots that have

been demonstrated in the demo performed in WP9. These details include the visualization of a specific weather forecasting time step for the temperature and relative humidity at 2m for each pilot region computed by WRF_2km@CMCC, which are two of the essential variables for fire risk assessment. The weather forecasts have been calculated at a range of 72h, with the main features of the configuration used for the operative runs summarized below:

- Vertical coordinates: *sigma-pressure* with 60 vertical levels
- Time Integration Scheme: 3$^{rd}$-order Runge-Kutta (RK)
- Horizontal discretization: Arakawa C staggered grid
- Spatial Integration scheme: 6° order centered difference
- Latitude-longitude grid: Lambert conformal projection
- Time frequency: 1 hour
- Horizontal spatial resolution: ~2km

The semi-operational runs for each of the three pilots were performed on the JUNO supercomputer of the CMCC Supercomputing Center (SCC). JUNO hosts 12.240 processor cores and 170 dual processor nodes, with each node featuring a memory of 512 GB.

### *3.4.1 Apulia-Italy Pilot*

The Italian region of Apulia was the first territory to be considered as a pilot site within the framework of this task. A major motivation to focus on Apulia, apart from the promising results obtained in the past with the WRF model over this domain, lies in the presence of the Gargano Promontory, which features forests prone to wildfires. Figure 23 shows the weather forecast of the 2m-temperature (left) and relative humidity at 2m (right) computed with WRF_2km@CMCC for the 28/07/2024 at 00:00:00 UTC+2, with the initialization performed on the 27/07/2024 at 00:00:00 UTC+2, over the Apulia pilot domain. In the horizontal, the domain consists of 210x210 grid points. The runs are performed allocating one node on JUNO with 32 processors. Each 72h-forecast takes up to 6.5h.

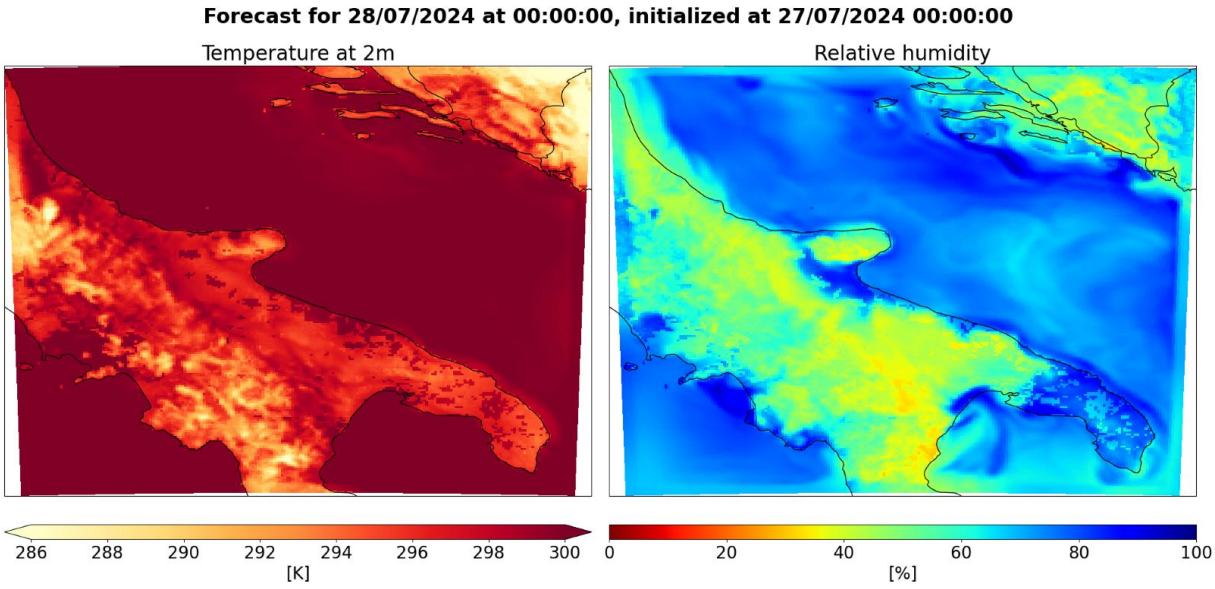


**Figure 23.** WRF_2km@CMCC forecast for the 28/07/2024 at 00:00:00 UTC+2 (initialized 27/07/2024 at 00:00:00 UTC+2) of the temperature at 2m (left) and relative humidity at 2m (right) over the numerical domain of the Apulia pilot.

### *3.4.2 Tepilora-Italy Pilot*

Tepilora describes a regional natural park located in north-eastern Sardinia, Italy. Figure 24 displays the output of WRF_2km@CMCC of the forecast time step on 28/07/2024 at 00:00:00 UTC+2 for 2m-temperature and 2m relative humidity over the numerical domain of this pilot. Here again, the forecast run was initialized at 27/07/2024 at 00:00:00 UTC+2. In this case, the pilot domain includes 3.840 grid points (60 W-E x 64 N-S). Owing the smaller size of the domain with respect to the Apulia pilot, again one computation node was allocated, but with just 18 processors. The entire 72h-runs have a duration of around 45min.

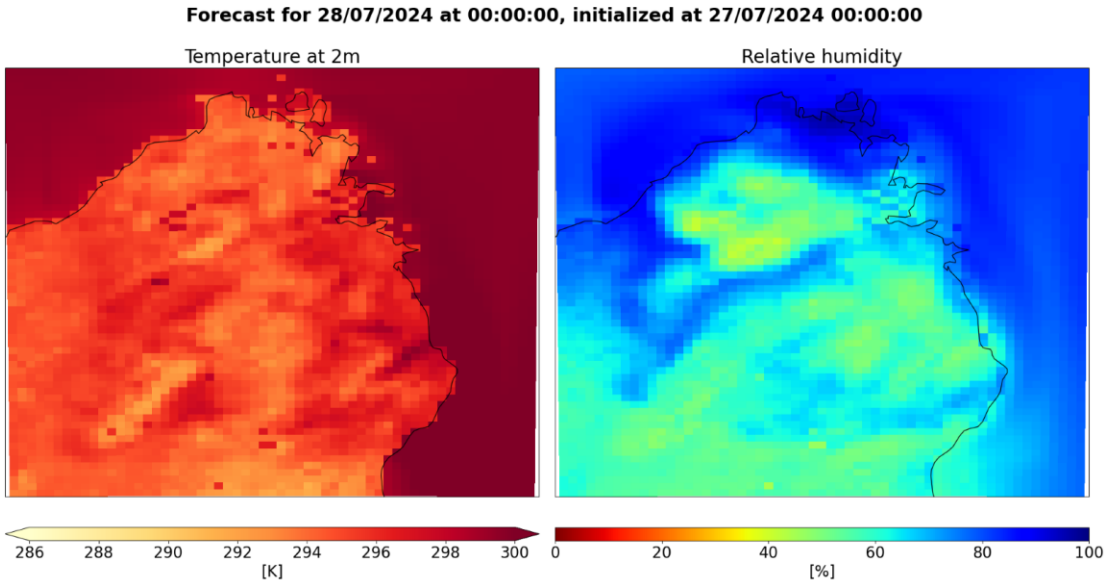


**Figure 24.** WRF_2km@CMCC forecast for the 28/07/2024 at 00:00:00 UTC+2 (initialized 27/07/2024 at 00:00:00 UTC+2) of the temperature at 2m (left) and relative humidity at 2m (right) over the numerical domain of the Tepilora pilot.

### *3.4.3 Cova de Beira-Portugal Pilot*

The third pilot site is the Portuguese subregion of Cova da Beira. Figure 25 visualizes the forecast step on 28/07/2024 at 00:00:00 UTC+2 of the 2m-temperature and the 2m relative humidity over the respective simulation domain. As in the previous two examples, the corresponding forecast was initialized at 27/07/2024 at 00:00:00 UTC+2. The domain comprises 137 grid points in west-east and 180 grid points in north-south direction. As the extent of the domain is somewhat similar to the Tepilora pilot, the same allocation of resources is used to perform the forecasts (i.e. 1 node with 18 processors). One 72h forecast run takes approximately 4h.

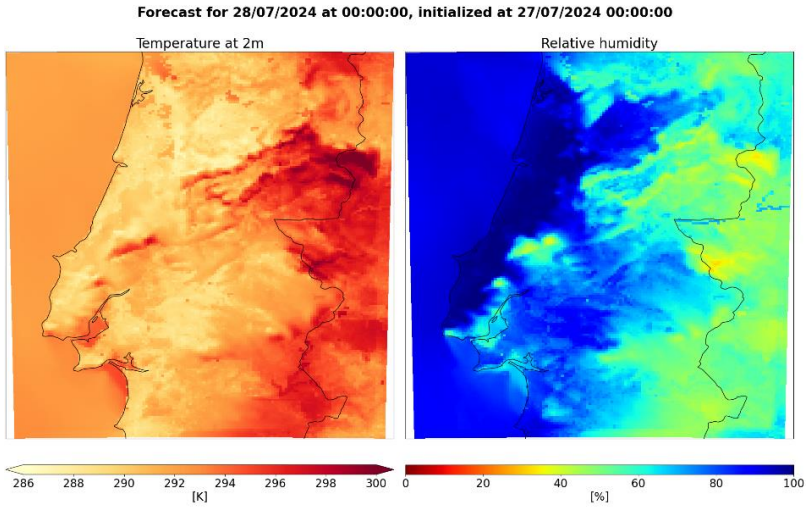


**Figure 25.** WRF_2km@CMCC forecast for the 28/07/2024 at 00:00:00 UTC+2 (initialized 27/07/2024 at 00:00:00 UTC+2) of the temperature at 2m (left) and relative humidity at 2m (right) over the numerical domain of the Portugal pilot.

# 4 Data collection, aggregation, and pre-processing of in-situ devices

## 4.1 Tool description

The task includes the development of a distributed data architecture and software protocols to communicate with various in-situ data collection devices, such as IoT devices, EMDCs, and FCCs. IoT devices developed for UP4a and UP9b detect fire/smoke and assess air quality, respectively, by utilizing ML algorithms and collecting sensor data from the field. EMDCs (UP4b) collect and process images/videos to identify fire occurrences and serve as gateways for data transmission to the SAL. FCCs (UP10) are employed to support commanders with vital services such as the Data Ingestion Pipeline (DIP), the SILVANUS dashboard (UP11), and data synchronization with the cloud, which is important in cases that lack network connectivity or cannot afford data transmission latencies. By including these devices and protocols the project facilitates the handling of fire emergency events by enabling local data processing, storage, and visualization. Some of these technologies and protocols were tested in the 2023 pilot activities (see D9.3). In upcoming pilot exercises, we will conduct more thorough tests on these technologies and protocols, while focusing on those that have not yet been tested in the field.

## 4.2 Innovations and updates

### *4.2.1 Fire detection from IoT devices*

Several updates and improvements have been made during the M19-M36 period both on hardware and software regarding fire detection from IoT devices (UP4a). These improvements address key challenges encountered during pilot tests and exercises, such as the need for reliable data transmission, efficient cooling, and fast data processing, making sure the IoT is responsive in several operational scenarios. An overall description of the device, the produced data format and data ingestion pipeline can be found in D4.1, while details on the development and training of the fire and smoke detection models can be found in D4.2 and D5.1, respectively.

A major update is the redesigning of the case (see Figure 26), which is now larger and features better air-flow management for cooling down the device's components. The new design makes the case more ergonomic and enables improved hardware organization. One example of these updates is the inclusion of a placeholder for the Global Positioning System (GPS) antenna, which makes it more stable and ensures there are no dangling wires, minimising the chances they get damaged by forest animals and harsh weather conditions.


**Figure 26.** (Left) New case for our IoT, (Right) Previous IoT case.

Another important update is the successful integration with SAL, enabling data transmission from IoT devices (and IoT Gateways), their ingestion in the SILVANUS cloud and/or FCCs, and finally, their visualization on the SILVANUS dashboard (see Figure 27). In the dashboard end users can view archived IoT data, but most importantly get the status of the deployed IoTs and receive notifications for any detected fire or smoke events from the IoT captured images. This addresses real-time data monitoring and response.

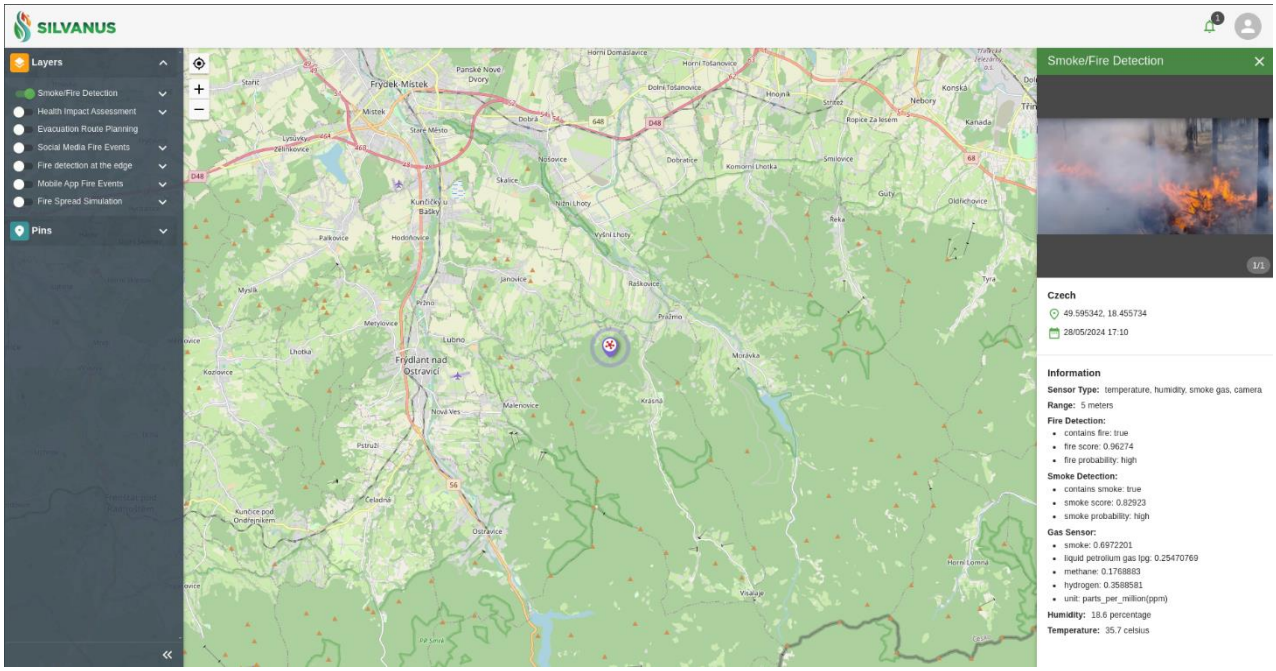


**Figure 27.** Example of SILVANUS dashboard for Smoke/fire detection (IoT) layer**.**

Furthermore, the IoT's software was updated to further improve its functionalities. In detail, the code was optimized by applying concurrent programming techniques, thus separating the data capturing, processing, and transmission operations. This improvement significantly reduced the time required for the device to alert all responsible personnel about fire-related events addressing the need for rapid response in emergency situations. In addition, leveraging images collected from pilots the performance of the fire/smoke detection ML algorithms were further increased by applying finetuning techniques.

Additionally, the synergy between IoT devices and EMDCs was tested in two different scenarios. The first one, was to use EMDCs as a gateway. In other words, the IoT devices deployed in the field transmitted data to the gateway (EMDC), which was then responsible for the propagation of that data to SAL. In the second scenario, the EMDC was used as another IoT device, that would be deployed on a fire engine or in a building in the area of interest (e.g., lookout tower), to ensure it has adequate power supply. A camera sensor, such as a Universal Serial Bus (USB) or Internet Protocol (IP) camera, was connected to the EMDC to collect images that were later analysed by the detection ML algorithms for any potential fire or smoke outbreaks. These updates enhance connectivity and data flow, important to effectively monitor the field for possible fire/smoke events in heterogeneous scenarios.

### *4.2.2 Fire detection from edge devices*

The tool detects fire and smoke using near real time footage or photos taken by drones or cameras.

This tool provides a fire and smoke detection using high-end devices for the videos or photos taken from any source. It can analyse the images in near real time (less than 1 sec). The objective of the tool is to provide a very reliable detection of fire and smoke during the phases of prevention (A) or fight against the fire (B). With this automatic detection, the operator doesn't need to check manually and continuously all cameras. The system can raise an alert based on the automatic detection to make the operator aware of the possible danger whenever a fire/smoke is detected. In this sense, is a tool for supporting the decision and helping the operator by providing evidence (in the form of images with fire or smoke detected) of what is really happening. A detailed description of the tool can be found in D5.1.

During the last period, the algorithm used has been retrained using new images of fire and smoke in forest areas. Also, we tested the full integration in remote and corrected the issues of connectivity with Kafka and the dashboard (format of the JavaScript Object Notation (JSON) and general connectivity issues with Kafka).

### *4.2.3 IoT for air quality assessment*

Making use of a network of IoT devices, both portable and stationary, that are equipped with various gas and particle sensors, this module effectively monitors the spatial evolution of pollutant levels. The data collected from these in-situ observations is transmitted wirelessly to a remote server, which is tasked with implementing the European Air Quality Index (EAQI) methodology to assess air quality, and subsequently sharing the inferred outcomes via suitable RESTful APIs and SILVANUS infrastructure to relevant stakeholders. The architecture of the Air Quality Assessment (AQA) system (UP9b) is demonstrated in detail in D5.3.

The proposed system was enriched throughout the M19-M36 period to further support the decision-making process related to fire management. More precisely, health-related messages containing guidelines for both the general population and vulnerable groups complements the deduced level of air quality index. These sensitive groups encompass adults and children suffering from respiratory issues as well as adults with heart conditions. For instance, the characterization of the ambient air quality as ‘Poor’ is accompanied, adopting EAQI recommendations, by the messages ‘Reduce intense activities outdoors, if you experience symptoms such as sore eyes, a cough or sore throat.’ and ‘Reduce physical activities, particularly outdoors, especially if you experience symptoms.’ addressing the general population and vulnerable groups, respectively. In addition, a list of key performance indicators pertaining to relative risk has been attached to the outcomes of this system, based on the recommendations provided by the Health Risks of Air Pollution in Europe (HRAPIE) project regarding concentration–response functions for key pollutants. These indicators provide a quantitative basis for estimating the health burden associated with exposure to emissions from wildfires over both short-term and long-term periods. A comprehensive description of the aforementioned AQA system updates is presented in D5.5.

### *4.2.4 Forward Command Centers*

The FCCs are key components of the SILVANUS platform, enabling data processing and decision-making directly at forest fire sites (i.e., at the edge). The FCC is designed and built as a customized replica of the SILVANUS platform, ensuring that its functionalities are available at incident sites. An FCC consists of the following components:

- **Kubernetes Cluster**: A single Kubernetes cluster to host system components and the decision support system's UPs.
- **SAL:** Enables ingested data to be stored and handled directly on the FCC.
- **Message Queue Instances**: Utilizing RabbitMQ to enable north-south and east-west data stream communication.
- **Retrieval Mechanisms**: Direct and claim check pattern retrieval mechanisms, where the former is suitable for lightweight data and the latter is better suited for large data files.
- **Query Interface**: Allows retrieval of data at rest using either of the retrieval methods.
- **DIP:** Enables data preprocessing from SILVANUS data sources as well as external data sources.
- **DSS's UPs**: Play a crucial role in forest fire management. The FCC allows the DSS to operate without needing to access the SILVANUS cloud.

#### *4.2.4.1 FCC East-West Communication*

The goal of the East-West Communication API is to ensure secure communication with the platform's peripheral systems, particularly among the various FCCs within the same pilot site. The east-west API is implemented leveraging the federation mechanisms provided by RabbitMQ.

Federation in RabbitMQ allows the connection of queues and exchanges across different RabbitMQ brokers, facilitating the creation of distributed and scalable systems. Unlike clustering, federation enables flexible connections between independent brokers, ideal for scenarios such as those in SILVANUS pilots. The main advantages of federation include scalability, to distribute the workload across multiple brokers; flexibility, to allow connecting brokers across different networks or data centers; reliability, for allowing continuation of service even if one broker fails; efficiency, reducing latency in message delivery.

To implement federation in RabbitMQ, Federation Links are configured between brokers, specifying the destination broker's URI, credentials, and the queues or exchanges to be federated. Each RabbitMQ instance has its own "Master.exchange" connected to internal fanout exchanges like "air-quality" and "evacuation-routes." These "Master.exchanges" are federated with each other, facilitating message distribution among the different RabbitMQ instances. Opting for a single federated exchange simplifies management, providing a single control point for message flow, with the possibility to further refine internal routing logic using bindings with specific routing keys according to specific needs.

The FCC can be used to ingest data collected during drone flights into the SILVANUS Cloud. In the first period (till M18) we developed a process to request computation of drone flight routes (computed using algorithms developed in T4.6) using an EmerPoll system deployment at the FCC (see Figure 28). As an input a monitored area and flight latitude together with the number of available drones is inserted. As the output the optimal flight route plans are calculated using different available services for a specified number of drones. The individual routes need to be subsequently loaded to individual drones (resp. to drone pilots). The related work was also published in [24].

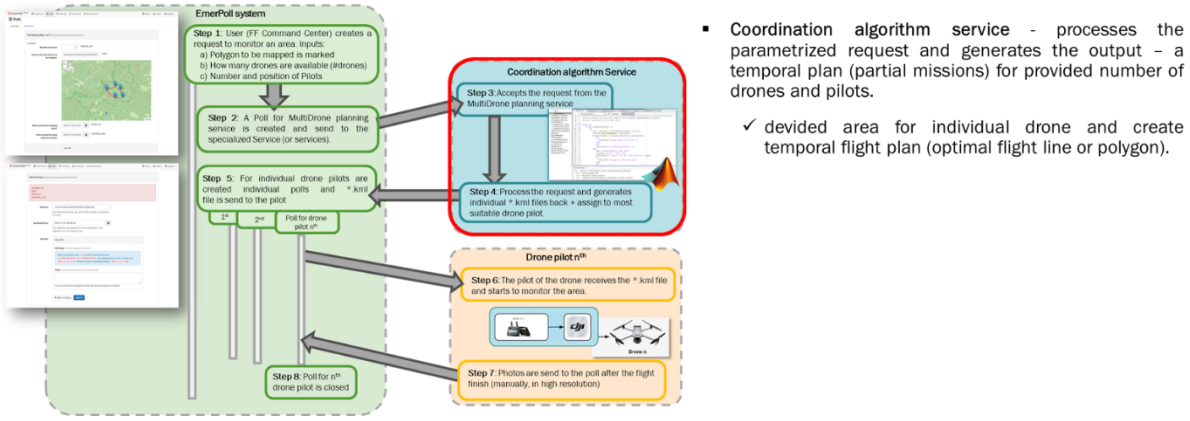


**Figure 28.** Process of computing drone flight routes at the FCC.

In the second period (M19-M36) the implementation of the approach as well as the in-field demonstration preparations were carried out. The main innovations and improvements achieved are:

- the development of the DroneReport poll template,
- an FCC-capable deployment of the EmerPoll subsystem,
- service for semi-automatic drone image stitching and
- implementation of the DroneReport and images ingestion into SAL.

The unique feature of the approach is collection and aggregation of images from many various drones (and drone types) in a uniform and integrated way. Images are primarily stored in the EmerPoll image storage. Selected images can be ingested into SAL and be passed to other components for processing, such as for fire detection (described in the previous section). Recently, a dedicated map layer is being implemented and integrated in the Dashboard to display the drone flight paths as well as the images taken by drones.

## 4.3 Scientific results and drawbacks

### *4.3.1 Fire detection from IoT devices*

UP4a - Fire detection from IoT devices, has participated in several pilots/tabletop exercises including Croatia, France, Gargano - Italy (tabletop exercise), Greece (tabletop exercise), Australia, and Czech, mentioned in chronological order. The status of UP4a related KPIs, during trial period 1, was listed in D9.3. In total, seven tests were completed in the four field exercises (Croatia, France, Australia and Czech) and the Gargano tabletop exercise, where material was collected offline by an EMDC from 2 different smoke sources. In Australia, two different tests were completed, one with a static IoT device and a moving one, were an IoT was mounted to a UGV.

An important accomplishment is that we achieved for the fire detection a false alarm rate of 5%, significantly below the 15% threshold. Nevertheless, we have more work to do for the smoke detection false alarm rate as it exceeds slightly the desired limit. Despite this, we have surpassed the expected 70% on true positive rates for both the fire and smoke with 90% and 93%, respectively. However, the missing rates for fire and smoke detection are 10% and 7%, which are above the desirable limit. This indicates a need to refine our algorithms further. Though it is anticipated that the fire detection algorithm will outperform that of smoke one, because of the phenomena form. That is fire is more easily distinguished because of its color and shape, compared to smoke which can be easily confused with mist or clouds.

Regarding the time required to identify an ignition and notify relevant personnel, the IoT performs within the expected time range. With measurements taken at 6 frames per second, the total time varies from 20.5 to 46 seconds (depending on the network connectivity), which is under the 1-minute mark. More tests will

be conducted in the upcoming periods using different network types to have a more thorough analysis of the time needed to process and send the collected data.

Despite the successful deployment and testing of the IoT device in several field exercises, some limitations still need to be addressed. The sensor measurement capabilities of the IoT are limited by design (a few meters), so the device primarily uses them to collect microclimate information rather than for fire/smoke predictions. The device therefore relies on the camera input to provide fire/smoke event predictions, which can be affected by lighting conditions and/or device positioning. For example, placing the device in direct sunlight will compromise the camera's functionality and therefore the ML algorithms' results. In addition, the camera can be obscured by dense vegetation, which will limit its detection range. It is therefore crucial to strategically place the devices in the area of interest, considering the terrain characteristics, to ensure maximum camera coverage. Additionally, the device relies on Wi-Fi or cellular networks for data transmission, which could lead to communication issues with the SILVANUS platform in their absence or if the bandwidth is limited. Moreover, the devices have limited battery life and lack autonomous charging capabilities, which restricts their operational duration in the field, although we are considering incorporating solar panels to enable them to be recharged. Another concern is that the current case is not waterproof, which could lead to damage during heavy rain and compromise the device's functionality. Furthermore, the current case material is not fireproof either, but the final product will be made from materials that are both waterproof and heat-resistant to some extent.

UP4a has demonstrated strong capabilities in the early detection of fire and smoke events, leading to minimizing wildfire casualties, cost savings from forest fires, forest conservation, and protection of wildlife and forest animal habitats. However, there is still work to be done to address the highlighted limitations. Addressing these challenges is crucial for improving the reliability and performance of the IoT devices across various environmental conditions.

### *4.3.2 Fire detection from edge devices*

The module has been used by real pilots in the field providing a very good level of detection. Also, it has demonstrated its potential for near real time analysis. As a potential drawback or limitation, we need to test further and quantify the number of false positives (that is, false fire/smoke detections) and false negatives (that is, real fire and smoke not detected). As a secondary drawback, the necessity of a high-end device (computer) needed for the analysis must be mentioned.

### *4.3.3 IoT for air quality assessment*

The real-time monitoring of the concentrations of harmful emissions released into the atmosphere during a wildfire and the consequent characterization of the ambient air quality and the compilation of the list of the relative risk indicators are valuable information for decision-makers within an integrated fire management system. Nevertheless, given that the portable IoT devices lack a fixed position and are instead transported by individuals, careful consideration should be given to the dimensions and weight of each constituent due to their impact on the personnel's mobility and endurance. In addition, some emerging challenges of the proposed system pertain to its effects on communication and energy costs. The utilization of Hypertext Transfer Protocol (HTTP) proves to be costly in terms of communication resources, leading to increased network overhead and higher mobile data usage compared to alternative communication protocols in existing literature, despite the convenient framework offered by REST API. Furthermore, the portable solution's energy consumption is elevated because of employing Raspberry Pi, in addition to 4 Inter-Integrated Circuit (I2C) sensors, 1 Universal Asynchronous Receiver-Transmitter (UART) sensor, and 1 UART 4G HAT, as opposed to a simple Micro-Controller Unit (MCU). A more comprehensive analysis of results and drawbacks is presented in D5.5.

### *4.3.4 Forward Command Centers*

The current version of the FCC has many functionalities crucial for wildfire management. However, several improvements can be made to the FCC design and implementation to make it more suitable for real-time deployment. The following are some of these improvements:

- **Dynamic Management and Enforcement for Tenant Isolation**: Currently, tenant isolation regarding data streams (i.e., over the message bus) requires manual configuration by the message bus admin. We propose an upgraded mechanism where data access requests and responses occur only between data producers and consumers, streamlining the process and enhancing security.
- **Salvaging Data with Incomplete or Incorrect Metadata**: In the current version, the SAL discards ingested data with issues related to their metadata, which can significantly impact real deployment if data sources are misconfigured. The proposed enhancement aims to reduce discarded ingested data by automatically fixing improperly attached metadata.
- **Hosting Proprietary Software Locally**: Implementing a mechanism to host proprietary software on the FCC will enable it to operate independently without needing internet access. This will enhance the FCC's reliability and functionality in remote or isolated locations.

The East-West Communication API has been implemented through Rabbit MQ Federation, which involves some specific advantages in the context of the SILVANUS platform. In particular, when connectivity is lost, messages are still queued locally, and they are consumed as soon as connectivity is back. This aspect is a clear advantage if compared with the alternative clustering approach, as federation supports failover, while instead clustering would better support use-cases in which high volume of data transfer is required and latency is minimized.

Another aspect that needs to be considered is that federation replicates messages across all RabbitMQ nodes, which on the one side ensures failover, but on the other side this adds additional latency, as replication of data across nodes requires additional time to be achieved.

The national legislative and safety requirements for pilots and drone flights (both piloted and unmanned) are a major limiting factor for deploying the developed solution for in its full potential. For instance, the pilot should always have its drone in direct visibility and needs to directly control its operations. The other potential drawbacks of the proposed solution are:

- Collection of images from drones is available only after the drone landing.
- Need to recompute routes if one of the drones becomes non-operational or unavailable or if the area becomes a non-flight zone.
- Trade-off between image file sizes - its transfer times from an FCC and time required to stitch images together. In case of larger high-quality files, a relatively high-speed data connection to ingest images from FCC into the SILVANUS Cloud is required. Long times required to stich images (for instance flight 2-UISAV: 37 min. for 140 images with 80% overlay).

The results of all the test flights are shown in Table 5.

**Table 5.** Process of computing drone flight routes at the FCC.

| Flight | # of UAV | # of images taken | Flight Time | Merging process status | Merging process time (fast-orthophoto) | Merging process time (high resolution) |
|---|---|---|---|---|---|---|
| flight 1-UISAV | 1 | 119 | 21m61s | done | 14m33s | - |
| flight 1-THALES | 1 | 119 | - | done | 15m9s | - |
| flight 2 -UISAV | 2 (120m) | 65 | 11m47s | not all | 5m11s | 17m6s |
| flight 2-UISAV | 2 (115m) | 75 | 13m49s | done | 6m1s | 20m58s |
| flight 2-UISAV | 2 | 65+75 | - | done | - | 37m56s |

## 4.4 Demonstration report

### *4.4.1 Fire detection from IoT devices*

The participation in several pilots resulted in the collection of feedback from various stakeholders and the continuous advancement of the fire/smoke detection IoT devices. Below is a timeline of the trials, challenges, achievements, improvements, and outcomes for each pilot where the IoT was tested.

#### *4.4.1.1 Croatia Pilot - 2023*

**Field testing:** First pilot to test the IoT in-situ. At the time only the fire detection algorithm was deployed on the IoT along with the Global System for Mobile Communications (GSM)/ Long-Term Evolution (LTE) adapter, and temperature/humidity sensors. The device successfully captured data, transmitted them to Catalink's server and visualised on a simple UI for validation. At the time SAL and the SILVANUS dashboard were under development and could not be tested, therefore temporary substitutes were used. The device efficiently sent data using the on-board SIM card and a wired connection to a MESH in the Sky (UP12) node, as shown in Figure 29. The device recorded 53 image-related events with GPS coordinates and sensor measurements. Additionally, 44 short videos were collected that were later used to extract more frames for the evaluation of the smoke detection algorithm.

**Challenges:** Due to heavy rain in the days prior, the firewood was soaked, resulting in only heavy smoke and very small flames, which the fire detection algorithm could not detect. This issue, combined with stakeholder feedback emphasizing that early smoke detection is crucial, highlighted the need for an additional smoke detection algorithm.

**Figure 29**. UP4a IoT device connected to the MESH in the Sky node.

#### *4.4.1.2 France Pilot - 2023*

**Field testing:** The focus of this pilot was to test the newly added smoke sensor and detection algorithm, along with the use of the EMDC as an IoT Gateway. In detail, two IoT devices were placed on trees monitoring the area the smoke would be released (see Figure 30) and an EMDC near the command centre. Data collected from the IoT devices were sent to the Gateway, which in turn transmitted data to SAL for storing. Additionally, mocks up of the SILVANUS dashboard and the IoT layer (among others) were shown to the end-users for the collection of feedback on their functionality. The devices captured 428 smoke-related events consisting of images, GPS coordinates, and sensor readings.

**Challenges:** At the time, the devices were not assigned static IPs from the SILVANUS Virtual Private Network (VPN), making communication and data exchange more challenging. To address this, a temporary VPN was set up between the IoT devices and the Gateway, but this solution made the setup less stable. As a result, some data were stored only locally in the Gateway and not transmitted to SAL. Additionally, the dense vegetation and mist-like smoke caused the smoke detection algorithm to struggle in recognizing smoke in some instances. The algorithm's low performance highlighted the absence of such challenging scenarios in the training data and underscored the importance of proper IoT device placement. To address these issues, data collected from the pilot were used to improve the algorithms, and acquiring information about the pilot location beforehand will aid in better device positioning. Lastly, with the addition of the new components to the IoT, the case became too small, and some components were exposed making the device more fragile. This emphasized the need for the designing of a new device that would better fit the new sensors.


**Figure 30.** UP4a enhanced with smoke detection algorithm.

#### *4.4.1.3* *Italy Tabletop - 2023*

Prior to this tabletop exercise, the EMDCs were setup with a USB camera (which acted as a stationary IoT device) in Gargano to collect material, such as videos, of two fire simulations using a smoke grenade and lightning torch, respectively. Later, that material was used to test both the fire and smoke detection algorithms on varying smoke densities and fire behaviours. These results were later presented to the participating stakeholders for feedback and the initiation of a discussion of how the EMDCs (with the detection algorithms) and/or IoTs could be leveraged to improve the monitoring and protection of the national park.

#### *4.4.1.4* *Greece Tabletop - 2023*

During this tabletop exercise the IoT was presented to the local stakeholders through a simulated wildfire scenario, where the functionalities and benefits of using such devices were presented, along with results from the previous pilot activities. The **collected feedback was positive** and highlighted the effectiveness and need of such fire/smoke detecting devices for the early identification of fire outbreaks and their imminent suppression.

#### *4.4.1.5* *Australia Pilot - 2023*

**Field testing:** The main objectives of this pilot was to test the live transmission of data from the devices to SAL and the synergy of UP4a and UP5a, for the detection of smoke from UGVs. During the pilot, one static and one moving IoT device were tested in variations of the smoke detection task. For the smoke detection while on a moving object, an IoT was mounted on a UGV, as shown in Figure 31, where the first was checking for any smoke outbreaks while the second was mapping the field. For the case of the static IoT, a smoke machine was used to produce smoke at different densities, to validate the performance of the smoke detection algorithm. A total of 131 events were send to SAL and 160 short videos stored locally for fine-tuning the IoT's algorithms.

**Challenges:** During the experiments, the IoT struggled to detect low-density smoke. Post-pilot analysis revealed the issue was not with the detection algorithm but with the device's programming. The device was set to capture a new image only after the previous one was successfully transmitted, causing it to miss the

low-density smoke while waiting for transmission to complete. To overcome this issue concurrent programming methods were applied to reduce the response time of the device.

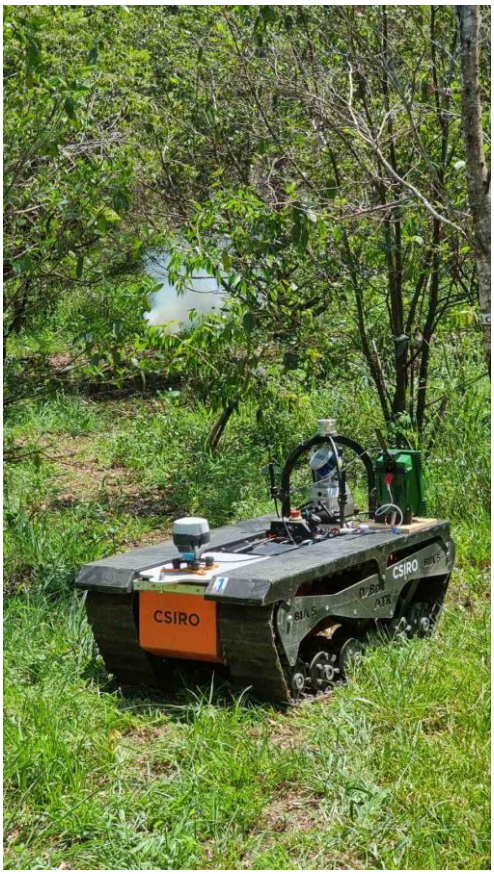


**Figure 31.** UP4a mounted on CSIRO's UGV.

#### *4.4.1.6 Czech Pilot - 2024*

**Field testing:** This was a crucial pilot as it would be the debut of the SILVANUS dashboard and the testing of the full data pipeline, i.e., starting from the data transmission, moving to their storing in SAL and finally their visualisation to the end-users. Regarding the IoTs, two devices were placed on trees, at two different locations, to test their detection capabilities for fire and smoke in varying landscapes and lighting conditions. Furthermore, two different connectivity tests were conducted, a wireless connection to the satellite network provided by MESH in the Sky (UP12) nodes, this time deployed on a drone, and with a cellular network from the devices' SIM card. These tests helped to measure response times when propagating data to SAL and to validate event visualisation in the SILVANUS dashboard (Figure 32). In this pilot the redesigned case was introduced (as presented in Section 4.2.1) and the improved code (exploitation of concurrent programming) was assessed. In total 184 events were captured and 168 of them were successfully transmitted to SAL. Additionally, 182 videos stored locally for fine-tuning the algorithms.

**Challenges:** The transparency of the flames and lighting conditions affected the fire detection algorithm, causing it to miss the fire events on some occasions. However, the collected data will be used to further enhance the algorithms. Additionally, after this discovery, augmentation techniques, such as altering image brightness, will be applied to improve the algorithm's robustness in similar situations.

Through these in-situ experiments and offline exercises, we have achieved significant advancements, including the improvement of the fire and smoke detection algorithms, upgraded case for the IoT, successful field tests in various conditions, and enhanced code execution time. The IoT devices will be further tested in Portugal, Croatia and Greece pilots, while they will be presented remotely in the remaining pilot sites through the respective layer in the SILVANUS dashboard. During these pilots the IoTs and

detection algorithms capabilities will be further tested, through different scenarios (including their communication with the FCC), while simultaneously focusing on overcoming the aforementioned limitations. Lastly, UP4a will contribute data to UP9h – Integrated Data Insights (see D5.3 and D5.4) for the generation of possible threat alerts to the end users, to take informed decisions.

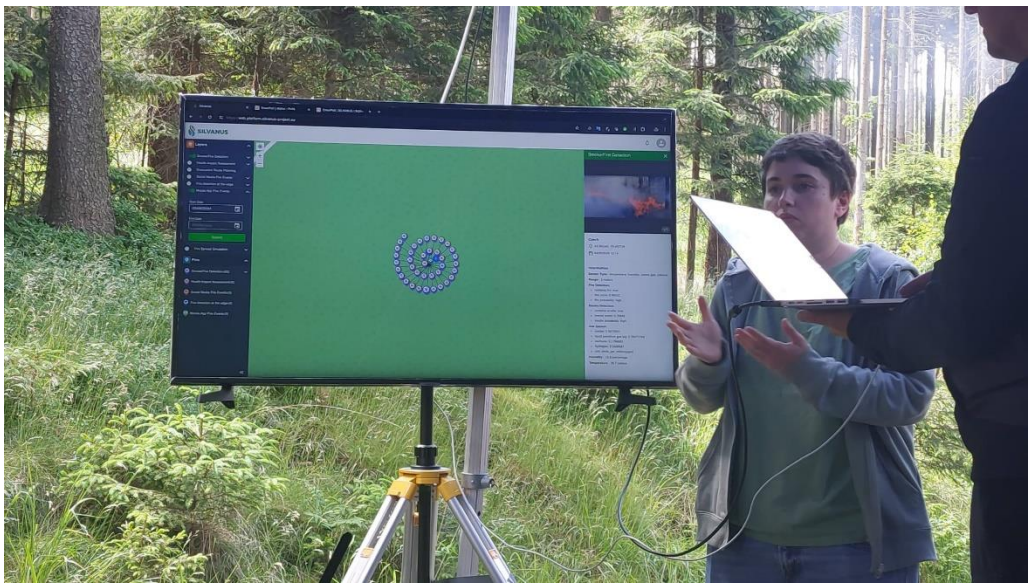

**Figure 32.** Visualizing UP4a layer on dashboard.

### *4.4.2 Fire detection from edge devices*

The module has been used in a real field pilot in Czech Republic (Ostrava area). It will be demonstrated during 2024 in other pilots. Confirmed so far, the tool will be used in Portugal and Italy (in both venues) presential and Romania and Greece in remote. In any case, the list of pilots where the tool will be present is still under negotiation and may still change.

Previously the tool was demonstrated in a tabletop exercise in Evia (Greece). At that time, the tool was not integrated with the dashboard.

During Czech pilot, the module was able to detect fire and smoke using videos and photos (taken on site) and even a live cam (computer webcam) with high reliability fulfilling the KPIs. In this regard, video was analysed at 14 frames per second (soft real time). There was only one false positive detection measured in the full video. False negatives rate was not measured in the video, but none was detected in the photos.

During this pilot, we full integration of the Up4b with the SAL and the dashboard was demonstrated. Participants were able to see the images of fire detected and the alerts raised in the dashboard.

### *4.4.3 IoT for air quality assessment*

The proposed Health Impact Assessment module was presented during both a tabletop exercise, which took place in Evia (Greece), and a pilot held in Krásná (Czech Republic) as demonstrated in Figure 33 and Figure 34. In both events, the proposed system received favourable responses/comments from the stakeholders who were involved. There was a notable level of interest shown towards IoT devices that are specifically adopted for the purpose of monitoring the ambient air quality and assessing the potential health impacts on firefighters following a wildfire outbreak. The inclusion of portable equipment was particularly valued, given that firefighters could conveniently affix it to their vehicles and transport it effortlessly. A comprehensive pilot testing scheme is currently in progress, encompassing the implementation of this system in the field settings across a range of wildfire incident scenarios.

*a) Equipment installation* *b) Equipment operation*

**Figure 33.** UP9b and MESH in the Sky integration.

It is necessary to underscore that synergies were developed with SILVANUS partners throughout the pilot's activities. Illustrative instances encompass our integration with RINI's MESH in the Sky (UP12) Figure 33, and the installation of portable solution on an UGV (3MON) as shown in Figure 34. In this way, the UGV (UP5a) can proceed first to monitor the air quality. If the air quality is deemed good, the firefighters can then move in and operate at the fire scene. Our involvement in forthcoming pilots includes the placement of the equipment onto an UAV - drone followed by the assessment of its functionality within an FCC (UP10).

**Figure 34.** Installation of UP9b on an UGV.

### *4.4.4 Forward Command Centers*

East-West Communication first demonstration

East-West Communication has been tested in a first demonstration scenario which involved the interaction with the EMER POLL application. The demonstration involved the usage of two VMs simulating each of them an EMDC instance.

The scheme of this demonstration is provided in Figure 35. The Fire Report module is designed for the detailed collection of fire data. Users can report new fires by providing comprehensive information such as location, description, and the status of the fire. The CE App is an application dedicated to end-users for submitting fire reports to the EMER POLL system. EMER POLL is the central report management system, responsible for collecting and organizing information sent by users.

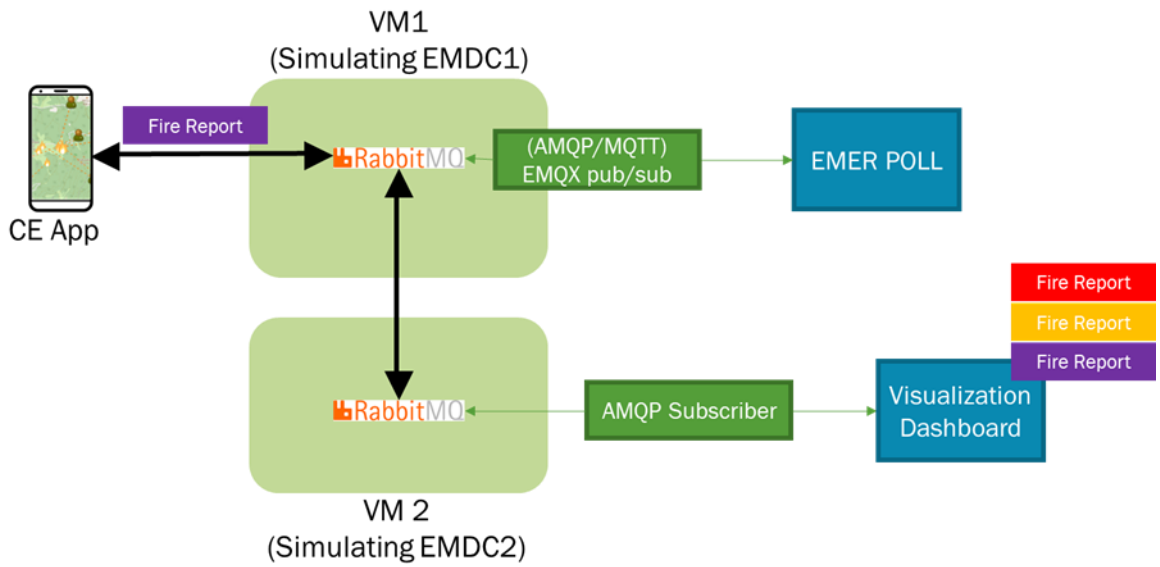


**Figure 35.** Architectural scheme of the East-West API first demonstration.

A second VM is connected to the first one by means of a RabbitMQ federation, which allows to transfer the fire report messages and to visualize it on a UI dashboard hosted on EMDC2, providing redundancy and ensuring operational continuity in case of failures or overloads of EMDC1. The distribution of the workload between EMDC1 and EMDC2 allows the system to handle a high volume of data and operations, improving overall efficiency and reducing the risk of service interruptions. The dashboard provides users with a visual overview of the status of fires and received reports, presenting detailed information through an intuitive graphical interface. Operators can continuously monitor the development of emergency situations, viewing GPS coordinates of fires, user-provided descriptions, and other relevant data.

In this scenario, an app user spots a fire and uses the CE App to send a detailed report. The app allows the user to input crucial information such as their current location and the location of the fire, a detailed description of the event, and the status of the fire. Once the report is completed, the app sends the message via the Message Queue Telemetry Transport (MQTT) protocol to a specific channel configured on the EMDC1. The EMDC1 receives the message and transmits it to various subscribers, including EMDC2 and EMER POLL. EMER POLL captures the message and initiates a poll to gather further details about the fire, while EMDC2 serves as a backup, ensuring redundancy and operational continuity. The report is displayed in real-time on the EMDC2 dashboard, where operators can monitor the fire status, view GPS coordinates, and read the notes provided by the user.

This demonstration focused on the implementation and verification of using different communication channels for each pilot site, testing the interaction between RabbitMQ instances installed on EMDC1 and EMDC2 connected via RabbitMQ Federation. The main objective of this demonstration was to evaluate the effectiveness of distributed communication and centralized management of fire reports in both real and simulated environments. During the demonstration, a dedicated MQTT communication channel was used to send and receive messages related to fire reports. This approach allowed us to isolate and monitor the performance of each channel, ensuring that data was transmitted reliably and promptly. The interaction between RabbitMQ instances was closely monitored to evaluate latency, message transmission speed, and overall system reliability. The federation configuration allows for bidirectional and asynchronous communication, enabling the data centers to continuously synchronize data and manage reports in a coordinated manner.

# 5 Social media sensing and concept extraction

The Social Media Sensing framework is a comprehensive system designed and developed to collect, analyse, and process social media posts related to fire incidents to support the involved stakeholders in decision making during a fire incident. This framework integrates various components that enrich the data and detect relevant fire events, which are then stored in a MongoDB and made accessible for end-users through visual interface.

## 5.1 Tool description

The process, as illustrated in Figure 36, begins with data collection, primarily utilizing a Twitter Crawler (described in Section 5.2.1.1) that connects to the Twitter API [25]. This crawler is responsible for formulating complex queries to retrieve tweets that are potentially related to fire incidents. Once the data is collected, the retrieved posts are stored in a MongoDB database in JSON format. These posts are then subjected to further analysis by different components within the Social Media Analysis Toolkit (described in Section 5.2.2), which operate through individual API calls. The results of these analyses are appended to the original JSON posts, enriching them with additional insights.

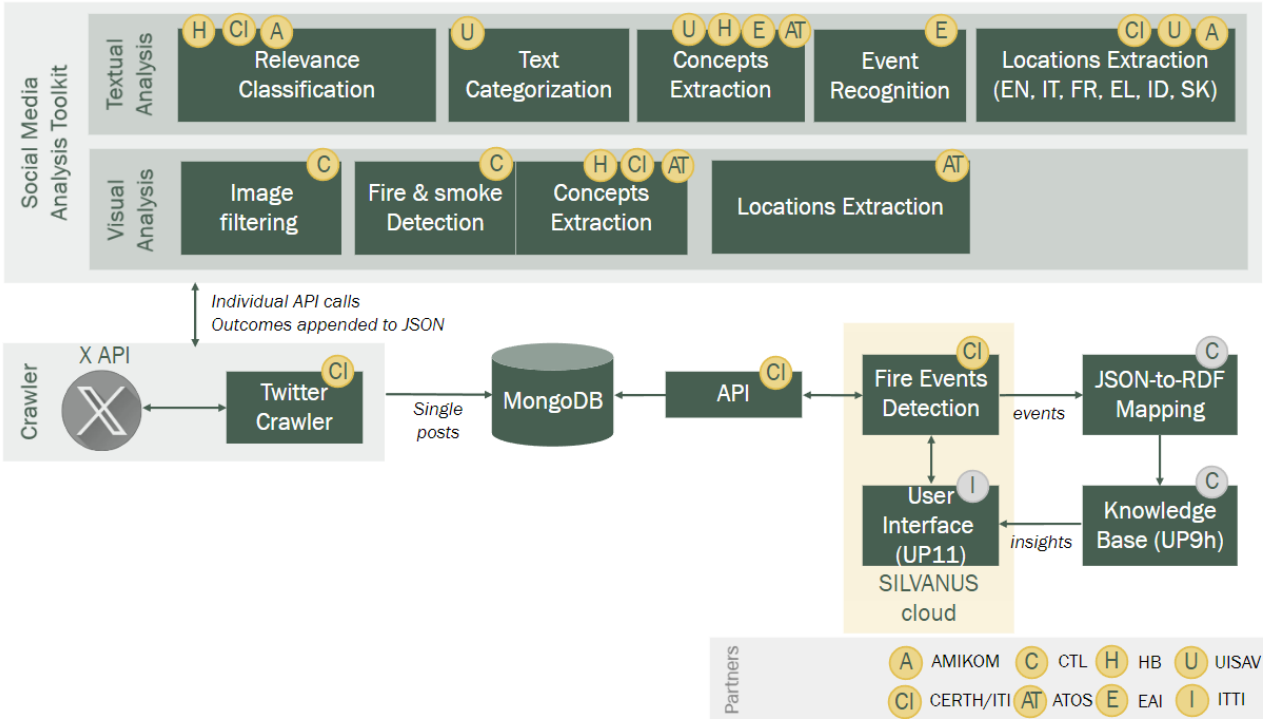


**Figure 36.** Social media sensing Framework.

As shown in Figure 36, the Social Media Analysis Toolkit comprises both textual and visual analysis modules. The textual analysis includes several key modules:

- **Relevance Classification:** This module assesses whether the post is pertinent to fire incidents.
- **Text Categorization:** It organizes posts into specific categories based on their content.
- **Concepts Extraction:** This module identifies and extracts significant entities such as names, organizations, and locations from the text.
- **Event Recognition:** It detects and identifies important events mentioned in the text.
- **Locations Extraction:** This module automatically extracts geographical information from the text, supporting multiple languages, including English, Italian, French, Greek, Indonesian, and Slovak.

In parallel with textual analysis, the visual analysis involves the following modules:

- **Image Filtering:** This module filters out irrelevant images from the posts.
- **Fire & Smoke Detection**: It analyses images to detect signs of fire or smoke.
- **Concepts Extraction:** This module annotates images with labels from a pool of concepts related to fire events, enabling the retrieval of high-level information from them
- **Locations Extraction:** It extracts location data directly from the images, providing additional geographical context to the post.

Once the posts have been enriched with additional information from the Social Media Analysis Toolkit, they are periodically analysed by the Fire Events Detection module (see Section 5.2.2.3). This module scans the collection of posts to detect any fire-related events. When an event is detected, it is sent to a message queue within the SAL, where it can be consumed by the SILVANUS UI to be visualized as pins on a map (see Section 5.4).

## 5.2 Innovations and updates

### *5.2.1 Social Media Crawling*

The Social Media Crawling module plays a crucial role in monitoring and analysing real-time data from various social media platforms. Its main goal is to deliver a continuous and dependable stream of data, with a particular focus on posts related to wildfires based on pre-defined search criteria. By utilizing the functionalities of different social media APIs, this module ensures prompt and efficient data collection, which is vital for disaster management and response.

This section provides an update of the X Crawler that has been presented in D4.2 Section 4.1, describing the changes and the new features added in the crawler.

#### *5.2.1.1 Implementation of the X Crawler*

X is a widely used platform [26] with millions of unique users who engage in various activities, including disaster reporting like fire incidents. To leverage this, CERTH has developed a crawler that consumes posts in near real-time from the X API, providing a crucial tool for monitoring and analysing real-time data. This section outlines the crawler's key components and functionalities. The crawler's main goal is to retrieve social media posts related to wildfires, ensuring a continuous and dependable stream of data. Among its key features are the ability to monitor X's public stream in real-time, filtering mechanisms that focus on specific keywords or user accounts, and the capability to manage large volumes of data. Additionally, the crawler operates within ethical guidelines and complies with X's terms of service, ensuring that all data is lawfully acquired.

X has revised its pricing policy for the X API, transitioning it from a free service to a service with a monthly subscription fee. Most of the rate limits have seen significant reductions compared to the previous free version of the Twitter API, and the available functionalities have been noticeably scaled down. The Free tier no longer allows reading posts, and the Basic tier, which costs $100 per month, only permits reading up to 10,000 posts per month and does not include access to real-time data or short-term historical data. Access to these functionalities is now restricted to the Pro tier, which comes with a hefty price tag of $5,000 per month. Given the significant reduction in rate limits and functionalities across the different tiers of the current X (formerly Twitter) API, the Basic version emerges as a practical choice for a system designed to collect X posts for wildfire detection. Despite the limitations, such as the restricted rate limit of 10,000 posts per month and the absence of real-time data access, the Basic tier still offers essential features like the ability to write and read posts, which are critical for continuously monitoring social media activity.

The decision to choose the Basic version is driven by a balance between cost and functionality. At $100 per month, it provides a feasible solution for maintaining a stream of relevant data without the prohibitive costs

associated with the Pro version. While the Pro version offers more extensive capabilities, its high cost makes it less viable.

In response to the recent changes in X's API policies, CERTH has adjusted its approach to a monitoring-based crawling strategy. This shift is designed to improve the efficiency of post retrieval while mitigating the impact of rate limits. By periodically monitoring the number of posts and concentrating crawling efforts during periods of significant activity spikes, CERTH can more effectively capture data related to specific search queries. This approach optimizes the collection of relevant posts.

The crawler utilizes the Recent Tweet Counts [27] endpoint to periodically monitor social media activity, looking for significant increases in the number of X posts since the last check. When such an increase is detected, it suggests a potential event, prompting the retrieval of posts generated within that specific time frame. Conversely, if no significant rise in post volume is observed, the crawler stops collecting posts during that period. This method ensures that relevant posts are maximized while staying within the platform's rate limits.

During times of heightened activity, the crawler collects posts by querying the Recent Search endpoint [28] of the X API at 30-minute intervals. These intervals are governed by the platform's rate limit, which permits calls to the Recent Search endpoint only every 30 minutes. The queries are carefully crafted with the search criteria keywords, phrases, and user accounts relevant to wildfires. The crawler then retrieves data that meets the predefined criteria set in the Recent Search endpoint. After retrieving the relevant posts, the X Crawler sends the data to the Social Media Analysis Toolkit. This module processes and analyses the posts to extract useful information. It comprises both textual and visual analysis modules. Finally, after analysis, the posts and their extracted data are stored in a MongoDB database in JSON format. An overview of the crawling architecture is depicted in Figure 37.

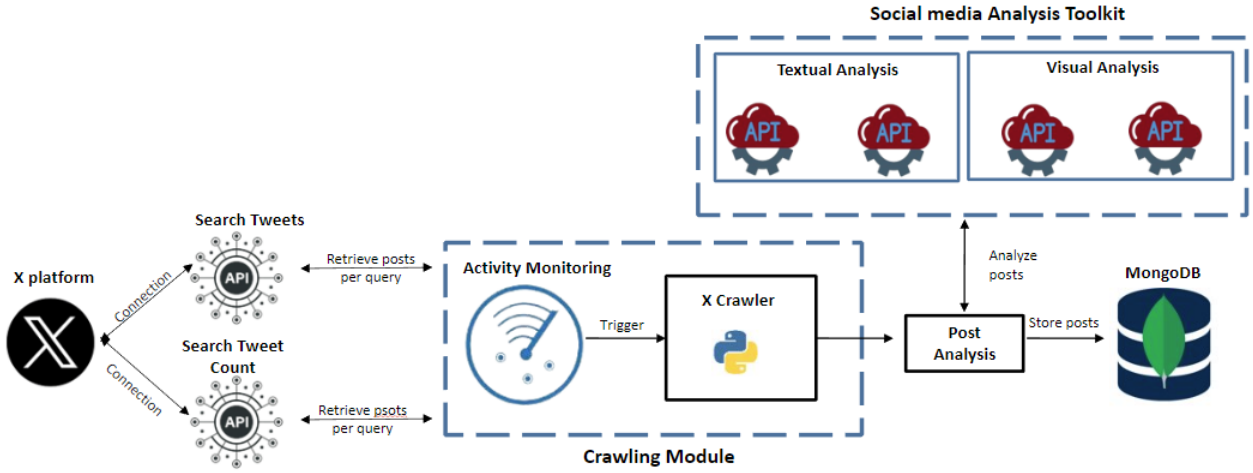


**Figure 37.** Crawling Architecture.

*5.2.1.2 Monitoring and anomaly detection*

An analysis was conducted on a dataset derived from the social media platform X, focusing on historical data collected between June 2019 and August 2022. This dataset consists of 7,990 posts about wildfires in Greece.

The aim of the study was to develop and evaluate different methodologies for ongoing monitoring of social media activity, especially during times of increased posting activity. The research had two key objectives: firstly, to pinpoint and visually highlight periods of significant activity on the platform; and secondly, to identify the most effective technique for real-time monitoring and detecting anomalies in anticipation of

future events. The following methodologies employed in the analysis, offering a comparison of their approaches, results, and effectiveness in identifying the best strategy for continuous monitoring.

#### 5.2.1.2.1 K-means

For anomaly detection within social media data, we employed the K-Means [29] clustering algorithm to analyse patterns in tweet counts. The data were grouped into 30-minute intervals, with the tweet counts aggregated within each interval to form a time series. This time series served as the input for the K-Means clustering algorithm, which was tasked with identifying clusters of tweet activity over time.

K-Means clustering partitions the data into a predefined number of clusters based on similarity. Each data point, representing a 30-minute tweet count, is assigned to the nearest cluster centre. The algorithm iteratively refines these clusters until the optimal partitioning is achieved.

As illustrated in the accompanying Figure 38, the tweet counts are visualized as coloured dots, with each colour representing a different cluster. The red crosses denote the centres of these clusters, representing the average tweet activity within each group. This visualization reveals distinct periods of tweet activity:

- **Clusters of High Activity:** These are often associated with significant events or topics that generate increased interest on social media.
- **Consistent Activity:** Some clusters capture periods of steady tweet activity, indicating ongoing discussions or sustained interest in certain topics.
- **Outliers:** The clustering also highlights periods that deviate from typical patterns, which could signal unusual or noteworthy events.

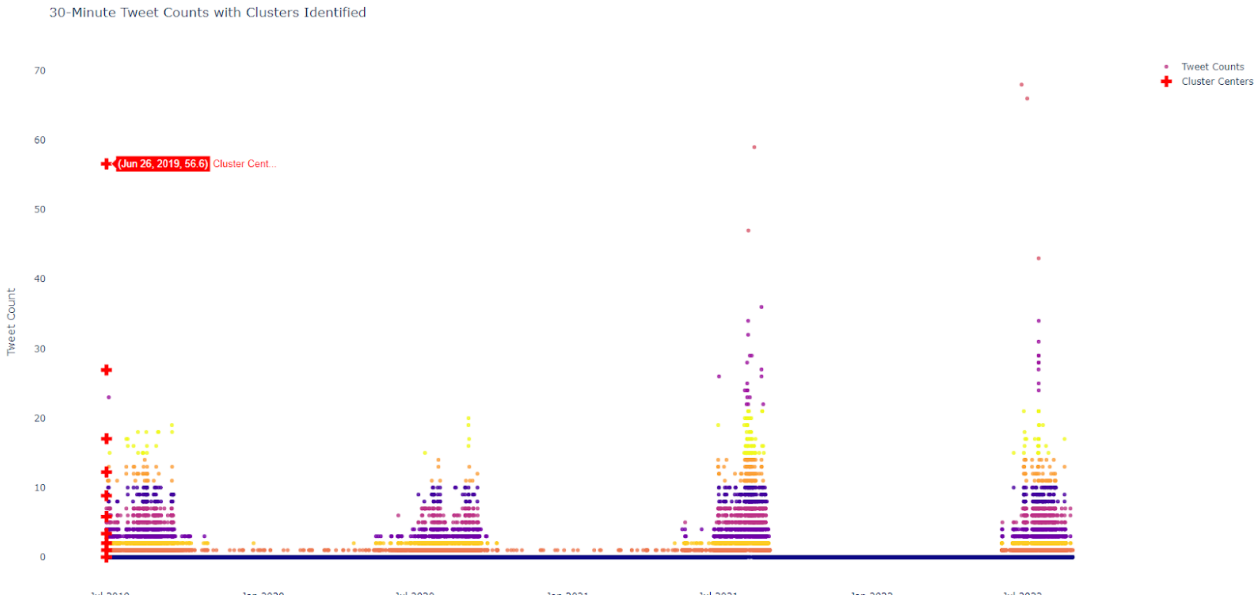


**Figure 38.** Results for K-Means method.

K-Means clustering proved to be a valuable tool in monitoring social media activity. By identifying clusters of high tweet counts, we can detect periods of significant public interest.

#### 5.2.1.2.2 Isolation forest

In the isolation forest algorithm [30], the process began by aggregating data into specific time intervals, 30-minute periods, where post counts were summed within each interval. The result was a time series of post counts, which served as the input for the anomaly detection algorithm.

In the time series of aggregated post counts, as illustrated in Figure 39, anomalies are highlighted with red markers. This visual representation makes it easy to identify periods of unusually high posting activity. What

sets this method apart is its ability to distinguish true anomalies from regular cyclical patterns, such as daily or weekly fluctuations. This capability significantly reduces the likelihood of false positives, ensuring that the identified anomalies genuinely reflect significant deviations from the norm. This distinction helps reduce the number of false positives and ensures that the identified anomalies are truly significant deviations from the norm.

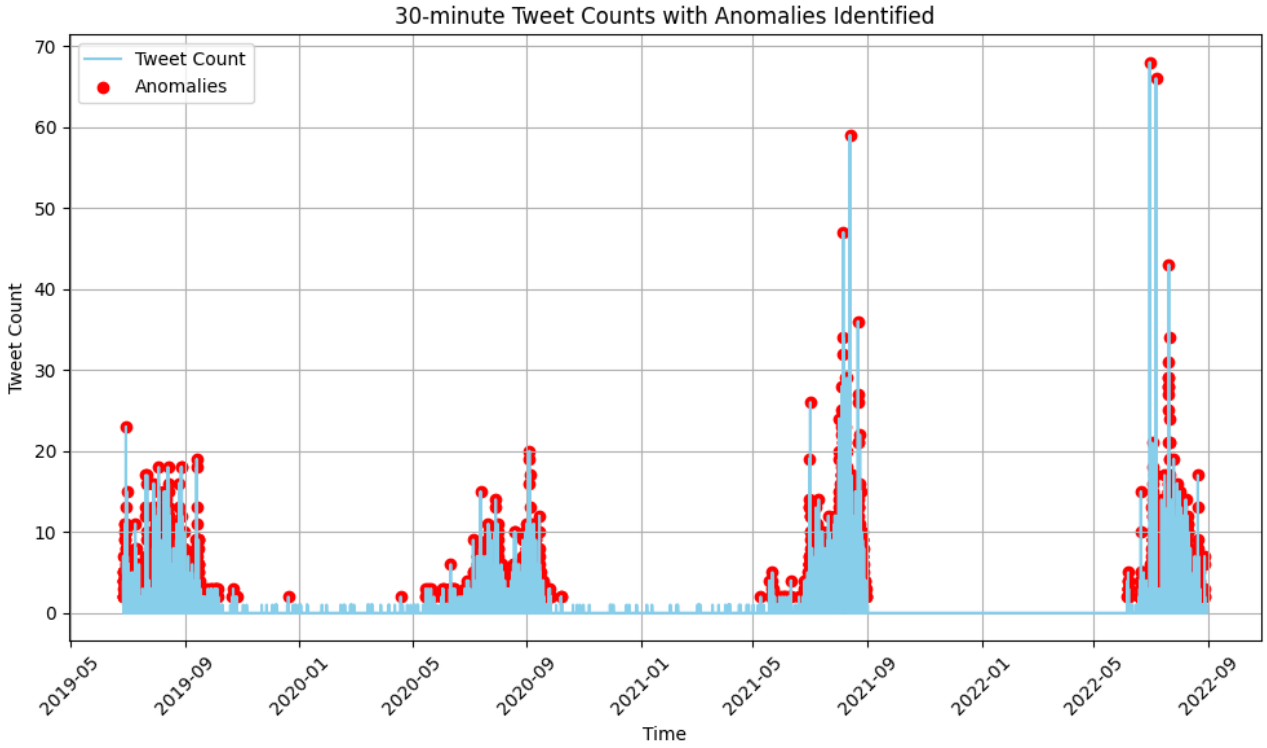


**Figure 39.** Results for Isolation Forest method.

#### 5.2.1.2.3 Long Short-Term Memory

The Long Short-Term Memory (LSTM)-based [31] approach was employed to detect anomalies in the tweet count data over time. The data were first aggregated into 30-minutes intervals, resulting in a time series where each point represents the total number of tweets within a given hour. This time series was then used as input for the LSTM model, which is particularly suited for capturing and predicting sequential patterns in data.

As shown in the Figure 40, tweet counts are represented by the light blue line, while anomalies are indicated by red markers. These anomalies appear where the actual tweet counts differ significantly from what the LSTM model predicted. Unlike the Isolation Forest method, which detects anomalies based on how much they stand out from the overall data, the LSTM approach focuses on finding irregularities in the sequence of tweet activity over time.

LSTM can lead to a higher frequency of detected anomalies, as seen in Figure 40. While this can be useful for identifying all potential deviations, it also requires careful interpretation to distinguish between truly significant anomalies and minor fluctuations in tweet activity. This potential for over-detection underscores the importance of setting appropriate thresholds to balance the detection of significant events against the risk of false positives.

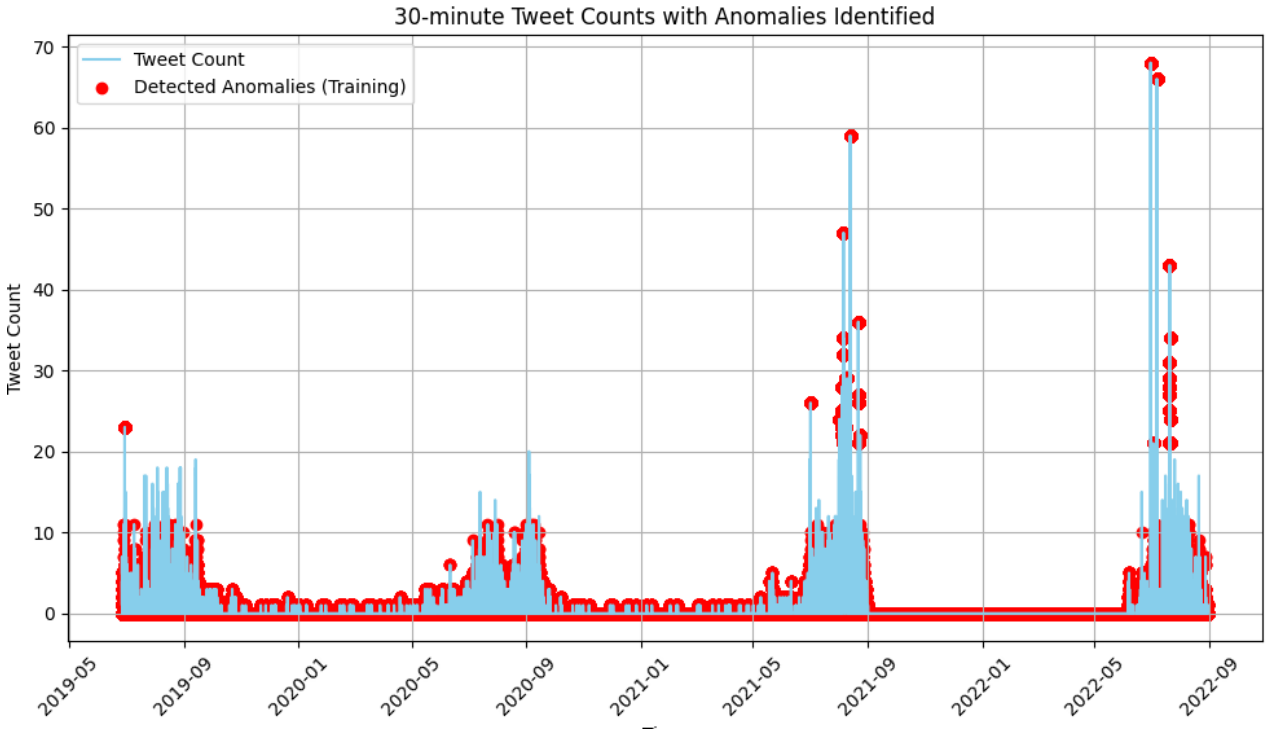


**Figure 40.** Results for LSTM method.

#### 5.2.1.2.4 Moving Average

The Moving Average method [32] was used to detect anomalies in tweet counts over time. In this approach, a moving average was calculated over the time series to establish a baseline of typical tweet activity. The data were aggregated into 30-minute intervals, creating a time series that captures the tweet counts within each period. This time series served as the foundation for the anomaly detection process.

As shown in the Figure 41 the tweet counts are depicted by the light blue line, while anomalies are marked with red dots. These anomalies represent periods where tweet activity significantly deviated from the expected pattern, as determined by the moving average. Unlike more complex methods like LSTM or Isolation Forest, the Moving Average method is straightforward and easy to implement. It effectively identifies spikes or drops in activity that deviate from the overall trend.

One of the strengths of the Moving Average method is its simplicity and transparency. It provides a clear and easily interpretable way to detect anomalies, making it ideal for scenarios where a quick and understandable analysis is required. However, its reliance on smoothing can sometimes overlook more subtle or abrupt changes in the data, which might be captured by more sophisticated methods.

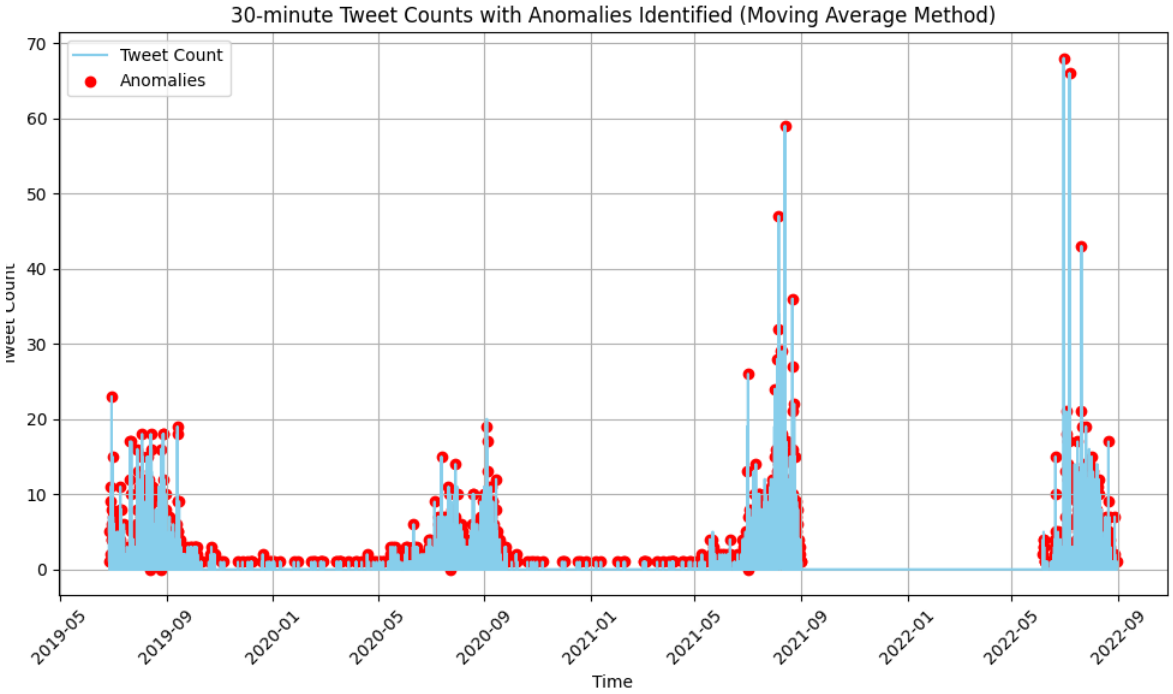


**Figure 41.** Results for Moving Average.

#### 5.2.1.2.5 Comparison of different methods

After evaluating the strengths and weaknesses of each method across four key criteria—Robustness, Outlier Detection, Reduced False Positives, and Flexibility—a comparison table was created. These criteria were chosen to ensure the reliability and effectiveness of the anomaly detection methods in handling diverse and unpredictable social media data. Robustness and Outlier Detection are essential for accurately identifying significant deviations, while Reduced False Positives and Flexibility ensure that the method is both precise and adaptable to various monitoring scenarios. Together, they provide a balanced assessment of each method's capability to detect meaningful anomalies in real-time.

Based on the comparison in Table 6, the **Isolation Forest** method is the most suitable choice for anomaly detection. Its strengths include:

- **Robustness:** Isolation Forest is highly effective across various data patterns, including both dense and sparse regions, making it well-suited for the diverse nature of social media activity.
- **Outlier Detection:** Its focus on detecting outliers ensures that significant deviations, which could indicate noteworthy events, are identified without being overwhelmed by regular fluctuations.
- **Reduced False Positives:** The method's ability to distinguish between true anomalies and cyclical patterns significantly reduces the likelihood of false positives, ensuring that detected anomalies genuinely reflect important deviations.
- **Flexibility:** While computationally intensive, the algorithm's flexibility and effectiveness in unsupervised anomaly detection make it a powerful tool for continuous monitoring.

**Table 6.** Comparison of Anomaly Detection Methods.

| Criteria | K-Means | Isolation Forest | LSTM | Moving Average |
|---|---|---|---|---|
| Robustness | Moderate: Effectiveness depends on cluster selection and data regularity. | High: Works well across diverse data patterns, handling both dense and sparse data. | High: Captures patterns and detects subtle anomalies over time. | Moderate: Handles general trends well but may miss subtle or abrupt changes. |
| Outlier Detection | Limited: Focuses on grouping similar patterns rather than detecting outliers. | Excellent: Specifically designed for outlier detection without needing labelled data. | High: Effective in detecting both isolated spikes and irregular sequences. | Moderate: Identifies deviations from the moving average but may miss more complex outliers. |
| Reduced False Positives | Moderate: Depends on the choice of clusters, may miss or incorrectly identify anomalies. | High: Effective at distinguishing true anomalies from regular patterns, minimizing false positives. | Moderate: Sensitive to minor deviations, which can lead to a higher rate of false positives. | Moderate: Straightforward, but its simplicity might lead to overlooking or misclassifying anomalies. |
| Flexibility | Moderate: Requires careful selection of the number of clusters to be effective. | High: Adaptable to different types of data with tuning of parameters like contamination. | Moderate; Powerful but requires complex tuning and interpretation to avoid over-detection. | High: Easy to implement and interpret, ideal for quick, basic analyses. |

In conclusion, Isolation Forest provides the best balance of accuracy, robustness, and practicality for the specific task of monitoring tweet activity related to wildfires. Its ability to identify true anomalies while minimizing false positives ensures that critical events are detected promptly and accurately, making it the optimal choice for ongoing monitoring.

#### *5.2.1.3 X Crawling*

The X data that was collected and stored spans the period from the beginning of the crawl on December 23, 2021, until September 02, 2024. During this period, approximately 1 million fire-related tweets were amassed across five pilot projects: P01, P03, P07, P09, and P10.

As seen in Table 7, English emerged as the dominant language in all pilots, gathering the largest number of tweets, which was anticipated given its global usage. Italian is the second most prevalent language, with around 30,000 tweets, while Greek follows with approximately 12,000 tweets. Notably, Indonesian Twitter users do not seem to actively engage in fire-related discussions or reporting, as shown by the lack of collected tweets. On the other hand, both Italian and Greek Twitter users show considerable activity surrounding fire incidents, indicating that fire-related issues are of significant concern within these regions.

A key insight from this data collection is the limited geolocation information provided directly by Twitter. Only about 0.1% to 0.8% of tweets have geolocation data embedded from Twitter's JSON output. However, by leveraging the CERTH Localization web service, the percentage of geotagged tweets dramatically increases, ranging from 24.2% to 49.5% depending on the pilot. This enhancement in geolocation data is invaluable, especially when combined with external data sources like satellite imagery. It underscores the importance of CERTH's service, as Twitter itself lacks a substantial number of geotagged tweets, making third-party localization crucial for extracting geospatial insights from social media data.

Moreover, between 5.7% and 10.8% of the tweets collected contain at least one image. While this may seem like a small portion of the overall dataset, the visual content within these tweets can provide rich and detailed information. These images offer an additional layer of data that can be utilized in comprehensive visual analyses through tools such as the Social Media Analysis Toolkit, adding depth to the overall understanding of fire-related incidents captured in the tweets.

**Table 7.** Pilots and tweets collected for each pilot.

| Pilot | Language | Collected Tweets | Location from Twitter | Location from CERTH's Localization | Contain Image |
|---|---|---|---|---|---|
| P01 | English | 47,036 | 141 (0.3%) | 12,111 (25.8%) | 1,270 (9.4%) |
| | Italian | 26,457 | 53 (0.2%) | 13,963 (52.8%) | 1,055 (4.0%) |
| P03 | English | 247,506 | 1,485 (0.6%) | 68,352 (27.6%) | 20,243 (8.2%) |
| | French | 209 | 1 (0.1%) | 79 (37.8%) | 9 (4.3%) |
| P07 | English | 134,173 | 1,208 (0.9%) | 31,931 (23.8%) | 9,231 (6.9%) |
| | Greek | 46,301 | 69 (0.1%) | 18,748 (40.5%) | 2,006 (4.3%) |
| P09 | English | 28,137 | 56 (0.2%) | 9,927 (35.3%) | 2,078 (7.4%) |
| P10 | English | 39,180 | 78 (0.2%) | 10,274 (26.2%) | 1,506 (9.1%) |
| | Indonesian | 0 | 0 (0.0%) | 0 (0.0%) | 0 (0.0%) |

### *5.2.2 Social Media Analysis Toolkit*

This section provides an update on the analysis modules in the Social Media Analysis Toolkit, as described in D4.2. These modules are organized based on the type of input they process—either textual content (see Section 5.2.2.1) or visual content (see Section 5.2.2.2).

Social media posts collected from platforms like Twitter, Facebook, and web crawlers are fed into the various components of the Social Media Analysis Toolkit through individual API calls. These modules then analyse both the visual and textual content of each post, enriching the post's metadata with deeper insights. This enhanced metadata is crucial for the Fire Events Detection module, which uses it to identify fire-related events and provide actionable information to users of the SILVANUS Dashboard.

#### *5.2.2.1 Textual analysis of social media content*

This subsection provides an update of the textual analysis modules provided by the Social Media Analysis Toolkit in D4.2. The analysis methods employed by these modules are as follows:

I. Text Categorization - Relevance Classification - Concepts Extraction: Involves categorizing social media posts based on their relevance to the topic of interest and classifying them into different categories based on their content.
II. Event Recognition: Identifies significant events mentioned in social media posts, e.g., as fires.
III. Location Extraction: Extract location information from social media posts text.

#### 5.2.2.1.1 Text Categorization - Concepts Extraction

HB has experimented with and successfully completed a new web service module for the UP3 pipeline called 'TweetAnalyzer for Information Fusion' to further develop textual analysis, complementing the existing 'Relevance Classification and Text Categorization' plus 'Textual Concepts Extraction' modules for text categorization and text mining in the Social Media Analysis Toolkit (D4.2). The module computes the heated nature of messages by their semantic intensity [33], [34] vs. social importance [35] measured by sentiment analysis vs. centrality values, with the results as new metadata forwarded to the ontology and the knowledge graph. By a proactive stance, it thereby connects T4.4 with T5.2, utilizing CERTH's test dataset of tweets, and is hosted by courtesy of CERTH.

The TweetAnalyzer is a tool which does statistical analysis of tweets by means of sentence embeddings (SBERT) [36], dimensionality reduction (t-SNE) [37]), sentiment analysis (SentiArt) [38], centrality analysis (Page Rank) [39], clustering (k-means) [40], and various other methods. The functions of this tool are fourfold:

- Isolate and highlight phrases and word pairs (bigrams, typically adjective-noun phrases) within tweet texts, so that they can be used for further analysis.
- Perform analysis of tweets, both with and without the context of isolated phrases or bigrams, expressing results as coordinate values.
- Offer a method to investigate which coordinates can logically be grouped together by k-means clustering, and to investigate which grouping would be optimal by silhouette analysis [41].
- Provide a method through which the 2-d vs. 3-d coordinates can be plotted to allow for visual inspection of the results.

The tool does this by offering a backend API, which returns responses in JSON format. The API is Python-based (FastAPI) [42], employs Docker for isolation and containerization, and uses Poetry for Python package versioning [43]. It employs Git for version control and is documented using Swagger [44].

After originally having carried out two sets of relevance classification experiments on altogether 5 social media datasets in English, Greek, Italian and Spanish, all provided by CERTH, HB reused established text processing workflow combining preprocessing, sentence embedding, dimensionality reduction, clustering of semantic content, and visualization as reported in D4.2. Starting with M19, in more detail we explored our strategy called Track B there, addressing the dynamics of evolving semantic content. Underlying that line of thought was the observation of topic outbursts [45]. After having also checked out two other methods, Top2vec [46] and TopSBM [47] [48], we stated that Topic Modelling (TM) by Latent Dirichlet Allocation (LDA) [49] could reliably identify wildfires as the central concept in the test data. We foresaw that by extracting changing compositions of index terms in document sets characterized by some progressive feature such as geographical location or timestamps, one could characterize event progress as a series of topic outbursts by Dynamic Topic Modelling (DTM) [50]. A version of this idea is also implemented as part of the BERTopic toolkit we used for testing [51].

Next, we focused on the spatiotemporal nature of the clustering of semantic content, i.e. how to handle the problem of incoming social media messages with evolving topicality such as the changing intensity of a fire event vs. its changing perception in the public eye. Key to our line of thought was to develop a method which can be reused both in a firefighting context and its political fallout, such as opinionated debates about responsibilities typical of Phase B. Also, it was important that the tool should be able to monitor good news expected in Phase C too, not just bad one's characteristic for disaster management. By respective language fine-tuning, our module was conceived to help not just firefighting but other forms of natural and humanitarian disasters as well. The respective round of experiments until M30 was solely using the timestamped 4K tweets dataset provided by CERTH.

The scientific novelties underlying HB's web service module development efforts are listed in the next section. In terms of development, within the context of the project period, three major releases of the TweetAnalyzer were published:

- Release v0.1 (24-07-03) drafted the basic structure of the module, offering Dockerization, version control, and package management for its requirements. Functionalities/deliverables provided were an endpoint to measure SentiArt scores for a tweet message (implemented as sum of its parts) as well as a t-SNE endpoint.
- Release v0.2 (24-07-12) expanded the set of endpoints with new ones for bigram isolation and ranking of word pairs, as well as a Page Rank implementation for a list of tweets. Furthermore, some optimizations to SentiArt were introduced to speed it up.
- Release v0.3 (24-08-26) added plotting endpoints for scatter, contour, and arrow diagrams of the t-SNE-produced coordinates and changed the code structure to a tidier version using Python mixins.

#### 5.2.2.1.2 UISAV Facebook Post Analyser module

The Named Entity Recognition (NER) model of the Facebook Post Analyser has been enhanced to recognize more classes (extended from 3 to 7), fine-tuned and evaluated. Description with examples of the recognizable Named Entity (NE) classes by the new model is provided in Table 8. A new dataset has been prepared using a manually annotated subset of 1,869 Facebook posts, selected from a total of 2,355 posts collected before the deprecation of the Facebook Groups API and the removal of the ability for group admins to install apps in groups, even with admin or developer roles on the app.

**Table 8.** NE classes recognized by the new enhanced NER model of the Facebook Post Analyser.

| Class | Description | Examples |
|---|---|---|
| keyword | action regarding the fire (verb, noun, adjective) | inhale, extinguish, burn, fire, flame, smoke |
| time | time of the event | 11:15, today, yesterday, before nine |
| unit | unit which was used during the response | volunteer fire brigades, mountain rescue service, soldiers, policemen, HaZZ |
| fire_area | what was on fire (noun) | grass, hall, flat, house, chimney, garage, cathedral, straw |
| equipment | what equipment were used during the response | dron, vehicle, rank, helicopter, pump, jet |
| location | location of the event | cities, villages, hills, factory areas, countries |

The dataset was annotated by 3 annotators, their annotations were merged with Cohen's kappa set to 0.7 and aligned by a fourth annotator. The distribution of the entity classes in the dataset is provided in Table 9.

Ratio of 80%, 10% and 10% was used to create training, validation and test datasets (i.e., 1884, 235, 236).

The fine tuning involved experimenting with different batch sizes, number of epochs, max. number of tokens, optimizers, learning rate schedulers to finally settle to 8 posts in a batch, 512 tokens, Adaptive Moment Estimation (ADAM) optimizer, Polynomial learning rate scheduler.

**Table 9.** Distribution of entity classes within the dataset.

| Class | Counts |
|---|---|
| fire_area | 949 (586) |
| keyword | 1,749 (182) |
| time | 623 (468) |
| unit | 1,099 (536) |
| equipment | 159 (228) |
| location | 1,452 (680) |
| other | 24,822 |
| all words | 33,534 |

Another hyper parameter that was examined was the minimal length of the post (number of words divided by whitespace) and inclusion of posts without entities of interest. Subsequently, the number of relevant posts has changed accordingly leading to train 4 different models:

- All posts (Model 1): 2,355
- All posts with length at least 3 words (Model 2): 2,217
- Posts with length at least 3 words and with entities only (Model 3): 1,783
- Posts with entities only (Model 4): 1,869

The models were evaluated over the test dataset. Detailed evaluation results are presented in Table 10, while the overall evaluation is presented in Table 11.

**Table 10.** Precision, Recall and F1 statistics for individual classes and models.

| | Precision | | | | Recall | | | | F1 | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Class/Model** | **M1** | **M2** | **M3** | **M4** | **M1** | **M2** | **M3** | **M4** | **M1** | **M2** | **M3** | **M4** |
| Location | **0.89** | 0.85 | 0.89 | 0.81 | **0.92** | 0.90 | 0.89 | 0.86 | **0.90** | 0.88 | 0.89 | 0.83 |
| Fire area | 0.73 | 0.77 | 0.70 | **0.81** | 0.76 | 0.74 | 0.78 | **0.81** | 0.74 | 0.75 | 0.74 | **0.81** |
| Unit | **0.79** | 0.64 | 0.70 | 0.74 | **0.83** | 0.74 | 0.72 | 0.81 | **0.81** | 0.69 | 0.71 | 0.77 |
| Equipment | 0.69 | 0.55 | **0.76** | 0.63 | 0.69 | 0.67 | **0.76** | 0.61 | 0.69 | 0.60 | **0.76** | 0.62 |
| Keyword | **0.88** | 0.88 | 0.87 | 0.88 | 0.84 | **0.90** | 0.86 | 0.82 | 0.86 | **0.89** | 0.87 | 0.85 |
| Time | **0.85** | 0.80 | 0.76 | 0.77 | 0.76 | **0.87** | 0.83 | 0.78 | 0.80 | **0.83** | 0.79 | 0.77 |

**Table 11.** Overall performance of individual models.

| | Posts with entities only | All posts |
|---|---|---|
| All posts | Precision: 0.79770<br>Recall: 0.81104<br>F1: 0.80431<br>Accuracy: 0.90250 | Precision: **0.83304**<br>Recall: 0.83598<br>F1: **0.83451**<br>Accuracy: **0.93727** |
| Posts with at least 3 words | Precision: 0.80000<br>Recall: 0.82572<br>F1: 0.81266<br>Accuracy: 0.90559 | Precision: 0.79595<br>Recall: **0.83969**<br>F1: 0.81724<br>Accuracy: 0.91890 |

Looking at the F1 statistics, Model1 and Model2 are the best for 2 classes. Location and Unit for Model 1 and Keyword and Time for Model2. The Equipment class is the hardest one to predict. This is because the number of the entities in this class is small (159). In our opinion the Model 1 could be considered as the "best" one, because of its robustness. To build this model, the largest dataset (all posts) was used. Also, the Precision is good compared to other models (4 classes). But on the other hand, the numbers are not as good as it was expected, when it comes to F1 statistics. Especially numbers of Equipment (7% worse than best result) and Fire area (7% worse than best result) classes. If other classes of Model 1 are compared, F1 statistics are close to the best ones.

It can be seen from the overall results that the model trained on all the posts (Model 1) gives the best results, regarding precision, F1 and accuracy. The only statistics in which it was not superior is Recall, where the model trained on posts with length of at least 3 words (Model 2) gives the best performance. Please note that Recall of Model 1 is close to that of Model 2. The reason for this is because the one- or two-word(s) posts are non-informative. They consist mainly of keywords – like "fire arise", "date", "time", or they are using emoticons for fire and fire truck and so on. But in the end, all the models are almost on the same level, when comparing the F1 statistics.

#### *5.2.2.1.3* CERTH's Relevance estimation module

Social media data is emerging as a vital tool for real-time situational awareness in natural disaster scenarios, including wildfires. Suitably filtering Twitter text data can improve practical response by authorities. CERTH developed a relevance estimation module that employs AI and Natural Language Processing (NLP) to determine the relevancy of a tweet to a wildfire event. Numerous NLP techniques have proven effective in binary text classification tasks. However, over the past decade, Deep Neural Networks and Transfer Learning have increasingly come to dominate the NLP field – largely motivated by the introduction of Transformer architectures and the subsequent development of models based on these architectures, such as Bidirectional Encoder Representations from Transformers (BERT).

In the context of SILVANUS, we approach Relevance Estimation as a supervised classification task. In this case, the datasets used are binary-labelled (0 or 1), signifying an irrelevant or relevant tweet, respectively. To illustrate the module's functional, we consider two arbitrary sentences containing keywords from the data collection phase. An example of a relevant sentence could be, "The fire is burning the whole forest right now!", whereas an irrelevant sentence might be, "Arson's rap is on fire!". The model's output should manifest as a probability mass function for the outcomes 0 and 1, denoting the probability of a tweet being irrelevant and its complementary probability of being relevant.

In summary, the primary objective of this process is to filter posts from X, ensuring that users receive content relevant to the specific topic of interest, namely, wildfire incidents. This study will focus on three different languages, English, Italian and Greek.

In this study we first evaluate the performance of the statistical Born Classifier (BR) to establish an ML baseline before exploring transfer learning with transformer models. This approach is motivated by two primary factors: first, the limited number of comprehensive studies on disaster relevance estimation, particularly given the inherent data collection biases in social media analysis; and second, the diversity of languages, which can impact the performance of transfer learning models due to differences in training datasets. The BR will be assessed using two NLP representation techniques: Term Frequency-Inverse Document Frequency (TF-IDF) and Bag of Words (BoW). Subsequently, our analysis will shift to the implementation of DL methods, using open-source BERT models and their variations as the basis for transfer learning. The application of the BR will be implemented using scikit-learn library [52] in Python.

Fine-tuning is a transfer learning technique where a model pre-trained on a large dataset is further trained on a smaller, specialized dataset to adapt its knowledge for specific tasks. This process enhances the

model's performance on the new data by adjusting its learned features to better suit the nuances of the targeted domain or task.

Table 12 presents the annotated data used in our experiments for Greek, Italian, and English. For English, the FireXPosts dataset from the "Fire Detection for Emergency Responders using X" by [53] was utilized. This dataset also includes a Greek subset of 582 posts, which was used to supplement the previously collected 4,880 Greek posts. All posts were gathered using CERTH's crawler, and the annotation process was supported by expert end users.

**Table 12.** Number of relevant and irrelevant tweets per language.

| **Languages** | **Number of tweets** | | |
|---|---|---|---|
| | **Relevant** | **Irrelevant** | **Total Number** |
| Greek | ~ 2190 (40 %) | ~ 3280 (60 %) | ~ 5470 (100%) |
| Italian | ~ 580 (38 %) | ~ 950 (62%) | ~ 1530 (100%) |
| English | ~ 620 (44%) | ~ 780 (56 %) | ~ 1400 (41%) |

Table 13 presents all the models evaluated on each dataset, including the BR, and highlights the best-performing ones for each language.

**Table 13.** Performance metrics for all models.

| **Language** | **Model** | **Accuracy** | **Recall** | **Precision** | **F1-score** | **Training Time(s)** |
|---|---|---|---|---|---|---|
| Greek | BoW & BR | 0.8492 | 0.8756 | 0.7739 | 0.8216 | 0.11 |
| | TF-IDF & BR | 0.8547 | 0.8594 | 0.7919 | 0.8243 | 0.10 |
| | 'nlpaueb/bert-base-greek-uncased-v1'<br>8 epochs: {0: 2e-6, 7: 2e-7} | 0.8770 | 0.8880 | 0.8208 | 0.8529 | 391.53 |
| | 'EftychiaKarav/DistilGREEK-BERT'<br>13 epochs: {0: 2e-5, 12: 2e-6} | 0.8688 | 0.8563 | 0.8245 | 0.8396 | 310.00 |
| | **'ClassCat/roberta-small-greek'<br>5 epochs: {0: 2e-5, 2: 2e-7}** | **0.8806** | **0.8606** | **0.8451** | **0.8527** | **87.55** |
| Italian | BoW & BR | 0.9216 | 0.9652 | 0.8473 | 0.9024 | 0.02 |
| | TF-IDF & BR | 0.9477 | 0.9826 | 0.8898 | 0.9339 | 0.02 |
| | 'dbmdz/bert-base-italian-cased'<br>11 epochs: {0: 2e-6, 9: 2e-8} | 0.9288 | 0.9572 | 0.8700 | 0.9104 | 144.73 |
| | 'osiria/distilbert-base-italian-cased'<br>16 epochs: {0: 2e-6, 8: 2e-7} | 0.9340 | 0.9419 | 0.8911 | 0.9157 | 112.03 |
| | **'osiria/roberta-base-italian'<br>16 epochs: {0: 2e-6, 8: 2e-7}** | **0.9523** | **0.9264** | **0.9264** | **0.9383** | **201.76** |
| English | BoW & BR | 0.7964 | 0.7244 | 0.8070 | 0.7635 | 0.05 |
| | TF-IDF & BR | 0.7857 | 0.6850 | 0.8131 | 0.7436 | 0.05 |
| | 'google-bert/bert-base-cased'<br>4 epochs: {0: 2e-5, 1: 2e-6, 2: 2e-7} | 0.8310 | 0.8264 | 0.7989 | 0.8113 | 44.66 |

| | | | | | | |
|---|---|---|---|---|---|---|
| | 'distilbert/distilbert-base-cased'<br>8 epochs: {0: 2e-6} | 0.8164 | 0.7918 | 0.7918 | 0.7904 | 141.50 |
| | **'FacebookAI/roberta-base'<br>4 epochs: {0: 2e-5, 2: 2e-7}** | **0.8386** | **0.8565** | **0.7951** | **0.8239** | **154.00** |

#### *5.2.2.2 Visual analysis of social media content*

This subsection provides the updates of the visual analysis modules provided by the Social Media Analysis Toolkit in Deliverable D4.2. The methods employed by these modules are:

I. Fire and Smoke Detection in Images: This method utilizes computer vision techniques to identify flames or smoke in images (see Section 5.2.2.2.1).
II. Visual Concept Extraction: Uses object recognition algorithms to extract meaningful visual concepts and objects from images, such as buildings, vehicles, and people. The extracted information is stored in the metadata of the post (see D4.2, Section 5.2.2).
III. Location Extraction: Extracts location information from images of a social media post (see D4.2, Section 5.2.3).

##### 5.2.2.2.1 Fire and Smoke Detection in Images

The fire and smoke detection algorithms used for the Social Media Sensing Toolkit, that were discussed in Chapter 4 of this deliverable, are identical to those deployed in the fire detection IoT devices (UP4a). For the purposes of this task, the selected algorithms were deployed on CTL's servers and a web service was created (using REST API) to allow CERTH to use them on demand. In detail, the API receives HTTP POST requests with a list of image URLs, contained in social media posts, to be checked for fire/smoke occurrences. In turn, the API responds with the fire and smoke scores per image in the JSON format of the social media analysis modules described in D4.2 Section 6.1. An example of the API call, using URL client (cURL), is shown in Figure 42 and the corresponding response in Figure 43. It is important to note that if a link is invalid (e.g., the image is no longer available) it will be ignored by the detection service and scores will only be provided for the remaining images, as shown in the example provided. To maintain near real-time performance for the toolkit, the image-based fire/smoke detection service aims to provide responses in under a minute, depending on network bandwidth and the number of query images. Since the service typically operates over Wi-Fi or Ethernet, network bandwidth is usually not a concern, leaving the number of images as the primary factor. Based on the conducted tests, typically a social media post contains 0-5 images, the API response time for such cases is between 20-40 seconds, well within the one-minute threshold.

```
curl --request POST 'http://fire-detection.catalink.eu:1414/social_media' \
--header 'Content-Type: application/json' \
--data-raw '{
"urls":["https://pbs.twimg.com/media/FQCpFwRaIAQIni2.png", "https://pbs.twimg.com/
    media/FQDK24ZagAcgJiY.jpg", "https://pbs.twimg.com/media/FQDMkX1XIAUsD_z.jpg"]
}'
```

**Figure 42.** Example of image-based fire/smoke detection API call, using cURL command.

```
{
  "analysis_concepts": {
    "module_name":"CTL_concept_analysis_visual",
    "analysis_type":"Visual",
    "concepts": [
      "Fire",
      "Smoke"
    ],
    "metadata": [
      {
        "image_url":"https://pbs.twimg.com/media/FQCpFwRaIAQIni2.png",
        "concepts_data": [
          {
            "concept":"Fire",
            "score":0.042772
          },
          {
            "concept":"Smoke",
            "score":0.999005
          }
        ]
      },
      {
        "image_url":"https://pbs.twimg.com/media/FQDK24ZagAcgJiY.jpg",
        "concepts_data": [
          {
            "concept":"Fire",
            "score":0.001108
          },
          {
            "concept":"Smoke",
            "score":0.981263
          }
        ]
      }
    ]
  }
}
```


**Figure 43.** Example response of the image-based fire/smoke detection API.

As mentioned in Section 5.2.1.1 of D4.2, the fire and smoke detection algorithms have been trained on a robust and diverse image dataset, which was further enriched with examples from SILVANUS pilot activities. The challenges encountered during the live testing of the IoT devices, such as poor lighting conditions, are less common in this context due to the human factor, which ensures images are captured under ample light. This reduces the likelihood of missing a fire event in an image. However, visual verification alone is not sufficient to categorise a social media post as relevant to an ongoing wildfire, as demonstrated in [53]. For example, a post might refer to a wildfire awareness campaign that uses images from past incidents to emphasise its message, rather than indicating a current event. While visual verification is important to make a post more trustworthy, it must be paired with textual analysis to ensure accurate identification of relevant content. This prevents overwhelming the SILVANUS platform users with unrelated posts and thus alerts.

#### *5.2.2.3 Fire Events Detection*

The Fire Event Detection module plays an important role in monitoring and identifying potential wildfire incidents by leveraging social media data. This system combines temporal and spatial clustering technique to group relevant posts, enabling the detection of fire events across various geographic regions in near real-time. By analysing both the timestamp and location of posts, the module ensures that the fire events are based on the most current and relevant information available.

Clustering is a critical component of the fire event detection process, allowing the system to group social media posts based on their temporal and spatial proximity. This process is essential for detecting and analysing potential fire incidents across various geographic regions using social media data. The workflow begins with the retrieval of relevant data from MongoDB collections, each representing distinct data sources collected by the crawlers (see Section 5.2.1.1). A key aspect of this workflow is dynamically determining the date range for analysis, focusing on data from the previous month. This strategy ensures that the system operates with the most current information, enhancing its ability to respond promptly to real-time events. By using up-to-date data, the system increases its sensitivity to detecting and analysing potential fire incidents, thereby improving overall situational awareness and responsiveness.

After the data is retrieved, the system formulates queries on MongoDB based on the predefined time range and the relevance of each post. This step ensures that only posts potentially related to fire events are considered for further analysis. By filtering posts based on relevance—such as specific features or locations indicative of fire-related content—the system reduces computational overhead and concentrates resources on the most pertinent data. This focused approach enhances the system's ability to accurately identify and respond to actual fire events, ensuring that alerts and actions are based on credible and timely social media data.

Clustering is a crucial phase where the system groups post according to their temporal and spatial proximity to identify potential fire incidents. This is accomplished using the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm [54], configured with parameters like the minimum number of posts per event and the time window size. For example, setting the "minimum posts per event" parameter to 10 ensures that only clusters with at least 10 posts are considered significant enough to indicate a possible fire, reducing false positives and improving the reliability of detection. The "time window size" parameter determines how long posts are aggregated for clustering, which affects the granularity of event detection.

The clustering process further organizes posts into weekly segments based on their chronological order, making the data more manageable for analysis. Within each time segment, the system evaluates the geographic and temporal proximity of posts. The spatial proximity check ensures that posts within a 10-kilometer distance are grouped together, while the temporal check ensures that posts within the same batch period are considered. Posts that meet both criteria are clustered together, representing potential fire events. These clusters, which contain a specified minimum number of posts, are then formatted into JSON for further analysis.

Throughout the detection process, the system monitors its performance, logging execution times for clustering and keyword extraction to ensure efficiency and timely insights into emerging fire events. Implemented in Python, the system operates via a Flask-based RESTful API, allowing it to receive HTTP GET requests and initiate the detection process every 15 minutes. This configuration ensures that the system can be remotely triggered and promptly respond to new data inputs.

Once fire events are identified, they are stored along with their associated metadata in MongoDB, making them readily accessible for further use within components of the SILVANUS platform, including visualization through SAL.

To assess the fire detection module's performance the fire event detection algorithm was run without any relevance filtering, meaning all detected posts were considered, regardless of their specific connection to wildfires. To verify the accuracy of these detected events, a cross-referencing process was undertaken. This involved comparing the detected events' details—such as keywords like "wildfire," "fire," "blaze," and the associated dates and locations—with reliable information from news reports and official sources available on the internet. This step was crucial for confirming that the events identified were indeed significant wildfire incidents. By checking against published news articles and official reports, we ensured that each event's classification as a real wildfire was based on credible external validation.

The initial execution of this method resulted in the identification of 51 events. Out of these, 42 were confirmed as actual wildfire incidents, while 9 were categorized as non-wildfire events or warnings. As a result, the initial accuracy of detecting real wildfire events based on keywords and context was approximately 82%.

To further improve accuracy, a Relevance Estimation analysis was applied to filter out non-relevant posts. This refinement reduced the number of detected events from 51 to 42. Among the 9 events that were filtered out, 8 were initially classified as non-wildfire events or warnings (False Positives), and 1 was a real wildfire event (False Negative).

Consequently, the filtered accuracy, which accounted for the relevance of the posts, improved to approximately 93.2%. This increase in accuracy reflected a more precise identification of real wildfire events and a significant reduction in false positives.

The results presented in Table 14 show that the fire event detection module exhibited strong performance, achieving an initial precision of 82% and a recall of 100% without filtering. After the relevance estimation algorithm was incorporated, precision improved significantly to 98%, while recall decreased slightly to around 95%. The initial unfiltered method was more comprehensive in detecting events but had a higher rate of false positives. In contrast, the filtered method improved precision but at the expense of missing a few actual wildfire events. The relevance estimation algorithm effectively refined the detection process, though further optimization and fine-tuning could enhance the module's performance, ensuring more accurate detection of real wildfire events while maintaining high precision.

**Table 14.** Fire event detection evaluation.

| Statistic | Initial Method | Filtered Method |
|---|---|---|
| True Positives (Real Wildfire Events) | 42 | 41 |
| False Positives (Non-Wildfire Events) | 9 | 1 |
| False Negatives (Non-Detected Real Wildfire Events) | - | 2 |
| True Negatives | - | - |
| Total Events | 61 | 42 |
| Precision | 0.82 | 0.98 |
| Recall | 1.00 | 0.95 |
| Accuracy | ≈ 82% | ≈ 93.2% |

#### 5.2.2.4 *Integration with SAL*

This section discusses the integration of the Fire Event Detection Module with other components within the SILVANUS platform, utilizing the SAL. Figure 44 illustrates the flow architecture of this integration.

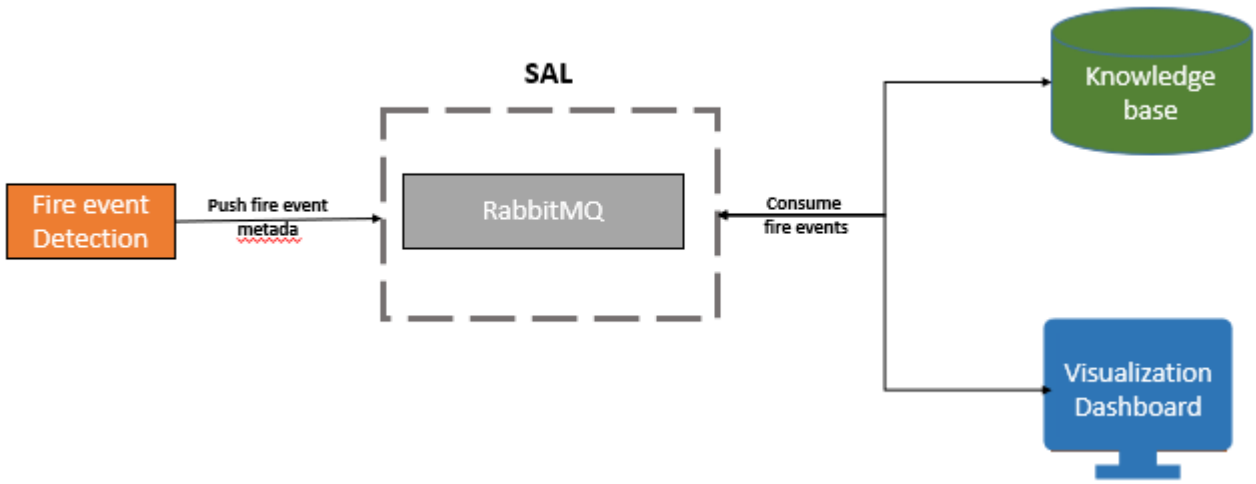


**Figure 44.** Integration flow of the Fire Event Detection Module within the SILVANUS platform

The Fire Event Detection Module is integrated into the SILVANUS platform through the SAL, ensuring smooth communication with other modules. The module detects fire events from social media platforms and formats them into a JSON structure. This JSON contains critical metadata, including the event's location, time, and additional relevant details.

Once a fire event is detected, it is forwarded to the SAL, where it enters a RabbitMQ message queue. The SAL serves as an intermediary that allows for the reliable delivery of fire event data to the consuming modules.

The SILVANUS Visualization Dashboard is designed to consume fire events directly from the RabbitMQ queue within the SAL. Upon receiving a fire event, the dashboard processes the JSON data, extracting necessary information to visualize the event. This visualization is typically displayed as a pop-up on a map interface, allowing users to quickly identify and assess the location and details of the fire event. Simultaneously, the SILVANUS Knowledge Base also consumes fire events from the RabbitMQ queue. The Knowledge Base's role involves transforming the fire event data into Resource Description Framework (RDF) format, which is a structured and semantic representation of the data. After this transformation, the RDF data is stored in the Knowledge Base's database, making it accessible for further analysis, reporting, and integration with other components of the SILVANUS platform [47].

#### 5.2.2.5 *Fire Events stored to knowledge database*

This section offers a concise update on the integration of social media posts into the Semantic Knowledge Base (SemKB) for subsequent utilization by other SILVANUS UP, such as UP9h. The current transformation to RDF and the population of the SemKB are conducted directly through the SAL, where CERTH transmits the data. CTL subscribes to the respective queue, consumes the data, and performs the data transformation as detailed comprehensively in D5.4, thereby populating the graph.

| | | | |
|---|---|---|---|
| 813 | :b461e603e7304f979147c50c360bfb77 | rdf:type | owl:NamedIndividual |
| 814 | :b461e603e7304f979147c50c360bfb77 | rdf:type | :RaspberryPiDetection |
| 815 | :bde1a7f9fe614489a6e2d03e5107e227 | rdf:type | owl:NamedIndividual |
| 816 | :bde1a7f9fe614489a6e2d03e5107e227 | rdf:type | :Measure |
| 817 | :bde1a7f9fe614489a6e2d03e5107e227 | rdf:type | :MeasureInDegrees |
| 818 | :bde1a7f9fe614489a6e2d03e5107e227 | rdf:type | :SensoryData |
| 819 | :bde1a7f9fe614489a6e2d03e5107e227 | rdf:type | :TemperatureSensoryData |
| 820 | :causesFire | rdf:type | owl:ObjectProperty |
| 821 | :causesFire | rdf:type | owl:ReflexiveProperty |
| 822 | :causesIncident | rdf:type | owl:ObjectProperty |
| 823 | :containsDetection | rdf:type | owl:ObjectProperty |
| 824 | :contributesToBuildupIndex | rdf:type | owl:ObjectProperty |
| 825 | :contributesToFireWeatherIndex | rdf:type | owl:ObjectProperty |
| 826 | :contributesToInitialSpreadIndex | rdf:type | owl:ObjectProperty |
| 827 | :d902925135f0474586b6f14c949c03bb | rdf:type | owl:NamedIndividual |
| 828 | :d902925135f0474586b6f14c949c03bb | rdf:type | :Coordinates |
| 829 | :da4ae9db2dad4d62b5970ef5b016753e | rdf:type | owl:NamedIndividual |
| 830 | :da4ae9db2dad4d62b5970ef5b016753e | rdf:type | :Geometry |
| 831 | :dc6b8a093227448ca34cfeaf19731963 | rdf:type | owl:NamedIndividual |

**Figure 45.** RDF Types and Object Properties.

The backbone of a semantic knowledge graph is its RDF structure, which defines how data is stored, categorized, and linked. Figure 45 is part of a table listing RDF types and object properties associated with various entities within the SemKB. This detailed tabulation includes the RDF type (e.g., owl:NamedIndividual, SensoryData, TemperatureSensoryData) and the associated properties (e.g., causesFire, containsDetection). This information underpins the ontology-driven structure of the knowledge base (for more ontology related details refer to D3.1), ensuring that all data is systematically categorized and linked according to well-defined semantic rules.

A chord diagram, as can been seen in Figure 46 provides a compelling overview of the interdependencies between various classes within the knowledge base. Each segment of the diagram represents a different class, such as Concept, GeoLocation, and RaspberryPiDetection. The arcs that connect these segments illustrate the flow of data and the relationships between different classes. This visualization not only captures the richness of the graph but also highlights the critical connections that enable the system to function as an integrated whole. It serves as a visual testament to the intricate design and thoughtful organization of the data architecture within the SILVANUS project.

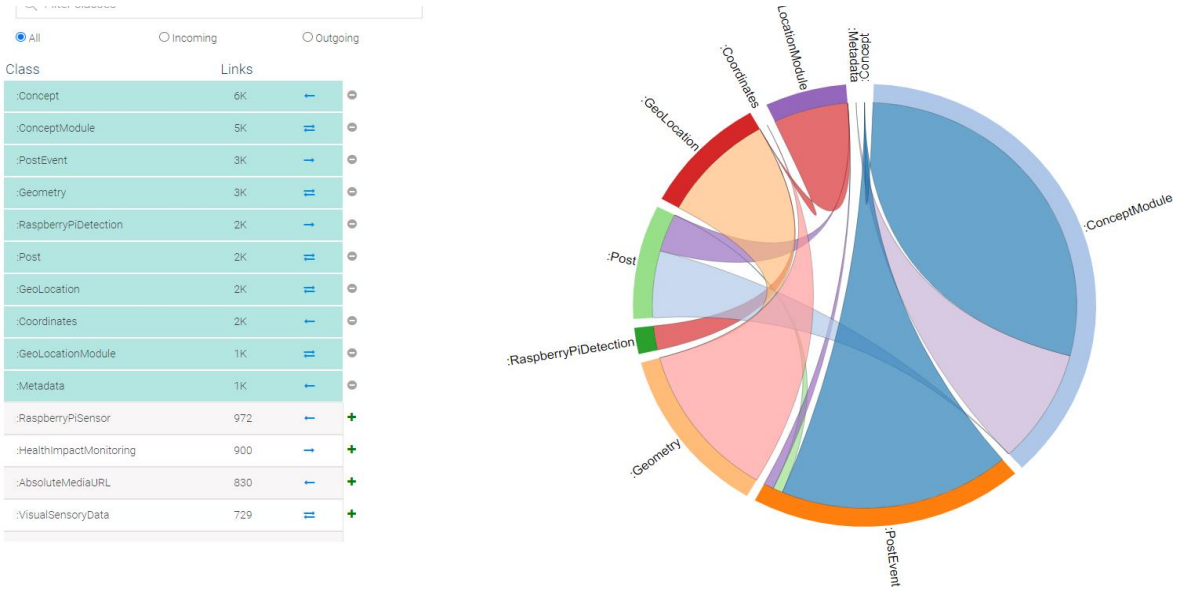


**Figure 46.** Visualizing Interdependencies within the Knowledge Base.

Figure 47 displays a node-link diagram representing a specific social media post in the knowledge graph. The central node corresponds to the unique identifier for the post (ID: 1602746121675218944). It is classified as an owl:NamedIndividual of type Post. Surrounding nodes include various attributes and related entities, such as the language (el), platform (twitter), and associated media (e.g., image). This visualization provides a clear depiction of how individual posts are represented and interconnected within the knowledge graph, highlighting the relationships between different data attributes and modules.

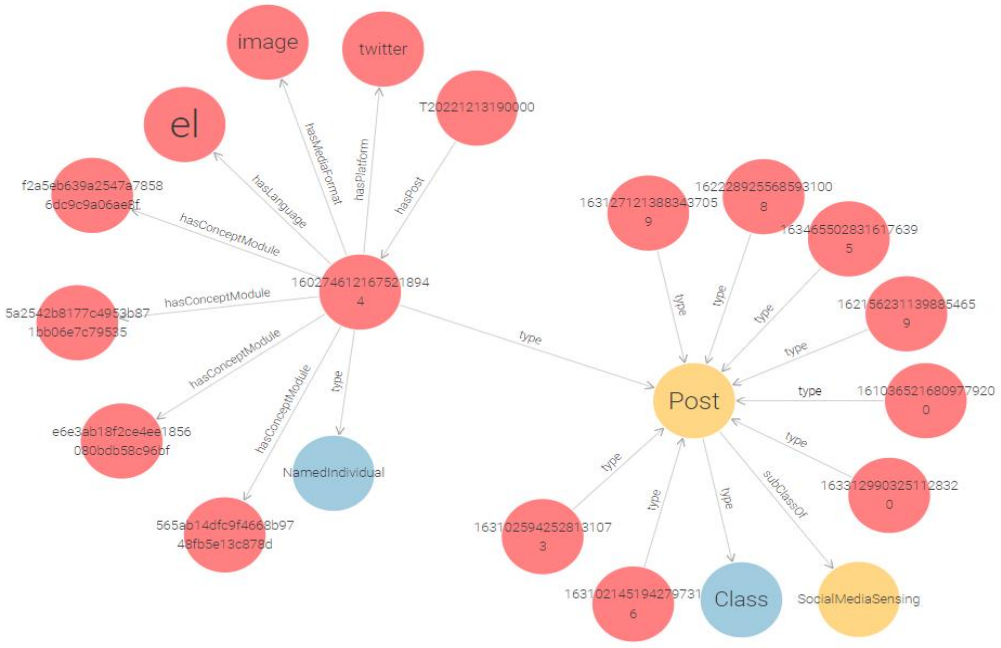


**Figure 47.** Visualizing the social sensing ontological network.

SemKB builds a network of interconnected entities instead of storing isolated pieces of information. Figure 48 showcases the broader network of connections surrounding the same social media post. In this expanded view, the post is linked to a range of semantic classes, such as Post, Class, and SocialMediaSensing. Each of these classes represents a broader category of knowledge, which helps in categorizing and inferring additional information from the data. This network of connections not only facilitates richer querying but also supports advanced inferencing, enabling the system to derive new insights from existing data.

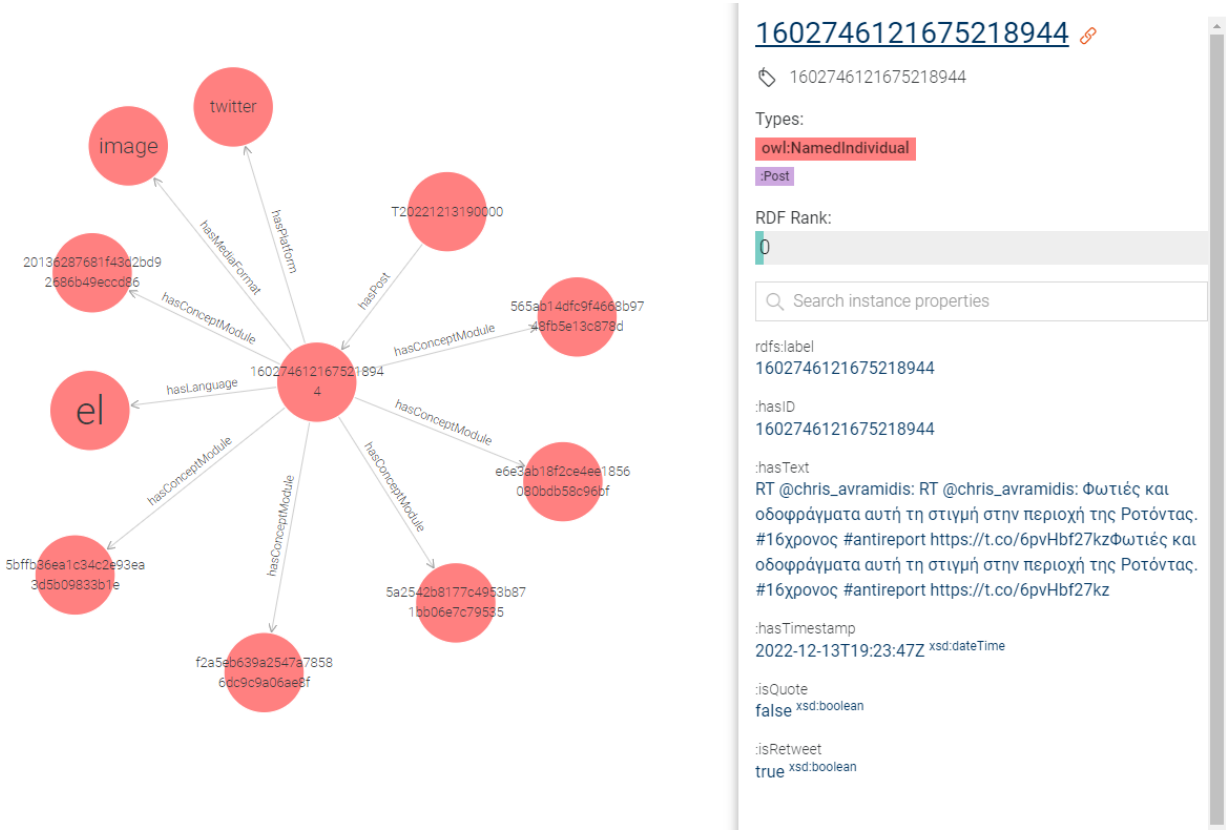


**Figure 48.** Visualizing social media in the form of graph database.

Integrating social media data into the SemKB requires preserving the relationships and context that ensure a comprehensive understanding of the content. This is achieved through the structured application of the SILVANUS ontology, which provides a robust framework for maintaining these critical connections. Figure 48 again illustrates how a specific tweet is represented within the knowledge graph. The central node corresponds to the tweet, identified by a unique ID (1602746121675218944), and acts as a hub linking various aspects of the post. This node is classified as an owl:NamedIndividual under the Post class, indicating its role as an individual social media entry. Surrounding this central node are other nodes that represent different attributes of the tweet, such as the language ("el" for Greek), the platform (Twitter), and any associated media like images. This diagram provides a clear example of how part of the tweet is saved in the form of RDF. It highlights the interconnected nature of the data, enabling a more detailed semantic understanding of the content within the knowledge graph.

## 5.3 Scientific results and drawbacks

### *5.3.1 Social media Crawling*

The analysis done in the Section 5.2.1.2 summarises various methods and algorithms used for ongoing monitoring of social media activity, especially during times of increased posting activity. It was found that

Isolation Forest provides the best balance of accuracy, robustness, and practicality for the specific task of monitoring tweet activity related to wildfires. Here we outline some limitations of the Isolation Forest method based on the analysis.

The main technique involved using the Isolation Forest algorithm to identify anomalies in the post count data. Isolation Forest is effective at detecting outliers without needing labelled anomaly data. The algorithm was configured with a contamination parameter that controlled the expected proportion of anomalies in the dataset. Anomalies were identified based on the timestamps provided by the algorithm's output.

One of the primary strengths of this method lies in its robustness, particularly in detecting outliers across diverse post patterns. Whether dealing with dense or sparse regions of the dataset, the method proved effective in identifying a wide range of anomalies. However, it is worth noting that the approach is computationally intensive and requires careful tuning of the contamination parameter to achieve optimal results.

### *5.3.2 Textual Analysis*

HB combined the following, beyond state-of-the-art scientific ideas and insights for tool design:

- We have employed sentence embedding based on Large Language Models (LLMs), more specifically the SentenceTransformers framework (see [36]), to base our content representation for ML on a viable semantic framework.
- Respectively, the dimension-reduced content representation vectors were visualized in scatterplots together with superimposed contour maps based on Gaussian Kernel Density Estimation (KDE), PageRank scores, and sentiment scores (Figure 49). This was achieved by tagging tweets with their sentiment vs. social importance values as computed by sentiment analysis and centrality analysis. The resulting 3-dimensional contour maps can help the visual identification of sentiments or perceived importance in topical groups (clusters) of content.
- Tool development (Figure 50) to demonstrate the above resulted in a web service module processing text messages as a combination of semantics, its primary factor, and perceived emotional/social relevance as its associated, secondary aspect. This combination resulted in information fusion. The fusion concept can be applied to other document genres and other media, i.e., our field approach to semantic content is both novel and expected to be generic.
- By adding new topical subsets of adjective-noun bigrams, the tool can be adapted to other areas of disaster management, i.e., its scope is expandable.

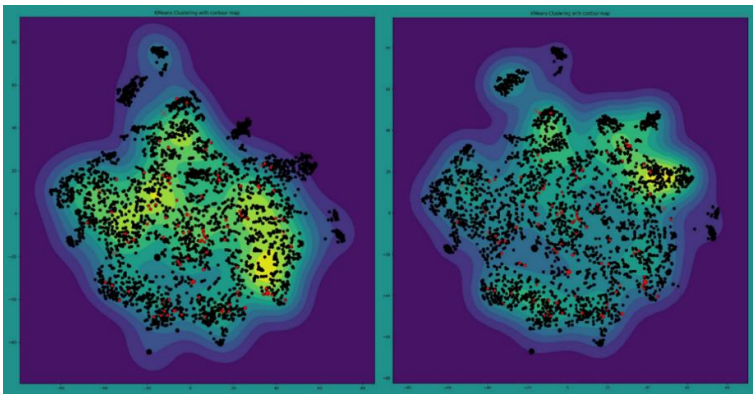

**Figure 49.** (left) Semantic intensity contour map of 4K tweets with 117 tweets, in red, indexed by high sentiment value phrases (bigrams) for comparison. Elevation corresponds to the heated nature of messaging. (right) Social importance contour map of 4K tweets. Elevation corresponds to the perceived importance (e.g., number of retweets) of messages.

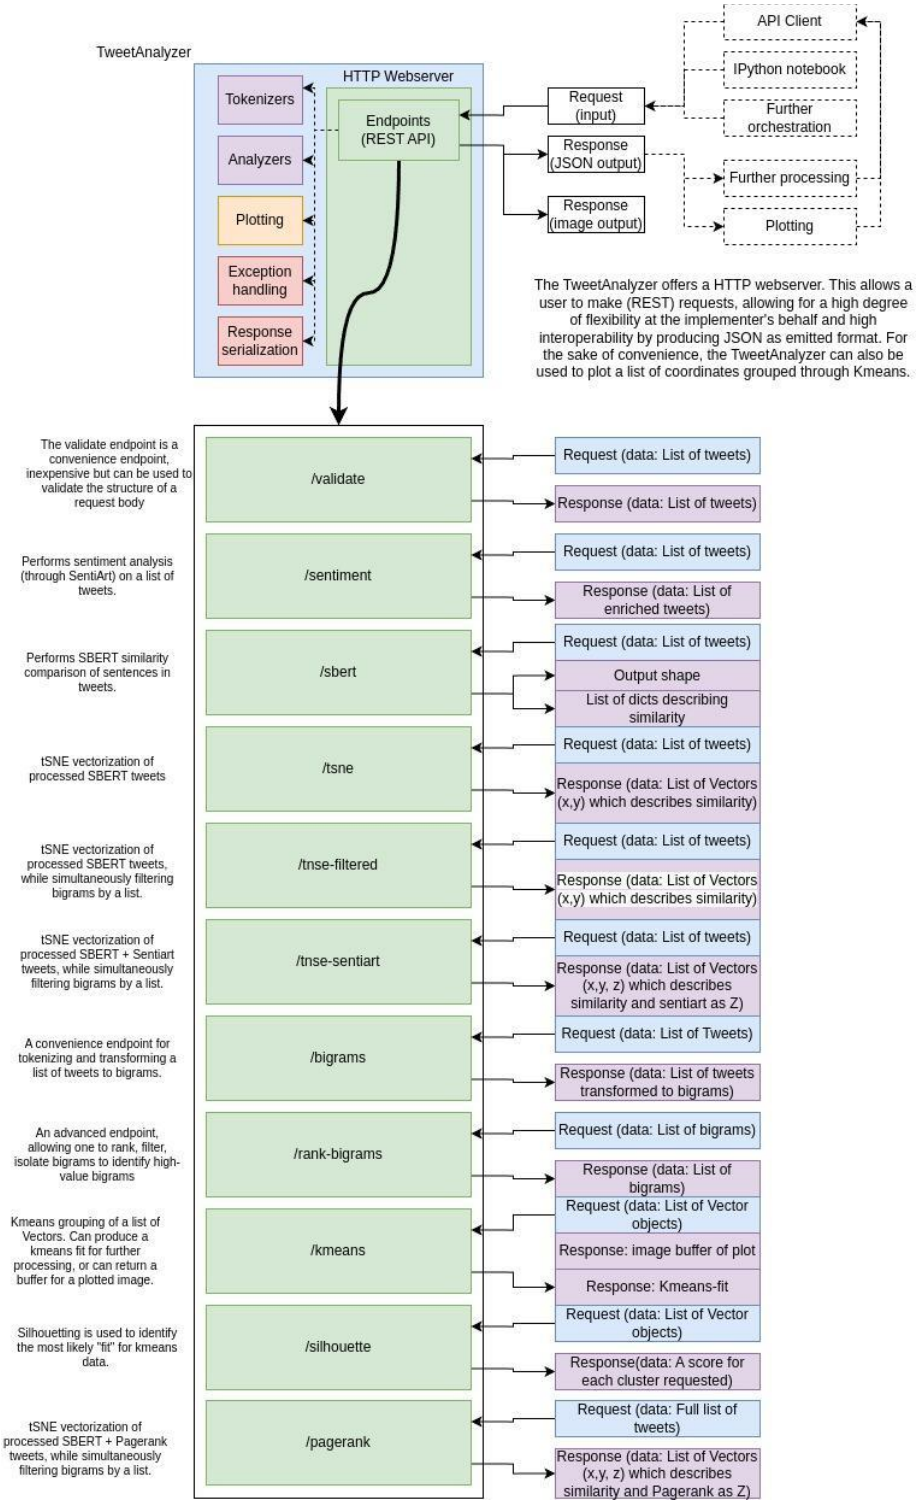


**Figure 50.** Workflow of the 'TweetAnalyzer for Information Fusion' web service module.

Known scientific drawbacks include the following:

- In a real-time environment social media sensing receives messages one by one in a stream, not in bulk as in a dataset. This means that for their clustering on the fly, some established categorial basis must pre-exist. We considered but did not follow up the handling of this research problem.
- Scalability and expandability tests were unfinished. In the future, the interpretation of cluster topics based on maximizing their semantic coherence would be crucial for the understanding of evolving fire events.
- Due to reasons of use restriction from X (Twitter), the tweets dataset provided by CERTH for experimentation proved to be too small with its language insufficient to achieve results beyond an indicative proof-of-concept level.
- The language of the dataset was not focusing on firefighting as a critical process with well-defined goals, but more on political topics. Therefore, a coupling of messages about an evolving fire event with its respective incoming physical signals could not be accomplished but would be possible with better and scalable data if within reach in WP4. Applied e.g. to professional firefighting communication, the interpretation of event related messages could become more accurate, a potential dashboard component.
- As shown in Figure 51, which contains four sets of coloured locations indicating tweets containing bigrams with high sentiment intensity accumulated over four updates, the red, blue, green and magenta dots do not sufficiently overlap with content attractors in the PageRank-based "landscape" of tweets. We experienced the same situation in the sentiment-based fusion landscape. This finding calls for the calibration of adjective-noun bigrams better tailored to firefighting scenarios, something that could be easily remedied with ontology and knowledge graph (KG) construction partners. Such conceptual fine-tuning would extract the most influential text words responsible for the 'sinks' with the most attractive 'force' in a fusion landscape.

Further, in terms of the capacities of our new web service module, the following must be noticed:

- Due to the number of requirements used, the resulting Docker is of a hefty 6 GB size at submission date.
- The default export format is JSON-based, complicating import in Comma Separated Value (CSV) based projects.
- The interface is a backend JSON API (albeit with a web-based Swagger interface), requiring implementers to know REST principles.
- The module is designed to analyse tweet data by taking in a list of tweets; in its present form it will not cater well to individual tweets evaluated vs. the "existing" base.
- The algorithms used require high amounts of computational power.

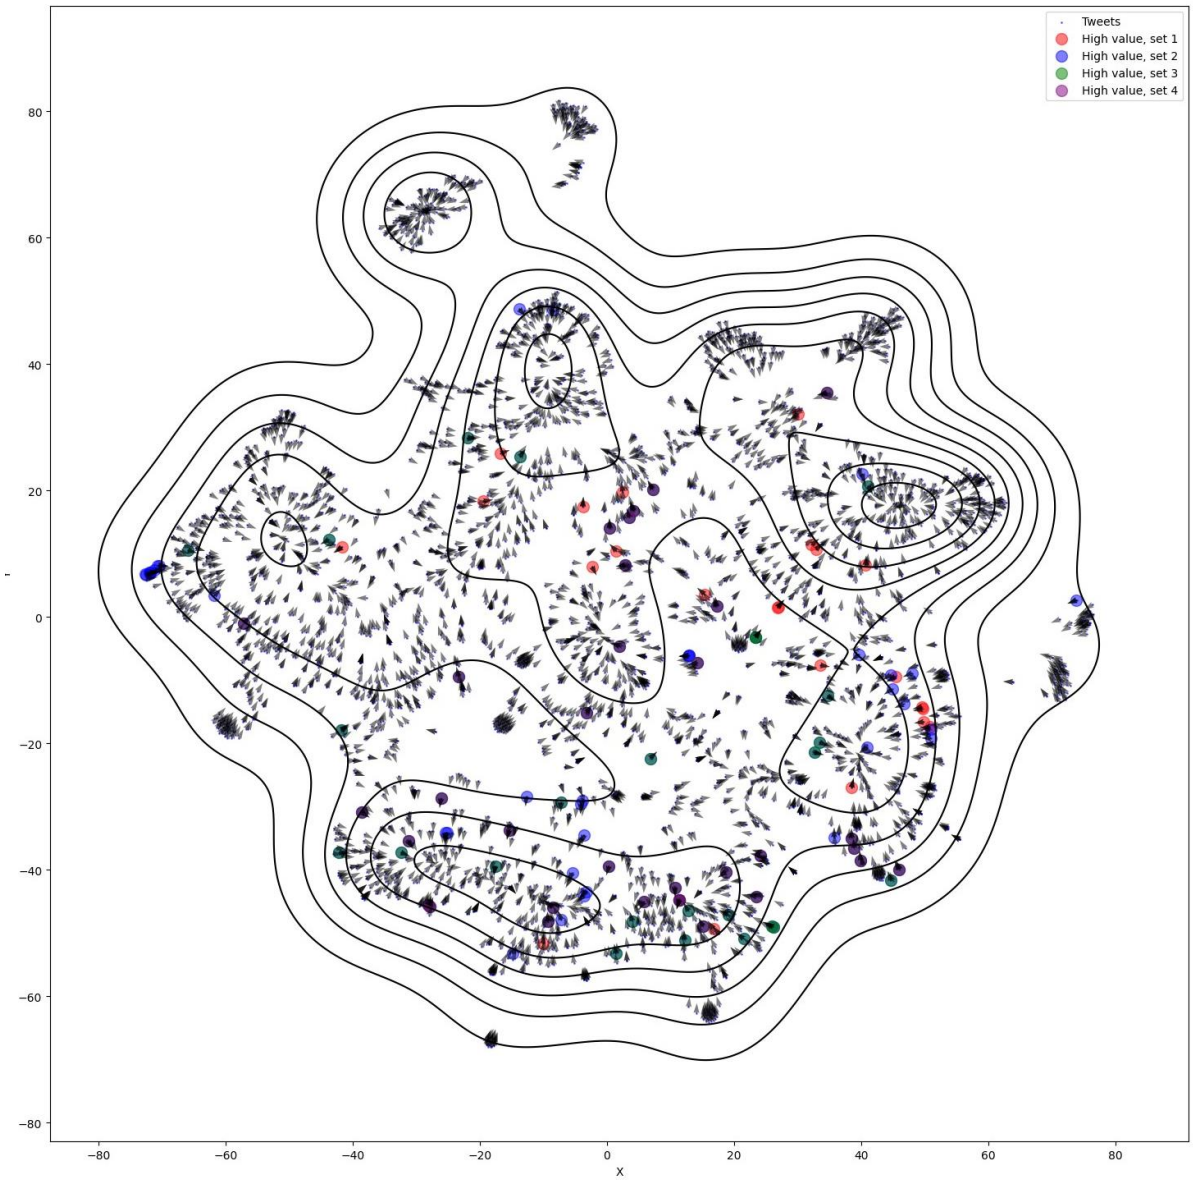


**Figure 51.** PageRank-based "landscape" of tweets, expansion phase No 4 (at 4 K): the sentiment intensity of marked up tweets does not overlap with the attractor structure**.**

### 5.4 Demonstration report

The SILVANUS dashboard (Figure 52) is a user-friendly tool designed for visualizing information about wildfire events. This platform integrates various data sources and layers, enabling users to monitor, analyse, and respond to wildfire threats effectively.

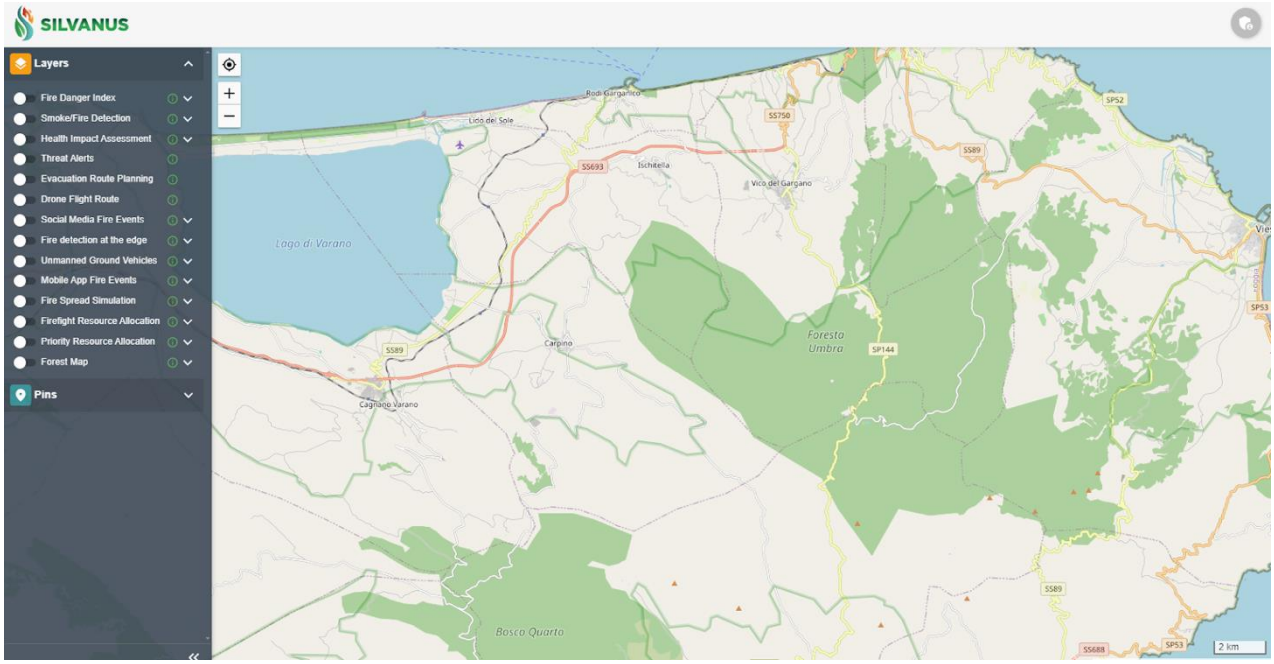


**Figure 52.** SILVANUS dashboard.

This dashboard includes a Map Section, where fire events and related data are visually represented. The map provides a geographical view of the region, showing various elements such as roads, forests, and other key landmarks. The core functionality of the map is to consume the new fire event that exist in SAL message queue and display them in the map as markers. These markers indicate locations where wildfires have been detected or where related events are taking place.

The Layer Section on the left side of the dashboard is crucial for customizing the map's display. This section allows users to filter and select various layers that represent different types of data related to wildfires. For instance, users can choose to display the FDI, which highlights areas with different levels of fire risk, or toggle on Smoke/Fire Detection, which shows places where active fires or smoke have been identified. Other layers might include Evacuation Route Planning, Fire Spread Simulation, or Social Media Fire Events, each providing specific insights that help in understanding and managing the wildfire situation. By activating or deactivating these layers, users can tailor the map view to focus on the most relevant information, enhancing their ability to make informed decisions.

When a user clicks on a pin on the map while actively using the Social Media Fire Events layer (Figure 53), a new section pops up, offering detailed information about the selected fire event. This information is typically presented in a panel on the right side of the screen. The details provided can include an image related to the fire, which might be a photo uploaded by a user or captured by a monitoring device. Additionally, text descriptions often accompany the image, providing context such as the severity of the fire, actions being taken, or any relevant updates. The event's exact location is also displayed, usually with coordinates and a map marker, allowing for precise identification of the affected area.

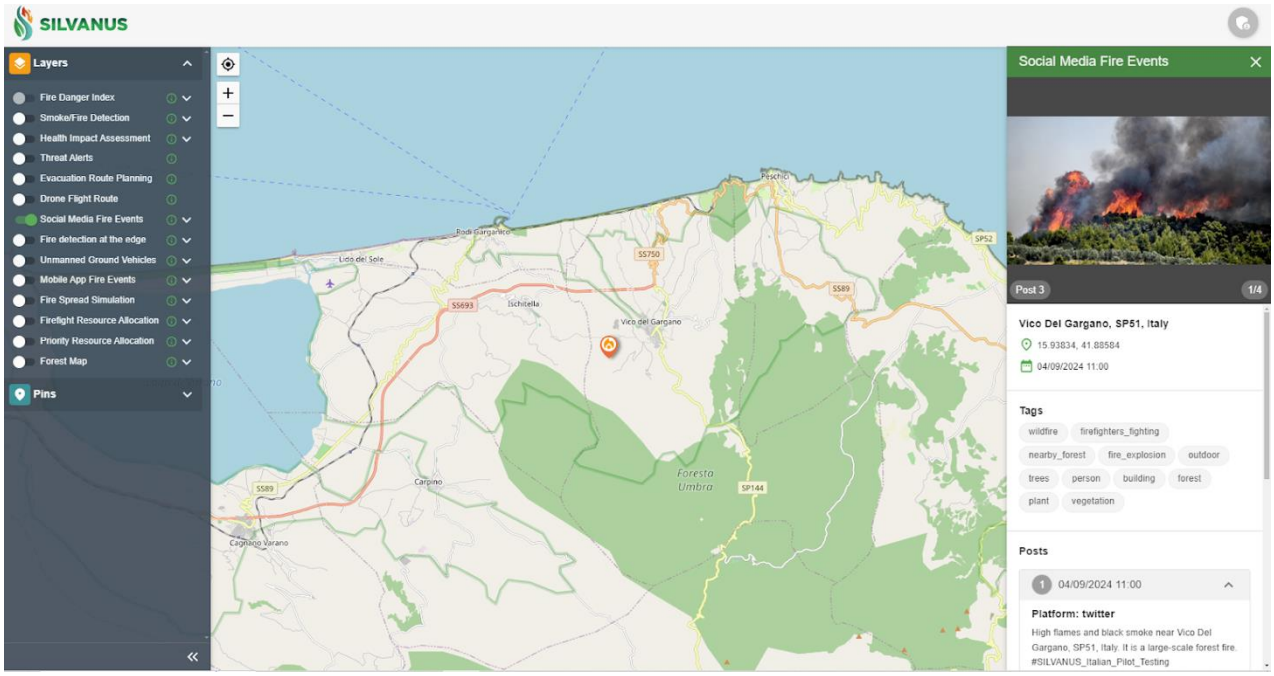


**Figure 53.** Visualization example of a fire event in SILVANUS dashboard.

Furthermore, this panel provides a timestamp, indicating when the fire was detected or when the information was last updated. This is crucial for understanding the progression of the event and for coordinating timely responses. Tags are also included, which categorize the event with keywords like "wildfire," "firefighters fighting," or "nearby forest," helping users quickly understand the nature of the event immediately.

Overall, the SILVANUS dashboard serves as a powerful tool for visualizing wildfire event. By combining an interactive map with customizable data layers and detailed event information, it supports both the monitoring of ongoing wildfires, and the strategic planning required to mitigate their impact. This integration of various data points into a single, accessible platform makes it an essential resource for anyone involved in wildfire management and response.

# 6 UGV monitoring for wildfire behaviour

## 6.1 Tool description

The goal of this system is to create automated uncrewed ground robots capable of gathering crucial data on wildfire behaviour. In this reporting period, CSIRO concentrated its R&D on the development of a multi-robot navigation using automatic map merging and place recognition. This capability allows single operators to oversee multiple vehicles and visualise them on the same navigation map.

In addition, the system incorporated mobile manipulation to allow a humidity sensor to extend to the ground to obtain ground-level soil humidity readings. This is particularly useful for the mitigation capability, where the readings can calibrate wider soil moisture maps, and feed into analysis of fire spread risk prior to a fire.

## 6.2 Innovations and updates

In this reporting period CSIRO continued the development and assessment of UGVs to acquire information about forest terrain. The focus was on creating new planning systems to increase the UGVs' ability to autonomously navigate intricate landscapes and create a terrain map of the forest region, including the use of humidity sensors and improved navigation offroad. Compared to the previous period, CSIRO expanded testing and deployment of UGVs, using, in addition to the BIA5 ATR [55] the Boston Dynamics Spot legged robots. All tests were done at the Australian test pilot site, where all partners met in November 2023. The results have shown successful UGV navigation in unstructured terrain, such as grassy areas amid trees, covering significant distances. The PB5 objective was fully met during this reporting phase, with further details available in Task 4.5 - UGV monitoring for wildfire behaviour.

In the demos presented by CSIRO, the ground navigation vehicle exhibited various functionalities, including autonomous navigation to and from the simulated wildfire frontlines by self-exploring the forest terrain. It autonomously transmitted data to a base station and then to a SILVANUS server via a REST interface, which was developed in collaboration with partner Dell. Additionally, the vehicle autonomously estimated foliage density, tree density and canopy coverage indices, and autonomously localized itself based on prior maps. The use of prior maps for navigation is not always available, but when they are, they were shown to be a valuable source of prior information for improved navigation.

In addition to autonomous navigation, we developed multi-agent capabilities using new machine-learning based place recognition methods. These allowed us to deploy multiple UGVs simultaneously, sharing each other’s map information. This was demonstrated in the Australian pilot using one ATR and two-legged Spot vehicles, as seen in Figure 54. An example of the ATR autonomously navigating in a forested environment is showing in Figure 55.

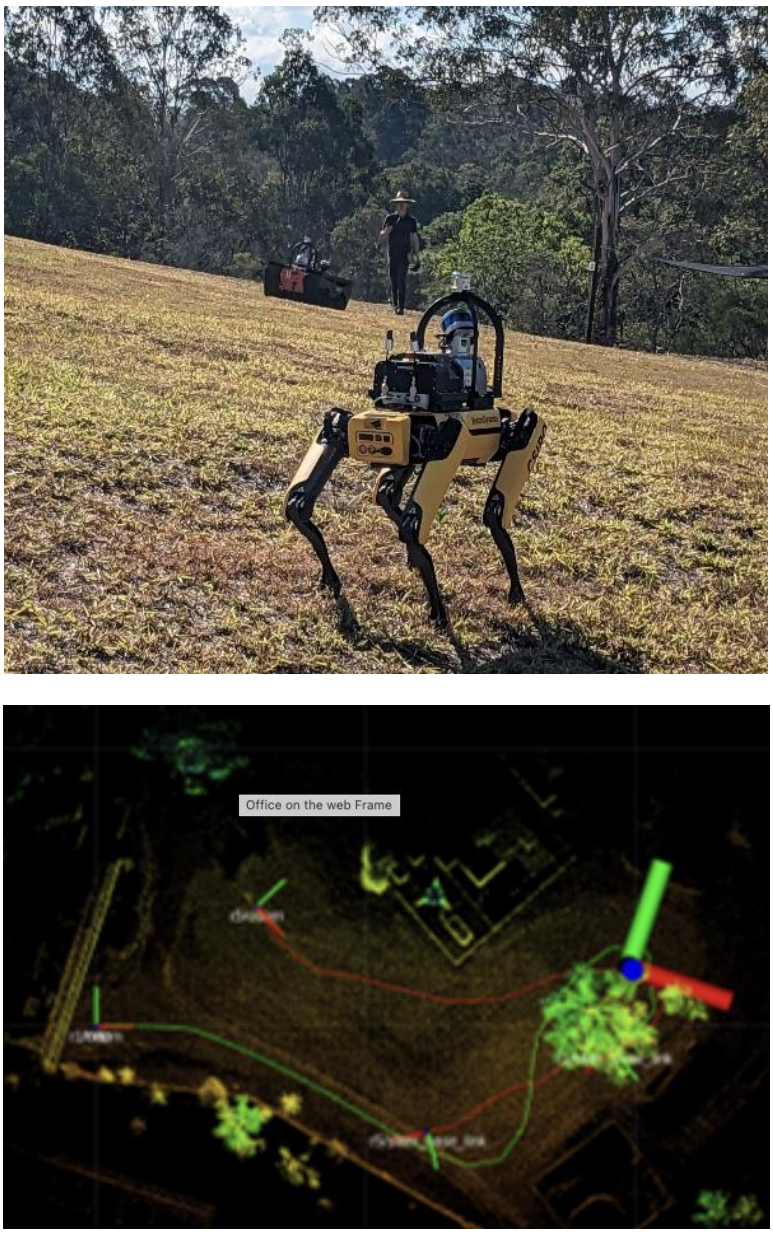


**Figure 54.** Multiple UGVs demonstrating online map merging at the Australia pilot.

**Figure 55.** The CSIRO UGV platform autonomously navigating in complex, natural terrain in the presence of (artificial) smoke, during a simulated mid-fire scouting operation.

**Summary of updates:**

- Robot navigation through dense forest environments in the presence of grass and undergrowth. Undergrowth detection has been an important new piece of work, which allows the robot to penetrate the forest without treating grass and plants as an immovable obstacle.
- Semantic segmentation of forest terrains based on ML methods. This technology provides a semantic label to every observed point from the vehicle's scanning lidar. Labels include ground, mud, water, grass, wood and foliage. This is demonstrated in Figure 56.
- Improved real-time forest analysis from lidar, specifically: tree density, canopy density and canopy coverage. We also have algorithms to estimate the full tree geometry from lidar.
- A reporting pipeline to package and send the canopy analysis, robot pose, and camera images back to the base station using RF comms nodes, and then up to the SILVANUS platform. This included the ability to query the SILVANUS platform and display the information on a satellite view of the fire front area. This is demonstrated in Figure 57.
- Mobile manipulation to allow a humidity sensor to extend to the ground to obtain ground-level soil humidity readings. This is particularly useful for the phase A mitigation capability, where the readings can calibrate wider soil moisture maps, and feed into analysis of fire spread risk prior to a fire.

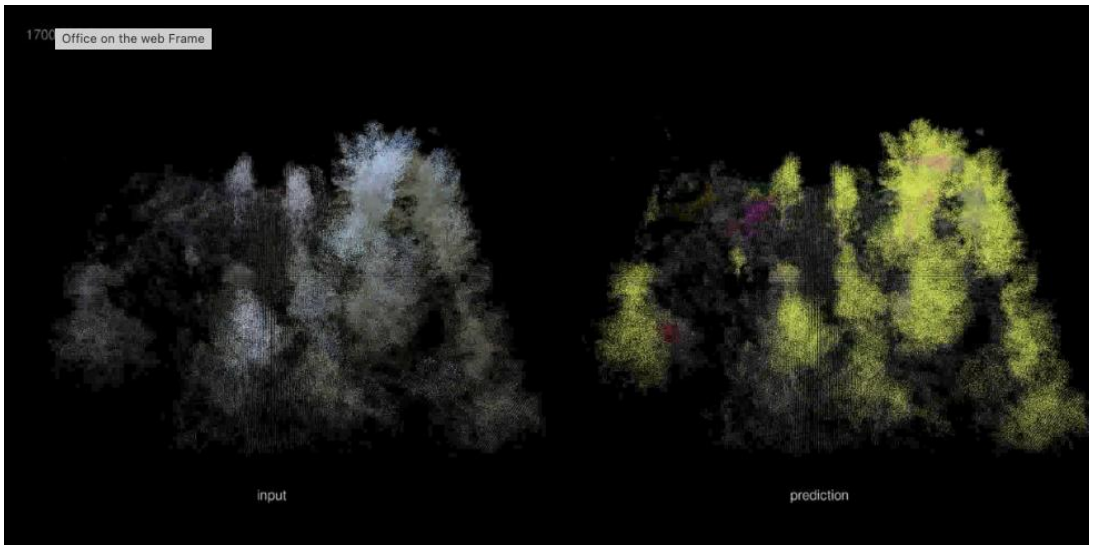


**Figure 56.** Semantic classification showing coloured point cloud (left) and predicted class labels (right). Data were collected during demonstrations at the Australia pilot.

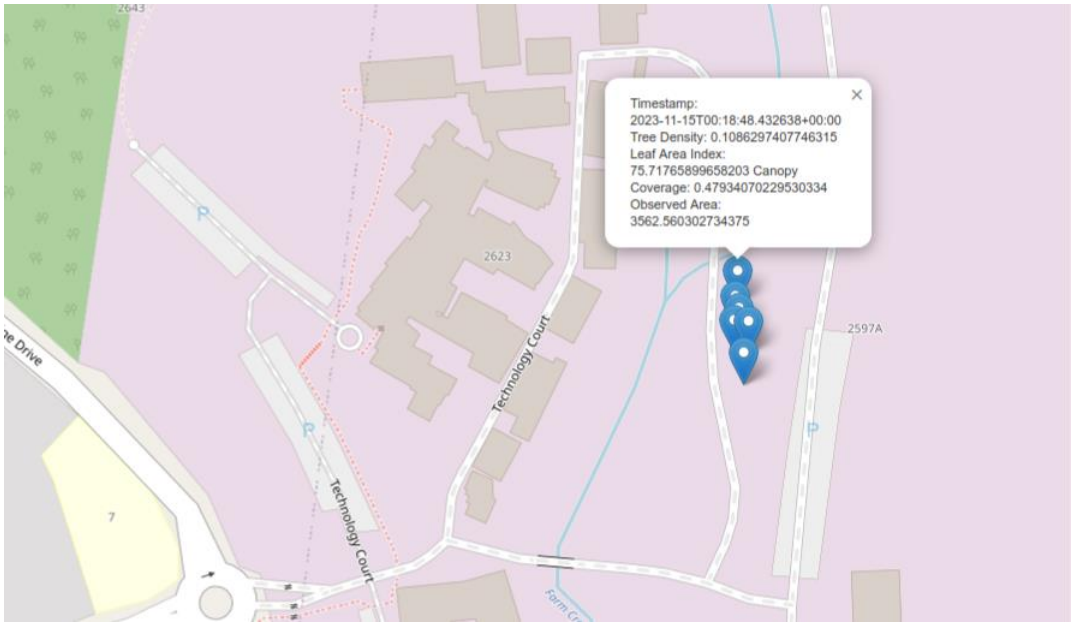


**Figure 57.** Data recorded by the robot during an autonomous exploration mission at the Australia pilot. Records were automatically uploaded to the SILVANUS platform and can be retrieved and plotted on an aerial map view. Information includes location, time, and *forest analytics results.*

Several of these developments began prior to M18, and so the recent work has included a large increase in field tests to road-test these behaviours in realistic scenarios. This has included testing the behaviours in the presence of smoke, moist ground, and undergrowth; and on two different classes of vehicle: tracked and legged.

We also collaborated with CTL to allow their smoke detector to be placed on the moving vehicle and report the detection of smoke in the vehicle's forward direction, directly to the SILVANUS platform.

### 6.3 Scientific results and drawbacks

We published papers on the topic of forest navigation which include:

- ForestTrav: 3D LiDAR-only Forest traversability estimation for autonomous ground vehicles [56].
- Deep robust multi-robot re-localisation in natural environments [57].

Navigation in cluttered forest environments is an extremely complicated robotics problem and operation is slow. The speed issue is something that needs to be addressed in the future. Negotiating terrain and obstacles (sticks, branches, etc) also poses a challenge, slowing down the operation speed.

### 6.4 Demonstration report

D4.5 for CSIRO was showcased at the Australian Pilot in November 2023, which demonstrated the technology described above. The results above were extracted from this pilot demonstration, which was presented through live demos to all partners presents.

# 7 UAV deployment for remote sensing and identification of wildfire

## 7.1 Tool description

Area surveillance using unmanned remote devices is useful in wildfire management, as it allows to detect fire starts, monitor existing fires or map threatened surroundings without endangering human operators. For such missions, UAVs are equipped with sensing devices to collect informative footage for firefighters and other incident managers. The sensors are typically regular, high-resolution RGB cameras but can also include thermal cameras, Lidars or other advanced sensing technologies.

An algorithmic toolbox to procure optimal flight plans for area surveillance by UAVs has been developed. The algorithmic module generates optimal sweeping trajectories based on several sensing parameters, such as camera shooting angle, as well as flight parameters like altitude and turning radius.

The role of this module within the broader Integrated Fire Management Approach is illustrated in Figure 58. It plays a crucial part during the vigilance phase, by either monitoring an ongoing incident or by detecting new fire starts.

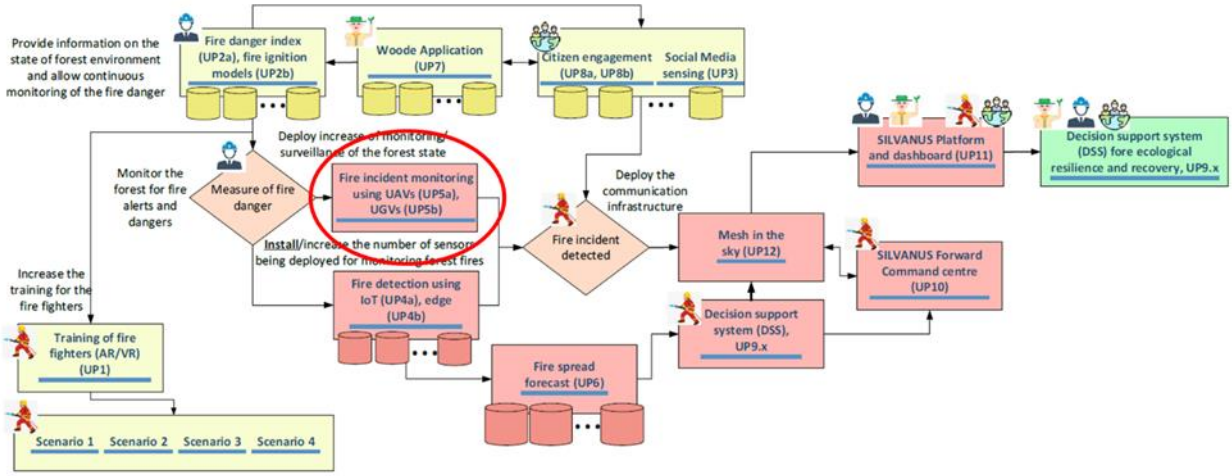


**Figure 58.** Place of trajectory optimization module within the integrated Fire Management approach.

## 7.2 Innovations and updates

Beyond sheer trajectory optimization improvements, much of the later research and development in T4.6 work carried out on the trajectory optimization module focused on subarea definition, which is how you decompose the overall area to monitor. In particular, work was done, as illustrated below on how to handle zone definitions with topological holes in them or to accommodate “no-fly” zones constraints in the coordination algorithm services and EmerPoll system (see Figure 28). The unique feature of the approach is the deviation of area as we can see in the Czech pilot space. In Figure 59, the green polygon represents the monitored area, and the red polygon represents the no-fly zone. The results of the planned two paths for drones are shown in Figure 60. For illustration, we tried the algorithm in the Romanian pilot area. The monitoring area (white polygon) was intentionally enlarged (green polygon) so that we could use more drones. The results of the planned five paths for drones are shown in Figure 61. The result of the distribution of flight paths when applying the no-fly zone can be seen in Figure 62.

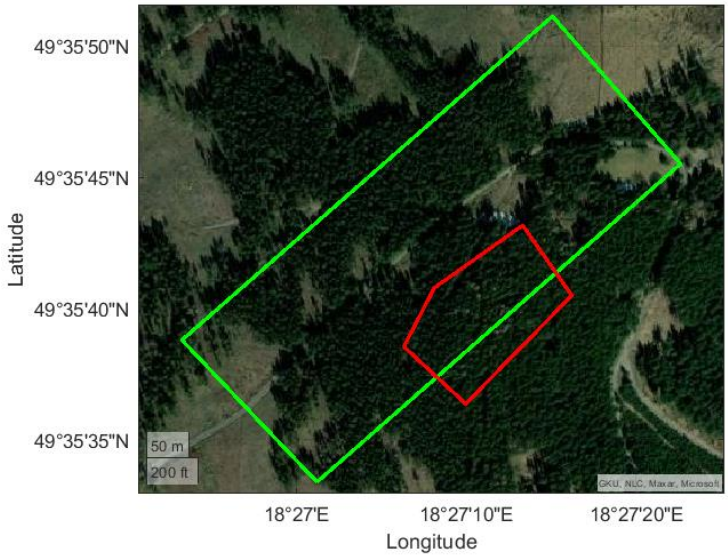


**Figure 59.** Monitored area for Czech pilot, green polygon represents the monitored area, and red polygon represents no-fly zone.

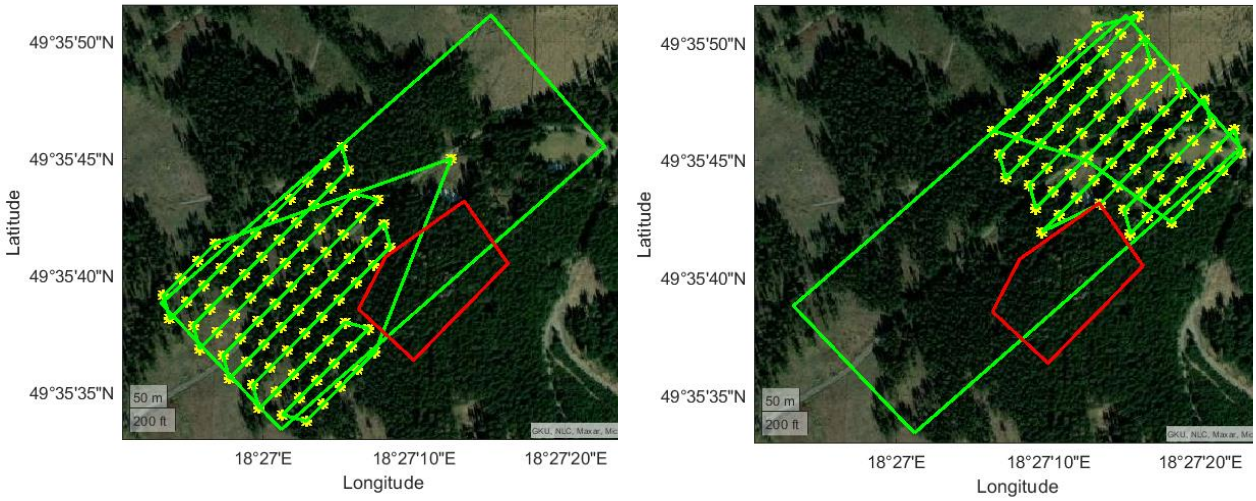


**Figure 60.** Results of drone path generation for two drones in Czech pilot.

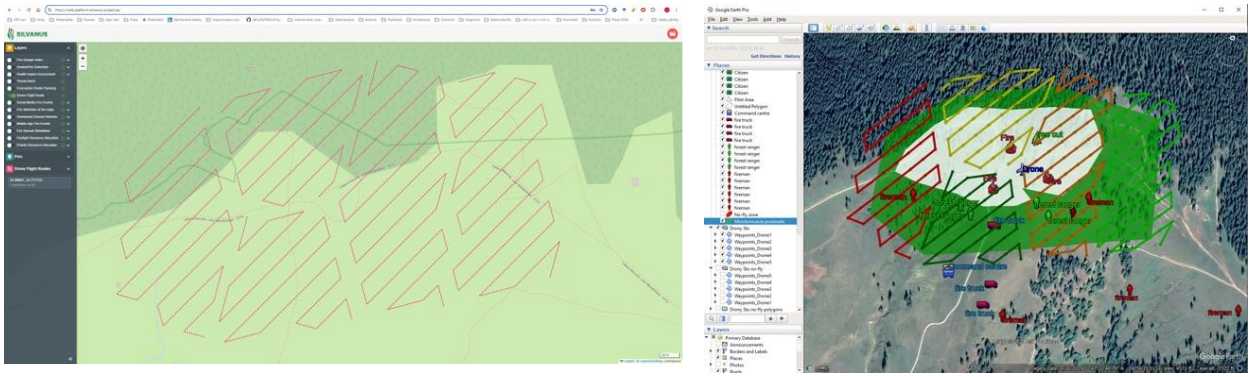

**Figure 61.** Results of drone path generation for five drones in Romania pilot. Dashboard visualization and visualization in Google earth pro.

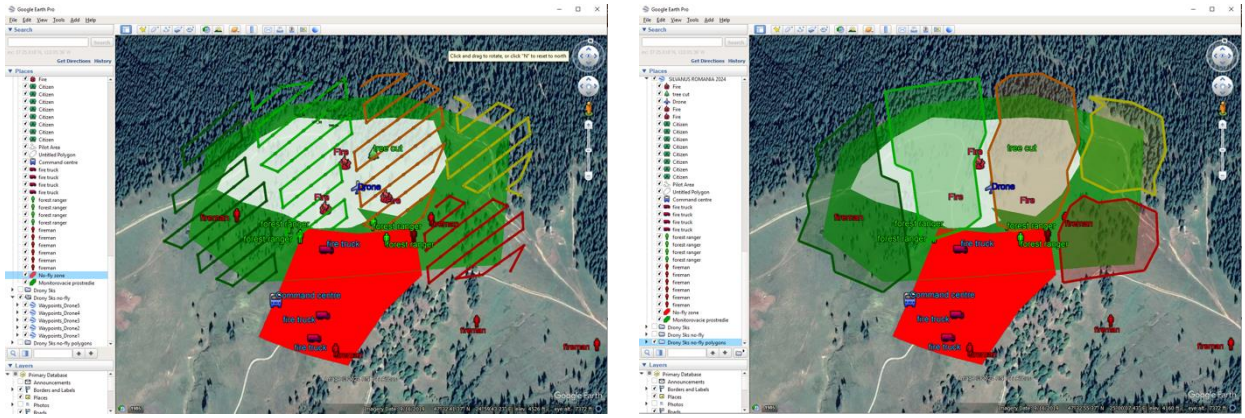

**Figure 62.** Results of drone path generation for five drones in Romania pilot with non-fly zone. Individual path or polygon.

### 7.3 Scientific results and drawbacks

The output of the tool is a fully functional coordination and trajectory optimization module for UAV fleets. This component effectively solves the underlying mathematical optimization problems to generate full mission plans, including area coverage and optimal path computation in a sufficiently rapid manner to facilitate agile “re-planning” of surveillance missions. In addition to the primary tabletop simulation-based, which utilizes real data, several real-world test flights were conducted to confirm that this technology can be effectively employed under suitable natural conditions. Ongoing work as part of WP9 and upcoming pilot projects aims to enhance the module’s robustness and ease of use while also addressing decentralized coordination as a longer-term perspective.

### 7.4 Demonstration report

Besides the numerous demonstrations conducted during project meetings and reviews, the module was utilized to compute and demonstrate trajectories during the tabletop exercises of the Italian Gargano Park pilot and the Greek pilot in fall of 2023. More recently, several actual test flights were carried out in the field during the 2024 edition of the Czech pilot in Krasna.

More specifically, the drone’s image collection and aggregation capabilities as well as ingestion to SAL from the FCC were tested on-site (Figure 63 a). Multiple trajectories were generated to monitor the selected incident site using two different algorithms by UISAV and TRT. The image below shows the results of two optimally planned drone routes (Figure 63 b) which were computed at the FCC. The drones had to operate at different altitudes (levels) (120 and 110 m above sea level) due to national safety requirements. The second photo (Figure 63 c) shows a drone controller with loaded data generated and prepared in the prior step. Based on this data two parallel drone missions were launched using two distinct drones. The final image (Figure 63 d) shows the result of stitching process of partial images collected and aggregated from the planned drone flights.

*a) Preparations for the mission at the pilot site (UISAV, TRT, 3MON, HZSCR).*

*b) Trajectories of optimal planned drone routes in KML-format precomputed and prepared for drone pilots.*

*c) KML-formatted drone route prepared and loaded to the drone controller before the flight.*

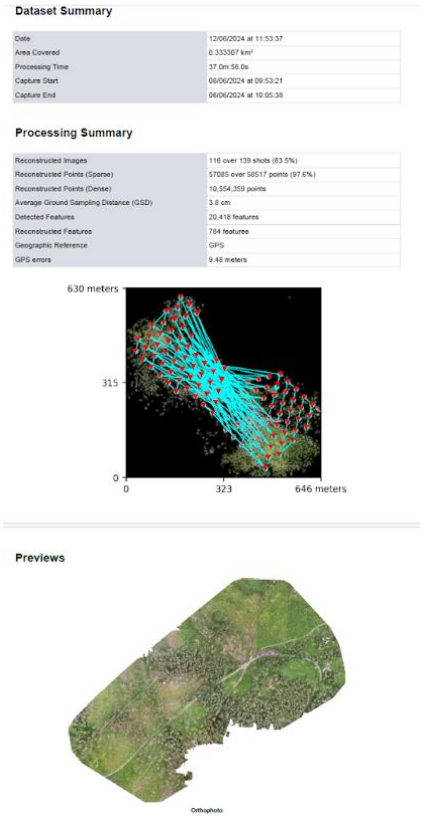


*d) Result of the drone image stitching from two independent drone flights.*

**Figure 63.** UAV drone flight data collection preparation, execution and aggregation at the Czech pilot site.

# 8 Conclusion

This deliverable summarizes the final implementation status at M36 of the various SILVANUS components for the ingestion, aggregation and pre-processing of heterogeneous data sources, including Satellite EOs, weather and climate data, in-situ IoT devices, social media sensing, UGVs and UAVs, and their integration into the SILVANUS platform. These integrated components provide advanced capabilities for fire detection. The deliverable also outlines the limitation and drawbacks encountered during the development and testing phases, as well as how these capabilities support the demonstration activities in WP9. These components are considered the backbone of the SILVANUS application and services and will be fully utilized in the SILVANUS Pilot sites in the coming months. This demonstrative phase of the SILVANUS system's capabilities will be extensively tested to gather feedback on the various developed components from users and stakeholders.

# 9 References


[1] E.-M. C. o. C. C. (. Foundation), “geolake,” [Online]. Available: https://github.com/CMCC-Foundation/geolake. [Accessed 25 09 2024].

[2] “geokube,” [Online]. Available: https://zenodo.org/records/10597965. [Accessed 19 09 2024].

[3] “Prefect,” Prefect, 13 09 2024. [Online]. Available: https://www.prefect.io/. [Accessed 13 09 2024].

[4] S. Kondylatos, I. Prapas, . M. Ronco, I. Papoutsis, G. Camps-Valls,, M. Piles, . M.-Á. Fernández-Torres and N. Carvalhais, “Wildfire danger prediction and understanding with deep learning.,” *eophysical Research Letters,* vol. 49, no. 17, 2022.

[5] “NASA MODIS,” [Online]. Available: https://lpdaac.usgs.gov/products/mod11a1v006/. [Accessed 19 09 2024].

[6] “The European and Global Drought Observatories,” [Online]. Available: https://drought.emergency.copernicus.eu/. [Accessed 19 09 2024].

[7] G. I. O. Środowiska, “Corine Land Cover - CLC,” [Online]. Available: https://clc.gios.gov.pl/index.php/9-gorne-menu/clc-informacje-ogolne/58-klasyfikacja-clc-2. [Accessed 25 09 2024].

[8] “WorldPop,” [Online]. Available: https://www.worldpop.org/. [Accessed 25 09 2024].

[9] “xarray.DataArray.interp_like,” [Online]. Available: https://docs.xarray.dev/en/latest/generated/xarray.DataArray.interp_like.html. [Accessed 19 09 2024].

[10] W. Skamarock, J. Klemp, J. Dudhia, D. Gill and Z. Liu, “A Description of the Advanced Research WRF Model Version 4.3,” Open Sky, 2021.

[11] R. Cornes, . G. van der Schrier, E. J. M. van den Besselaar and P. Jones, “An Ensemble Version of the E-OBS Temperature and Precipitation Data Sets,” *Journal of Geophysical Research: Atmospheres,* no. 123, p. 9391–9409., 2018.

[12] M. Raffa, A. Reder, G. F. Marras, M. Mancini, G. Scipione, M. Santini and P. Mercogliano, “VHR-REA_IT Dataset: Very High Resolution Dynamical Downscaling of ERA5 Reanalysis over Italy by COSMO-CLM,” *Data,* vol. 6, p. 8, 2021.

[13] M. Raffa, M. Adinolfi, . A. Reder, . G. F. Marras, M. Mancini, G. Scipione, M. Santini and P. Mercogliano, “Very High Resolution Projections over Italy under different CMIP5 IPCC scenarios,” *Scientific Data,* vol. 10, p. 238, 2023.

[14] S. Materia, Á. Muñoz, M. d. C. Álvarez-Castro, S. Mason, F. Vitart and S. Gualdi, “Multimodel Subseasonal Forecasts of Spring Cold Spells: Potential Value for the Hazelnut Agribusiness,” *Weather and Forecasting,* vol. 35, no. 1, pp. 237-254, 2020.

[15] P. Bourgault, D. Huard, T. J. Smith, T. Logan, A. Aoun, J. Lavoie, É. Dupuis, G. Rondeau-Genesse, R. Alegre, C. Barnes, A. B. Laperrière, S. Biner, D. Caron, C. Ehbrecht and Fyk, “xclim: xarray-based climate data analytics,” *The Open Journal,* vol. 8, no. 85, p. 5415, 2023.

[16] ECMWF, “Monitoring European climate using surface observations,” [Online]. Available: https://surfobs.climate.copernicus.eu/surfobs.php. [Accessed 25 09 2024].

[17] V. Wagner, “Structure of the Canadian forest fire weather index,” *Environment Canada, Forestry Service Ottawa, ON, Canada,* vol. 1333, 1974.

[18] M. McElhinny, J. Beckers, C. Hanes, M. Flannigan and P. Jain, “A high-resolution reanalysis of global fire weather from 1979 to 2018 – overwintering the Drought Code,” *Earth System Science Data,* vol. 12, no. 3, pp. 1823--1833, 2020.

[19] N. Perez-Zanon, L.-P. Caron, M. d. C. Alvarez-Castro, L. Batté, S. Corti, M. Dominguez, F. Fabiano, S. Gualdi, J. von Hardenberg, L. Lledó, N. Manubens, P. Marson, S. Materia and Torr, “CSTools R package bringing state-of-the-arts postprocessing methods to seasonal-to-decadal forecast users,” EGU General Assembly 2020, 2020.

[20] L.-P. Caron, N. Pérez-Zanón, C. Alvarez-Castro, L. Batté, S. Corti, M. Dominguez, S. Gualdi, J. von Hardenberg, L. Lledó, N. Manubens, P. Marson, S. Materia, E. Sánchez and V. Scha, “CSTools: a new R package for the calibration, combination, downscaling and analysis of seasonal forecasts,” 2019.

[21] V. Lucarini, D. Faranda, . A. C. Moreira Freitas, J. M. Freitas, . T. Kuna, M. Holland, . M. Nicol, . M. Todd and S. Vaienti, Extremes and recurrence in dynamical systems, John Wiley & Sons, 2016.

[22] D. Fernanda, G. Messori, M. C. Alvarez-Castro and P. Yiou, “Dynamical properties and extremes of Northern Hemisphere climate fields over the past 60 years.,” *Nonlinear Processes in Geophysics,* vol. 24, no. 4, pp. 713-725, 2017.

[23] P. De Luca , G. Messori, F. Pons and D. Faranda, “Dynamical Systems Theory Sheds New Light on Compound Climate Extremes in Europe and Eastern North America,” *Quarterly Journal of the Royal Meteorological Society,* vol. 146, p. 1636–1650, 2020.

[24] J. Zelenka, T. Kasanický, E. Gatial, Z. Balogh, A. Majlingová, Y. Brodrechtova, S. Kalinovská, R. Rehák, Y. Semet and G. Boussu, “Coordination of Drones Swarm for Wildfires Monitoring.,” in *Proceedings of the 20th ISCRAM Conference* , Omaha, Nebraska, 2023.

[25] “X Developer Platform,” 2024. [Online]. Available: https://developer.x.com/en/docs/x-api. [Accessed 2024].

[26] “Exploding Topics,” 2024. [Online]. Available: https://explodingtopics.com/blog/top-social-media-platforms. [Accessed 02 09 2024].

[27] “X Developer Platform,” [Online]. Available: https://developer.x.com/en/docs/x-api/tweets/counts/api-reference/get-tweets-counts-recent.

[28] “X Developer Platform,” [Online]. Available: https://developer.x.com/en/docs/x-api/tweets/search/api-reference/get-tweets-search-recent#:~:text=The%20recent%20search%20endpoint%20returns,the%20Project%2Dlevel%20Tweet%20cap.

[29] S. Na, L. Xumin and G. Yong, “Research on k-means Clustering Algorithm: An Improved k-means Clustering Algorithm,” in *Third International Symposium on Intelligent Information Technology and Security Informatics*, Jian, China, 2010.

[30] J. Lesouple, B. Cédric, M. Spigai and J.-Y. Tourneret, “Generalized isolation forest for anomaly detection.,” *Pattern Recognition Letters,* vol. 149, pp. 109-119, 2021.

[31] T. Ergen and S. S. Kozat, “Unsupervised Anomaly Detection With LSTM Neural Networks,” *IEEE Transactions on Neural Networks and Learning Systems,* vol. 31, no. 8, pp. 3127-3141, 2019.

[32] S. Hansun, “A new approach of moving average method in time series analysis,” in *2013 Conference on New Media Studies*, Tangerang, Indonesia, 2013.

[33] G. de Melo and M. Bansal, “Good, Great, Excellent: Global Inference of Semantic Intensities,” *Transactions of the Association for Computational Linguistics,* vol. 1, pp. 279-290, 2013.

[34] M. A. Aina Garí Soler, “MultiSem at SemEval-2020 Task 3: Fine-tuning {BERT} for Lexical meaning,” *CoRR,* vol. abs/2007.12432, 2020.

[35] D. F. Gleich, “PageRank Beyond the Web,” *SIAM Review,* vol. 57, no. 3, pp. 321-363, 2015.

[36] N. R. I. Gurevych, “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks,” *arXiv,* vol. abs/1908.10084, 2019.

[37] G. C. Linderman and S. Steinerberger, “Clustering with t-SNE, Provably,” *SIAM Journal on Mathematics of Data Science,* vol. 1, no. 2, pp. 313-332, 2019.

[38] A. Jacobs, “Sentiment Analysis for Words and Fiction Characters From the Perspective of Computational (Neuro-)Poetics,” *Frontiers in Robotics and AI,* vol. 6, 2019.

[39] S. Brin and P. Lawrence, “The anatomy of a large-scale hypertextual web search engine.,” *Computer networks and ISDN systems,* vol. 39, pp. 107-117, 1998.

[40] J. MacQueen, “A Test for Suboptimal Actions in Markovian Decision Problems,” *INFORMS,* vol. 15, no. 3, pp. 559-561, 1967.

[41] P. J. Rousseeuw, “Silhouettes: A graphical aid to the interpretation and validation of cluster analysis,” *Journal of Computational and Applied Mathematics,* vol. 20, pp. 53-65, 1987.

[42] “FastAPI,” [Online]. Available: https://fastapi.tiangolo.com/. [Accessed 15 09 2024].

[43] “Poetry,” [Online]. Available: https://python-poetry.org/. [Accessed 15 09 2024].

[44] “Swagger,” [Online]. Available: https://swagger.io/. [Accessed 2924 15 09].

[45] K. Jon, “An impossibility theorem for clustering,” *Advances in neural information processing systems,* vol. 15, 2002.

[46] D. Angelov, “Top2vec: Distributed representations of topics.,” *arXiv,* p. arXiv:2008.09470, 2020.

[47] “TopSBM,” [Online]. Available: https://topsbm.github.io/. [Accessed 15 09 2024].

[48] M. Gerlach, . T. d. P. Peixoto and . E. Altmann, “A network approach to topic models,” *Science Advances,* vol. 4, no. 7, p. eaaq1360, 2018.

[49] D. Blei, A. Ng and M. Jordan, “Latent dirichlet allocation.,” *Journal of machine Learning research,* vol. 3, pp. 993-1022, 2003.

[50] D. Blei and J. Lafferty, “Dynamic topic models." Proceedings of the 23rd international conference on Machine learning,” in *Proceedings of the 23rd international conference on Machine learning*, 2006.

[51] M. Grootendorst, “BERTopic: Neural topic modeling with a class-based TF-IDF procedure.,” *arXiv,* p. arXiv:2203.05794, 2022.

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

[53] E. Michail, A. Bozas, D. Stefanopoulos and S. Paspalakis, “Incorporating Social Media Sensing and Computer Vision Technologies to Support Wildfire Monitoring.,” in *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium. IEEE*, 2024.

[54] D. Deng, “Research on anomaly detection method based on DBSCAN clustering algorithm.,” in *5th International Conference on Information Science, Computer Technology and Transportation (ISCTT). IEEE*, 2020.

[55] “BIA5’s latest ATR platform integration,” [Online]. Available: https://bia5.com/bia5s-latest-atr-platform-integration/. [Accessed 25 09 2024].

[56] F. A. Ruetz, N. Lawrance, E. Hernández, P. V. K. Borges and T. Peynot, “ForestTrav: 3D LiDAR-Only Forest Traversability Estimation for Autonomous Ground Vehicles,” *IEEE Access,* vol. 12, pp. 37192-37206, 2024.

[57] M. Ramezani, E. Griffiths, M. Haghighat, A. Pitt and P. Moghadam, “Deep Robust Multi-Robot Re-localisation in Natural Environments,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.