


TRUSTED
EXTREMELY PRECISE
MAPPING AND PREDICTION
FOR EMERGENCY
MANAGEMENT

Deliverable D3.1: First report on algorithms for extreme data analytics







Project Information

Project acronym:	TEMA
Project full title:	Trusted Extremely Precise Mapping and Prediction for Emergency Management
Call identifier:	HORIZON-CL4-2022-DATA-01
Type of action:	HORIZON Research and Innovation Actions
Start date:	1 December 2022
End date:	30 November 2026
Grant agreement no:	101093003

D3.1 – First report on algorithms for extreme data analytics			
Executive Summary:	<p>Deliverable D3.1 “First report on algorithms for extreme data analytics,” is the first Deliverable of Work-package 3 (WP3) within the TEMA project. This document encapsulates the research achievements of Tasks T3.1, T3.2, and T3.3 over the first 18 months (M1-M18) of the project. See the Section “Executive Summary”.</p>		
WP:	3		
Author(s):	See table below for a full list of authors		
Editor:	Ioannis Pitas, Nikolaos Marios Militsis (AUTH)		
Leading Partner:	Aristotle University of Thessaloniki (AUTH)		
Participating Partners:	All		
Version:		Status:	Final
Deliverable Type:	R — Document, report	Dissemination Level:	PU - Public
Official Submission Date:	31 May 2024	Actual Submission Date:	3 June 2024







Disclaimer

This document contains material, which is the copyright of certain TEMA contractors, and may not be reproduced or copied without permission. All TEMA consortium partners have agreed to the full publication of this document if not declared “Confidential”. The commercial use of any information contained in this document may require a license from the proprietor of that information. The reproduction of this document or of parts of it requires an agreement with the proprietor of that information.

The TEMA consortium consists of the following partners:

No.	Partner Organisation Name	Partner Organisation Short Name	Country
1	ARISTOTELIO PANEPISTIMIO THESSALONIKIS	AUTH	GR
2	DEUTSCHES ZENTRUM FÜR LUFT – UND RAUMFAHRT EV	DLR	DE
3	ENGINEERING - INGEGNERIA INFORMATICA SPA	ENG	IT
4	ATOS IT SOLUTIONS AND SERVICES IBERIA SL	ATOS IT	ES
4.1	ATOS SPAIN SA	ATOS SP	ES
5	UNIVERSIDAD DE SEVILLA	USE	ES
6	TECNOSYLVA SL	TSYL	ES
7	NORTHDOCKS GMBH	ND	DE
8	PARIS-LODRON-UNIVERSITÄT SALZBURG	PLUS	AT
9	THE LISBON COUNCIL FOR ECONOMIC COMPETITIVENESS ASBL	LC	BE
10	LATITUDO 40 SRL	LAT40	IT
11	NELEN & SCHUURMANS TECHNOLOGY BV	NS	NL
12	FRAUNHOFER GESELLSCHAFT ZUR FORDERUNG DER ANGEWANDTEN FORSCHUNG EV	FHHI	DE
13	UNIVERSITÀ DEGLI STUDI DI MESSINA	UNIME	IT
14	KAJAANIN AMMATTIKORKEAKOULU OY	KAMK	FI
15	KAJAANIN KAUPUNKI	KAJ	FI
16	KENTRO MELETON ASFALIAS	KEMEA	GR
17	DIMOS MANTOUDIYOU - LIMNIS - AGIAS ANNAS	D.MALLIAN	GR
18	REGIONE AUTONOMA DELLA SARDEGNA*RAS	RAS	IT
19	BAYERISCHES ROTES KREUZ	BRK	DE







Document Revision History

Version	Description	Contributions
0.1	ToC	AUTH
0.2	First complete draft	AUTH, ATOS, DLR-DFD, DLR-KN, FHHI, PLUS, UNIME
0.3	First complete draft	AUTH
1.0	Internal review	USE
1.1	Final version	AUTH
1.2	Submission to EC	AUTH
1.3	Length reduction and re-submission to EC	AUTH

Authors/Contributors

Name	Partner
Ioannis Pitas	AUTH
Nikolaos Marios Militsis	AUTH
Dimitrios Papaioannou	AUTH
Michael Siavrakas	AUTH
Matthaios D. Tzimas	AUTH
Ioanna Koroni	AUTH
Filippos Kitsos	AUTH
Anestis Kaimakamidis	AUTH
Eleonor Fragachan	ATOS
Ricard Munne Caldes	ATOS
Monika Friedemann	DLR-DFD
Marc Wieland	DLR-DFD
Michael Nolde	DLR-DFD
Martin Mühlbauer	DLR-DFD
Dmitriy Shutin	DLR-KN
Victor Prieto Ruiz	DLR-KN
Ekkehard Schnoor	FHHI
Jawher Said	FHHI
Joost van Dijk	NS
Nicolette Volp	NS
Joep Grispén	NS






Name	Partner
Bernd Resch	PLUS
David Hanny	PLUS
Sebastian Schmidt	PLUS
Merve Keskin	PLUS
Lorenzo Carnevale	UNIME
Massimo Villari	UNIME
Antonio Celesti	UNIME
Maria Fazio	UNIME
Antonino Galletta	UNIME
Armando Ruggeri	UNIME

Reviewers

Name	Organisation
J. Ramiro Martínez de Dios	USE









Table of Contents


Table of Contents.....	1-5
List of Figures.....	1-7
List of Tables.....	1-9
List of Terms and Abbreviations	1-10
Executive Summary	1-11
1 Introduction.....	1-12
1.1 Purpose and scope of the document.....	1-12
1.2 Structure of the document.....	1-12
2 Summary of the work carried out.....	2-13
2.1 Objectives.....	2-13
2.2 Summary of the work carried out with respect to the objectives.....	2-14
3 Explainable and robust analytics.....	3-21
3.1 Introduction.....	3-21
3.2 Generic XAI.....	3-23
3.3 XAI on Diffusion Models image generation.....	3-26
3.4 AI robustness	3-28
3.4.1 Decentralized Inference for Forest Fire Classification.....	3-32
4 Real-time semantic visual analysis and remote sensing.....	4-33
4.1 Introduction.....	4-33
4.2 AI algorithms for visual data analysis	4-34
4.2.1 Fire detection	4-35
4.2.2 Fire region segmentation	4-39
4.2.3 Forest Fire Classification.....	4-41
4.2.4 Burnt forest region segmentation.....	4-42
4.2.5 Flood region segmentation	4-44
4.2.6 Person and car detection in flooded areas.....	4-45





4.2.7	Privacy protection.....	4-47
4.2.8	Transformer architecture acceleration.....	4-49
4.2.9	Synthetic data generation.....	4-52
4.2.10	Person re-identification	4-53
4.3	Methods for satellite/SAR data analysis	4-54
4.3.1	Satellite-based flood detection and assessment.....	4-54
4.3.2	Satellite-based fire detection and assessment	4-58
4.4	Multimodal analysis for the construction of 3D smoke concentration maps.....	4-59
5	Social media and text semantic analysis.....	5-61
5.1	Introduction.....	5-61
5.2	Semantic topic modelling.....	5-61
5.2.1	Multilinguality handling	5-62
5.2.2	Post relevance classification/assessment	5-63
5.3	Sentiment analysis for short texts.....	5-65
5.3.1	AI-based text sentiment analysis	5-65
5.3.2	Consensus-based labelling	5-68
5.3.3	AI-based text emotion analysis.....	5-69
5.4	Spatial hot spot analysis.....	5-70
5.5	Contrastive image-language models	5-71
6	TEMA Core: Parallel and Distributed System.....	6-73
6.1	Kubernetes Cluster.....	6-74
6.2	Federation.....	6-74
6.3	Networking.....	6-75
6.4	VM for Easy Cluster Joining	6-75
6.5	Deployment of Containerized Applications in a Kubernetes Cluster.....	6-76
6.6	Management Kubernetes Cluster with Rancher.....	6-76
6.7	Deployment and Integration of Containerised Application on a K3s Cluster	6-78
7	Conclusion	7-79
8	References	8-80





A. Appendix: Related publications and technical reports..... 8-94

List of Figures

Figure 1: (a) Speed of local XAI explaining individual predictions. (b) Speed of global XAI explaining global model strategies for EfficientNet_b0..... 3-22

Figure 2: Prototype-based explanations for a fire-classification model based on PCX: two prominent concepts of “smoke” and “fire” are most crucial for the class prediction. The appearance of similar concepts suggests a high degree of redundancy within the trained neural network (PURE)..... 3-24

Figure 3: Forest Fire synthetic image generated by ATOS through diffusion models, and the corresponding DAAM heatmap for the prompt “fire”..... 3-28

Figure 4: Examples of (a) forest fire image, (b) forest fire image with FGSM adversarial attack and (c) forest fire image with FGSM+ADAR adversarial attacks. 3-31

Figure 5: Fire detection inference of RT-DETR trained with AUTHs size-balanced $L1$ loss. 4-36

Figure 6: Comparison between AUTH loss and $L1$. Specifically, this is an example in which AUTH has two DNN prediction-target sets that have the same IoU and one set is the scaled version of the other one. The method defines a reference set with error equal to $L1_{ref}$ and gradually increases the height and width of the scaled box by the factor s . The scaled box has loss $L1_{scaled}$ (ours), is computed with respect to the scale factor s and ensure smaller contribution of larger object to the DNN loss function. 4-37

Figure 7: Examples of segmented forest fires using the I2I-CNN [PAP2021]. 4-39

Figure 8: Example of burnt forest region segmentation on sample from the burnt region segmentation partition of BLAZE dataset using the CNN-I2I. (a) raw image, (b) burnt area segmentation mask..... 4-42


Figure 9: Examples of flood segmentation using the PSPnet. [ZHA2017]..... 4-44

Figure 10: Examples of (a) person and (b) car detection in flooded areas..... 4-46

Figure 11: Example of face detection and blurring by AUTH..... 4-48

Figure 12: (a) Raw image sample from celebA dataset containing human face. (b) Image sample where the human face is totally hidden based on PAR, yet recognizable by a chosen face recognizer..... 4-49

Figure 13: The proposed SESA module. By enforcing a circulant structure onto the projection matrices of query, key, and value subspaces, it significantly reduces the spatial complexity from $O(D^2)$ to $O(D)$ for each respective learnable matrices Wq , Wk , and Wv . By leveraging the Discrete Fourier Transform (DFT) matrix and Fast Fourier Transform (FFT), AUTH decreases the theoretical computational complexity from $O(D^2)$ to $O(D \log(D))$ 4-51






Figure 14: Forest Fire (a) and Flood (b) Images from a drone POV generated by ATOS through diffusion models. 4-53

Figure 15: Reference dataset used for training, validation and test of satellite-based flood detection and assessment. 4-55

Figure 16: Comparison of ONNX model inference throughput for different model precisions. Measured on a NVIDIA RTX A4000 GPU using ONNX-Runtime "CUDAExecutionProvider" across a range of batch sizes..... 4-56

Figure 17: Results from satellite-based monitoring of the Greece floods 2023 with examples showing inundated areas on 2023-09-09 (A) and flood duration from 2023-09-01 until 2023-10-04 (B) around the city of Larissa..... 4-57

Figure 18: Sample from the generated Donnie Creek wildfire dataset (derived from Sentinel-3). The burn severity attribute is used for visualization..... 4-59

Figure 19: Data collected in a wind tunnel to validate trace gas dispersion models under simplified uniform flow. Fig. a) represents the measured smoke plume after reaching equilibrium, and b) the model [HIN2024]..... 4-60

Figure 20: Illustration of the proposed method, AnchorBERT. Given an input sentence, a base encoder is leveraged to extract the contextualized word representation. Next, a class anchor is used for each emotion in order to enrich the text’s representation. To this end a multi-headed attention module is employed. Finally, a feed forward neural network projects the enriched representation into the binary task at hand. AUTH illustrates its model both during training (left) and inference (right). 5-67


Figure 21: The TEMA Core is a complex infrastructure that includes many actors (e.g., admin-users, end-users, technology owners, third-parties technology owners). The main components are the Big Data Storage and the Parallel Processing units. The infrastructure is designed to allow a federated massive computation of natural disasters data. Moreover, it meets the need of low latency thanks to the use of the Cloud-Edge Continuum paradigm, where the computation may migrate from the cloud to edge resources and vice versa. 6-73


Figure 23: A virtual private network (VPN) enables easy and secure access to the Kubernetes cluster. Nodes connect to the VPN and receive an IP address with which each node is visible within the private network..... 6-75

Figure 25: Deployment process in Kubernetes: The deployer reads the deployment manifest and applies the configuration to the Kubernetes cluster. Once confirmed, Kubernetes starts the application pods on selected worker nodes. 6-76

Figure 26: Rancher Cluster Dashboard showing an overview of the cluster's status and capacity. Key metrics include the total resources, number of nodes, and deployments. The capacity section details the usage of pods, CPU, and memory. 6-77

Figure 27: Rancher Cluster Nodes view displaying all nodes in the cluster. The dashboard shows whether each node is a master or worker, and provides details on the percentage usage of CPU, RAM, and pods, as well as the operating system running on each node. 6-78





List of Tables

Table 1: Evaluating of robustification methods at white box attack settings on the fire classification partition of BLAZE dataset. 3-30

Table 2: Robustness Against Adversarial Attacks on MNIST dataset [DEN2012]. 3-31

Table 3: Forest fire classification accuracy on BLAZE dataset (fire classification partition) achieved by one agent assuming that all agents are honest. 3-33

Table 4: Fire detection evaluation results on the Forest-Fire dataset. 4-38

Table 5: Evaluation of state-of-the-art region segmentation architectures on the Flame dataset [SHA2020]. 4-40

Table 6: Classification accuracy and speed of seven DNN classifiers for fire classification on the fire classification partition of BLAZE dataset. 4-41

Table 7: Evaluation of multiple real-time semantic segmentation architectures on the burnt region segmentation partition BLAZE dataset. 4-43

Table 8: YOLOv6-S and YOLOv6-L performance and efficiency on Nvidia GTX 1080. 4-47

Table 9: Accuracy of several CNN architectures on MNIST dataset where parts of the images containing sensitive information such as human faces or car plates are totally hidden. ResNet trained with PAR maintains its performance even though the images are partly occluded. 4-49

Table 10: Comparison of state-of-the-art real-time object detectors on COCO val 2017 [LIN2014]. 4-51


Table 11: Performance comparison of top-performing DNN models for semantic segmentation of water bodies from Sentinel-1 (radar) and Sentinel-2 (multi-spectral) satellite images. 4-54

Table 12: Performance metrics for 60 clusters computed on the TweetEval [BAR2020] dataset averaged over a total of 10 runs when UMAP was used for dimensionality reduction for the JTS variants. The sentiment statistics marked with * were computed independently from the topic model and k denotes the number of keywords used for the computation of TC and TD. The best scores for each metric are highlighted in bold [HAN2024] 5-63

Table 13: Comparison of results for different relevance classification models (P = precision, R = recall, F1 = F1 score, GS = Gaussian Score; NB: Naive Bayes, RF: Random Forest, SVM = Support Vector Machine, CNN = COnvolutional Neural Network, BERT = Bidirectional Encoder Representations from Transformers) [BLO2024]. The best scores for each metric are highlighted in bold. 5-64

Table 14: F1 scores for sentiment classification tasks in different languages based on the work of Barbieri et al [BAR2022]. The highest value for each language is highlighted in bold. 5-66

Table 15: F1 scores for sentiment detection in CovidEMO. The highest value for each sentiment is highlighted in bold. 5-67






Table 16: F1 scores for sentiment detection in HurricaneEMO. The highest value for each sentiment is highlighted in bold..... 5-67


Table 17: F1 scores for sentiment detection in CancerEMO. The highest value for each sentiment is highlighted in bold..... 5-68


Table 18: F1 scores for sentiment detection in NLPCC 2018. The highest value for each sentiment is highlighted in bold..... 5-68

Table 19: Accuracy scores for sentiment detection in AUTH’s Mastodon dataset. The highest value for each sentiment is highlighted in bold. 5-68

List of Terms and Abbreviations

Abbreviation	Description
TEMA	Trusted extremely precise mapping and prediction for emergency management
AI	Artificial Intelligence
XAI	Explainable Artificial Intelligence
NDM	Natural Disaster Management
WP	Work Package
DNN	Deep Neural Network
CNN	Convolutional Neural Network
Roi	Region of Interest
IoU	Intersection over Union






Executive Summary


Deliverable D3.1 “First report on algorithms for extreme data analytics,” is the first Deliverable of Work-package 3 (WP3) within the TEMA project. This document encapsulates the research achievements of Tasks T3.1, T3.2, and T3.3 over the first 18 months (M1-M18) of the project. These tasks focus on addressing the challenges in extreme data analytics for NDM by leveraging heterogeneous data sources including edge devices and sensors (e.g., drones, wind sensors, stream flow gauges etc.), satellite images, geospatial data, meteorological data, and geosocial media. These data sources are heterogeneous, voluminous, frequently updated, complex, multilingual, dispersed, sparse, and extreme in nature. Towards this end, novel real-time semantic visual analysis and remote sensing techniques have been developed to accurately detect fires, floods and humans in danger. These techniques leverage advanced deep learning methods, namely deep image classification, deep object detection and semantic region segmentation using Deep Neural Networks (DNNs) applied to both RGB and satellite images. Furthermore, the trustworthiness and robustness of these methods are enhanced through the application of eXplainable Artificial Intelligence (XAI) techniques and AI robustness mechanisms. Additionally, AI algorithms for semantic analysis of social media and textual content have been developed, ensuring comprehensive and reliable data analysis across various platforms such as Twitter (X).

The research under Task T3.1 has been directed towards developing novel trustworthy AI algorithms for DNNs, addressing TEMA objective **OA1**, which aims to enhance the trustworthiness of extreme data analysis algorithms. Task T3.2 has concentrated on Deep Learning Architectures for semantic visual analysis, particularly in RGB images captured by drones as well as in Satellite and Synthetic Aperture Radar (SAR) data processing, thereby contributing to TEMA objectives **OA2** and **OA3** to improve the accuracy and responsiveness of data analysis algorithms. Task T3.3 has explored Deep Learning methods for analysing geosocial media, news posts, and textual content, also supporting objective **OA2**.

Key achievements of this research period include the publication of 14 papers (4 in journals and 10 in conference proceedings) and the submission of 18 technical reports to various academic conferences and journals for publication. Additionally, the development of 9 TEMA technological components has been instrumental in enhancing the functionalities of the TEMA platform, as outlined in Deliverable D2.2. A Federated Cloud-Edge Platform for executing extreme data analytics algorithms was also developed, beginning with the design and implementation of the TEMA Core, the central unit that will manage and coordinate all TEMA technological components.

The techniques devised in Task T3.1 will be further refined and integrated into diverse AI algorithms across other TEMA tasks, with the explainable AI outputs specifically feeding into Task T5.3 focused on augmented reality and rapid visualization. Outputs from Task T3.2 are set to be enhanced and primarily utilized in Task T4.3, which deals with information fusion. Similarly, the





methodologies from Task T3.3 will undergo further refinement and be applied in both Task T4.3 and Task T4.4, the latter concentrating on data-fusion-based decision support and process triggering.

This deliverable not only summarizes the main research outputs and technical challenges associated with designing innovative algorithms for extreme data analytics, but can also serve as a crucial reference for big data analytics and Natural Disaster Management (NDM) researchers. As a public document, it plays a vital role in disseminating TEMA findings to the broader scientific community.

1 Introduction


1.1 Purpose and scope of the document


Deliverable D3.1 “First report on algorithms for extreme data analytics” is the first Deliverable of the third Work-package (WP3) of the TEMA project. The main purpose of this document is to report the initial research results of Tasks T3.1 “Explainable and robust analytics”, T3.2 “Real-time semantic visual analysis and remote sensing” and T3.3 “Social media and text semantic analysis” between M1-M18. More specifically the research effort was focused on the following areas:

- Novel trustworthy AI algorithms for DNNs.
- Deep learning architectures for semantic visual analysis.
- Novel methods for satellite/SAR data analysis and remote sensing.
- Deep learning architectures for semantic analysis of geosocial media/news posts and textual content.
- Distributed semantic analysis across the edge-to-cloud continuum.

1.2 Structure of the document

The remainder of the document is structured as follows. Section 2 summarizes the main research efforts and key outputs, with respect to TEMA objectives. Section 3 describes the development of novel trustworthy AI algorithms for DNNs. Section 4 analytically describes the development novel semantic visual analysis and remote sensing AI algorithms for heterogeneous data modalities. Section 5 describes the development of novel semantic analysis algorithms for geosocial media/news posts and textual content. Section 6 describes in detail the TEMA Core, hosting and coordinating the TEMA technological components. Finally, conclusions are drawn in Section 7.






2 Summary of the work carried out


2.1 Objectives

TEMA envisions addressing the challenges in extreme data analytics for NDM like regional floods, flash floods, and wildfires by leveraging heterogeneous data sources including edge devices and sensors (e.g., drones, wind sensors, stream flow gauges etc.), satellite images, geospatial data, meteorological data, and geosocial media. These data sources are heterogeneous, voluminous, frequently updated, complex, multilingual, dispersed, sparse, and extreme in nature. The main objective of TEMA WP3 "Trustworthy Federated Analytics," is to develop novel methods for accurate and fast/real-time semantic data analytics using inputs from multiple modalities. This includes creating algorithms for trustworthy AI, federated analytics, and DNN data scarcity mitigation. This will involve distributed semantic analysis across the edge-to-cloud continuum to minimize latency, using AI and DNNs for trustworthy and flexible heterogeneous data analytics. Fast computations are crucial for decision-making in emergencies, supported by the vast cloud computing resources, spanning the entire edge-to-cloud continuum to meet the demands of extreme data analytics.

The specific TEMA objectives linked to Tasks T3.1, T3.2, and T3.3 are derived from existing challenges in extreme data analytics and from the socially relevant use-cases, i.e., regional floods, flash floods, and forest fires/wildfires. They are presented below, along with accompanying Key Performance Indicators (KPIs) and Target Values (TVs) as defined in Section 1.1.1 of Part B of TEMA's Description of the Action (DoA). Task T3.1 contributes to TEMA objective **OA1** "Increase trustworthiness of extreme data analysis algorithms". Task T3.2 contributes to TEMA objectives **OA2** "Increase accuracy of extreme data analysis algorithms" and **OA3** "Increase responsiveness/speed of extreme data analysis algorithms". Task T3.3 contributes to TEMA objectives **OA2** "Increase accuracy of extreme data analysis algorithms".

Section 2.2 summarizes the Research and Development (R&D) activities conducted under Tasks T3.1, T3.2, and T3.3 by TEMA partners, aligned with TEMA objectives in their natural order. The section begins with a concise description of R&D on novel semantic visual analysis and remote sensing AI algorithms tailored for heterogeneous data modalities. Following this, it describes the progress achieved in making these algorithms trustworthy by incorporating eXplainable AI (XAI) techniques and developing robust defenses against attacks, such as deliberate arson. The discussion then shifts to the development of cutting-edge semantic analysis algorithms for geosocial media, news posts, and textual content. Lastly, the section highlights the design and implementation of the TEMA Core, a central unit devised to host and orchestrate all TEMA technologies, which are primarily developed within the WP3, WP4, and WP5 work packages.





2.2 Summary of the work carried out with respect to the objectives

AI algorithms for visual data analysis

Fire detection


AUTH has introduced an innovative training methodology tailored for RT-DETR, a real-time detection transformer optimized specifically for fire detection scenarios. AUTH approach addresses the challenge of imbalanced ground-truth object sizes by implementing a size-based weighting mechanism for the standard $L1$ loss used in training object detection algorithms. This modification significantly improves DNN training efficiency and enhances fire detection performance, particularly for datasets with objects fires enclosed in varying size bounding boxes. This involves the following improvements with regards to the fulfilment of KPIs:


- **Object Detection Accuracy:** With this AUTH new training methodology, RT-DETR achieves a remarkable mean Average Precision (mAP) increase of **+6.2%** compared to state-of-the-art approaches. This advancement exceeds the KPI value for "**Object Detection Accuracy**" under objective **OA2** "Increase Accuracy of Extreme Data Analysis Algorithms".
- **Speed and Responsiveness:** RT-DETR maintains high-speed visual analysis capabilities, being ideal for real-time fire detection applications. Thus, it satisfies TEMA objective **OA3** "Increase responsiveness/speed of extreme data analysis algorithms" and much exceeds the related the "**Visual analysis speed**" KPI. Specifically, RD-DETR processes input images at **108 FPS** (Frame per Second) on a T4 GPU with TensorRT FP16, which is much faster than real-time video analysis (25-30 FPS).

ATOS created a synthetic forest fire dataset through the application of diffusion models in order to face the data scarcity regarding natural disasters. This development complies with the **OA2** "Increase Accuracy of Extreme Data Analysis Algorithms" objective by increasing the amount and quality of training data for the fire detector, which includes varied scenarios that are hardly to come by on any existing dataset (different weather, seasons, and times of day).

Person and car detection in flooded areas

AUTH has utilized and optimized YOLOv6 models specifically for the critical task of identifying and localizing humans and vehicles in flooded regions. These models are designed to efficiently detect and localize objects in high-resolution images, ensuring accurate identification of humans being in danger. AUTH YOLOv6 object detection models were trained in novel flood image datasets. They satisfy the objective **OA3** "Increase Responsiveness/Speed of Extreme Data





Analysis Algorithms" by delivering rapid and accurate person/car detection in flooded regions at **33 FPS** on a GeForce GTX 1080 GPU which exceeds the "**Visual analysis speed**" KPI.

Flood region segmentation


AUTH introduced a novel flood region segmentation dataset, containing annotated images sourced from diverse datasets. It serves as a benchmark for evaluating state-of-the-art flood region segmentation DNN models. Leveraging the ST++ self-training method with the PSPnet architecture, the novel AUTH flood region segmentation methods enhance the region segmentation accuracy and efficiency. This method utilizes pseudo-labeling for unlabeled flood images, aligning with the TEMA objective of increasing the accuracy of extreme data analysis algorithms. The dataset enables comprehensive DNN model training and evaluation, while the AUTH methodology demonstrates promising improvements towards reaching the target value of 5% for the "**Semantic segmentation accuracy**" KPI, achieving a **3.5%** mean Intersection over Union (mIoU) increase over existing benchmarks. It also satisfies the objective **OA2** "Increase accuracy of extreme data analysis algorithms". Additionally, AUTH solution processes images at a speed of **79 FPS** on a GeForce GTX 1080 GPU, which is much faster than "**Visual analysis speed**" KPI of objective **OA3** "Increase responsiveness/speed of extreme data analysis algorithms".


Burnt forest region segmentation

AUTH created a novel annotated burnt forest region segmentation dataset through gathering RGB images from public drone videos and annotating them meticulously to support DNN region segmentation systems. Through rigorous region segmentation performance evaluation, AUTH identified CNN-I2I region segmentation architecture as the top performer, boasting a real-time inference speed of **145 Frames Per Second (FPS)**, which much exceeds the "**Visual analysis speed**" KPI of **OA3** "Increase responsiveness/speed of extreme data analysis algorithms". This method also delivered an **8.56%** mIoU improvement over the state-of-the-art approach, surpassing the "**Semantic segmentation accuracy**" KPI of objective **OA2** "Increase accuracy of extreme data analysis algorithms". These findings highlight the effectiveness of using the AUTH dataset in fine-tuning the CNN-I2I architecture for real-time burnt forest region segmentation.

Visual privacy preservation

GDPR provisions for personal privacy protection do apply to visual data collected during natural disasters. This is an issue that is often overlooked by NDM authorities. AUTH introduced a novel method called Privacy via Adversarial Reprogramming (PAR), which utilizes Adversarial






Reprogramming strategies to tackle the vital challenge of maintaining visual privacy in DNN models. PAR works by reprogramming the target DNN model (e.g., a face recognizer) to retain the DNN model functionality while hiding sensitive information (e.g., the face region) in the source image. Through careful adjustment of the DNN model parameters, PAR achieves remarkable outcomes, maintaining the accuracy of the model at an impressive **97%**. This degree of privacy preservation ensures that the DNN model remains useful for its intended purposes while preventing unauthorized access to sensitive personal visual data (facial images). Even though the areas of RGB images that contain sensitive information, such as human faces or car plates are completely hidden, the proposed PAR method does not contradict with the KPI "**Image Recognition Accuracy**" of TEMA objective "**OA2** "Increase accuracy of extreme data analysis algorithms", as the DNN model performance is maintained.


Satellite data analysis

Satellite-based flood detection and assessment

DLR uses a modular processing system for surface water monitoring, automating satellite data search, pre-processing, analysis, and dissemination over specified areas of interest. They employ Sentinel-1 (radar) and Sentinel-2 (multi-spectral) satellite images, processed with pre-trained Deep Neural Network (DNN) models for semantic segmentation to identify water areas. Time-series analysis is utilized to derive additional information about permanent water bodies and flood duration. TEMA extends this approach by DLR, developing and testing various DNN models specifically for water segmentation.

This technology complies with the TEMA objectives **OA2** "Increase accuracy of extreme data analysis algorithms" and **OA3** "Increase responsiveness/speed of extreme data analysis algorithms". The initial version of satellite-based flood detection and assessment technology has notably enhanced **flood segmentation accuracy (OA2)** and **visual analysis speed (OA3)**. Compared to previous DLR-developed Convolutional Neural Networks (CNN) models for flood mapping, TEMA latest advances show a **10%** increase in segmentation accuracy (IoU) for Sentinel-1 (radar) and a **4%** increase for Sentinel-2 (multi-spectral) image analysis. In addition, the inference throughput has increased by **10 MP/s** compared to earlier models. This technology can analyze data in less than 10 seconds on a consumer-grade GPU (excluding data download, pre-processing, and result publishing), which is significantly faster than manual or semi-automated methods used by the Copernicus EMS.






Satellite-based fire detection and assessment


DLR operates a flexible processing chain for automatically deriving burn areas from optical satellite images, compatible with various optical sensors that include near-infrared and red bands. This setup reduces the time between overpasses and enables near-real time result distribution. The system primarily uses mid-resolution data, updating results four times daily (from Sentinel-3 A/B and Aqua/Terra MODIS sensors). This development aligns with TEMA's objectives **OA2** "Increase accuracy of extreme data analysis algorithms" and **OA3** "Increase responsiveness/speed of extreme data analysis algorithms". Compared to other systems like Copernicus EMS, JRC EFFIS, and NASA MCD64A1, the **spatial accuracy and completeness of burnt area detection** show a significant improvement of around **23%** with respect to the Intersection over Union (IoU) performance measure.

Explainable Artificial Intelligence (XAI)

FHHI contributes to the aspect of explainability for AI, in particular for sensitive applications, such as natural disaster management. Among the results detailed in **Section 3.2**, the following key contributions are particularly noteworthy. To automate XAI and reduce the amount of human labour (required for interpreting and validating the explanations), FHHI has developed new explainability methods based on prototypes, building up on FHHI's previous work on concept-based explanations. Moving beyond traditional attribution maps, prototype-based methods model the distributions of latent relevancies as Gaussian Mixtures Models, each for the different classes involved. This allows to define prototypes (samples closest to the means) of (sub-)classes (global XAI), enabling comparison with individual samples (local XAI) with applications for outlier detection. FHHI has recently begun exploring new research avenues, in close collaboration with AUTH and DLR. Together, they are working on developing models and creating explainability methods for new architectures, such as YOLO-v6 object detector. The prototype-based methods developed by FHHI are compliant with the "**Speed of global XAI explaining global strategies**" KPI TEMA objective **OA1** "Increase trustworthiness of extreme data analysis algorithms", as the time required to compute global explanations does not exceed the time consumed to compute all local explanations by more than an order of magnitude. FHHI's prototype-methods demonstrate promising progress towards reaching the target value of "**Speed of local XAI explaining individual predictions**" KPI as they currently exceed the time consumed for the prediction itself by a factor of 6. FHHI's research on explainability may not only help for providing XAI-enhanced tools for the end-user, but also to help in improving the models under development.

ATOS has studied the explainability of diffusion models applied on Task 3.5. These explanation methods allow users to understand which pixels are being affected during the diffusion process for a specific word or concept on the input prompt. For every concept, a heatmap is





generated explaining which regions of the image correspond to that concept. With these explanations users can understand what the diffusion model is conveying while also getting a segmentation from every concept, which can be used for image labelling.

AI robustness


DNN robustness against adversarial attacks


AUTH has conducted extensive research to advance in DNN robustness methodologies. A significant proposal involves transitioning from the widely used One-versus-All (OvA) classification scheme to the innovative One-versus-One (OvO) approach. This novel AUTH method integrates the OvO DNN classification strategy with geometrically inspired DNN training criteria. This adaptation enhances the robustness and accuracy of DNN classification models, showcased in the fire and smoke recognition task under adversarial attacks, and align with TEMA objective **OA2** "Increase accuracy of extreme data analysis algorithms", with a specific focus on enhancing KPI "**Image Recognition Accuracy**". Notably, the proposed AUTH method significantly improves the **Correct Classification Rate (CCR)** under the adversarial attack known as Basic Iterative Method (BIM), showcasing a significant increase of **13.21%**.

Moreover, AUTH has introduced the novel adversarial defense strategy known as "Adversarial Defence via Adversarial Reprogramming." This defence method is designed to fortify DNN systems against adversarial attacks, that subtly manipulate input data to deceive DNNs, thereby bolstering the trustworthiness and reliability in safety critical scenarios. These advances underscore AUTH commitment to improve AI system robustness and performance, particularly in critical NDM applications, where AI system accuracy and resilience are of paramount importance.

Decentralised DNN inference for forest fire classification

In natural disaster emergencies, centralised DNN decision making systems can be compromised by attacks, e.g., in the case of deliberate arson. Decentralized inference addresses these challenges by distributing data processing across multiple nodes on the edge, ensuring faster, more reliable analysis, and maintaining operational continuity even when parts of the DNN network are compromised through attacks. AUTH has developed a decentralized DNN inference framework that operates in two distinct stages. First, through the Individualized Agent Model Selection Process (IAMSP), each autonomous DNN agent engages in a DNN model selection technique by consulting neighbouring DNN agents. This process allows each DNN agent, particularly the poor-performing DNNs, to optimize their recognition performance based on collective insights. Notably, AUTH framework delivers significant performance gains. For instance, when applying this framework to the EfficientNet B4 architecture for the fire recognition task, its






accuracy significantly improves by **16.69%** compared to its stand-alone operation. This achievement surpasses the targeted “Image recognition accuracy” KPI of objective OA2 “Increase accuracy of extreme data analysis algorithms”. Secondly, AUTH introduces the Quality of Inference (Qol), a pioneering Byzantine Fault-Tolerant (BFT) consensus protocol. Qol serves as a universal inference rule accepted by all participating DNN agents in the system, ensuring robust and consistent decision-making, in the presence of attacks, e.g., malicious or compromised DNN agents. While Qol ensures the robustness of decentralized DNN inference, other protocols, e.g., ones based on majority voting completely fail to counter such attacks. Both novel AUTH approaches showcase the ability to ensure NDM decision making robustness through decentralized inference systems.


Social media and text semantic analysis

Natural disasters create a user response in social media, e.g., Twitter (X) text posts. PLUS developed a framework for the simultaneous analysis of semantics, sentiment, space, and time in social media posts, resulting in two journal articles (one under review) and a conference article on disaster management applications. This approach uniquely identifies local semantic topics relevant to disaster contexts, incorporating sentiment, spatial, and temporal information into multimodal clusters. Compared to existing joint topic-sentiment models, the framework significantly improves semantic topic quality (measured by Normalised Positive Pointwise Mutual Information) and sentiment classification accuracy, aligning with objective **OA2** “Increase accuracy of extreme data analysis algorithms”, specifically **enhancing sentiment analysis accuracy by more than 10% over the SOTA and achieving more than double the topic quality score**. Large Language Models (LLMs) like Llama-2 were used to generate readable topic labels and identify information critical to emergency responders. Different variants of the framework were explored, visualising the trade-off between runtime and empirical performance. This supports **OA3** “Increase responsiveness/speed of extreme data analysis algorithms”, targeting **real-time social media analysis**.

PLUS also advanced active learning methods for identifying disaster-related social media posts, outperforming keyword filtering and other approaches on datasets from CrisisLex, the 2021 Ahr Valley, and the 2023 Chile wildfires, further supporting **OA2**. This work will be detailed in an upcoming conference article and a published journal article on relevance classification of flood-related Tweets.

Concurrently, ATOS IT developed software for generating image captions by comparing state-of-the-art Image-to-Text models, including OpenFlamingo, Blip2, Monkey, and LLava, using Twitter images from the Ahrthal floods provided by PLUS. This development supports both **OA2** and **OA3**. ATOS SP assisted in these efforts.






These advancements by PLUS and ATOS significantly enhance both the **accuracy (OA2)** and **speed (OA3)** of extreme data analysis algorithms in social media and disaster management contexts.


Text Sentiment analysis

Towards text sentiment analysis, AUTH proposed an innovative approach reintroducing sentiment interrelation, while simultaneously addressing multiple independent binary sentiment classification problems. By leveraging sentiment representatives defined as anchors within the feature space of each binary classifier, this method simplifies the sentiment representations based on their similarity to these anchors. This strategy effectively combines the strengths of different methodologies, resulting in significant advances. Furthermore, as TEMA should perform multilingual text sentiment analysis, AUTH created an annotated dataset of MASTODON posts to run post sentiment analysis. AUTH text sentiment progress is particularly notable on this dataset it achieved a **9.3%** increase in **sentiment analysis accuracy for social media posts**, compared to existing state-of-the-art methods. This accomplishment is very close to the KPI of TEMA objective **OA2** "Increase accuracy of extreme data analysis algorithms".

Federated Cloud-Edge Platform

The UNIME contribution is aimed at accomplishing a Federated Cloud-Edge Platform for the execution of extreme data analytics algorithms. The initial focus is on the design and then on the implementation of the so-called TEMA Core, the central unit able to host and orchestrate all TEMA technologies, which are developed primarily in WP3, WP4 and WP5 work packages. UNIME creates a microservices architecture based on Docker containers, whereas the orchestration of the microservices is performed using Kubernetes technology. TEMA Core is based on distributed IT resource nodes deployed on TEMA partners execution facilities, such as UNIME, AUTH, USE and PLUS. The overall TEMA network infrastructure is fully trusted leveraging virtual private network (VPN) connections, protected through 509v3 Certificates.






3 Explainable and robust analytics

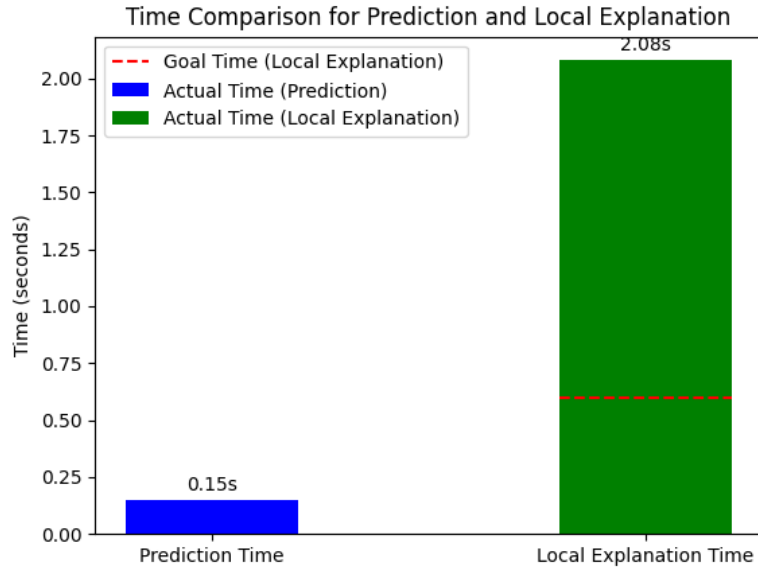

3.1 Introduction

In the rapidly evolving field of AI, the development of robust and explainable algorithms has become paramount to ensure trust and reliability in AI systems, in particular for high-stakes applications such as natural disaster management. Task T3.1 "Explainable and Robust Analytics", spans from M7 to M30 and involves a collaborative effort among FHHI, AUTH, and ATOS, with additional links provided by UNIME. This task is dedicated to advancing the state of the art in DNNs from the viewpoint of explainability and robustness. FHHI will lead with a focus on generic XAI methods, ATOS will specialize in XAI for visual data, AUTH will work on AI robustness, and UNIME will facilitate connections to Task T3.4. This is laid out in more detail in the following sections. This task is closely related to TEMA objective **OA1** "Increase trustworthiness of extreme data analysis algorithms" while ensuring that the time required to generate human interpretable local and global explanations must not exceed the time consumed for the prediction itself by a factor of 4 and an order of magnitude respectively.

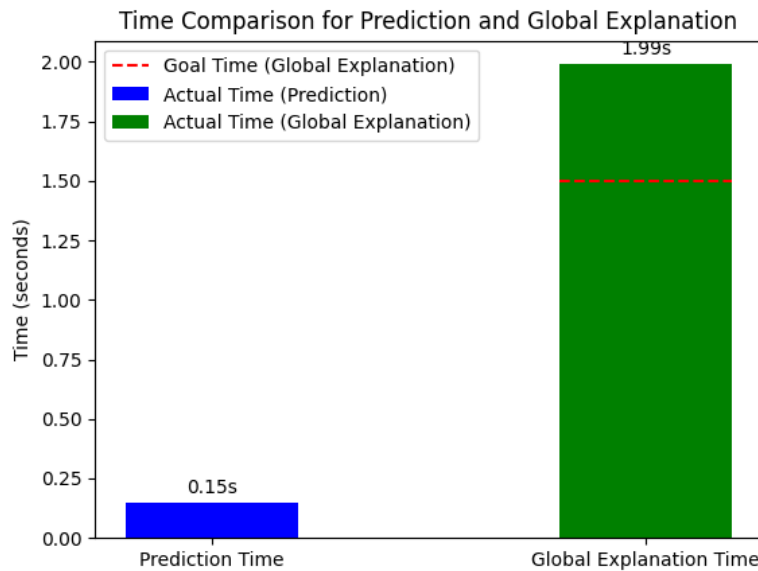
- For local XAI, FHHI computes a heatmap based on layer wise relevance propagation (LRP) for the predicted class as well as conditional heatmaps for the 5 most dominant concepts in time that currently exceeds the time for the prediction itself by a factor of 14.
- For global XAI, for a precomputed Gaussian mixture model (GMM) to cluster the latent relevancies, providing prototypes (samples closest to the means of the cluster). Then, by swiping through the entire dataset FHHI searches for reference samples for the top concepts appearing in the prototypes. The time needed for generating the global explanations exceeds the time needed for the prediction itself by a factor of 13.

Figure 1 depicts the time needed for generating local and global predictions in comparison to the time needed for the prediction itself for EfficientNet_b0. While currently not yet meeting the KPIs, this has to be understood as worst-case experiments. By tailoring explanations to the end-users needs, code optimization, as well as dimensionality reduction techniques (e.g. based on spectral clustering opposed to GMM-clustering in a high-dimensional latent space) FHHI expects further improvements towards reaching the KPIs during the next reporting periods.







(a)



(b)

Figure 1: (a) Speed of local XAI explaining individual predictions. (b) Speed of global XAI explaining global model strategies for EfficientNet_b0.






3.2 Generic XAI

Deep Neural Networks are prone to learning and relying on spurious correlations present in the training data. In safety-critical applications and high-stakes decision-making, including natural disaster management, such reliance can lead to fatal consequences. Post-hoc methods offer a solution to mitigate these issues without the need for computationally expensive retraining. However, they can harm the model performance by globally shifting the distribution of latent features.

Advances beyond SOTA

To tackle the aforementioned challenge, FHFI proposes a reactive approach conditioned on model-derived knowledge and XAI insights [BAR2024]. FHFI's method, R-CLArC, builds upon P-CLArC to effectively counteract model overcorrection. [DRE2024a] presents a novel method for model correction on the concept level that explicitly reduces model sensitivity towards biases via gradient penalization. When modeling biases by Concept Activation Vectors (CAVs), the importance of choosing robust directions is highlighted, as traditional regression-based approaches such as Support Vector Machines (SVMs) tend to result in diverging directions. [DRE2024b] models the distributions of latent relevances as Gaussian Mixtures Models, each for the different classes involved. This allows to define prototypes (samples closest to the means) of (sub-)classes (global XAI), enabling comparison with individual samples (local XAI) with applications for outlier detection. **Figure 2** depicts the correlating features in the fire classification partition BLAZE dataset developed by AUTH, using two prototypes of an EfficientNet_b0 in its penultimate layer. For each prototype, the relevant concepts and their corresponding relevance scores (%) are shown. In the spirit of mechanistic interpretability, [DRE2024c] deals with disentangling polysemanticity (neurons simultaneously encoding of more than one concept) by decomposing into multiple monosemantic "virtual" neurons, to increase latent interpretability, especially regarding safety critical tasks. [BLE2024] investigates the question of explaining predictive uncertainty, i.e. why a model 'doubts', which has been scarcely studied. It reveals that predictive uncertainty is dominated by second-order effects, involving single features or product interactions between them.



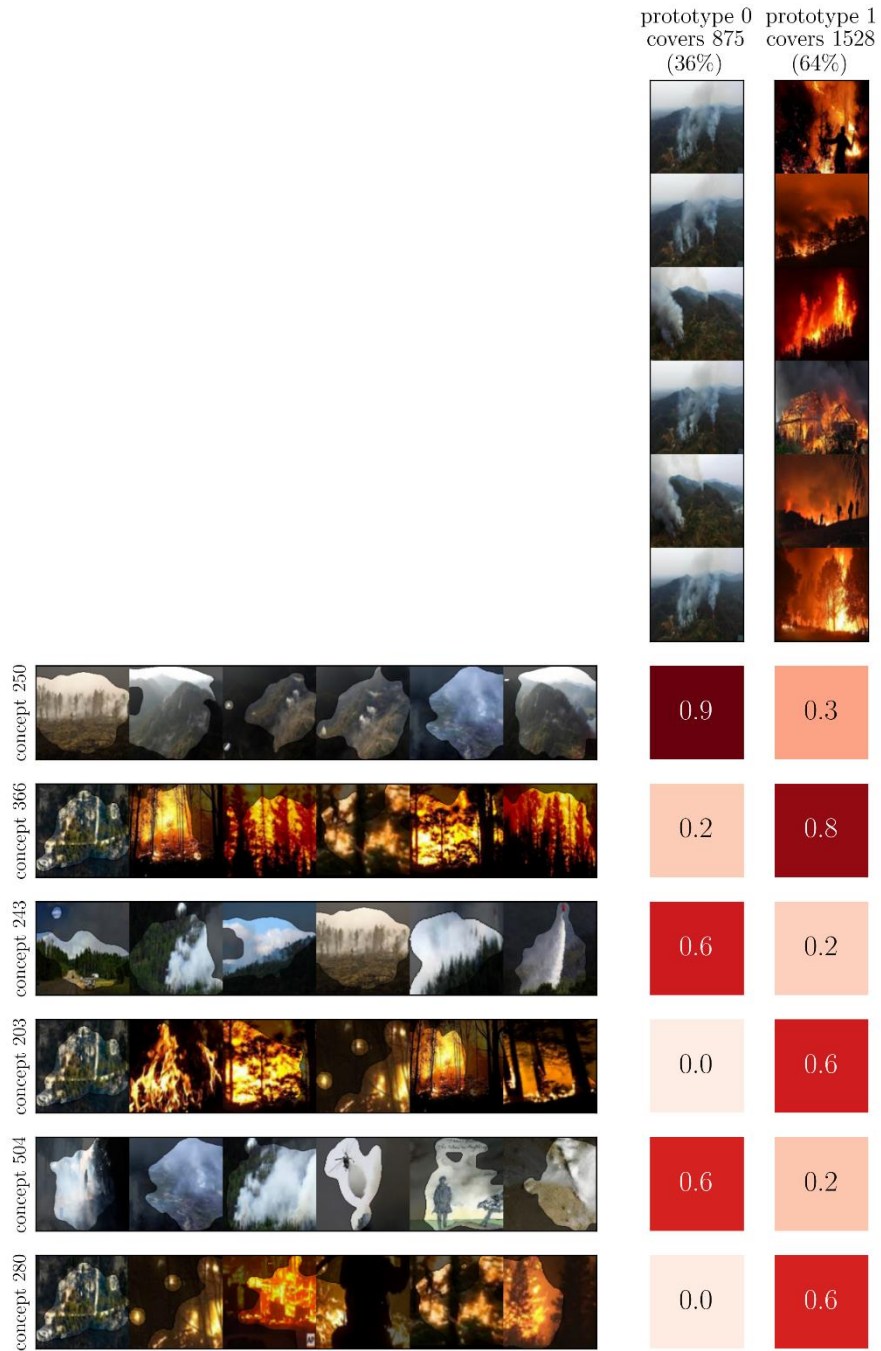





Figure 2: Prototype-based explanations for a fire-classification model based on PCX: two prominent concepts of “smoke” and “fire” are most crucial for the class prediction. The appearance of similar concepts suggests a high degree of redundancy within the trained neural network (PURE).







One of the unsolved challenges in XAI is determining how to most reliably estimate the quality of an explanation method in the absence of ground truth explanation labels. Competing evaluation methods, or "quality estimators," often yield conflicting rankings, complicating the selection of the best-performing explanation method. [HED2023] proposed MetaQuantus, a novel framework that conducts a meta-evaluation of various quality estimators in XAI. Released under an open-source license, FHFI's work serves as a development tool for XAI and Machine Learning (ML) practitioners, facilitating the verification and benchmarking of newly developed quality estimators within specific explainability contexts. Usefulness to human users, possibly from a non-scientific background such as staff of rescue and emergency agencies in natural disaster management, is an important aspect of XAI in practice, and a way to compare different XAI methods. In [DAW2023], three XAI methods (Prototypical Part Network, Occlusion, and Layer-wise Relevance Propagation) are compared with regard to their usefulness of explainability methods for end-users.

The TEMA project is a big data project, which revolves around tackling machine learning tasks that involve handling vast amounts of data that is not only voluminous and rapidly changing but also diverse in nature, applied to natural disaster management. While there is a plethora of research in the field of explainable AI that focuses primarily on computer vision tasks or natural language processing, FHFI aims to extend explainability to other cases from the variety of possible data sources. For Task T3.1 with regard to generic XAI methods, FHFI has notably investigated time series data (such as audio signals) from an explainability point of view. [VIEL2024] introduces a virtual inspection layer that transforms time series data into an interpretable representation, allowing relevance attributions to be propagated through local XAI methods. The proposed method, DFT-LRP, applies the discrete Fourier transform and Layer-wise Relevance Propagation (LRP) to enhance interpretability and reveal differences in classification strategies and correlations in various time series classification tasks (such as speech). [BEC2024] explores post-hoc explanations for DNNs in audio by using Layer-wise Relevance Propagation (LRP) to identify relevant features and introducing audible heatmaps for superior interpretability. It introduces AudioMNIST, a novel Open-Source audio dataset and derive hypotheses about feature selection, tested through manipulations of the input data. [FRO2023] leverages XAI to understand classification strategies in DNNs for audio event detection. They compare different input representations (waveform vs. spectrogram) using Layer-wise Relevance Propagation to uncover representation-dependent decision strategies and make informed decisions about model robustness and alignment with human requirements. Another preprint dealing with yet a different data source is [TIN2024].

Other work on generic XAI research performed by FHFI includes the following. [WEB2024] paper introduces Layer-wise Feedback Propagation (LFP), a new training approach for neural-network-like predictors that uses Layer-wise Relevance Propagation (LRP) to assign rewards to connections based on their contributions to solving a task. LFP distributes rewards throughout the





model without relying on gradients, strengthening structures with positive feedback and reducing the influence of structures with negative feedback. Data attribution is the problem of finding the training datapoints that are responsible for a given decision.


To address methodological caveats for the empirical interpretation of Model Parameter Randomisation Tests (MPRT), [HED2024] proposes Smooth MPRT and Efficient MPRT. The former minimises the impact that noise has on the evaluation results through sampling and the latter circumvents the need for biased similarity measurements by re-interpreting the test through the explanation's rise in complexity. DualView [YOL2024] proposes an efficient and scalable alternative to methods of data attribution. It uses a surrogate interpretable model to generate explanations hundreds of thousands of times faster, while producing explanations of comparable quality to other methods, across many different criteria. While currently this is only established for classification, a future extension could include a generalization to models for segmentation/object detection tasks.


3.3 XAI on Diffusion Models image generation

Given the specific nature of TEMA project, the availability of visual data pertaining to natural disasters is notably lacking, and in some instances, virtually non-existent. To address this challenge, one of the primary strategies TEMA project involves the generation of synthetic visual data through the utilization of diffusion models. However, it is imperative that all synthetic data generated be meticulously labeled to ensure its utility for the models outlined in the following sections.

SOTA


Zero-shot detection models offer a promising avenue for open-vocabulary detections without the prerequisite of prior class information, aligning seamlessly with ATOS's imperative to label images in scenarios characterized by data scarcity. At present, the pinnacle of achievement in zero-shot detections resides with GroundingDINO, which has a significantly high Average Precision (AP) of 52.5% on zero-shot tasks, as evidenced by its performance on the COCO dataset [LIU2023]. Throughout its research endeavors, ATOS tested GroundingDINO on a batch of synthetic images depicting forest fires. However, the accuracy of fire detections within these images yielded false positives and some inaccuracies, necessitating manual correction of labels to ensure the integrity of ATOS's dataset.








Advances beyond SOTA

In order to augment the quality of the generated labels, ATOS proposes a novel approach centered on the application of explainability to diffusion models. By leveraging Diffusion Attentive Attribution Maps (DAAM) [TAN2022], ATOS gains invaluable insights into how different words within the input prompt of the diffusion model influence distinct regions within the generated image. By harnessing the maps generated through DAAM during the diffusion process, ATOS can integrate this invaluable information with the labels provided by GroundingDINO, thereby elevating the accuracy of ATOS’s labeling process. **Figure 3** depicts a forest fire Image generated through a diffusion model, and the corresponding DAAM heatmap for the prompt “fire”.



(a)





(b)


Figure 3: Forest Fire synthetic image generated by ATOS through diffusion models, and the corresponding DAAM heatmap for the prompt “fire”.


3.4 AI robustness

In NDM, the integrity of the DNNs performance for semantic extreme analysis is vital. Adversarial attacks (e.g., in case of deliberate arson), which can subtly manipulate input data to deceive DNNs, pose a significant threat to the accuracy of these systems. Ensuring robustness helps maintain the trustworthiness and effectiveness of automated responses, thereby enhancing the overall resilience and responsiveness of disaster management efforts, ultimately safeguarding lives and property during critical events.

SOTA

An NDM AI system may face attacks in the case of deliberate disaster causes, e.g., arsons. Therefore, AI system robustness must be studied and defense systems must be developed. In the dynamic landscape of Machine Learning, the emergence of Adversarial Attacks [SZA2013] poses a critical challenge to the reliability and security of deployed DNN models. Adversarial attacks involve the intentional manipulation of input data to deceive Machine Learning models by adding small, imperceptible to human observers, perturbations, leading to DNN misclassifications [SZA2013]. As these attacks become increasingly sophisticated, the need for robust defense






strategies [MUS2019] becomes paramount to ensure the trustworthiness of artificial intelligence systems across various applications, e.g., for accurate disaster response coordination.

Advances beyond SOTA

In order to tackle the previously outlined challenges, a thorough examination of the DNN model decision boundaries, as established within an adversarial defense strategy, becomes a necessity. To address this issue, AUTH proposes a novel adoption of the One-versus-One (OvO) classification scheme, instead of the commonly employed One-versus-All (OvA). Specifically, AUTH introduces a new approach that combines the OvO DNN classification strategy with geometrically inspired criteria. The working principle is based on the fact that it is harder to fool many classifiers at the same time, than just one. Instead of creating multiple binary classifiers, this approach leverages the within class relationships, thus shaping a decision space that is naturally more robust to adversarial attacks.

AUTH has applied these robust DNN methods on the fire classification partition of BLAZE dataset to showcase that OvO strategy can counter attacks to forest fire classifiers. AUTH reports the results obtained at white-box settings on the the fire classification partition of BLAZE dataset. This is a challenging setting in which the attacker has full knowledge of the DNN model architecture and the parameters of the defender. The training procedure is divided into two modes: a "from scratch" (FS) training procedure, where the model, along with the respective defense strategy, is trained for 400 epochs; and a pretrained strategy, in which a SoftMax-only (SM) training step is applied for 200 epochs before applying the defense and training it for an additional 200 epochs using the SoftMax-only as pretrained step. **Table 1** shows the performance of various OvO methods on the fire classification partition of BLAZE dataset in comparison to the ones achieved by OvA vanilla approaches.

Specifically, AUTH reports the classification accuracy obtained on clean samples and the one obtained after performing the FGSM [GOO2014], BIM [KUR2018] and MIM [DON2018] adversarial attacks. PCL and HCP stand for Prototype Confirmation Loss and Hyperspherical Class Prototype methods respectively. The OvO HCP method achieves a fire classification accuracy increase of **13.21%** and **8.42%** in the case of white- box attacks, namely in BIM and in MIM attacks, respectively. In comparison with the state-of-the-art method mentioned above, a classification accuracy increase greater than the TEMA 5% KPI is achieved for the classification task.






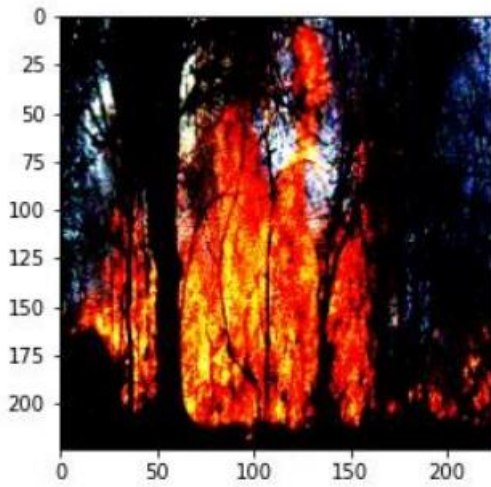



Table 1: Evaluating of robustification methods at white box attack settings on the fire classification partition of BLAZE dataset.

Model	Clean DNN Accuracy (%)	White-Box Accuracy (%)		
		FGSM	BIM	MIM
-	-	FGSM	BIM	MIM
OvA FS	73.25	2.55	0.00	0.00
OvA AT (SM + PGD) [MAD2017]	68.72	18.12	0.39	0.00
OvA PCL [MUS2019]	72.81	4.98	0.00	0.00
OvA HCP [MYG2022]	72.05	42.69	26.61	27.44
OvO FS	68.92	15.44	20.10	22.53
OvO HCP	73.58	44.48	39.82	35.86

Attackers can try to alter NDM data, e.g., forest fire images and force NDM AI systems to fail. An example of such an attack is showing in **Figure 4 a, b**. Even in perceivable forest fire image perturbations can render a DNN fire recognizer useless. In such a setting, AUTH proposes a novel adversarial defense method, namely Adversarial Defense via Adversarial Reprogramming (ADAR). The Adversarial Reprogramming method tries to find a single adversarial perturbation vector (program) that is applied to all inputs to a model in order to cause it to perform a new task chosen by the adversary. AUTH treats adversarial image classification as the new task for which AUTH creates an adversarial program. **Figure 4c** depicts the attacks of a forest fire image 4b, where an ADAR defense perturbation has been added. It is evident that the adversarial attacked image 4b and the defended image 4c contain perturbations that cannot be detected by humans, yet they impact classification DNN Models. Specifically, the ADAR-defended image 4c can now corrected be classified as forest fire image by a DNN forest fire classifier.






(c)


Figure 4: Examples of (a) forest fire image, (b) forest fire image with FGSM adversarial attack and (c) forest fire image with FGSM+ADAR adversarial attacks.

This novel method has a dual advantage over adversarial training or other existing adversarial defense methods. First, it can be applied to any DNN model (e.g., a pretrained or adversarially one or one trained on another task). Secondly, it does not affect the DNN model parameters at all. Rather, it is added as an extra input on top of the input data. **Table 2** presents the classification accuracies of: a) baseline ResNet50 DNN model, b) ADAR, c) ADAR combined with adversarial training and d) adversarial defense methods on adversarial test data samples of MNIST dataset, created with FGSM for an ϵ noise range of 0.1 – 0.5. **Table 2** shows that ADAR combined with adversarial training achieves higher classification accuracy on the MNIST dataset for attacks noise $\epsilon = 0.1$ and $\epsilon = 0.5$. This work is ongoing to explore its full potential for NDM AI system defense against attacks.

Table 2: Robustness Against Adversarial Attacks on MNIST dataset [DEN2012].

Method/ Noise	$\epsilon = 0.0$			$\epsilon = 0.1$			$\epsilon = 0.3$			$\epsilon = 0.5$		
Attack type	FGS M	BIM	MIM	FGS M	BIM	MIM	FGSM	BIM	MIM	FGSM	BIM	MIM
ResNet50	98.74	98.74	98.74	92.45	91.53	90.8	32.53	17.56	16.61	19.43	5.29	5.87
Adv. Train	94.68	94.68	94.68	99.15	98.49	98.67	58.05	70.83	55.86	29.38	28.46	23.07
ADAR	94.14	94.14	94.14	91.72	90.76	90.63	78.92	66.84	64.49	55.0	23.38	26.07
ADAR + Adv. Train	87.72	87.72	87.72	89.86	88.0	88.04	84.78	76.45	74.67	72.86	54.05	51.8





Random Smooth	98.7	98.7	98.7	94.13	91.38	90.95	46.69	14.2	11.97	25.48	0.94	1.05
Feature Squeeze	99.0	99.0	99.0	93.68	92.50	91.77	50.42	21.74	19.96	28.39	1.76	2.16
Defensive Distill	96.05	96.05	96.05	88.07	87.18	86.49	50.8	23.34	22.14	31.72	8.86	8.74

3.4.1 Decentralized Inference for Forest Fire Classification


SOTA


In the realm of natural disaster management (NDM), multi-agent systems can play a pivotal role in robust decision-making for disaster identification and management [DOM2011]. For example, such DNN agents can reside on several drones monitoring a forest-fire in a distributed way. In the field of Distributed Deep Neural Network (D-DNN) learning, such edge agents can transmit their results to the cloud for centralized DNN inference [ZHO2019]. However, this centralized solution is not robust as it has a central possible failure point. In the case of decentralized inference on the edge, e.g., using majority voting [MAL2022], DNN agents are susceptible to attacks forcing one or more of them to behave erroneously or even maliciously. Existing SOTA aggregation solutions (e.g., majority voting or weighted average) among the DNN agents [KIT1998] cannot provide robust decentralized DNN inference.

In simple terms, the strategies outlined above, e.g., majority voting, predominantly depend on the mutual trust among participating agents, presuming their consistent and reliable operation under all circumstances. However, this reliance poses a considerable risk of failure due to potential factors such as system crashes, computational errors, or malicious attacks. In such scenarios, certain agents may present a façade of normalcy, while clandestinely engaging in activities aimed at subverting and compromising the integrity of the NDM AI system [CAS2002]. So far, no robust decentralized DNN inference architectures have been proposed for NDM tasks.

Advances beyond SOTA

Motivated by these challenges, AUTH introduces a robust decentralized DNN inference framework operating in two modes. Firstly, AUTH proposes an Individualized Agent Model Selection Process (IAMSP), where a DNN agent-to-agent model selection technique is applied, enabling each autonomous DNN agent to enhance its performance through consultation with neighboring DNN agents. Once a DNN agent uses IMSP to decide which DNN agent to consult, it can use Weighted Average or Majority Voting to aggregate their inferences. Through experiments focused on forest-fire classification using 7 DNN agents, AUTH observes a notable enhancement in the classification accuracy of poor-behaving individual DNN agents. Specifically, in the case of a





poor-performing EfficientNet B4 agent, its individual fire classification accuracy is boosted from **61.90%** to **78.69%** , achieving an accuracy increase of **16.69%** compared to its stand-alone operation (without peer-consultation). This increase is much greater than the 5% KPI for the classification task.

Secondly, AUTH proposes the Quality of Inference (QoI), a novel Byzantine Fault Tolerant (BFT) consensus protocol, designed to serve as a universally accepted inference rule for all participating DNN agents within the DNN system, even if some of them behave maliciously. In such a setting, DNN inference aggregation methods fail completely to help DNN agents reach a consensus, since they lack fault-tolerant properties. On the opposite, the QoI protocol ensures that DNN agents do reach consensus despite attacks. To this end, AUTH compares the QoI performance with typically centralized aggregation rules such as Weighted-Average and simple Majority Voting [KIT1998] as shown on **Table 3**. The proposed consensus protocol results in an overall accuracy score surpassing those obtained through typical centralized DNN ensemble aggregation techniques.


Table 3: Forest fire classification accuracy on BLAZE dataset (fire classification partition) achieved by one agent assuming that all agents are honest.


Method	Centralized Voting Rules		QoI Consensus Protocol
	Weight Average	Majority Voting	Average
-			
DNN aggregation	84.65	83.66	85.64
IMASP (Avg condition)	-	-	85.51

4 Real-time semantic visual analysis and remote sensing

4.1 Introduction

Real-time semantic visual analysis and remote sensing play an indispensable role in NDM, providing critical insights and timely information for effective response and mitigation efforts. The ability to swiftly and accurately analyze visual data from various sources, such as satellite imagery and video feeds, enables emergency responders to assess the extent of damage, identify affected areas, and prioritize resources efficiently. The task T3.2 “Real-time Semantic Visual Analysis and







Remote Sensing” focus on developing advanced AI algorithms for semantic visual analysis and remote sensing, tailored for heterogeneous data modalities, as described in Subsection “Visual data analysis and remote sensing” of Section 1.2.2.A of Part B of TEMA’s Description of the Action (DoA). AUTH leads the task for forest fire and flood image analysis, concentrating on image and video instance segmentation, object detection, combating DNN overfitting, privacy preservation algorithms, and accelerating DNNs. DLR-DFD contributes by creating innovative methods for analyzing satellite and SAR data and engaging in R&D activities centered on remote sensing. DRL-KN focus on multimodal analysis to construct 3D maps of smoke.

In this context, the collaborative efforts of AUTH and DLR are instrumental in advancing the state-of-the-art in real-time semantic visual analysis and remote sensing. AUTH focus on image and video instance segmentation, object detection, and combating DNN overfitting directly contributes to enhancing the accuracy and speed of analysis in disaster scenarios. Moreover, AUTH work on privacy preservation algorithms ensures the ethical handling of sensitive information during crisis response operations. DLR expertise in satellite and SAR data analysis is particularly valuable in providing comprehensive and high-resolution imagery of disaster-affected regions. Their novel methodologies and research efforts in remote sensing technology significantly enhance the capabilities of disaster management authorities to assess damages, monitor environmental changes, and plan effective response strategies. ATOS, USE, NS, FHFI and UNIME provide links to other WPs and tasks.

4.2 AI algorithms for visual data analysis

Current methods for semantic image and video analysis encounter challenges in achieving a balance between speed and accuracy, particularly when handling extensive datasets [NAE2022]. Commonly, these Machine Learning tasks rely on DNNs, encompassing object detection, semantic or instance segmentation, and image or video classification. In an NDM setting, the following items (objects) are the classical visual analysis targets: a) fire, smoke, burnt region (in the case of forest fires), b) flood water (in the case of floods), c) human related objects, namely persons or cars in danger and d) infrastructure objects e.g., houses in danger. There are specialized solutions tailored for emergency response and prevention, such as visual fire detection [MAJ2022], smoke segmentation or detection [CA02022], and flood detection or recognition [HER2022, MAD2021]. In the recent past, most NDM DNNs used were CNNs variants. However, the scarcity of relevant training datasets coupled with the highly dynamic conditions in such datasets, including variations in lighting, aberrations, occlusions, movements, and zoom levels, pose significant challenges for data-intensive DNNs, consequently limiting their accuracy. As a result, the current state-of-the-art in visual analysis for real-time extreme data analysis during emergencies is still immature, as it necessitates new solutions that offer good compromise between speed, accuracy, and throughput






[NAE2022]. Newer DNN architectures, e.g., Transformer Networks [VAS2017] provide promising DNN solutions in an NDM setting. Newer CNN-based architectures are quite promising as well.

4.2.1 Fire detection

SOTA

In the initial phase of TEMA project, the most recent DNN YOLO v5 [JOC2023], with its variants (small, medium, large), was employed as the state-of-the-art architecture for object detection. Specifically, in fire detection frameworks [YAR2023, WU2023], YOLOv5 was utilized as the primary detector due to its superior performance in real-time detections, as evidenced by its testing on various fire and smoke detection datasets [VEN20232]. However, recently, Real-Time Detection Transformers (RT-DETR) [ZHA2023] have outperformed the YOLO family models on the COCO dataset, showcasing remarkable classification and localization capabilities. Despite this, RT-DETR has not been as extensively tested in fire detection scenarios as other DNN models [YAR2023], such as YOLO family [VEN2022], the Vision Transformer (ViT) [SHA2021], and the Detection Transformer [WAN2023]. RT-DETR [ZHA2023] is a DETR [CAR2020] based Transformer DNN model that is trained with a classification loss for the confidence score and two regression losses for the bounding box coordinates, namely the L_1 loss and the L_{GIoU} [REZ2019] loss, which is based on the Intersection over Union (IoU) measure. In contrast to L_{GIoU} , L_1 is not invariant to the size of the bounding boxes resulting in different L_1 error for larger DNN predicted or ground truth Regions of Interest (Rols), despite the high similarity with smaller Rols. This issue which is created from the imbalance between ground truth Rols of different size becomes more pronounced in architectures such as RT-DETR and DN-DETR [LI2022] that employ the denoising mechanism. This mechanism [LI2022] reduces the instability of bipartite matching in DETR-based models by adding noisy queries near a corresponding ground-truth Rol (bounding box) resulting in a much higher number of training targets. These shortcomings were addressed by AUTH in a novel way. An example of fire detection inference of RT-DETR trained with AUTH size-balanced L_1 loss is depicted on **Figure 5**.







Figure 5: Fire detection inference of RT-DETR trained with AUTHs size-balanced L_1 loss.

Advances beyond SOTA

In scenarios involving ground truth fire Rols of varying scales within the same image as fire hot spots nearby a wildfire, larger boxes exhibit higher L_1 errors due to their larger distances in the pixel domain [REZ2019]. This sensitivity results in disagreements between the L_1 and the L_{GIoU} loss, which impact the training of the neural network, especially when denoising is applied, as the number of these incidents increases. To mitigate this problem in scenarios with multiple fire targets within an image, AUTH adjusts the contribution of each prediction-target error to the L_1 loss based on its corresponding ground truth Rol (bounding box) size. More specifically, AUTH assigns greater weight to predictions with smaller bounding box sizes using the softmax activation function. Like L_1 , where each weight is equal to $1/N$, the AUTH chosen weights sum to one, promoting more stable DNN training across the entire dataset. As illustrated in **Figure 6**, AUTH's loss mechanism ensures that larger objects contribute less compared to the traditional $L1$ loss. This approach potentially enhances training, especially in scenarios with multiple ground truths per image.



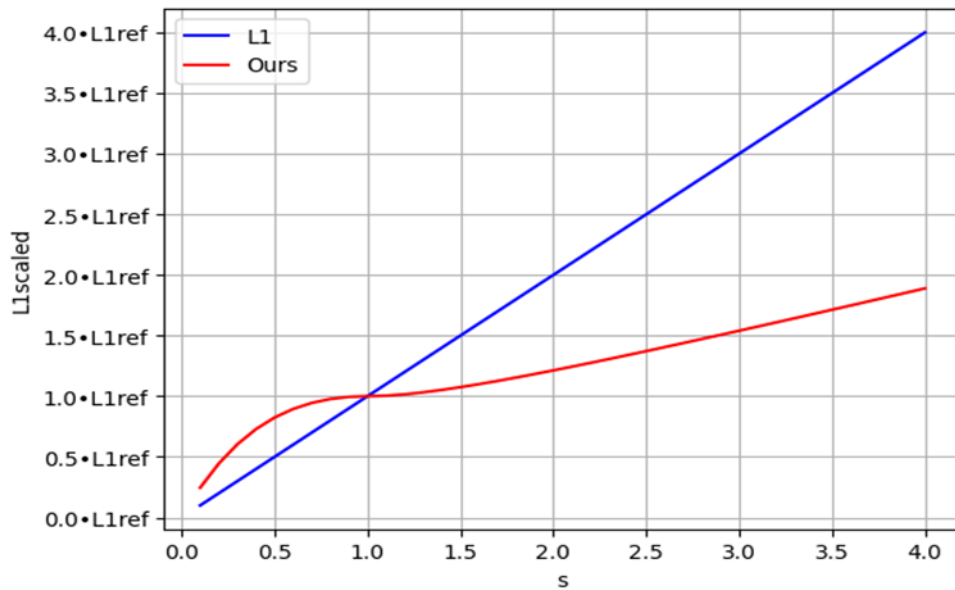





Figure 6: Comparison between AUTH loss and L_1 . Specifically, this is an example in which AUTH has two DNN prediction-target sets that have the same IoU and one set is the scaled version of the other one. The method defines a reference set with error equal to L_{1ref} and gradually increases the height and width of the scaled box by the factor s . The scaled box has loss $L_{1scaled}$ (ours), is computed with respect to the scale factor s and ensure smaller contribution of larger object to the DNN loss function.

Novel fire detection methods utilize the mean Average Precision (mAP) metric to evaluate their performance on detecting objects/fires, which awards/penalizes object/fire bounding box predictions based on their alignment with the corresponding ground-truth boxes. In most objects such as cars, numerous “children” objects that belong to different classes (e.g., car wheel, car window) collectively contribute in creating the “parent” object (car). Consequently, each “parent” object corresponds to exactly one ground-truth bounding box. However, in the case of objects like fire, “children” objects belong to the same class as the “parent” object (fire), which creates uncertainty regarding whether each “child” is, in fact, a “parent” object. This uncommon property of fire image instances introduces uncertainty for both human annotators and DNNs/CNNs concerning the number of fire Rols (bounding boxes) required to represent a fire image region (object) accurately. To tackle this, AUTH proposes a new evaluation measure for fire detection, namely Image-level mean Average Precision (ImAP) [TZI2023]. Instead of looking at each predicted bounding box separately, ImAP evaluates the fire detection models on their ability to predict fire object bounding boxes in the whole image.

This approach has been implemented on RT-DETR with a ResNet50 backbone and tested across three datasets: two data sets focusing on fire detection (Forest-Fire, Fire-Rob) and one data set focusing on fire-smoke detection(D-Fire). The Forest-Fire dataset, designed exclusively for fire







detection, was created by merging three distinct datasets: Dfire [VEN2022], Jhope [ICH2022], and Crossican [TOU2017]. Crossican, initially a forest fire segmentation dataset, was transformed into a fire detection dataset through image processing techniques. Jhope, sourced from Roboflow, includes a wide variety of fire types. Smoke annotations were removed from Dfire to repurpose it solely for fire detection. Additionally, AUTH developed a specialized test set primarily focused on forest fires and wildfires, aimed at aiding fire detector training. **Table 4** depicts the fire detection evaluation results on the Forest-Fire dataset. The mean Average Precision (mAP) measure of the fire detection on the Forest-Fire dataset, show that RT-DETR achieves a superior fire detection performance of **61.71%** using the AUTHs modified loss called RTDETR-Lsb surpassing the YOLOv5 (small, medium, large) by **8.81%**, **6.71%**, and **6.21%** respectively, and the original RT-DETR by **1.4%**. This technological component will be incorporated into the TEMA platform (as detailed in D2.2) within the context of WP6.

Table 4: Fire detection evaluation results on the Forest-Fire dataset.

MODEL	mAP (KPI-METRIC)	mAP50	ImAP	ImAP50
YOLOV5-S [JOC2023]	52.9	82.1	59.78	90.62
YOLOV5-M [JOC2023]	55.0	83.5	61.96	90.5
YOLOV5-L [JOC2023]	55.5 (<i>SOTA</i>)	84.2	61.53	90.51
RT-DETR [ZHA2023]	60.43	85.99	69.01	94.07
RT-DETR-Lsb (OURS)	61.71 (+6.21%)	86.19	69.42	94.67

Based on **Table 4**, RT-DETR trained with the AUTH size-balanced L_1 loss surpasses the previous state-of-the-art fire detection Deep Neural Network, YOLOv5-L, by **6.21%** using the mAP evaluation measure. This advancement exceeds the expected value for "**Object Detection Accuracy**" KPI under objective **OA2** "Increase Accuracy of Extreme Data Analysis Algorithms". Additionally, RT-DETR demonstrates real-time performance, with an inference speed of 108 FPS on a T4 GPU using TensorRT FP16. Thus, it is compliant with the TEMA objective **OA3** "Increase responsiveness/speed of extreme data analysis algorithms". as it greatly exceeds the "**Visual analysis speed**" KPI.





4.2.2 Fire region segmentation

SOTA

Limited research has been performed on CNN-based fire region segmentation, e.g. using the FLAME dataset [GUA2022], that was focused on region segmentation accuracy. It employed the Feature Pyramid Network to export region heatmaps. DeepLabV3+ was also employed for semantic forest fire segmentation on UAV RGB images [LEE2023, QUR2023]. However, these methods use Conditional Random Field (CRF) components. Therefore, they have high computational complexity and their Frames Per Second (FPS) rate is much less than the real-time rate of 25-30 FPS. As far as the AUTH team knows, no comprehensive comparative study of CNN-based forest fire region segmentation algorithms has been performed so far. Furthermore, the employed region segmentation measures can not be readily explainable by the NDM authorities. AUTH addresses the abovementioned shortcomings in a novel way.


Advances beyond SOTA


AUTH conducted a comprehensive comparative analysis of algorithms for forest fire image segmentation. Specifically, the CNN-based BiSeNet [YU2018], I2I-CNN [PAP2021], and PID-Net [XU2023] region segmentation algorithms were applied on Unmanned Aerial Vehicle (UAV) fire images for segmenting fire regions. Examples of segmented forest fires using the I2I-CNN [PAP2021] are depicted on **Figure 7**.



Figure 7: Examples of segmented forest fires using the I2I-CNN [PAP2021].

Fire region segmentation Intersection over Union (IoU) is the most widely used region segmentation metric [REI2022]. Mean IoU is the most common and efficient measure in both object







detection and semantic segmentation. However, NDM authorities must be able to interpret fire region segmentation output in an objective and explainable way. Therefore, AUTH proposes the implementation of novel region segmentation measures that are tailored specifically for forest fire segmentation. They are designed to augment the interpretability of DNN outputs without the need of ground truth fire images, furnishing end-users with enhanced insights into DNN model performance, beyond the one offered by conventional measures, such as the mean IoU.

Since a forest fire is more dangerous if it has multiple sources (regions) that cover big areas and are spatially dispersed, the following novel fire segmentation metrics are proposed to accurately describe the fire region extent in a more explainable, context-rich form: a) D_n is the number of predicted fire region instances, b) D_a measures the average fire region area (in pixels) and c) D_s measures the spatial dispersion of the fire region instances. By offering additional contextual information regarding forest fires, these measures empower stakeholders to make more informed decisions regarding forest fire management. The evaluation of several state-of-the-art region segmentation architectures on the Flame dataset with respect to Mean IoU, D_n , D_a , and D_s is presented on **Table 5**. When using the proposed new forest fire segmentation metrics, the I2I-CNN-R18 region segmentation method is the clear winner. This technological component will be incorporated into the TEMA platform (as detailed in D2.2) within the context of WP6.

Table 5: Evaluation of state-of-the-art region segmentation architectures on the Flame dataset [SHA2020].

CNN Architecture	Input Data	Evaluation metric			
		mIoU	DN	DA	DS
PID-Net-R18	RGB	0.9146	0.3006	0.1503	0.0149
	RGB+HSV	0.9113	0.3082	0.1853	0.0166
	RGBS	0.9129	0.269	0.148	0.0132
I2I-CNN-R18	RGB	0.89532	0.702	0.326	0.055
	RGB+HSV	0.82677	0.961	0.453	0.082
	RGBS	0.82959	0.989	0.514	0.082
BiSeNet-R18	RGB	0.90482	0.414	0.205	0.019
	RGB+HSV	0.90781	0.353	0.19	0.019
	RGBS	0.90495	0.372	0.196	0.02
BiSeNet-R101	RGB	0.87358	0.45	0.247	0.026
	RGB+HSV	0.86541	0.456	0.254	0.027
	RGBS	0.81785	0.628	0.39	0.037





4.2.3 Forest Fire Classification

SOTA

Image classification is a classical Machine Learning task that is related to object detection and semantic segmentation. Currently, the SOTA in wildfire image classification is an XceptionNet, which was trained and evaluated on the FLAME [SHA2021] dataset. It achieved **76.23%** classification accuracy on the Fire/No-Fire binary classification task.


Advances beyond SOTA


Besides working on forest fire detection and segmentation, AUTH worked on forest fire classification, as classification is an independent Machine Learning task from the other two. To this end AUTH employed the fire classification partition of the BLAZE dataset (see **Section 4.2.4**). Three evaluation protocols take place. For the first one, AUTH makes use of the first three classes, to assess the burnt severity of an area for comparative purposes. For the second one, AUTH makes use of all 5 classes, namely: 'burnt region', 'half-burnt region', 'unburnt region', 'fire image', and 'smoke image'. For the third one, AUTH merges the 'fire image' and 'smoke image' classes under one label called 'fire image', since in many cases fire and smoke images bear the same semantic significance for firefighters. The classification accuracy results of various newly trained DNN classifiers, along with the inference speed measured in FPS, can be seen on **Table 6**.

Table 6: Classification accuracy and speed of seven DNN classifiers for fire classification on the fire classification partition of BLAZE dataset.

Model	Evaluation Protocol 1 (%)	Evaluation Protocol 2 (%)	Evaluation Protocol 3 (%)	FPS
InceptionNetV3 [SZE2015]	54.36	59.22	65.92	109
ResNet101 [HE2015]	65.55	61.01	68.92	100
ResNet50	68.01	69.43	76.90	196
EfficientNetB0 [TAN2019]	77.74	78.94	82.32	140
EfficientNetB1	77.41	82.71	85.32	103
EfficientNetB2	78.97	80.03	84.49	106
EfficientNetB3	81.21	80.03	83.85	91

It is clearly seen that several DNN classifiers achieve very good fire classification accuracy (around **80%**) while operating faster than real-time. Moreover, we fine-tuned our pretrained EfficientNet





B1 from Evaluation Protocol 3 on the FLAME dataset, and achieved a classification accuracy of **92.28%**, which is over **14%** in comparison with the above-mentioned SOTA (**76.23%**), more than the **5%** TEMA KPI classification accuracy increase.

4.2.4 Burnt forest region segmentation SOTA

Currently, the task of segmenting burnt regions on drone images did not receive significant attention in the literature. There is only one publicly available dataset, the Burned Area UAV Dataset [RIB2023], researchers that was used to train and assess the performance of a UNet [RON2015] architecture, which achieves a Mean Intersection over UNion accuracy of 66.26% on a new dataset created by AUTH.


Advances beyond SOTA


To fill this research gap, AUTH has developed a new burnt region dataset and employed six region segmentation DNNs for training and evaluation. This dataset contains both burnt forest (primarily) and cityscape images. An example of burnt region segmentation on a sample from the burnt region segmentation partition of BLAZE dataset using the CNN-I2I is depicted in **Figure 8**.



Figure 8: Example of burnt forest region segmentation on sample from the burnt region segmentation partition of BLAZE dataset using the CNN-I2I. (a) raw image, (b) burnt area segmentation mask.

DNN based forest region segmentation methods usually require a sufficient amount of properly annotated data to achieve high segmentation performance. Due to the scarcity of such





ground truth data, AUTH prioritized the creation of a burnt region segmentation dataset called BLAZE [All2024]. AUTH has gathered 5.5k RGB images depicting scenes before, during, and after a wildfire, creating the BLAZE dataset. This dataset is suitable for fire classification and burnt region segmentation tasks. Specifically, all images are assigned to 5 classes: 'Burnt region', 'Half-Burnt region', 'Unburnt region', 'Fire', and 'Smoke' images.


The BLAZE dataset is divided into two partitions. The fire classification subset, which includes all 5.5k images assigned to the five classes. The burnt region segmentation subset, which includes 1028 annotated images, where each pixel is labeled as either 'burnt region' or 'unburnt region'. The annotations are manually created burnt region masks that are overlaid on images, separating burnt from unburnt regions. The dataset was split into non-overlapping training and test data sets.


Furthermore, AUTH used the burnt region segmentation partition of BLAZE dataset to train and the evaluate six real-time region segmentation methods. The Mean Intersection over Union (MIoU), as well as the Intersection over Union (IoU) region segmentation results are shown in **Table 7**.

Table 7: Evaluation of multiple real-time semantic segmentation architectures on the burnt region segmentation partition BLAZE dataset.

Model	MIoU(%)	Burnt region IoU(%)	Unburnt region IoU(%)
CNN-I2I [PAP2021]	74.82	75.45	74.19
CABiNet [KUM2021]	74.58	75.14	74.02
PIDNet-S [XU2023]	73.76	75.07	72.45
PIDNet-M	71.27	72.24	70.29
PIDNet-L	68.23	69.33	67.12
UNet++ [ZHO2018]	63.51	66.96	60.07

The best results were obtained by the CNN-I2I method. AUTH then measured the inference speed of the CNN-I2I method, which was found to be **145 FPS**. The CNN-I2I method achieved an accuracy increase of **8.56%** in comparison with the state-of-the-art UNet method, an increase that is greater than the **5%** KPI for the segmentation task. In addition, the **145 FPS** inference speed means that the method operates much faster than real-time (**25-30 FPS**).





4.2.5 Flood region segmentation

SOTA

Flood region segmentation, is the task of segmenting flood water pixels on an input flood image. In the recent years, few related publications appeared in the literature that analyze drone images. Consequently, publicly available datasets suitable for training DNN models, especially those designed for Unmanned Aerial Vehicles (UAVs), are scarce. For instance, a recent approach [HER2022] opts to train DNN models using FloodNet dataset [RAH2021], a UAV imagery dataset that depicts post-disaster scenes. Therefore, it is not the ideal choice for a flood monitoring solution.

In the initial phase of TEMA project, AUTH evaluated several benchmark Semantic Segmentation DNN architectures, which could support real-time flood region segmentation when operating in commercial Nvidia GPUs. After extensive experimentation, AUTH shifted its focus to PSPnet [ZHA2017], a popular region segmentation architecture that adopts a pyramid pooling module to capture global region context information. It has a Resnet50 [HE2016] backbone and outperformed other DNNs that were specifically designed for real-time segmentation, while managing to maintain real-time inference capabilities.. It scored **82.94%** mIoU on AUTH flood region segmentation test dataset. Examples of flood segmentation using the PSPnet. [ZHA2017] are depicted in **Figure 9**.






Figure 9: Examples of flood segmentation using the PSPnet. [ZHA2017].

Advances beyond SOTA

To overcome the flood region segmentation challenge, AUTH has curated a novel flood segmentation training dataset, comprising annotated images from diverse sources. In total, AUTH





gathered 548 annotated images from three different sources: 1) Kaggle - Flooded Area Segmentation, 2) Kaggle - Roadway Flooding dataset [SAZ2019] and 3) V-Floodnet [YON2023]. These images were split in a 3:1 ratio into training and validation sets. For flood segmentation DNN model testing, AUTH created a new flood segmentation test set by cropping, and manually annotating 567 video frames from a flash flood in Mandra, Greece.

To address the data scarcity problem, AUTH followed a different research approach as well. It exploited vast unlabelled flood image datasets available in the web, through semi-supervised training by employing the ST++ method[YAN2022]. This is a popular semi-supervised, self-training method which uses an already trained DNN model to predict pseudo-labels for unlabeled images. Then, during the final DNN training stage, it applies strong image perturbations to the unlabeled images, in order to prevent overfitting to the pseudolabels, since they may contain errors. This method can also be utilized in two distinct steps, to achieve even better results but with more training time. More specifically, AUTH chose to include unlabelled images from FloodImg DB [PAL2022], which is a large unlabeled flood dataset. Notably, AUTH flood region segmentation results are very promising; when choosing a 1:6 labeled-unlabeled image approximately ratio, AUTH flood region segmentation network outperforms the baselines ones and scores **86.55%** mIoU on AUTHs test set (~**3.5%** increase in performance). This technological component will be incorporated into the TEMA platform (as detailed in D2.2) within the context of WP6.


4.2.6 Person and car detection in flooded areas


SOTA

The integration of person/car detection in the TEMA platform stems from the critical need to precisely localize humans and vehicles that might be at risk during a flash or regional flood. The flood case particularity is that humans or cars may be partially submerged in flood water. Hence, their background has a particular visual appearance. Unfortunately, there are not any publicly available, annotated datasets dedicated for person or car detection in a flood context. A relevant study was published in [PAL2022]. However, the dataset labels were not published thus prohibiting AUTH to use it in its studies. To address the lack of training data, AUTH opted to use object detection models that were trained on other large benchmark datasets, like COCO [LIN2014]. Such datasets contain thousands of person and car instances. Therefore, a DNN model that was trained on them is expected to generalize better in different domains/contexts, as in the case of floods.

Advance beyond SOTA

In the beginning of TEMA project, AUTH chose to work with YOLOv6 model [LI2023], as it offers state-of-the-art performance for real-time object detection. Therefore, AUTH chose YOLOv6 small and large models and collected a test set to measure car and person detection performance





in flood scenarios. Such test set images depict flooded regions and submerged cars and/or people in close proximity or within the flood water. AUTH manually annotated three object classes (person, car, and truck). Examples of person and car detection in flooded areas are depicted on **Figure 10**.




(a)



(b)

Figure 10: Examples of (a) person and (b) car detection in flooded areas.

AUTH measured and compared the performance of Small and Large Yolov6 models. It is worth mentioning that these two chosen YOLOv6 models do not share the same architecture. The small YOLOv6 model is equipped with a single-branch architecture (during the inference stage), while the large YOLOv6 model was designed with a multi-branch architecture, achieving better performance [LI2023]. The results are presented in **Table 8**, when using two different mAP metrics (corresponding to different True Positive thresholds). The inference latency (ms) was measured using an Nvidia GTX commercial GPU.






Table 8: YOLOv6-S and YOLOv6-L performance and efficiency on Nvidia GTX 1080.

Model	mAP[0.5:0.95]	mAP[0.5]	Latency (ms)	FPS
YOLOv6-S	0.53	0.73	13.8	72.6
YOLOv6-L	0.61	0.80	30.3	33

The results are very promising, as both YOLO DNNs achieve very satisfactory performance faster than real-time. They can support the AUTH flood monitoring solution by precisely detecting both humans and vehicles. The final choice depends on the Tema HW platform characteristics. This technological component will be incorporated into the TEMA platform (as detailed in D2.2) within the context of WP6.


4.2.7 Privacy protection

SOTA

In the contemporary landscape of Artificial Intelligence (AI) applications, ethical concerns have become increasingly pronounced, particularly ones related to privacy protection. Such concerns arise also in the case of NDM, as collected visual data may depict humans. One of the primary ethical concerns in modern AI applications revolves around the use of facial recognition technology and its implications for privacy [CHA2020]. Therefore, various face blurring methods, as well as face de-identification methods [NOU2020, CHA2019, CHR2018] have been proposed in the literature.

Advances beyond SOTA

AUTH has developed a face detection and blurring technique that was incorporated in TEMA technological components as can be seen in **Figure 11**.







Figure 11: Example of face detection and blurring by AUTH.

This technological component will be incorporated into the TEMA platform (as detailed in D2.2) within the context of WP6. Furthermore, on the research front, AUTH proposes a pioneering approach named Privacy via Adversarial Reprogramming (PAR) [ELS2018], which harnesses the power of Adversarial Reprogramming techniques to address the critical issue of privacy preservation in machine learning models. PAR operates by reprogramming the target DNN model (e.g., the face classifier), thereby effectively concealing sensitive information while maintaining model functionality. For example, the adversarial perturbation noise (program) applied on facial image shown in **Figure 12** renders the facial image completely unintelligible by humans and face recognizers alike. However, a target DNN model can still recognize this image. So far AUTH has very promising results on MNSIST images and extends the PAR method for face recognition/detection tasks.







Figure 12: (a) Raw image sample from celebA dataset containing human face. (b) Image sample where the human face is totally hidden based on PAR, yet recognizable by a chosen face recognizer.


As depicted in **Table 9** the model trained with PAR maintains its performance even though that images are partly occluded due to the concealment of sensitive information.


Table 9: Accuracy of several CNN architectures on MNIST dataset where parts of the images containing sensitive information such as human faces or car plates are totally hidden. ResNet trained with PAR maintains its performance even though the images are partly occluded.

Model	Accuracy (%)
MLP	6.5
CNN	8.9
ResNet +PAR [HE2016]	97.89
AlexNet [KRI2017]	12.56
InceptionNet [SZE2015]	11.8

4.2.8 Transformer architecture acceleration SOTA

Transformers have emerged as a dominant force in Machine Learning, surpassing traditional DNN models like Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) in many tasks. Their unique self-attention mechanism allows for the parallel data sequence processing, enabling them to capture long-range dependencies more effectively than RNNs. Their ability to model complex data patterns surpasses that of CNNs. This makes them particularly well-






suited for a wide range of applications in Natural Language Processing and Computer Vision. They have already been used in TEMA for object detection, e.g., of forest fires. However, this strength comes at a cost: the computational demands of self-attention computation, especially in large models, present a significant bottleneck, necessitating research into more efficient computational methods. The computational complexity of the attention operation is quadratic with relation to the sequence length N .

While the research community has made strides in addressing the computational demands of transformers, their efforts have concentrated on optimizing the self-attention operation for processing larger sequence sizes N . These advances, albeit significant, overlook a crucial aspect: the reduction of parameters in projection matrices and the processing of input feature spaces of high dimensionality D .

Advances beyond SOTA

In response to this research gap, AUTH introduces a novel approach to Multi-Head Self-Attention (MHSA) computation. Specifically, AUTH proposes the Structured Efficient Self-Attention (SESA), a generic paradigm that eliminates the massive computation needs and memory requirements of MHSA. Inspired by the Johnson-Lindenstrauss (JL) lemma [JOH1984], an Adaptive Fast JL Transform (A-FJLT) is employed, that is parameterised by a single learnable circulant matrix. By enforcing a circulant structure to the projection matrices of Query, Key, and Value subspaces, it significantly reduces the spatial complexity from $O(D^2)$ to $O(D)$ for each respective learnable matrices \mathbf{W}_q , \mathbf{W}_k , and \mathbf{W}_v . Additionally, by leveraging the Discrete Fourier Transform (DFT) that can be computed through Fast Fourier Transform (FFT), AUTH decreases the MHSA computational complexity from $O(D^2)$ to $O(D \log |D|)$. This advance not only makes transformer models more suitable to systems with limited GPU resources, but also allows larger batch size during training, potentially leading to smoother loss function gradients and better Transformer model convergence. The proposed method is depicted on **Figure 13**. The proposed method can efficiently process high-dimensional feature vectors. This is a crucial factor for encoding richer information in complex computer vision tasks. It allows the elimination of **75%** of the learnable legacy MHSA parameters at a cost of a very slight sacrifices to object detection accuracy. Furthermore, the reduced Transformer parameter count also implies a lower overfitting risk. SESA properties are showcased on the demanding task of object detection on the COCO dataset [LIN2014], achieving comparable performance with its computationally intensive counterparts. AUTH showcases the efficiency of its proposed method in real-time object detection by integrating it with the RT-DETR-R50 model and compares it against state-of-the-art real-time object detectors. The results of this comparison are depicted on **Table 10**, using various mean Average Precision (mAP) levels. This table shows that the RT-DETR using the SESA method has **26.1** million fewer parameters than the best real-time object detector (YOLOv8-X) and slightly less mAP figures.



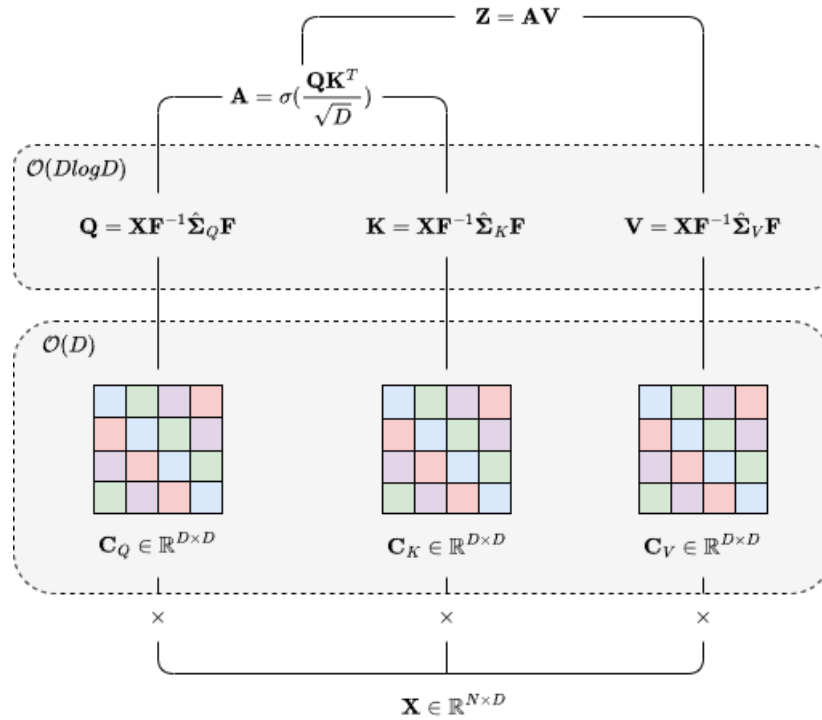





Figure 13: The proposed SESA module. By enforcing a circulant structure onto the projection matrices of query, key, and value subspaces, it significantly reduces the spatial complexity from $O(D^2)$ to $O(D)$ for each respective learnable matrices \mathbf{W}_q , \mathbf{W}_k , and \mathbf{W}_v . By leveraging the Discrete Fourier Transform (DFT) matrix and Fast Fourier Transform (FFT), AUTH decreases the theoretical computational complexity from $O(D^2)$ to $O(D \log(D))$.

Table 10: Comparison of state-of-the-art real-time object detectors on COCO val 2017 [LIN2014].

Model	Params(M)	mAP50	mAP50s	mAP50m	mAP50l
YOLOv7-L [CHI2022]	36.9	51.2	35.2	55.9	66.7
RT-DETR-R50 + E-C-SESA (proposed)	42.1	52.2	34.8	56.7	69.6
YOLOv8-L [JOC2023]	43.7	52.9	35.3	58.3	69.8
YOLOv5-L [JOC2022]	46.5	49.0	-	-	-
PPYOLOE-L [SHA2022]	52.2	51.4	31.4	55.3	66.1
YOLOv6-L [CHU2023]	59.6	52.8	34.4	58.1	70.1
YOLOv8-X [JOC2023]	68.2	53.9	35.7	59.3	70.7
YOLOv7-X [CHI2022]	71.3	52.9	36.9	57.7	68.6
YOLOv5-X [JOC2022]	86.7	50.7	-	-	-
PPYOLOE-X [SHA2022]	98.4	52.3	33.3	56.3	66.4







4.2.9 Synthetic data generation

As previously mentioned, the scarcity of training data poses a significant challenge for Machine Learning models tasked with addressing natural disaster scenarios. Recent advancements in diffusion models offer a promising solution by enabling the generation of synthetic images closely mirroring real-life scenes based on text prompts [WEI2024]. To confront this data scarcity head-on, ATOS is applying diffusion models to curate new datasets comprising synthetic images depicting natural disasters such as forest fires and floods, as observed from drone perspectives. These datasets encompass a diverse array of scenarios, encompassing variations in climate, time of day, and geographical locations, thereby enriching the training data available for ATOS's models. Examples of forest fire and flood images from a drone point of view (POV) generated through a diffusion model are depicted on **Figure 14**.



(a)






(b)


Figure 14: Forest Fire (a) and Flood (b) Images from a drone POV generated by ATOS through diffusion models.

4.2.10 Person re-identification

The person re-identification task integral to TEMA project necessitates the real-time processing of multiple video streams sourced from diverse origins to facilitate search and rescue operations during natural disasters. This multifaceted task is segmented into three distinct components: person detection, person tracking, and person ReID. The former two aspects must be executed across every video stream, while the latter entails querying the ReID process, which maintains a repository of identified individuals characterized by unique visual attributes.

Currently, the state-of-the-art (SOTA) for person detection in general-purpose scenes is YOLOv8 [REI2023]. In terms of tracking, BYTetrack [ZHA2022] emerges as the best approach in TEMA case given the highest FPS results, an essential criterion given TEMA requirement for real-time processing, especially at the edge. As for person ReID, ATOS employs a ResNet50 model trained on the Market1501 dataset to extract distinctive features unique to each individual.





4.3 Methods for satellite/SAR data analysis

4.3.1 Satellite-based flood detection and assessment

SOTA


Conventional rapid mapping methods that analyse satellite images to support situational awareness during disasters can be slow and labour-intensive [VOI2016]. Recent advances in deep learning and the availability of new large-scale remote sensing reference datasets have opened new possibilities for automated image analysis [BEN2022, BON2020, MAT2017].


Advances beyond SOTA

DLR deploys a modular processing chain for surface water monitoring that enables automatic satellite data search, pre-processing, analysis, and dissemination over predefined areas of interest [WIE2019]. Sentinel-1 (radar) and Sentinel-2 (multi-spectral) satellite images are analyzed using pre-trained DNN models for semantic segmentation to extract binary water masks. Time-series analysis is used to further derive secondary products about permanent water bodies [MAR2022] and flood duration [MAR2019]. Within TEMA, DLR further develops this method and have trained, validated, and tested various DNN models for water segmentation. DLR compared 16 different encoder and decoder (U-Net, DeepLab-V3+, FPN, PSPNet) combinations. **Table 11** shows the top-performing architecture (U-Net with EfficientNet-B4) and compares it against the previously best-performing network (U-Net with EfficientNet-B0) as of [WIE2023a], all networks have been trained independently with an AdamW optimizer, initial learning rate of 1e-3, weight decay of 1e-2 and a weighted combination of binary cross entropy and Lovász Hinge loss. Training is performed with 32-bit floating-point precision (FP32) using Pytorch.

Table 11: Performance comparison of top-performing DNN models for semantic segmentation of water bodies from Sentinel-1 (radar) and Sentinel-2 (multi-spectral) satellite images.

ID	Decoder	Encoder	Sentinel-1			Sentinel-2		
			IoU	Precision	Recall	IoU	Precision	Recall
1.2.1	U-Net	EfficientNet-B0	0.84	0.95	0.88	0.94	0.99	0.95
1.2.2	U-Net	EfficientNet-B4	0.94	0.98	0.97	0.98	0.99	0.99





The basis of these analyses formed a globally sampled and freely available reference dataset that is the result of previous work of the authors at DLR [WIE2023b]. Within TEMA, DLR has further extended the dataset with additional samples that cover particularly challenging environments. Each sample covers an area of 100 x 100 km and consists of triplets of Sentinel-1 and Sentinel-2 satellite images as well as a corresponding manually annotated water mask. The final dataset consists of 118 sample scenes (>160,000 sample tiles with 256 x 256 pixels) and covers a large variety of environmental and atmospheric conditions (**Figure 15**).

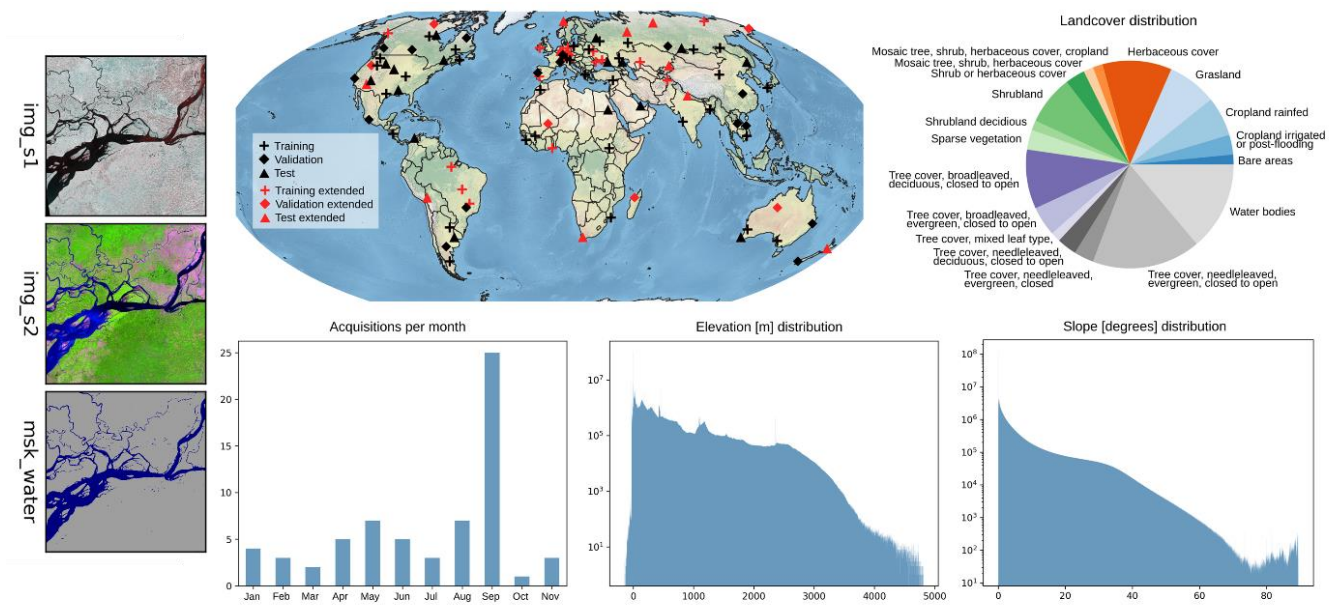



Figure 15: Reference dataset used for training, validation and test of satellite-based flood detection and assessment.

The trained models are deployed in the modular processing chain via the open ONNX (Open Neural Network Exchange) format, which is built to represent machine learning models in a standard across a variety of frameworks, tools, runtimes, and compilers. To improve inference speed in the production environment, DLR has converted the FP32 models to mixed-precision FP16 models, which effectively reduces model size and complexity without affecting the accuracy of the results. **Figure 16** shows a comparison of inference throughput across a range of batch sizes. With a throughput of 28 MP/s for the best performing model (FP16 at batch size 8), DLR is able to analyse a single Sentinel-2 image (covering an area of 100 x 100 km at a spatial resolution of 10 m) in less than 10 seconds on a consumer grade GPU (not including data download, pre-processing and publishing of results).








Figure 16: Comparison of ONNX model inference throughput for different model precisions. Measured on a NVIDIA RTX A4000 GPU using ONNX-Runtime "CUDAExecutionProvider" across a range of batch sizes.

DLR tested the preliminary version of “TFA-tech-08: Satellite-based flood detection and assessment” with above described DNN models in a “live-scenario” during the Thessaly floods in Greece September 2023 for which Copernicus EMS has been activated (<https://rapidmapping.emergency.copernicus.eu/EMSR692>). **Figure 17** shows results from the satellite-based flood monitoring with examples of the delineation of permanent and temporary flooded areas around the city of Larissa on 2023-09-09. DLR also computed a flood duration product [MAR2019] for the whole monitoring period (2023-09-01 until 2023-10-04). In total DLR processed 31 Sentinel-1 and 86 Sentinel-2 images for this disaster, covering an area of more than 35,000 km².



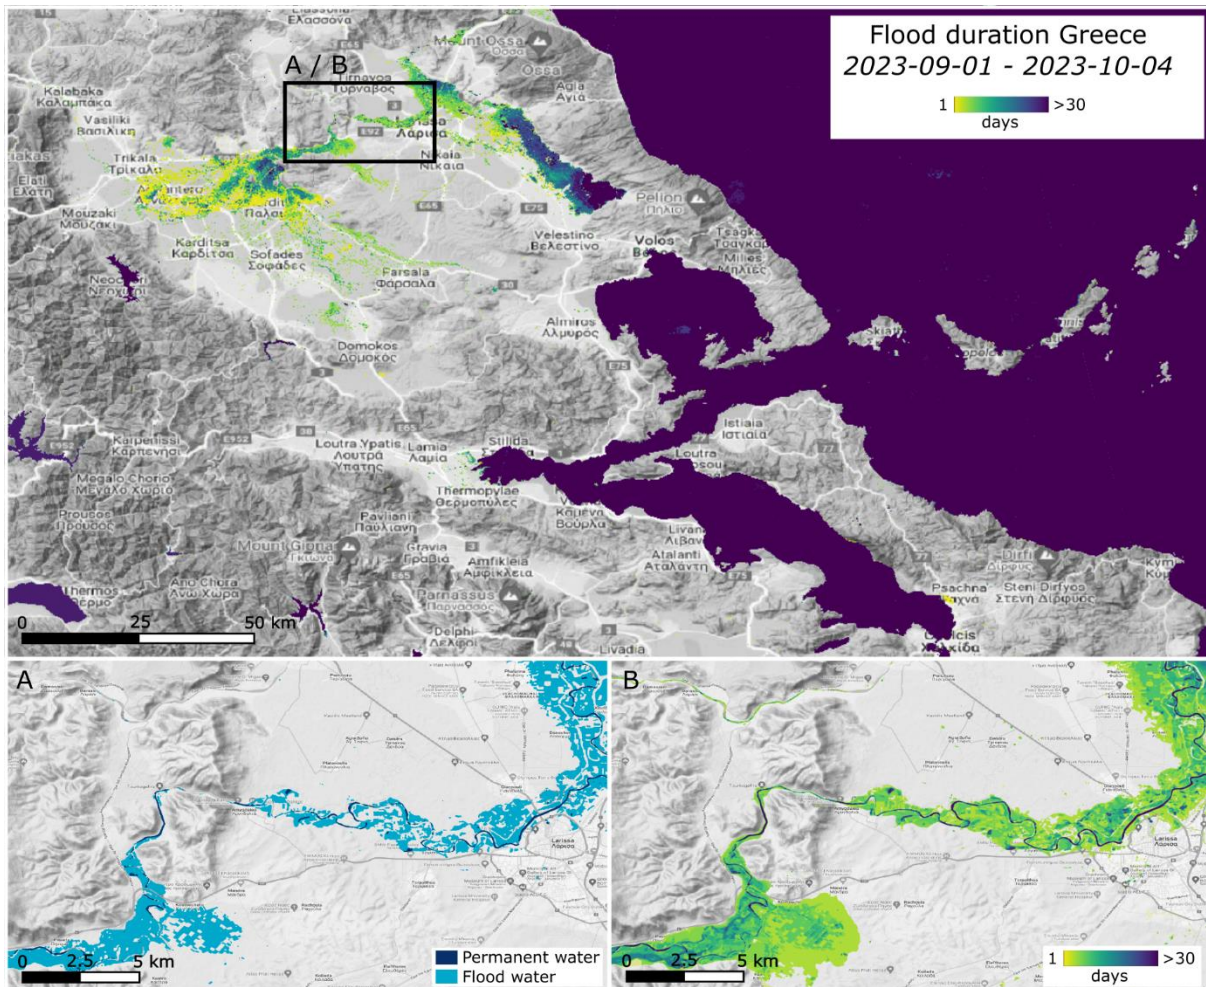





Figure 17: Results from satellite-based monitoring of the Greece floods 2023 with examples showing inundated areas on 2023-09-09 (A) and flood duration from 2023-09-01 until 2023-10-04 (B) around the city of Larissa.

Objectives and KPIs that are linked to this technology include objectives **OA2** “Increase accuracy of extreme data analysis algorithms” and **OA3** “Increase responsiveness / speed of extreme data analysis algorithms”. The preliminary version of “TFA-tech-08: Satellite-based flood detection and assessment” could already significantly increase flood segmentation accuracy (**OA2**) and speed of data analysis (**OA3**). Compared to previous CNN models that have been developed by DLR for flood mapping, the most recent advancements within TEMA lead to an increase of segmentation accuracy by **+10%** IoU for Sentinel-1 (radar) and **+4%** IoU for Sentinel-2 (multi-spectral) image analysis (**Table 11**). Moreover, inference throughput (as a measure of the speed of analysis) could be increased by **+10 MP/s** compared to previous models. With an approximate analysis time of less than 10 seconds on a consumer grade GPU (not including data download, pre-





processing and publishing of results), this is significantly faster than any manual or even semi-automated methods deployed for example by the Copernicus EMS.

Further developments within TEMA will focus on improving speed and efficiency of data access, pre-processing and dissemination for the satellite-derived water masks. Moreover, scientific studies and experiments related to object-detection in orthoimages from drone surveys will be conducted.


4.3.2 Satellite-based fire detection and assessment

SOTA

To this point, remote sensing based burnt area mapping is either done in a semi-automatic way, requiring the manual interaction of an expert in the field, or via a time-series analysis. The latter, however, requires imagery before and after a wildfire event over a timespan of at least several weeks, thus preventing the use of the results in time-critical applications. Additionally, available burnt area data usually only comprises the final fire perimeter, and the starting date. Valuable information, such as the progression of the burnt area over time and the burn severity is usually not included.

Advances beyond SOTA

DLR maintains and operates a modular processing chain for automated burn area derivation from optical satellite imagery. It is sensor agnostic, meaning it can utilize any given optical sensor as long as a band in the near-infrared and red domain is present, respectively. The significantly decreases the duration between overpasses and allows for a distribution of results in near-real time. The system is designed to primarily utilize mid-resolution data, allowing the results to be updated four times a day (2x Sentinel-3 A/B, 2x Aqua/Terra MODIS sensors). Additionally, high-resolution data (specifically Sentinel-2 MSI) is used whenever it gets available. The monitoring region comprises all of Europe, with additional areas of interest in typically fire-prone areas such as Canada, Chile, and Australia. The system is built upon a Graph Convolutional Network (GCN), which identifies fire-affected regions through a pixel-cluster based node classification. Pytorch Geometric is used for the training and inference process. **Figure 18** shows a sample result from a generated Sentinel-3/Sentinel-2 dataset covering the Canadian Donnie Creek region for summer 2023.








Figure 18: Sample from the generated Donnie Creek wildfire dataset (derived from Sentinel-3). The burn severity attribute is used for visualization.


This development is related to TEMA objectives **OA2** “Increase accuracy of extreme data analysis algorithms”) and **OA3** “Increase responsiveness / speed of extreme data analysis algorithms”. Continuous, large-scale wildfire data for Europe is available through the European Forest Fire Information System (EFFIS), which is based on a semi-automated way and is manually updated 2x per day (no updates are available outside office hours). No auxiliary information, such as burn severity, is available. The presented method is available globally, updated at least 4 times per day and features a wealth of relevant auxiliary information. Further activities in TEMA include the support of additional sensors to further increase the update rate.

4.4 Multimodal analysis for the construction of 3D smoke concentration maps

SOTA

In TEMA the key focus of the smoke and wind modeling is to understand the dispersion of airborne substances, particularly forest fire smoke, at initial as well as later stages of the fire. Naturally, an accurate representation of airflow is of key importance here: airflow is often the dominant factor that is responsible for spatial material transport. This is important since accurate airflow information is crucial for evacuations and response planning; during forest fires, real-time wind maps can support fire-fighting activities on the ground [BAT2019].





Existing smoke plume models, however, are highly focused on model accuracy. The assumptions of these models involve making them unsuited to real-time modelling. DLR directs the curious reader to [PET2022] for an in-depth explanation of the state-of-the-art algorithms for smoke plume dynamics.

Advances beyond SOTA

DLR aims to improve on the state-of-the-art solutions by providing more simplified model of the dynamics of smoke plumes, which is able to adapt to measurements taken from the fire. DLR plans on bringing configured hardware to validate and inform its models. TEMA objectives **OB1** (related to reduce latency) and **OC3** “Reduce mental load of end users”, both are achieved by such a model which adapts to data collected in the area.

Up until M18, DLR has made decisions about its hardware platform and the numerical models used to model the smoke plume. In great part because smoke-plume dynamics are a novel field of research and there exists little data to validate models, DLR is working first on building a framework to assess the accuracy of its models to achieve objective **OB2** “Increase Model-based prediction accuracy”. **Figure 19** shows an example of validation of trace gas dispersion models using wind tunnel data under simplified uniform flow conditions.

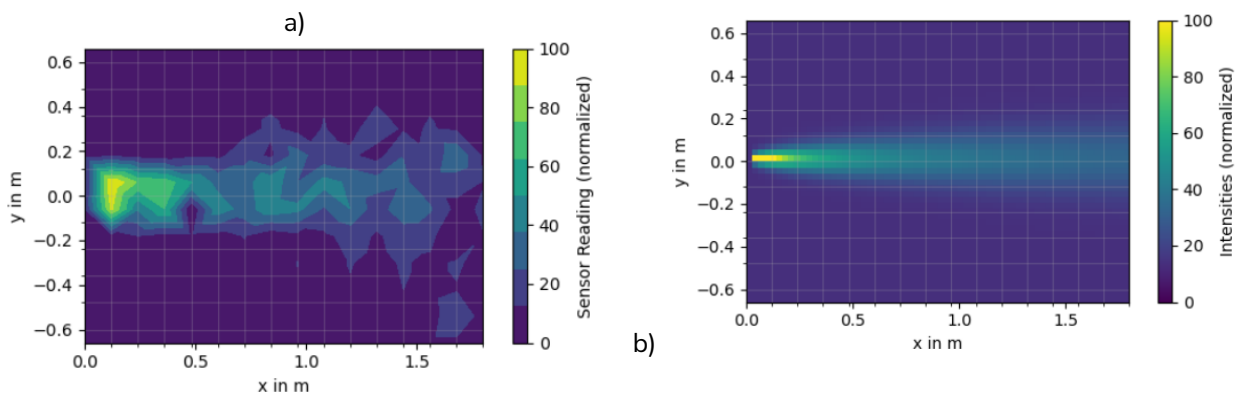




Figure 19: Data collected in a wind tunnel to validate trace gas dispersion models under simplified uniform flow. Fig. a) represents the measured smoke plume after reaching equilibrium, and b) the model [HIN2024].





5 Social media and text semantic analysis

5.1 Introduction


During the M7-M30 period task T3.3 “Social media and text semantic analysis” aims to develop advanced semantic analysis algorithms for geosocial media and news posts, as described in Subsection “Geosocial Media and News Analysis” of Section 1.2.2.A of Part B of TEMA’s Description of the Action (DoA). This task involves PLUS, AUTH, DLR, ATOS, ATOS SP, and FHHL. PLUS leads the task, focusing on topic identification, multilingual handling, and relevance classification. AUTH contributes by researching sentiment analysis for short texts, such as social media posts. ATOS investigates the use of CLIP methodology to combine images with text snippets. These efforts are expected to enhance the capability to analyze geosocial media and news content accurately and efficiently, supporting the overall project goals.


Social media analysis has become an established technique for aiding efforts in disaster management, stretching over all four phases *mitigation, preparedness, response, and recovery* of the disaster management cycle [ALB2011, PHE2021]. Geo-social media posts with a geographic reference can provide particularly actionable insights for disaster management as they allow for local insights. Historically, microblogging platforms such as X (formerly Twitter) have dominated this field of research due to their wide availability via Application Programming Interfaces (APIs). The vast amount of posts available on these platforms allow for the collection of statistically robust geographically referenced datasets [SER2023]. Commonly used techniques to analyze such data include semantic and sentiment analysis of texts along with spatial and temporal analysis methods. PLUS has developed and is working on methods touching upon all four modalities: Semantics, sentiments, space, and time.

5.2 Semantic topic modelling

SOTA

Topic modelling is a central technique to unveil latent semantic topics within large collections of documents in an unsupervised manner. More generally, it can be described as a method of content analysis for text which involves assigning one or more “topic codes” to each document. One of the main contributions to this field was made by [BLE2003] with Latent Dirichlet Allocation (LDA). It operates under a bag-of-words hypothesis and utilises probabilistic techniques to infer semantic topics. LDA’s generative concept has been extended to incorporate various other attributes such as sentiments, time and social media user behaviour. However, LDA neglects the relationships between individual words within a document. To address this limitation, newer embedding-based topic modelling approaches have emerged. For instance, [DIE2020] introduced






the Embedding Topic Model (ETM), which merges LDA's generative concept with pre-trained word and sentence embeddings. Additionally, [SIA2020] demonstrated the efficacy of clustering high-dimensional semantic embedding vectors as a potent topic modelling strategy, performing comparably to LDA. Further advancements were made by [ANG2020] with the proposal of Top2Vec which works on a clustering basis powered by Doc2Vec, UMAP, and HDBSCAN. Grootendorst [GRO2022] enhanced the technique, introducing BERTopic, which consistently outperformed traditional LDA and Top2Vec.


Advances beyond SOTA

PLUS aims to improve upon the SOTA by integrating additional modalities and modern language models such as Llama-2 and multilingual BERT models in the semantic topic modelling procedure. As a result, they developed a multimodal spatio-temporal joint topic-sentiment model that is capable of identifying geographically local sentiment-associated semantic topics with a temporal reference. The method is significantly more fine-grained than classic topic modelling techniques. The output was evaluated on a use-case basis where significantly more information useful to emergency responders could be extracted when compared to the semantic topic modelling techniques used so far. A conference publication regarding this development was accepted for AGILE 2024 and another journal article is currently under review. The technique also utilizes Llama-2 to extract human-readable topic labels and emergency-relevant information which has not been done before. The efforts align with **OA2**, “Increase accuracy of extreme data analysis algorithms” but require a qualitative evaluation as no comparable baseline models exist.

5.2.1 Multilinguality handling SOTA

Traditional topic models such as LDA are not particularly capable of handling mixtures of different languages due to the underlying bag-of-words hypothesis which disregards the meaning of words. Newer techniques such as BERTopic, in contrast, use Large Language Models (LLMs) for the computation of a numeric representation of the semantic meaning of each text. Consequently, they can be configured to use multilingual language models for processing. Such multilingual LLMs are usually trained for multilingual data using knowledge distillation [REI2020]. Texts with the same meaning in different languages are then mapped to an identical or similar numeric representation.





Advances beyond SOTA


PLUS developed an integrated topic modelling approach for the analysis of social media posts called the Joint Topic-Sentiment (JTS) model [HAN2024]. It is capable of producing sentiment-associated semantic topics, given posts in up to 100 languages. The underlying working principle is based on a clustering approach of semantic and sentiment embedding vectors. **Table 12** shows the outputs of PLUS's model compared to previous approaches for joint topic-sentiment modelling. PLUS's approach outperformed the adversarial approaches in terms of semantic topic coherence (TC), topic diversity (TD) and overall topic quality (TQ). In particular, the developed framework produced clusters with semantic topic quality scores of up to **0.23** while the best score among the previous approaches was **0.12**. The sentiment classification accuracy increased from **0.35** to **0.72** and the uniformity of sentiments within the clusters reached up to **0.9** in contrast to the baseline of **0.56**. Within TEMA, the method supports **OA2** both in terms of a significant increase in sentiment classification accuracy as well much better quantitative semantic topic quality.


Table 12: Performance metrics for 60 clusters computed on the TweetEval [BAR2020] dataset averaged over a total of 10 runs when UMAP was used for dimensionality reduction for the JTS variants. The sentiment statistics marked with * were computed independently from the topic model and k denotes the number of keywords used for the computation of TC and TD. The best scores for each metric are highlighted in bold [HAN2024] .

	$k = 10$			$k = 25$			S-EMR	SU	SC	DBI
	TC	TD	TQ	TC	TD	TQ				
JST	0.22	0.58	0.12	0.15	0.48	0.07	0.35	1.00	-	-
TSWE	0.22	0.50	0.11	0.16	0.40	0.07	0.35	1.00	-	-
BERTopic	0.18	0.69	0.12	0.12	0.46	0.05	0.72 *	0.56 *	-	-
JTS ^(k-means)	0.31	0.58	0.18	0.22	0.48	0.11	0.72	0.89	0.33	1.02
JTS ^(GSOM)	0.31	0.56	0.17	0.21	0.48	0.10	0.72	0.90	0.31	1.02
JTS ^(HDBSCAN)	0.32	0.73	0.23	0.19	0.60	0.11	0.72	0.77	-0.09	1.32

5.2.2 Post relevance classification/assessment SOTA

A major difficulty in relevance classification for geo-social media posts is the large amount of unstructured data. Traditionally, this was achieved through keyword-based filtering. Due to linguistic and grammatical diversity, however, this leads to blurring in the filtering, e.g., with regard to semantically ambiguous words. Several machine-learning-based approaches have also been proposed for relevance classification of social media content. For example, [DER2018] use a CNN as a first step to classify Tweets as informative or uninformative. In a second step, they sort the





informative messages into eight categories of actionability using a SVM. [MAD2019] investigate stacking a CNN and an Artificial Neural Network (ANN) to classify Tweets on Hurricane Harvey as informative or non-informative. [PAP2023] suggest using a GNN to combine textual information, imagery content and time for flood-related Tweets. However, most of these approaches refer to English-language Tweets, for which there is a large bias in NLP applications. This can cause problems if a language is morphologically more complex [HOV2021].


Advances beyond SOTA


To overcome this limitation, PLUS developed a machine learning algorithm for relevance classification of flood-related Tweets. PLUS compared the performance its approach - a fine-tuned Bidirectional Encoder Representations from Transformers (BERT) model - to more traditional, established models including Naïve Bayes (NB), Random Forest (RF), Support Vector Machine (SVM) and a tailored Convolutional Neural Network (CNN).

PLUS’s BERT-based approach was by far superior to all other approaches, leading other methods by 3-9 percentage points in the Gaussian score and 10-19 percentage points in the F1 score (**Table 13**). PLUS’s model achieved a performance on par with the SOTA. However, PLUS was able to prove that the misclassified texts were generally classified in semantically similar categories. The results therefore directly address topic identification accuracy in social media/news posts within **OA2** of TEMA.

Table 13: Comparison of results for different relevance classification models (*P* = precision, *R* = recall, *F1* = F1 score, *GS* = Gaussian Score; NB: Naive Bayes, RF: Random Forest, SVM = Support Vector Machine, CNN = COnvolutional Neural Network, BERT = Bidirectional Encoder Representations from Transformers) [BLO2024]. The best scores for each metric are highlighted in bold.

Model	Relevance Categories							
	1-Very Relevant				2-Rather Relevant			
	P	R	F1	GS	P	R	F1	GS
NB	0.39	0.32	0.35	0.58	0.38	0.47	0.42	0.74
RF	0.44	0.40	0.42	0.59	0.44	0.50	0.47	0.79
SVM	0.38	0.41	0.40	0.56	0.26	0.22	0.24	0.68
CNN	0.62	0.45	0.53	0.83	0.38	0.44	0.41	0.80
BERT	0.76	0.64	0.69	0.89	0.63	0.69	0.66	0.89
Model	3-Barely Relevant				4-Not Relevant			
	P	R	F1	GS	P	R	F1	GS
	NB	0.38	0.32	0.35	0.73	0.44	0.49	0.46
RF	0.35	0.36	0.36	0.74	0.56	0.51	0.53	0.78





SVM	0.17	0.18	0.18	0.69	0.30	0.31	0.31	0.67
CNN	0.50	0.68	0.58	0.84	0.64	0.47	0.54	0.87
BERT	0.72	0.64	0.68	0.90	0.73	0.86	0.79	0.9


5.3 Sentiment analysis for short texts


5.3.1 AI-based text sentiment analysis

SOTA

Sentiment analysis refers to the study of text author sentiments, as expressed within the text. This task has become particularly valuable for online social media posts, due to the variety of opinionated content posted on these platforms. Usually, texts are categorised into pre-defined classes such as negative/neutral/positive. Sentiment classification can occur at various levels, namely at document, sentence, or aspect ones. Text sentiment classification methods generally fall into three categories: lexicon-based techniques, which utilize pre-defined word or phrase dictionaries with associated sentiment scores, classical Machine Learning methods, such as Support Vector Machines (SVM) or random forests, for supervised learning, and Deep Learning methods, which rely on recurrent, Long Short-Term Memory (LSTM), or transformer-architecture neural networks.

The current state-of-the-art approaches use BERT-based Transformer encoders [DEV2019]. BERT consists of multiple layers, each containing a Multi-Head Self-Attention mechanism which enables the Transformer model to focus on different sentence parts simultaneously, capturing various context aspects. In recent years, researchers focus primarily on multi-label sentiment classification [ALH2021], [HE2018]. Perhaps the simplest approach to this end is by using a single multi-label classifier. This way sentiment interrelations are easily introduced to the text representation as the feature space is the same for all sentiments. Additional constraints can be added via sentiment model related loss functions [HE2018], [DEN2020]. However, Machine Learning models suffer greatly from the sentiment class data imbalances [HE2009] that are prominent in expressed sentiments [HEI2014] [ZEL2000], [TRA2015] due to their different occurrence frequencies. Another approach is to create separate binary DNN models for each sentiment, where the text representation does not consider intrinsic sentiment interrelationships. An obvious limitation of this approach is increased computational cost, since sentiment detection requires the simultaneous inference of multiple DNN for analyzing one text at a time. Nevertheless, it has been shown that such approaches often lead to impressive text sentiment analysis performance surpassing their multi-label classifier equivalents [DES2020]. Interestingly, even human annotators showcase a higher inter-annotator agreement, when asked to separately annotate each emotion [SOS2021].





BERT-based models can also be pre-trained on multiple languages, achieving high performance for multilingual learning tasks, including sentiment classification. **Table 14** shows the F1 scores for sentiment classification tasks in various languages, as reported by [BAR2022]. The multilingual XLM-RoBERTa and Twitter-XLM-RoBERTa models, fine-tuned with monolingual data, significantly outperformed FastText in the respective languages. **Table 14** depicts the mean recall of various models for sentiment classification on the English-language TweetEval dataset [BAR2020], as reported by Loureiro et al. [LOU2022]. The base RoBERTa model and Twitter-RoBERTa pre-trained on a large corpus of tweets achieved notably higher performance.


Table 14: F1 scores for sentiment classification tasks in different languages based on the work of Barbieri et al [BAR2022]. The highest value for each language is highlighted in bold.

	FastText	XLM-RoBERTa	Twitter-XLM-RoBERTa
Ar	45.98	63.56	67.67
En	50.85	68.18	66.89
Fr	54.82	71.98	68.19
De	59.56	73.61	76.13
Hi	37.08	36.60	40.29
It	54.65	71.47	70.91
Pt	55.05	67.11	75.98
Sp	50.06	65.87	68.52

Advances beyond SOTA

AUTH has proposed a novel approach to re-introduce sentiment interrelations, while maintaining multiple independent binary sentiment classification problems, by learning sentiment representatives for each class called sentiment anchors. They are defined in the feature space of each and every binary sentiment classifier. Thus the occurrence of each sentiment instance, can be simply represented by its similarity with each sentiment anchor. This way, AUTH combines the best of both worlds. Each binary classifier uses its own feature space; thus it has good sentiment classification performance and does not suffer from sentiment class imbalances. On the other hand, sentiment interrelations are introduced without referring to the outputs of the other classifiers in the pack, since the sentiment anchors have already been learned and stored in each independent sentiment model.

Figure 20 illustrates the proposed, AnchorBERT method. In this approach, a base encoder extracts contextualized word representations from an input text sentence. Class anchors are then used for each sentiment to enrich these representations via a multi-headed attention module. Finally, a feed-forward neural network projects the enriched representations into the binary task.



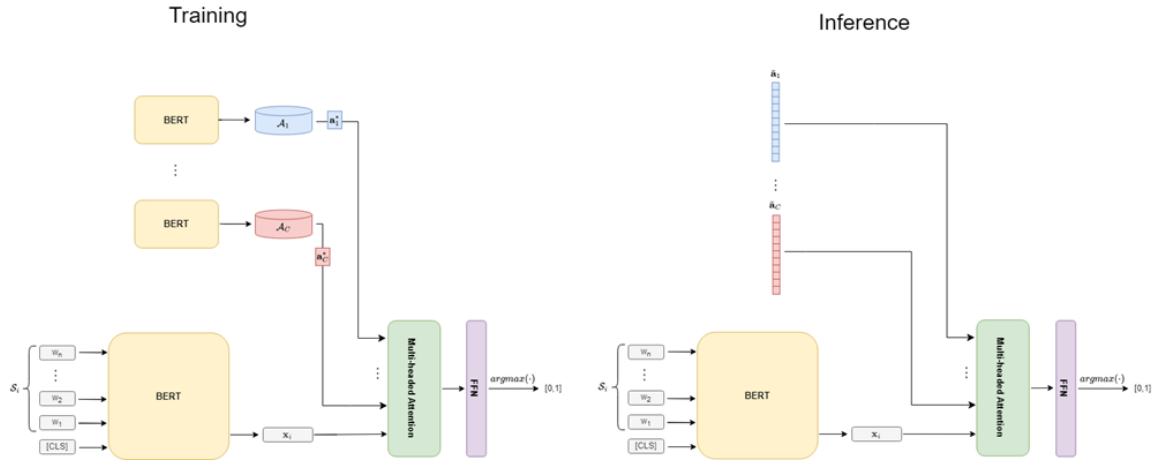



Figure 20: Illustration of the proposed method, AnchorBERT. Given an input sentence, a base encoder is leveraged to extract the contextualized word representation. Next, a class anchor is used for each emotion in order to enrich the text's representation. To this end a multi-headed attention module is employed. Finally, a feed forward neural network projects the enriched representation into the binary task at hand. AUTH illustrates its model both during training (left) and inference (right).


AUTH evaluates its model on four diverse datasets: CovidEMO [SOS2021], HurricaneEMO [DES2020], CancerEMO [SOS2020], and NLPCC 2018 [WAN2018]. Results are presented in **Table 15**, **Table 16**, **Table 17** and **Table 18** respectively.

Table 15: F1 scores for sentiment detection in CovidEMO. The highest value for each sentiment is highlighted in bold.

	ANG	ANT	DIS	FEA	JOY	SAD	SUR	TRU	AVG
AnchorBERT (ours)	0.759	0.584	0.652	0.66	0.747	0.735	0.637	0.553	0.666
Noisy-SSL	0.741	0.554	0.657	0.651	0.741	0.726	0.623	0.532	0.654
CT-BERT	0.735	0.577	0.629	0.644	0.725	0.717	0.617	0.520	0.644

Table 16: F1 scores for sentiment detection in HurricaneEMO. The highest value for each sentiment is highlighted in bold.

	ANG	ANT	DIS	FEA	JOY	SAD	SUR	TRU	AVG
BERTbase	0.676	0.750	0.668	0.683	0.540	0.585	0.557	0.674	0.641
AnchorBERT	0.714	0.749	0.670	0.679	0.557	0.615	0.564	0.691	0.655






Table 17: F1 scores for sentiment detection in CancerEMO. The highest value for each sentiment is highlighted in bold.

	ANG	ANT	DIS	FEA	JOY	SAD	SUR	TRU	AVG
AnchorBERT	0.691	0.845	0.729	0.715	0.831	0.719	0.735	0.652	0.74
BERT _{base}	0.678	0.825	0.72	0.718	0.83	0.724	0.698	0.658	0.731
eMLM	0.69	0.78	0.58	0.77	0.85	0.73	0.68	0.67	0.718

Table 18: F1 scores for sentiment detection in NLPCC 2018. The highest value for each sentiment is highlighted in bold.

	ANG	FEA	HAP	SAD	SUR	F1
AnchorBERT	0.629	0.415	0.777	0.578	0.468	0.573
BERT _{base}	0.613	0.395	0.753	0.531	0.418	0.542
DeepIntell	0.543	0.264	0.734	0.616	0.418	0.515

Text Sentiment Dataset creation

TEMA should process multilingual text data for sentiment analysis (at least English, German, Italian, Greek and Finnish) to cater the TEMA trial needs. To this end this dataset will be used in future multilingual text sentiment analysis experiments.

AUTH collected 766 publicly available “Mastodon” posts in Greek, regarding the Evros wildfires of 2023. Each post was annotated internally zero or more Plutchik-8 emotions. The dataset was split into training, validation, and test sets. Given the relatively small dataset size, experimentation with augmentation techniques yielded the best results, providing a **9%** increase in accuracy over baseline BERT. The accuracy scores for sentiment detection in AUTH Mastodon dataset are presented on **Table 19**.


Table 19: Accuracy scores for sentiment detection in AUTH's Mastodon dataset. The highest value for each sentiment is highlighted in bold.


	ANG	ANT	DIS	FEA	NEU	SAD	SUR	TRU	AVG
BERT _{base}	0.74	0.60	0.73	0.57	0.71	0.59	0.71	0.75	0.68
BERT _{aug}	0.77	0.71	0.77	0.73	0.84	0.74	0.79	0.79	0.77

5.3.2 Consensus-based labelling

SOTA

In order to train effective machine learning models, annotated examples must be provided. A well-annotated training data set is especially important for complex tasks such as language classification, e.g. in sentiment analysis. Historically, single-annotator strategies have often been





used for labelling complex data sets such as tweets, especially in combination with crowd working platforms such as Amazon Mechanical Turk [NAK2016,]. The issue with this approach that the quality of the data largely depends on the annotator, often leading to inconsistencies between different annotators [e.g. IMR2016]. A potential strategy to mitigate these effects to use a labelling strategy based on majority votes [e.g. EPI2010]. Recently, LLMs have also gained attention for assisting labelling using zero-shot classification. However, their performance is still nowhere near the effectiveness of human annotators [WAN2023].

Advances beyond SOTA

PLUS focused on a subcategory of sentiment analysis called emotion analysis, i.e. the assignment of concrete emotions to (parts of) texts. However, due to the complexity of human language, pure assignment at sentence or even text level is often not particularly useful, as a lot of information is lost in the process. A so-called aspect-based emotion analysis approach was pursued by PLUS for TEMA. Aspect terms (individual words or groups of words) to which an emotion relates in the text are identified. This emotion is then classified by a machine learning model.


Thus, a particular focus was placed on creating a high-quality training dataset for aspect-based emotion analysis. For this purpose, a detailed labelling process was carried out in which several human annotators had to manually classify tweets. There were three successive steps: First, they needed to recognize whether one or more emotions were explicitly or implicitly present in the tweet at hand. In a second step, the word or words to which the respective emotion referred had to be marked. In a final step, the emotion had to be classified. For this purpose, a distinction was made between four categories: joy, anger, sadness, and fear.


A labelling guide was created for labelling to ensure consistency between the human annotators. Nevertheless, a consensus decision was made for the model training. This related both to the agreement of the aspect terms and the emotion categories. The resulting training dataset is much more specific than the previous SOTA, especially with regard to the selection of aspect terms.

5.3.3 AI-based text emotion analysis

SOTA

Text emotion classification is a rapidly evolving field with several key approaches. Lexicon-based methods, pioneered by works like WordNet-Affect [STR2004] and NRC Word-Emotion Lexicon [MOH2013], rely on predefined dictionaries that map words to emotions. While these offer simplicity, they struggle with nuance and emotional intensity. Machine learning approaches address this by training algorithms on labelled data. Jain et al. [JAI2017] achieved high accuracy on





multilingual text using Naïve Bayes and so did Hasan et al. [HAS2019] using Support Vector Machines (SVMs). Deep learning techniques are also increasingly used for emotion classification due to their ability to handle complex relationships within text. The respective models [e.g. CHA2019] leverage Long Short-Term Memory (LSTM) networks to capture sentiment and semantic cues for improved emotion detection. Some advances also leverage transfer learning. Transfer learning allows researchers to reuse pre-trained models on new datasets, particularly valuable when dealing with limited labelled data in specific languages. For instance, Ahmad et al. [9] used transfer learning to achieve good performance in Hindi emotion detection by pre-training a model on English datasets.

In general, emotion analysis can be conducted on a document level, sentence level or aspect level. Much of the work conducted so far regarding emotion analysis for short text has been conducted on a document or sentence level [NAN2021].

Advances beyond SOTA


Using the training data described in **Section 5.3.1**, a model based on GRACE [LUO2020] was further developed in order to be able to carry out an aspect-based emotion analysis. Numerous hyperparameter specifications were tested in the model training in order to achieve the best possible fine-tuning. Due to the high specificity of the classification task, the resulting accuracies are relatively low compared to other NLP tasks. However, as there is no comparable SOTA (only for the less specific task of aspect-based sentiment analysis), they represent a very good first step. A master thesis was written on the topic of aspect-based emotion analysis by Christina Zorenböhmer.


For a sentence-level classification, existing SOTA models were used and fine-tuned with TEMA-specific training data. So far, no significant advance to SOTA was made. However, recent developments in deep learning will most likely lead to a large increase in model performance. Their use for the TEMA project still needs to be evaluated, mainly due to high usage costs.

5.4 Spatial hot spot analysis

SOTA

A spatial hotspot is generally an area with higher concentration of a certain events compared to the expected number when a random distribution of those events is assumed. Historically, hotspot detection has evolved from the study of point pattern analysis [CHA1995]. Practically, the density of points within a pre-defined area is compared against a model that assumes spatial randomness. If point density is significantly higher in one area, a spatial hotspot is





identified [GET1992]. Popular techniques to realize such analyses utilize spatial autocorrelation in the form of Moran's I or Geary's C as well as the G_i^* and LISA statistics [ANS1995]. Additionally, a visual representation can be achieved using thematic maps or kernel density estimation. These methods have frequently been used in the context of disaster management; however, they have been rarely evaluated for social media data specifically. Overall, spatial hotspot analysis relies mostly on statistical methods. However, machine learning based methods have also been proposed to predict hotspots in geographic space [ZHA2020].

Advances beyond SOTA


The TEMA project primarily uses spatial hot spot analyses to model the spatial distribution of social media posts. For this purpose, several implementations were developed that are based on a spatial aggregation of classified social media data. These can be, for example, tweets with identical emotions or posts with a disaster reference. Different approaches were compared for the aggregation, e.g. hexagonal grids. The classic Getis-Ord G_i^* method was evaluated as the most suitable for the project purposes. The high temporal resolution of many social media data also enables a spatiotemporal analysis that can make changes visible over different time periods.


As part of the TEMA project, hot spot analyses were also used to evaluate a potential acceleration of the satellite tasking process using social media data. A spatial aggregation of disaster-related social media data was carried out and compared with a baseline. If there was a significant increase in the proportion of disaster-related posts, an alert could be generated, which could enable early identification of significant areas and thus earlier data collection by high-resolution satellites or drones. Such a methodology does not yet exist in the standard routines of remote sensing data retrieval and processing, thus pushing the SOTA.

5.5 Contrastive image-language models

SOTA

Contrastive image-language models are relatively new, which is why there are not yet many studies using it. Image classification tasks, on the other hand, have received a lot of attention in the literature, also with regard to natural disasters. For the most part, CNN-based methodologies to identify images of floods posted on Twitter were developed [BAR2021, MAD2021a). [HAS2022] propose a methodology based on a CNN and transfer learning to perform sentiment analysis on disaster-related imagery. [LI2020] analyze which kind of image characteristics lead to higher user engagement on social media, using Google Cloud Vision to classify the image content.






However, in the media-effective, on-going discourse about ChatGPT, BLIP-2 has also already been proposed as a tool for automatic question answering for visual content [ZHU2023]. [LIU2023c] compare the performance of BLIP-2 against other similar models such as OpenFlamingo, LLaVa or MiniGPT4. [CHE2023] find that BLIP-2 outperforms Flamingo and BLIP considerably for video captioning. [JUH2023] propose a workflow to tag road segments from OpenStreetMap that employs BLIP-2, even though they still achieve a quite low accuracy.


Advances beyond SOTA

At the project's inception, BLIP2 [LI2023] stood out as the sole dependable image-to-text model, leveraging a CLIP image encoder coupled with a Large Language Model (LLM) to produce accurate image captions. ATOS's investigation delved into various branches of BLIP2 models, including those incorporating the Opt LLM (blip2-opt-2.7b, blip2-opt-6.7b) and those utilizing the Flan T5 LLM (blip2-flan-t5-xl, blip2-flan-t5-xxl). Across the board, these models provided similar results in generating brief descriptions and responding to general inquiries pertaining to input images.

Shortly thereafter, the introduction of OpenFlamingo [AWA2023] heralded a significant advancement. While still employing a CLIP image encoder, OpenFlamingo integrated the MPT-1B and MPT-7B LLMs. Notably, OpenFlamingo surpassed its predecessors by enabling more nuanced queries about input images and furnishing more elaborate descriptions compared to BLIP2.

More recently, LLaVA [LIU2023] has emerged as a game-changer in the field. Utilizing the Vicuna-7B and Vicuna-13B LLMs, LLaVA markedly enhances the quality and richness of image descriptions relative to BLIP2 and OpenFlamingo. Distinguished as a MultiModal Language Model, LLaVA not only excels in generating descriptions and responding to queries about images but also possesses the capability to infer specific details within images, such as words or numbers, and address follow-up inquiries.





6 TEMA Core: Parallel and Distributed System

The TEMA Core is a Cloud-Edge Continuum platform designed for the execution of extreme data analytics algorithms. UNIME worked on the architecture with the focus on the interplay between big data storage and parallel processing, orchestrated through a well-defined set of services and interfaces. The system employs a microservices approach, utilizing Docker containers to ensure modularity and flexibility. Kubernetes orchestrates these microservices, providing seamless deployment, scaling and management across the system. Alongside, central components contribute significantly to the TEMA ecosystem. An overview of the TEMA Core is shown in **Figure 21**.

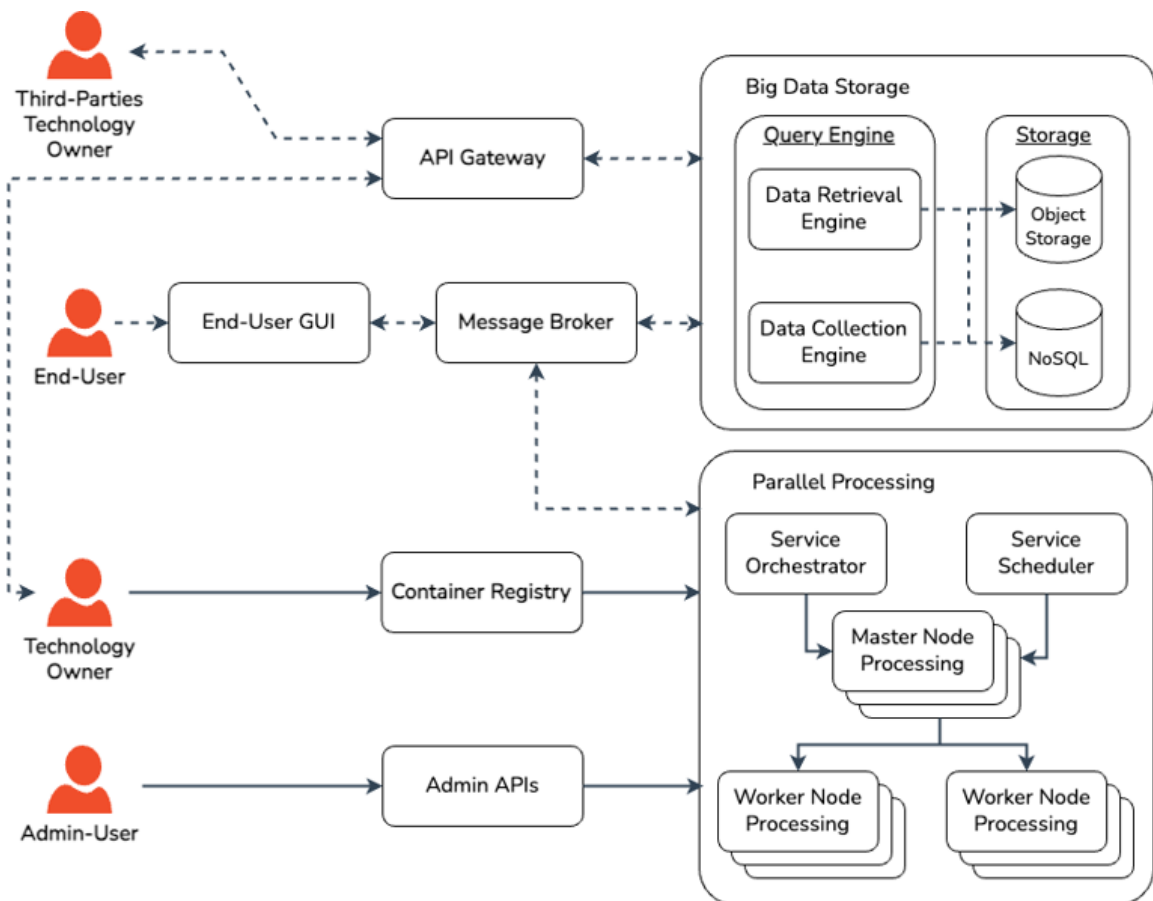




Figure 21: The TEMA Core is a complex infrastructure that includes many actors (e.g., admin-users, end-users, technology owners, third-parties technology owners). The main components are the Big Data Storage and the Parallel Processing units. The infrastructure is designed to allow a federated massive computation of natural disasters data. Moreover, it meets the need of low latency thanks to the use of the Cloud-Edge Continuum paradigm, where the computation may migrate from the cloud to edge resources and vice versa.





The distributed architecture of the TEMA platform adopts the federation paradigm, currently spanning multiple partner such as the University of Messina (UNIME), Aristotle University of Thessaloniki (AUTH), University of Seville (USE), and Paris Lodron University of Salzburg (PLUS).

Furthermore, within this distributed ecosystem, a Docker repository serves as the centralized hub for containerized application management. The cohesive infrastructure, orchestrated by Kubernetes, facilitates seamless coordination of workloads, ensures streamlined application deployment across the network, and efficiently manages computational resources across diverse geographic locations, optimizing both performance and scalability.

To facilitate the development of microservices with the Docker Engine, UNIME wrote the guidelines to dockerize TEMA applications, as described in the Appendix.

6.1 Kubernetes Cluster

A Kubernetes cluster consists of a set of worker machines, known as nodes, that run these applications within pods, which represent single instances of running processes. A pod can house multiple containers sharing the same network namespace, using container images that are created, stored in a registry, and then referenced in the pod.


Kubernetes operates on a distributed architecture with distinct roles for master and worker nodes. Master nodes manage the cluster, while worker nodes execute application workloads, providing necessary resources.


In this architecture, Kubernetes master nodes, currently located only at UNIME, assume the central governance role, overseeing resource allocation and workload distribution throughout the cluster. AUTH, USE and PLUS host Kubernetes worker nodes, leveraging their computational resources to efficiently execute containerized tasks.

6.2 Federation

Kubernetes objects are persistent entities representing the cluster's state, using labels as key/value pairs to specify attributes and enable efficient grouping and selection. Selectors filter objects by labels, facilitating domain and network overlay creation, allowing policies to govern microservice deployment on nodes with specific labels.

In the TEMA cluster, nodes are labeled with attributes like tier (cloud or edge), GPU presence, and the hosting partner's name, such as partner=AUTH for Aristotle University of





Thessaloniki. To ensure applications run exclusively on nodes hosted by AUTH, a Kubernetes node selector is used to specify the `partner=AUTH` label, confining tasks to the institution's nodes.

6.3 Networking

For secure access and easy management of the Kubernetes cluster, a virtual private network (VPN) is setup managed through a server located at UNIME infrastructure. The server is built using OpenVPN, which ensures a trusted and secure connection between nodes, safeguarded from X509v3 certificates able to protect all communications and to ensure data integrity and confidentiality.

Each node in the cluster has an OpenVPN client installed. The nodes are labelled to identify important attributes such as the type of tier (cloud or edge) and the partner project hosting the node.






Figure 22: A virtual private network (VPN) enables easy and secure access to the Kubernetes cluster. Nodes connect to the VPN and receive an IP address with which each node is visible within the private network.

6.4 VM for Easy Cluster Joining

An automated Click & Play Virtual Machine (VM) simplifies the onboarding process by seamlessly connecting to the TEMA VPN and automatically joining the Kubernetes cluster. This VM includes Docker, Git, and a local MongoDB instance pre-configured for development and testing. Users need to fulfill prerequisites such as downloading and installing Oracle VirtualBox to host the





VM, ensuring a stable internet connection for resource access and updates, and meeting hardware requirements like 8GB of RAM and a 4-core CPU on the host machine.

6.5 Deployment of Containerized Applications in a Kubernetes Cluster

Creating Kubernetes objects involves specifying their desired state and details in YAML manifest files. Deployments serve as Kubernetes objects for managing application releases and updates, abstracting pod management for efficient replica handling, scaling based on demand, and supporting version rollback for stability. The Deployment.yaml file outlines parameters such as replica count, container image, resource needs, and configuration data. Applying this manifest file to the cluster initiates pod creation and resource provisioning, illustrated in **Figure 23**, enabling Kubernetes to effectively manage containerized applications with robust scalability and adaptability.

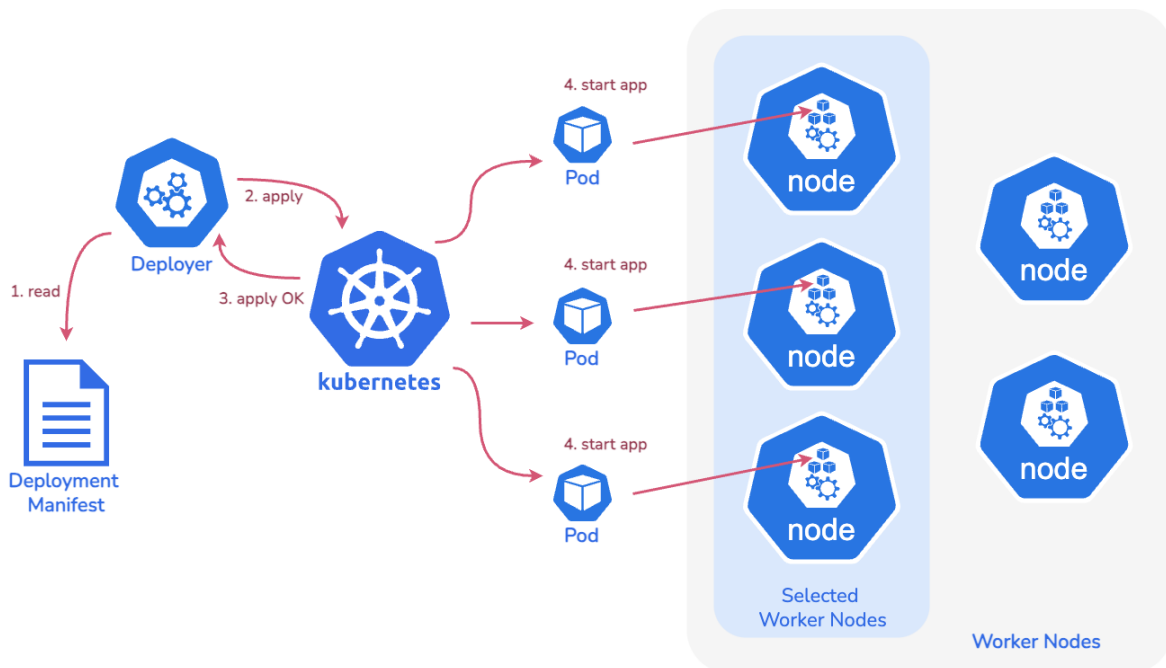




Figure 23: Deployment process in Kubernetes: The deployer reads the deployment manifest and applies the configuration to the Kubernetes cluster. Once confirmed, Kubernetes starts the application pods on selected worker nodes.

6.6 Management Kubernetes Cluster with Rancher

Rancher provides a user-friendly dashboard for administrators to manage Kubernetes clusters efficiently. It offers centralized oversight for deploying, scaling, and upgrading applications,





automating tasks like node provisioning and configuration management. The platform enhances productivity by integrating additional tools and services seamlessly. Administrators benefit from clear visualizations of cluster nodes and pod deployments, enabling real-time monitoring, proactive management, and troubleshooting within their Kubernetes infrastructure.

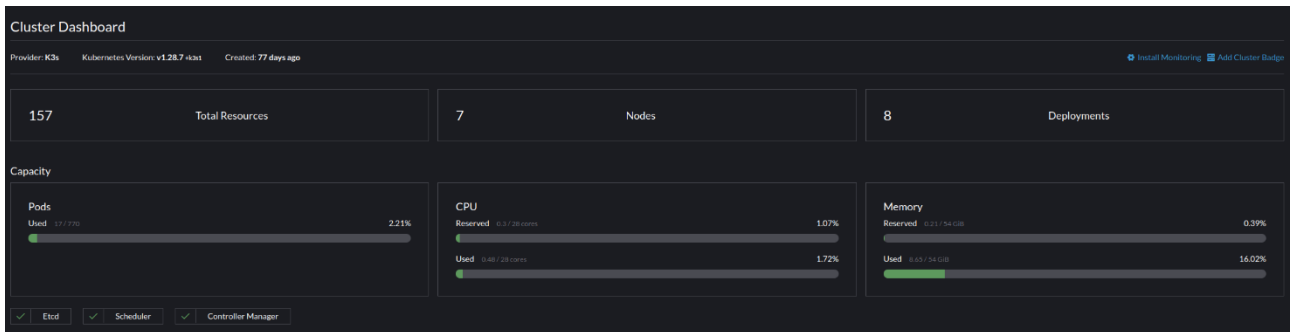




Figure 24: Rancher Cluster Dashboard showing an overview of the cluster's status and capacity. Key metrics include the total resources, number of nodes, and deployments. The capacity section details the usage of pods, CPU, and memory.

Furthermore, Rancher simplifies cluster management by automating tasks such as node provisioning, configuration management, and load balancing. Admin users can seamlessly integrate additional tools and services into their Kubernetes clusters, extending functionality and enhancing productivity.

Rancher offers a clear and concise visualization of cluster nodes, allowing administrators to easily view the status and health of each node within their Kubernetes infrastructure, along with their associated labels. This graphical representation enables quick identification of any issues or anomalies, facilitating proactive management and troubleshooting.






State	Name	Roles	Version	External/Internal IP	OS	CPU	RAM	Pods	Age
Active	auth-sgmr	Worker	v1.28.8+k3s1	~/10.8.0.37	Linux	1.1%	21%	0.91%	49 days
Active	tema-public	Worker	v1.28.7+k3s1	~/10.8.0.1	Linux	6.6%	55%	1.8%	77 days
Active	tema-rasp1	Worker	v1.29.4+k3s1	~/10.8.0.38	Linux	2.4%	37%	0.91%	14 days
Active	tema-rasp2	Worker	v1.29.4+k3s1	~/10.8.0.39	Linux	2.4%	98%	0.91%	14 days
Active	tema-vni1	Control Plane, Etcd	v1.28.7+k3s1	~/10.8.0.17	Linux	16%	66%	5.5%	77 days
Active	tema-vni2	Control Plane, Etcd	v1.28.7+k3s1	~/10.8.0.19	Linux	5.8%	21%	2.7%	77 days
Active	user-csapp	Worker	v1.29.4+k3s1	~/10.8.0.40	Linux	1.1%	21%	0.91%	7 days


Figure 25: Rancher Cluster Nodes view displaying all nodes in the cluster. The dashboard shows whether each node is a master or worker, and provides details on the percentage usage of CPU, RAM, and pods, as well as the operating system running on each node.

With Rancher's intuitive interface, administrators can visualize the deployment and status of pods across their Kubernetes clusters. This visual representation provides valuable insights into the distribution and performance of containerized applications, enabling administrators to efficiently monitor and manage workload deployment and scaling.

6.7 Deployment and Integration of Containerised Application on a K3s Cluster

A K3s cluster was established with three master nodes hosted by a partner institution to ensure high availability. K3s, recognized for its lightweight and efficient Kubernetes distribution, is ideal for resource-constrained environments. Each application deployed on the cluster required specific manifest files: deployment.yaml defined the application's desired state, clust-ip.yaml configured internal networking, and node-port.yaml exposed applications externally. These configurations facilitated seamless communication within and outside the cluster. Node selectors and labels were employed to target pods to specific nodes, optimizing resource utilization and operational efficiency during deployment.







7 Conclusion

Deliverable D3.1 “First report on algorithms for extreme data analytics” is the first Deliverable of the third Work-package (WP3) of the TEMA project. This document reports the initial research results of Tasks T3.1 “Explainable and robust analytics”, T3.2 “Real-time semantic visual analysis and remote sensing”, T3.3 “Social media and text semantic analysis” between M1-M18. The main outputs of the research carried out, were 14 publications (4 journal papers, 10 conference papers) and 18 technical reports, that were submitted to academic journals and conferences. Moreover, 9 technological components were developed, supporting the TEMA platform functionalities described in Deliverable D2.2. These TEMA technologies and the related ones developed in WP4 and WP5 will be interlinked and orchestrated through the developed TEMA Core. This document serves as a summary of the main research outputs and serves as a reference point for researchers summarizing the technical challenges of designing novel algorithms for extreme data analytics. Finally, as a public deliverable, it assists in the dissemination of the project’s results to the scientific community.


The techniques developed in T3.1 will be enhanced and integrated into various AI-based algorithms across different tasks within TEMA. Additionally, the XAI outputs from T3.1 will be utilized as input for Task T5.3 “Augmented Reality and rapid visualization”. The advancements from Task T3.2 will also undergo further refinement and primarily serve as inputs for Task T4.3 “Information fusion”. Similarly, the methods developed in Task T3.3 will be refined and utilized as inputs for both Task T4.3 “Information fusion” and T4.4 “Data-fusion-based decision support and process triggering”. Overall, TEMA technology integration will be performed in Task T6.2.






8 References

- [BAR2024] Dilyara Bareeva et al. "Reactive Model Correction: Mitigating Harm to Task-Relevant Features via Conditional Bias Suppression." arXiv preprint arXiv:2404.09601
- [DRE2024b] Maximilian Dreyer et al. "Understanding the (Extra-)Ordinary: Validating Deep Model Decisions with Prototypical Concept-based Explanations." arXiv preprint arXiv:2311.16681
- [DRE2024c] Maximilian Dreyer et al. "PURE: Turning Polysemantic Neurons Into Pure Features by Identifying Relevant Circuits." arXiv preprint arXiv:2404.06453
- [DRE2024a] Maximilian Dreyer et al. "From Hope to Safety: Unlearning Biases of Deep Models via Gradient Penalization in Latent Space." arXiv preprint arXiv:2308.09437
- [YOL2024] Galip Ümit Yolcu et al. "DualView: Data Attribution from the Dual Perspective." arXiv preprint arXiv:2402.12118
- [VIE2024] Johanna Vielhaben et al. "Explainable AI for time series via Virtual Inspection Layers." arXiv preprint arXiv:2303.06365
- [BLE2024] Florian Bley et al. "Explaining Predictive Uncertainty by Exposing Second-Order Effects." arXiv preprint arXiv:2401.17441
- [HED2024] Anna Hedström et al. "A Fresh Look at Sanity Checks for Saliency Maps." arXiv preprint arXiv:2405.02383
- [BEC2024] Sören Becker et al. "AudioMNIST: Exploring Explainable Artificial Intelligence for audio analysis on a simple benchmark." arXiv preprint arXiv:1807.03418
- [DAW2023] Karam Dawoud et al. "Human-Centered Evaluation of XAI Methods." arXiv:2310.07534
- [FRO2023] Annika Frommholz et al. "XAI-based Comparison of Audio Event Classifiers with different Input Representations." arXiv:2304.14019
- [WEB2024] Leander Weber et al. "Layer-wise feedback propagation." arXiv preprint arXiv:2308.12053
- [HED2023] Anna Hedström et al. "The Meta-Evaluation Problem in Explainable AI: Identifying Reliable Estimators with MetaQuantus." arXiv preprint arXiv:2302.07265
- [TIN2024] Christian Tinauer et al. "Explainable concept mappings of MRI: Revealing the mechanisms underlying deep learning-based brain disease classification." arXiv preprint arXiv:2404.10433





[LIU2023] Liu, Shilong, et al. "Grounding dino: Marrying dino with grounded pre-training for open-set object detection." arXiv preprint arXiv:2303.05499 (2023).

[TAN2022] Tang, Raphael, et al. "What the daam: Interpreting stable diffusion using cross attention." arXiv preprint arXiv:2210.04885 (2022).

[SZE2013] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2013). Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199.

[HAT2022] Hathaliya, Jigna J., Sudeep Tanwar, and Priyanka Sharma. "Adversarial learning techniques for security and privacy preservation: A comprehensive review." Security and Privacy 5.3 (2022): e209.

[MUS2019] A. Mustafa, S. Khan, M. Hayat, R. Goecke, J. Shen, L. Shao, Adversarial defense by restricting the hidden space of deep neural networks, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 3385–3394.

[DEN2012] Deng, L. (2012). The mnist database of handwritten digit images for machine learning research [best of the web]. IEEE signal processing magazine, 29(6), 141-142.

[GOO2014] I. Goodfellow, J. Shlens, C. Szegedy, Explaining and harnessing adversarial examples, arXiv 1412.6572 (12 2014).

[KUR2018] A. Kurakin, I. J. Goodfellow, S. Bengio, Adversarial examples in the physical world, in: Artificial intelligence safety and security, Chapman and Hall / CRC, 2018, pp. 99–112.


[DON2018] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, J. Li, Boosting adversarial attacks with momentum, in: 2018 IEEE / CVF Conference on Computer Vision and Pattern Recognition, 2018, pp. 9185–9193.


[MAD2017] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, A. Vladu, Towards deep learning models resistant to adversarial attacks, ArXiv abs/1706.06083 (2017)

[MYG2022] V. Mygdalis, I. Pitas, Hyperspherical class prototypes for adversarial robustness, Pattern Recognition 125 (2022) 108527

[NAE2019] Naeem, M., Jamal, T., Diaz-Martinez, J., Butt, S. A., Montesano, N., Tariq, M. I., ... & De-La-Hoz-Valdiris, E. (2022). Trends and future perspective challenges in big data. In Advances in Intelligent Data Analysis and Applications: Proceeding of the Sixth Euro-China Conference on Intelligent Data Analysis and Applications, 15–18 October 2019, Arad, Romania (pp. 309-325). Springer Singapore.

[MAJ2022] Majid, S., Alenezi, F., Masood, S., Ahmad, M., Gündüz, E. S., & Polat, K. (2022). Attention based CNN model for fire detection and localization in real-world images. Expert Systems with Applications, 189, 116114.





[HER2022] Hernandez, D., Cecilia, J. M., Cano, J. C., & Calafate, C. T. (2022). Flood detection using real-time image segmentation from unmanned aerial vehicles on edge-computing platform. *Remote Sens* 14: 223.

[MAD2021] Madichetty, S., Muthukumarasamy, S., & Jayadev, P. (2021). Multi-modal classification of Twitter data during disasters for humanitarian response. *Journal of ambient intelligence and humanized computing*, 12, 10223-10237.

[CAO2022] Cao, Y., Tang, Q., Xu, S., Li, F., & Lu, X. (2022). QuasiVSD: efficient dual-frame smoke detection. *Neural Computing and Applications*, 34(11), 8539-8550.

[WANG2023] Wang, Chien-Yao, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. "YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors." In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 7464-7475. 2023.

[YAR2023] Yar, Hikmat, Zulfiqar Ahmad Khan, Fath U. Min Ullah, Waseem Ullah, and Sung Wook Baik. "A modified YOLOv5 architecture for efficient fire detection in smart cities." *Expert Systems with Applications* 231 (2023): 120465.

[WU2022] Wu, Zongsheng, Ru Xue, and Hong Li. "Real-time video fire detection via modified YOLOv5 network model." *Fire Technology* 58, no. 4 (2022): 2377-2403.


[TZI2023] Tzimas, Matthaios Dimitrios, Christos Papaioannidis, Vasileios Mygdalis, and Ioannis Pitas. "Evaluating Deep Neural Network-based Fire Detection for Natural Disaster Management." In *Proceedings of the IEEE/ACM 10th International Conference on Big Data Computing, Applications and Technologies*, pp. 1-6. 2023.


[VEN2022] de Venâncio, Pedro Vinícius AB, Adriano C. Lisboa, and Adriano V. Barbosa. "An automatic fire detection system based on deep convolutional neural networks for low-power, resource-constrained devices." *Neural Computing and Applications* 34, no. 18 (2022): 15349-15368.

[ZHAO2023] Zhao, Yian, Wenyu Lv, Shangliang Xu, Jinman Wei, Guanzhong Wang, Qingqing Dang, Yi Liu, and Jie Chen. "Detrs beat yolos on real-time object detection." *arXiv preprint arXiv:2304.08069* (2023).

[CAR200] Carion, Nicolas, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. "End-to-end object detection with transformers." In *European conference on computer vision*, pp. 213-229. Cham: Springer International Publishing, 2020.

[WAN2023] Wang, Xinzhi, Mengyue Li, Mingke Gao, Quanyi Liu, Zhennan Li, and Luyao Kou. "Early smoke and flame detection based on transformer." *Journal of Safety Science and Resilience* 4, no. 3 (2023): 294-304.





[REZ2019] Rezatofighi, Hamid, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. "Generalized intersection over union: A metric and a loss for bounding box regression." In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 658-666. 2019.

[LI2022] Li, Feng, Hao Zhang, Shilong Liu, Jian Guo, Lionel M. Ni, and Lei Zhang. "Dn-detr: Accelerate detr training by introducing query denoising." In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 13619-13627. 2022.

[ICH2022] lichrak. 2022. jhope Dataset. <https://universe.roboflow.com/ichrak/jhope>. <https://universe.roboflow.com/ichrak/jhope> visited on 2023-10-12.

[TOU2017] Toulouse, Tom, Lucile Rossi, Antoine Campana, Turgay Celik, and Moulay A. Akhloufi. "Computer vision for wildfire research: An evolving image dataset for processing and analysis." Fire Safety Journal 92 (2017): 188-194.

[JOC2022] Glenn Jocher, Ayush Chaurasia, Alex Stoken, Jirka Borovec, NanoCode012, Yonghye Kwon, Kalen Michael, et al. 'Ultralytics/yolov5: V7.0 - Yolov5 SOTA Realtime Instance Segmentation'. Zenodo, 22 November 2022. <https://doi.org/10.5281/zenodo.7347926>.

[RIB2023] T. F. R. Ribeiro, F. Silva, J. Moreira and R. L. de C. Costa, "Burned area semantic segmentation: A novel dataset and evaluation using convolutional networks," ISPRS Journal of Photogrammetry and Remote Sensing, vol. 202, pp. 565-580, 2023.


[RON2015] O. Ronneberger, P. Fischer and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," CoRR, vol. abs/1505.04597, 2015.


[PAP2021] C. Papaioannidis, I. Mademlis and I. Pitas, "Autonomous UAV Safety by Visual Human Crowd Detection Using Multi-Task Deep Neural Networks," in 2021 IEEE International Conference on Robotics and Automation (ICRA), 2021.

[KUM2021] S. Kumaar, Y. Lyu, F. Nex and M. Y. Yang, "CABiNet: Efficient Context Aggregation Network for Low-Latency Semantic Segmentation," in 2021 IEEE International Conference on Robotics and Automation (ICRA), 2021.

[XU2023] J. Xu, Z. Xiong and S. P. Bhattacharyya, "PIDNet: A Real-time Semantic Segmentation Network Inspired by PID Controllers," in 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Los Alamitos, CA, USA, 2023.

[ZHO2018] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh and J. Liang, "UNet++: A Nested U-Net Architecture for Medical Image Segmentation," CoRR, vol. abs/1807.10165, 2018.





[GUA2022] Z. Guan, X. Miao, Y. Mu, Q. Sun, Q. Ye, and D. Gao, "Forest fire segmentation from aerial imagery data using an improved instance segmentation model," *Remote Sensing*, vol. 14, no. 13, 2022. [Online]. Available: <https://www.mdpi.com/2072-4292/14/13/3159>

[LEE2023] Y. J. Lee, H. G. Jung, and J. K. Suhr, "Semantic segmentation network slimming and edge deployment for real-time forest fire or flood monitoring systems using unmanned aerial vehicles," *Electronics*, vol. 12, no. 23, 2023. [Online]. Available: <https://www.mdpi.com/2079-9292/12/23/4795>

[QUR2023] S. Qurratulain, Z. Zheng, J. Xia, Y. Ma, and F. Zhou, "Deep learning instance segmentation framework for burnt area instances characterization," *International Journal of Applied Earth Observation and Geoinformation*, vol. 116, p. 103146, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S156984322200334X>

[YU2018] Yu, C., Wang, J., Peng, C., Gao, C., Yu, G., & Sang, N. (2018). Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 325-341).

[PAP2021] Papaioannidis, C., Mademlis, I., & Pitas, I. (2021, May). Autonomous UAV safety by visual human crowd detection using multi-task deep neural networks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 11074-11080). IEEE.

[XU2023] Xu, J., Xiong, Z., & Bhattacharyya, S. P. (2023). PIDNet: A real-time semantic segmentation network inspired by PID controllers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 19529-19539).


[REI2022] A. Reinke, M. D. Tizabi, C. Sudre, M. Eisenmann, T. Radsch, M. Baumgartner, L. Acion, M. Antonelli, T. Arbel, S. Bakas, P. Bankhead, A. Benis, M. J. Cardoso, V. Cheplygina, B. Cimini, G. Collins, K. Farahani, B. Glocker, P. Godau, and A. Noyan, "Common limitations of image processing metrics: A picture story," 04 2022.


[SHA2020] Shamsoshoara, A., Afghah, F., Razi, A., Zheng, L., Fulé, P., & Blasch, E. The flame dataset: Aerial imagery pile burn detection using drones (uavs), 2020. URL: <https://dx.doi.org/10.21227/qad6-r683>. doi, 10.

[ZHA2017] Zhao, H., Shi, J., Qi, X., Wang, X., & Jia, J. (2017). Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2881-2890).

[HE2016] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770-778).

[KRI2017] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2017). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6), 84-90.





[SZE2015] Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., ... & Rabinovich, A. (2015). Going deeper with convolutions. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 1-9).

[HER2022] Hernández, D., Cecilia, J. M., Cano, J. C., & Calafate, C. T. (2022). Flood detection using real-time image segmentation from unmanned aerial vehicles on edge-computing platform. Remote Sensing, 14(1), 223.

[LIAN2023] Liang, Y., Li, X., Tsai, B., Chen, Q., & Jafari, N. (2023). V-FloodNet: A video segmentation system for urban flood detection and quantification. Environmental Modelling & Software, 160, 105586.

[SAZ2019] Sazara, C., Cetin, M., & Iftekharuddin, K. M. (2019, October). Detecting floodwater on roadways from image data with handcrafted features and deep transfer learning. In 2019 IEEE intelligent transportation systems conference (ITSC) (pp. 804-809). IEEE.

[RAH2021] Rahnemoonfar, M., Chowdhury, T., Sarkar, A., Varshney, D., Yari, M., & Murphy, R. R. (2021). Floodnet: A high resolution aerial imagery dataset for post flood scene understanding. IEEE Access, 9, 89644-89654.

[YAN2022] Yang, L., Zhuo, W., Qi, L., Shi, Y., & Gao, Y. (2022). St++: Make self-training work better for semi-supervised semantic segmentation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 4268-4277).

[PAL2022] Pally, R. J., & Samadi, S. (2022). Application of image processing and convolutional neural networks for flood image classification and semantic segmentation. Environmental Modelling & Software, 148, 105285.


[LI2023] Li, C., Li, L., Geng, Y., Jiang, H., Cheng, M., Zhang, B., & Chu, X. (2023). Yolov6 v3. 0: A full-scale reloading. arXiv preprint arXiv:2301.05586


[LIN2014] Lin, T. Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., ... & Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13 (pp. 740-755). Springer International Publishing.

[SHA2023] Shahid, Mohammad, and Kai-Lung Hua. "Fire detection using transformer network." In Proceedings of the 2021 International Conference on Multimedia Retrieval, pp. 627-630. 2021.

[CHI2022] Chien-Yao Wang, Alexey Bochkovskiy, and Hong Yuan Mark Liao. "Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors", 2022.

[JOC2023] Jocher Glenn. "Yolov5 release v7.0.", <https://github.com/ultralytics/ultralytics/tree/main>, 2023





[JOC2022] Jocher Glenn. "Yolov5 release v7.0.", <https://github.com/ultralytics/yolov5/tree/v7.0>, 2022

[SHA2022] Shangliang Xu, Xinxin Wang, Wenyu Lv, Qinyao Chang, Cheng Cui, Kaipeng Deng, Guanzhong Wang, Qingqing Dang, Shengyu Wei, Yuning Du, and Baohua Lai. "Pp-yoloe: An evolved version of yolo.", 2022

[CHU2023] Chuyi Li, Lulu Li, Yifei Geng, Hongliang Jiang, Meng Cheng, Bo Zhang, Zaidan Ke, Xiaoming Xu, and Xiangxiang Chu. "Yolov6 v3.0: A full-scale reloading.", 2023.

[All2024] «<https://aiia.csd.auth.gr/blaze-fire-classification-segmentation-dataset/>».

[HE2015] K. He, X. Zhang, S. Ren, J. Sun, «Deep Residual Learning for Image Recognition,» CoRR, abs/1512.03385, 2015.

[SHA2021] A. Shamsoshoara, F. Afghah, A. Razi, L. Zheng, P. Z. Fulé, E. Blasch, «Aerial imagery pile burn detection using deep learning: The FLAME dataset,» Computer Networks, 193, p. 108001, 2021.

[SZE2015] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, «Rethinking the Inception Architecture for Computer Vision,» CoRR, abs/1512.00567, 2015.

[TAN2019] M. Tan, Q. V. Le, «EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks,» CoRR, abs/1905.11946, 2019.

[DOM2011] Domnori, Elton, Giacomo Cabri, and Letizia Leonardi. "Multi-agent approach for disaster management." 2011 International Conference on P2P, Parallel, Grid, Cloud and Internet Computing. IEEE, 2011.


[ZHO2019] Zhou, Zhi, et al. "Edge intelligence: Paving the last mile of artificial intelligence with edge computing." Proceedings of the IEEE 107.8 (2019): 1738-1762.


[MAL2022] Malka, May, et al. "Decentralized low-latency collaborative inference via ensembles on the edge." arXiv preprint arXiv:2206.03165 (2022).

[CAS2002] Castro, Miguel, and Barbara Liskov. "Practical byzantine fault tolerance and proactive recovery." ACM Transactions on Computer Systems (TOCS) 20.4 (2002): 398-461.

[KIT1998] Kittler, Josef, et al. "On combining classifiers." IEEE transactions on pattern analysis and machine intelligence 20.3 (1998): 226-239.

[TSI2024] Tsigos, K., Apostolidis, E., Baxevanakis, S., Papadopoulos, S., & Mezaris, V. (2024). Towards Quantitative Evaluation of Explainable AI Methods for Deepfake Detection. arXiv preprint arXiv:2404.18649.





[CHA2020] Chamikara, M. A. P., Bertók, P., Khalil, I., Liu, D., & Camtepe, S. (2020). Privacy preserving face recognition utilizing differential privacy. *Computers & Security*, 97, 101951.

[NOU2020] Nousi, P., Papadopoulos, S., Tefas, A., & Pitas, I. (2020). Deep autoencoders for attribute preserving face de-identification. *Signal Processing: Image Communication*, 81, 115699.

[CHA2020] E. Chatzikyriakidis, C. Papaioannidis and I.Pitas, "Adversarial Face De-Identification" in Proceedings of the IEEE International Conference on Image Processing (ICIP), Taipei, Taiwan, September 22-25, 2019.

[CHR2018] P. Chriskos, R. Zhelev, V. Mygdalis and I.Pitas, "Quality Preserving Face De-Identification Against Deep CNNs" in Proceedings of the IEEE International Workshop on Machine Learning for Signal Processing (MLSP), Aalborg, Denmark, September, 2018.

[ELS2018] Elsayed, G. F., Goodfellow, I., & Sohl-Dickstein, J. (2018, September). Adversarial Reprogramming of Neural Networks. In International Conference on Learning Representations.

[WEI2024] Wei, Yufeng. (2024). Review: Recent advances for the diffusion model. *Journal of Physics: Conference Series*. 2711. 012005. 10.1088/1742-6596/2711/1/012005.

[REI2023] Reis, Dillon, et al. "Real-time flying object detection with YOLOv8." arXiv preprint arXiv:2305.09972 (2023).


[ZHA2022] Zhang, Yifu, et al. "Bytetrack: Multi-object tracking by associating every detection box." European conference on computer vision. Cham: Springer Nature Switzerland, 2022.


[BEN2022] Bentivoglio, R., Isufi, E., Jonkman, S.N., Taormina, R., 2022. Deep learning methods for flood mapping: a review of existing applications and future research directions. *Hydrology and Earth System Sciences* 26, 4345–4378. <https://doi.org/10.5194/hess-26-4345-2022>

[BON2020] Bonafilia, D., Tellman, B., Anderson, T., Issenberg, E., 2020. Sen1Floods11: A georeferenced dataset to train and test deep learning flood algorithms for Sentinel-1, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 835–845. <https://doi.org/10.1109/CVPRW50498.2020.00113>

[MAR2019] Martinis, S., Bettinger, M., Wieland, M., Schläffer, S., Böhnke, C., Shakya, H., Nolde, M., Plank, S., Strunz, G., Hess, U., Sharma, V., Jangle, N., Milosch, O., 2019. An automatic system for near-real time flood extent and duration mapping based on Sentinel-1, Sentinel-2, and TerraSAR-X data. Milan.

[MAR2022] Martinis, S., Groth, S., Wieland, M., Knopp, L., Rättich, M., 2022. Towards a global seasonal and permanent reference water product from Sentinel-1/2 data for improved flood mapping. *Remote Sensing of Environment* 278, 113077. <https://doi.org/10.1016/j.rse.2022.113077>





[MAT2017] Mateo-Garcia, G., Gomez-Chova, L., Camps-Valls, G., 2017. Convolutional neural networks for multispectral image cloud masking, in: 2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS). IEEE, Fort Worth, TX, pp. 2255–2258. <https://doi.org/10.1109/IGARSS.2017.8127438>

[VOI2016] Voigt, S., Giulio-Tonolo, F., Lyons, J., Kučera, J., Jones, B., Schneiderhan, T., Platzack, G., Kaku, K., Hazarika, M.K., Czarán, L., Li, S., Pedersen, W., James, G.K., Proy, C., Muthike, D.M., Bequignon, J., Guha-Sapir, D., 2016. Global trends in satellite-based emergency mapping. *Science* 353, 247–252. <https://doi.org/10.1126/science.aad8728>

[WIE2023a] Wieland, M., Fichtner, F., Martinis, S., Groth, S., Krullikowski, C., Plank, S., Motagh, M., 2023a. S1S2-Water: A global dataset for semantic segmentation of water bodies from Sentinel-1 and Sentinel-2 data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 17, 1084–1099.

[WIE2023b] Wieland, M., Fichtner, F., Martinis, S., Groth, S., Krullikowski, C., Plank, S., Motagh, M., 2023b. S1S2-Water: A global dataset for semantic segmentation of water bodies from Sentinel-1 and Sentinel-2 satellite images (v1.0.0).


[WIE2019] Wieland, M., Martinis, S., 2019. A modular processing chain for automated flood monitoring from multi-spectral satellite data. *Remote Sensing* 11, 2330.


[BAT2019] Battiston, S., Friedemann, M., Gascón, D., Viseras, A., Cardil, A., Mendes, M. A. I., ... & Clandillon, S. (2019). HEIMDALL: A technological solution for multi-hazard management Support including wildfires.

[PET2022] Peterson, D. L., McCaffrey, S. M., & Patel-Weynand, T. (2022). *Wildland Fire Smoke in the United States: A Scientific Assessment* (p. 341). Springer Nature.

[NAK2016] Nakov, Preslav, Alan Ritter, Sara Rosenthal, Fabrizio Sebastiani, and Veselin Stoyanov. 2016. 'SemEval-2016 Task 4: Sentiment Analysis in Twitter'. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, edited by Steven Bethard, Marine Carpuat, Daniel Cer, David Jurgens, Preslav Nakov, and Torsten Zesch, 1–18. San Diego, California: Association for Computational Linguistics. <https://doi.org/10.18653/v1/S16-1001>.

[IMR2016] Imran, Muhammad, Prasenjit Mitra, and Carlos Castillo. 2016. Twitter as a Lifeline: Human-annotated Twitter Corpora for NLP of Crisis-related Messages. In *Proceedings of the 10th Language Resources and Evaluation Conference (LREC)*, 1638-1643. Portorož, Slovenia: European Language Resources Association (ELRA).





[IPE2010] Ipeirotis, Panagiotis G., Foster Provost, and Jing Wang. 2010. Quality management on Amazon Mechanical Turk. In Proceedings of the ACM SIGKDD Workshop on Human Computation, 64–67.

[WAN2023] Wang, Zhiqiang, Yiran Pang, and Yanbin Lin. 2023. 'Large Language Models Are Zero-Shot Text Classifiers'. arXiv. <https://doi.org/10.48550/arXiv.2312.01044>.

[STR2004] Strapparava, Carlo, and Alessandro Valitutti. 2004. 'WordNet Affect: An Affective Extension of WordNet'. In Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04), edited by Maria Teresa Lino, Maria Francisca Xavier, Fátima Ferreira, Rute Costa, and Raquel Silva, 1006-1010. Lisbon, Portugal: European Language Resources Association (ELRA). <http://www.lrec-conf.org/proceedings/lrec2004/pdf/369.pdf>

[MOH2013] Mohammad, Saif M., and Peter D. Turney. 2013. 'Crowdsourcing a Word-Emotion Association Lexicon'. arXiv. <https://doi.org/10.48550/arXiv.1308.6297>.

[JAI2017] Jain, Vinay Kumar, Shishir Kumar, and Steven Lawrence Fernandes. 2017. 'Extraction of Emotions from Multilingual Text Using Intelligent Text Processing and Computational Linguistics'. Journal of Computational Science 21 (July):316–26. doi: 10.1016/j.jocs.2017.01.010.

[HAS2019] Hasan, Maryam, Elke Rundensteiner, and Emmanuel Agu. 2019. 'Automatic Emotion Detection in Text Streams by Analyzing Twitter Data'. International Journal of Data Science and Analytics 7 (1): 35–51. doi: 10.1007/s41060-018-0096-z.


[CHA2019] Chatterjee, Ankush, Umang Gupta, Manoj Kumar Chinnakotla, Radhakrishnan Srikanth, Michel Galley, and Puneet Agrawal. 2019. 'Understanding Emotions in Text Using Deep Learning and Big Data'. Computers in Human Behavior 93 (April):309–17. doi: 10.1016/j.chb.2018.12.029.


[AHM2020] Ahmad, Zishan, Raghav Jindal, Asif Ekbal, and Pushpak Bhattacharyya. 2020. 'Borrow from Rich Cousin: Transfer Learning for Emotion Detection Using Cross Lingual Embedding'. Expert Systems with Applications 139 (January):112851. doi: 10.1016/j.eswa.2019.112851.

[LUO2020] Luo, Huaishao, Lei Ji, Tianrui Li, Daxin Jiang, and Nan Duan. 2020. 'GRACE: Gradient Harmonized and Cascaded Labeling for Aspect-Based Sentiment Analysis'. In Findings of the Association for Computational Linguistics: EMNLP 2020, 54–64. Online: Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.6.


[NAN2021] Nandwani, Pansy, and Rupali Verma. 2021. 'A Review on Sentiment Analysis and Emotion Detection from Text'. Social Network Analysis and Mining 11 (1): 81. doi: 10.1007/s13278-021-00776-6.


[CHAK1995] Chakravorty, Sanjoy. 1995. 'IDENTIFYING CRIME CLUSTERS: THE SPATIAL PRINCIPLES'.





- [GET1992] Getis, Arthur, and J. K. Ord. 1992. 'The Analysis of Spatial Association by Use of Distance Statistics'. *Geographical Analysis* 24 (3): 189–206.
- [ZHA2020] Zhang, Xu, Lin Liu, Luzi Xiao, and Jiakai Ji. 2020. 'Comparison of Machine Learning Algorithms for Predicting Crime Hotspots'. *IEEE Access* 8:181302–10.
- [ANS1995] Anselin, Luc. 1995. 'Local Indicators of Spatial Association—LISA'. *Geographical Analysis* 27 (2): 93–115.
- [MAD2021a] Madichetty, Sreenivasulu, and Sridevi Muthukumarasamy. 2021. 'A stacked convolutional neural network for detecting the resource tweets during a disaster'. *Multimedia Tools and Applications* 80: 3927–3949,
- [BAR2021] Barz, Björn, Kai Schröter, Ann-Christin Kra and Joachim Denzler. 2021. 'Finding relevant flood images on Twitter using content-based filters'. *ICPR International Workshops and Challenges* 12666.
- [LI2020] Li, Yiyi, and Ying Xie. 2020. 'Is a Picture Worth a Thousand Words? An Empirical Study of Image Content and Social Media Engagement'. *Journal of Marketing Research* 57 (1): 1–19.
- [HAS2022] Hassan, Syed Zohaib, Kashif Ahmad, Steven Hicks, Pål Halvorsen, Ala Al-Fuqaha, Nicola Conci and Michael Riegler. 2022. 'Visual Sentiment Analysis from Disaster Images in Social Media'. *Sensors* 22 (10): 3628.
- [LIU2023c] Liu, Yuliang, Zhang Li, Hongliang Li, Wenwen Yu, Mingxin Huang, Dezhi Peng, Mingyu Liu, Mingrui Chen, Chunyuan Li, Cheng-lin Liu, Lianwen Jin and Xiang Bai. 2023. 'On the Hidden Mystery of OCR in Large Multimodal Models'.
- [JUH2023] Juhász, Levente, Peter Mooney, Hartwig H. Hochmair, and Boyuan Guan. 2023. 'ChatGPT as a mapping assistant: A novel method to enrich maps with generative AI and content derived from street-level photographs'.
- [ZHU2023] Zhu, Deyao, Jun Chen, Kilichbek Haydarov, Xiaoqian Shen, Wenxuan Zhang and Mohamed Elhoseiny. 2023. 'ChatGPT Asks, BLIP-2 Answers: Automatic Questioning Towards Enriched Visual Descriptions'. 2023.
- [CHE2023] Chen, Aozhu, Ziyuan Wang, Chengbo Dong, Kaibin Tian, Ruixiang Zhao, Xun Liang, Zhanhui Khang and Xirong Li. 'ChinaOpen: A Dataset for Open-world Multimodal Learning'. 2023.
- [MAD2019] Madichetty, Sreenivasulu. and M. Sridevy. 2019. 'Detecting Informative Tweets during Disaster Using Deep Neural Networks'. *Proceedings of the 2019 11th International Conference on Communication Systems & Networks (COMSNETS), Bengaluru, India, 7–11 January 2019: 709–713.*





[DER2018] Derczynski, Leon, Kalina Bontcheva, Kenny Meesters, and Diana Maynard. 'Helping Crisis Responders Find the Informative Needle in the Tweet Haystack'. 2018. Proceedings of the 15th International Conference on Information Systems for Crisis Response and Management, Rochester, NY, USA, 20–23 May 2018.

[PAP2023] Papadimos, Thomas, Stelios Andreadis, Ilias Gialampoukidis, Stefanos Vrochidis, and Ioannis Kompatsiaris 'Flood-Related Multimedia Benchmark Evaluation: Challenges, Results and a Novel GNN Approach'. 2023. Sensors 23: 3767.

[HOV2021] Hovy, Dirk and Shrimai Prabhumoye. 'Five Sources of Bias in Natural Language Processing'. 2021. Lang. Linguist. Compass (15): e12432.

[ALB2011] R. Albtoush, R. Dobrescu, F. Ionescu, "A hierarchical model for emergency management systems," UPB Scientific Bulletin, vol. 73, no. 2, pp. 1-10, 2011.

[ANG2020] D. Angelov, "Top2Vec: Distributed Representations of Topics," arXiv, 2008.09470, 2020.

[BAR2020] F. Barbieri, J. Camacho-Collados, L. Espinosa Anke, L. Neves, "TweetEval: Unified Benchmark and Comparative Evaluation for Tweet Classification," Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 1644-1650, 2020.

[BAR2022] F. Barbieri, L. Espinosa Anke, J. Camacho-Collados, "XLM-T: Multilingual Language Models in Twitter for Sentiment Analysis and Beyond," Proceedings of the Thirteenth Language Resources and Evaluation Conference, pp. 258-266, 2022.

[BLE2003] D. M. Blei, A. Y. Ng, M. I. Jordan, "Latent dirichlet allocation," The Journal of Machine Learning Research, vol. 3, pp. 993-1022, 2003.


[DIE2020] A. B. Dieng, F. J. R. Ruiz, D. M. Blei, "Topic Modeling in Embedding Spaces," Transactions of the Association for Computational Linguistics, vol. 8, pp. 439-453, 2020.


[GRO2022] M. Grootendorst, "BERTopic: Neural topic modeling with a class-based TF-IDF procedure," arXiv, 2203.05794, 2022.

[HAN2024] D. Hanny, B. Resch, "Clustering-Based Joint Topic-Sentiment Modeling of Social Media Data: A Neural Networks Approach," Information, vol. 15, no. 4, Article 4, 2024.

[LOU2022] D. Loureiro, F. Barbieri, L. Neves, L. Espinosa Anke, J. Camacho-Collados, "TimeLMs: Diachronic Language Models from Twitter," Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pp. 251-260, 2022.

[PHE2021] J. Phengsuwan, T. Shah, N. B. Thekkummal, Z. Wen, R. Sun, D. Pullarkatt, H. Thirugnanam, M. V. Ramesh, G. Morgan, P. James, R. Ranjan, "Use of Social Media Data in Disaster Management: A Survey," Future Internet, vol. 13, no. 2, Article 2, 2021.





[REI2019] N. Reimers, I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 3980-3990, 2019.

[SER2023] H. N. Serere, B. Resch, C. R. Havas, "Enhanced geocoding precision for location inference of tweet text using spaCy, Nominatim and Google Maps. A comparative analysis of the influence of data selection," PLOS ONE, vol. 18, no. 3, e0282942, 2023.

[SIA2020] S. Sia, A. Dalmia, S. J. Mielke, "Tired of Topic Models? Clusters of Pretrained Word Embeddings Make for Fast and Good Topics too!," Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1728-1736, 2020.

[BLO2024] E. Blomeier, S. Schmidt, B. Resch, "Drowning in the Information Flood: Machine-Learning-Based Relevance Classification of Flood-Related Tweets for Disaster Management," Information, vol. 15, no. 3, Article 3, 2024.

[DEV2019] Kenton, Jacob Devlin Ming-Wei Chang, and Lee Kristina Toutanova. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding". In Proceedings of NAACL-HLT, pp. 4171-4186. 2019.


[ALH2021] Alhuzali, Hassan, and Sophia Ananiadou. "SpanEmo: Casting Multi-Label Emotion Classification as Span-Prediction". In Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, edited by Paola Merlo, Jorg Tiedemann, and Reut Tsarfaty, 1573–84. Online: Association for Computational Linguistics, 2021. <https://doi.org/10.18653/v1/2021.eacl-main.135>.


[HE2018] He, Huihui, and Rui Xia. "Joint Binary Neural Network for Multi-Label Learning with Applications to Emotion Classification". In Natural Language Processing and Chinese Computing: 7th CCF International Conference, NLPCC 2018, Hohhot, China, August 26--30, 2018, Proceedings, Part I 7, 250–59. Springer, 2018.

[DEN2020] Deng, Jiawen, and Fuji Ren. "Multi-Label Emotion Detection via Emotion-Specified Feature Extraction and Emotion Correlation Learning". IEEE Transactions on Affective Computing, 2020.

[HEI2014] Heiy, Jane E., and Jennifer S. Cheavens. "Back to Basics: A Naturalistic Assessment of the Experience and Regulation of Emotion". Emotion 14, no. 5 (2014): 878.

[HE2019] He, Haibo, and Eduardo A. Garcia. "Learning from Imbalanced Data". IEEE Transactions on Knowledge and Data Engineering 21, no. 9 (2009): 1263–84.





[ZEL2000] Zelenski, John M., and Randy J. Larsen. "The Distribution of Basic Emotions in Everyday Life: A State and Trait Perspective from Experience Sampling Data". *Journal of Research in Personality* 34, no. 2 (2000): 178–97.

[TRA2015] Trampe, Debra, Jordi Quoidbach, and Maxime Taquet. "Emotions in Everyday Life". *PLoS One* 10, no. 12 (2015): e0145450.

[DES2020] Desai, Shrey, Cornelia Caragea, and Junyi Jessy Li. "Detecting Perceived Emotions in Hurricane Disasters". *arXiv Preprint arXiv:2004.14299*, 2020.

[SOS2021] Sosea, Tiberiu, Chau Pham, Alexander Tekle, Cornelia Caragea, and Junyi Jessy Li. "Emotion Analysis and Detection during COVID-19". *arXiv Preprint arXiv:2107.11020*, 2021.

[SOS2020] Sosea, Tiberiu, and Cornelia Caragea. "Canceremo: A Dataset for Fine-Grained Emotion Detection". In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 8892–8904, 2020.


[WAN2018] Wang, Zhongqing, Shoushan Li, Fan Wu, Qingying Sun, and Guodong Zhou. "Overview of NLPCC 2018 Shared Task 1: Emotion Detection in Code-Switching Text". In *Natural Language Processing and Chinese Computing: 7th CCF International Conference, NLPCC 2018, Hohhot, China, August 26--30, 2018, Proceedings, Part II* 7, 429–33. Springer, 2018.


[LI2023] Li, Junnan, et al. "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models." *International conference on machine learning*. PMLR, 2023.

[AWA2023] Awadalla, Anas, et al. "Openflamingo: An open-source framework for training large autoregressive vision-language models." *arXiv preprint arXiv:2308.01390* (2023).

[LIU2023] Liu, Haotian, et al. "Improved baselines with visual instruction tuning." *arXiv preprint arXiv:2310.03744* (2023).


[JOH1984] William B. Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into hilbert space. *Contemporary mathematics*, 26:189–206, 1984






A. Appendix: Related publications and technical reports

No.	Title	Reference	Type
1	Understanding the (Extra-)Ordinary: Validating Deep Model Decisions with Prototypical Concept-based Explanations	[DRE2024b]	Article
2	PURE: Turning Polysemantic Neurons Into Pure Features by Identifying Relevant Circuits	[DRE2024c]	Article
3	From Hope to Safety: Unlearning Biases of Deep Models via Gradient Penalization in Latent Space	[DRE2024a]	Article
4	DualView: Data Attribution from the Dual Perspective	[YOL2024]	Article
5	Explainable AI for time series via Virtual Inspection Layers	[VIE2024]	Article
6	Explaining Predictive Uncertainty by Exposing Second-Order Effects	[BLE2024]	Article
7	A Fresh Look at Sanity Checks for Saliency Maps	[HED2024]	Article
8	Human-Centered Evaluation of XAI Methods	[DAW2023]	Article
9	AudioMNIST: Exploring Explainable Artificial Intelligence for audio analysis on a simple benchmark	[BEC2024]	Article
10	XAI-based Comparison of Audio Event Classifiers with different Input Representations	[FRO2023]	Article
11	Layer-Wise Feedback Propagation	[WEB2024]	Article
12	The Meta-Evaluation Problem in Explainable AI: Identifying Reliable Estimators with MetaQuantus	[HED2023]	Article
13	Explainable concept mappings of MRI: Revealing the mechanisms underlying deep learning-based brain disease classification	[TIN2024]	Article
14	Revisiting One-versus-One Classification for adversarial robustness		Technical Report
15	Adversarial Defense via Adversarial Reprogramming		Technical Report
16	Evaluating Deep Neural Network-based Fire Detection for Natural Disaster Management	[TZI2023]	Article
17	Enhancing Fire Detection with RT-DETR: A Size-Based Loss Weighting Approach		Technical Report
18	Forest Fire Segmentation Metrics and Algorithms		Technical Report
19	BLAZE: A Dataset for Wildfire and Burnt Area UAV Image Classification and Segmentation	[All2024]	Article





20	Real-time flood water segmentation with deep neural networks		Technical Report
21	Flooded house detection for natural disaster management		Technical Report
22	Forest Fire Image Classification Through Decentralized DNN Inference		Technical Report
23	Structured Efficient Self-Attention Showcased on DETR-based Detectors		Technical Report
24	Re-introducing Emotion Interrelations in Multilabel Text Emotion Detection		Technical Report
25	CORE SERVICES FOR MANAGING BIG DATA IN A NATURAL DISASTER		Technical Report
26	Data Operational Driven AI-based Architecture for Natural Disaster Management		Technical Report
27	Supporting the Natural Disaster Management Distributing Federated Intelligence over the Cloud-Edge Continuum: the TEMA Architecture		Technical Report
28	Reactive Model Correction: Mitigating Harm to Task-Relevant Features via Conditional Bias Suppression	[BAR2024]	Article
29	EDGEmergency: A Cloud-Edge Platform to Enable Pervasive Computing for Disaster Management		Technical Report

