MISSION 4
EDUCATION
RESEARCH

Attachment 2.2

ATTACHMENT 2.2

to the Technical - Scientific Report of the Implementing Entity
Reporting period no. 6 (months 21-27)

PRIN 2022 PNRR

Normative and Digital Solutions to Counter Threats during National Election Campaigns (RightNets)

Project Code: P2022MCYCK

CUP: D53D23022340001 - University of Macerata
PI - Prof. Giovanni Di Cosimo

CUP: D53D23022350001 - Sapienza University, Rome
PI - Prof. Francesco Bilancia

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
unIMC UNIVERSITÀ DI MACERATA

MISSION 4
EDUCATION
RESEARCH

Attachment 2.2

Normative and Digital Solutions to Counter Threats during National Election Campaigns (RightNets)

D6.2 Report on KPIs for monitoring digital election campaigns; D6.3 Report on the design of the computational experiments

Project Title	Normative and Digital Solutions to Counter Threats during National Election Campaigns
Project Acronym	RightNets
Call Identifier	PRIN 2022 PNRR - Decreto Direttoriale n. 1409 del 14-9-2022
Grant. no	P2022MCYCK
CUP	D53D23022340001
Start of Project	30 November 2023
Duration	24 months
Work Package	6
Abstract	The report addresses Key Performance Indicators (KPIs) for monitoring digital election campaigns and the design of analytical solutions for studying online political interactions. In this regard, a computational experiment was designed to empirically identify which characteristics of digital content most influence voter engagement during the 2024 European elections. As part of testing this pilot case with the proposed experimental methodology, a dataset was released and 364 Instagram images using deep learning classification techniques were analyzed. Two categories of KPIs were identified: engagement metrics (retweets, likes, shares, comments) aggregated into popularity levels, and predictive model evaluation metrics (accuracy, precision, recall, F1-score, AUC). Sixteen experimental configurations combining four neural feature extractors with four supervised classifiers were tested, with ViT-B16 and SVM achieving optimal performance.
Authors	Paolo Sernani (UNIMC)
PI	Giovanni Di Cosimo (UNIMC) – PI; Francesco Bilancia (Sapienza) – Substitute PI

The project *"Normative and Digital Solutions to Counter Threats during National Election Campaigns – RightNets"* is funded by the European Union – Next Generation EU, under the PRIN 2022 PNRR call, Mission 4, Component 2, Investment 1.1 (P2022MCYCK; CUP D53D23022340001). The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

RightNets

Report on KPIs for monitoring digital election campaigns (D6.2) and the design of the of the computational experiments (D6.3)

Paolo Sernani
Department of Law, University of Macerata, Italy, paolo.sernani@unimc.it

Introduction

This report documents deliverables D6.2 and D6.3 of the PRIN RightNets project, concerning respectively the Key Performance Indicators (KPIs) for monitoring digital election campaigns and the design of analytical solutions for studying online political interactions. During the development phase of the project, it was deemed scientifically more beneficial to focus activities on in-depth computational experimentation rather than on the creation of a web-based dashboard. This methodological choice is justified by the primary objectives of RightNets, which aim to understand the fundamental dynamics of digital election campaigns through the use of artificial intelligence and machine learning techniques.

The construction of a dashboard would require a solid understanding of the relevant metrics and behavioral patterns to be monitored. However, the context of the 2024 European elections represented a unique opportunity to empirically investigate which characteristics of digital content most influence voter engagement, an issue that remains open in scientific literature. The experiments conducted have therefore made it possible to identify the KPIs that are actually significant for monitoring online political communication, providing a solid empirical basis for any future developments in visualization and monitoring tools.

The experimental work focused on producing (and publicly uploading) the RightNets dataset, comprising 10,000 tweets, 411 Facebook posts, and 364 Instagram images collected during the 2024 European election campaign[1]. Through deep learning and automatic classification techniques, the visual and textual characteristics that determine the popularity levels of political content have been identified. This approach directly responds to the transparency and accountability objectives of the RightNets project, providing analytical tools to assess the effectiveness of communication strategies and identify potential anomalies in online interaction dynamics.

The rest of this document is composed of three sections. The first describes the KPIs identified, distinguishing between engagement metrics and predictive evaluation metrics. Next, the experimental protocol adopted is illustrated, specifying the deep learning architectures tested and the classifiers used. The final section discusses the future prospects for applying the methodology developed.

The Selected KPIs

The experiments conducted as part of the RightNets project required the identification and calculation of KPIs that met the objectives of transparency and accountability in digital election campaigns. The selected metrics fall into two main categories: engagement metrics and predictive model evaluation metrics.

[1] The dataset is described in P. Sernani, "Social media and electoral dynamics: A dataset of X and Facebook activity during the 2024 European elections," Data in Brief, vol. 59, p. 111407, 2025, doi: 10.1016/j.dib.2025.111407 and is available, in open-access, at https://github.com/rightnets/rightnets-social-media-activity-dataset

With regard to engagement metrics, the following data were collected for each social media platform. On X, the number of retweets, replies, likes, and mentions for each tweet were monitored, as well as Boolean attributes identifying the type of content. On Facebook, likes, shares, and comments for each post were tracked. On Instagram, likes, comments, and shares were recorded, with particular attention to the visual component of the content. These metrics were chosen based on their ability to quantitatively measure the level of interaction between voters and political content, which is central to assessing communication effectiveness and identifying potential anomalies in dissemination dynamics. These values were then aggregated into a composite indicator of total engagement, calculated as the sum of the interactions received by each post. This overall metric made it possible to categorize the content into three levels of popularity: low, with values between 0 and 1,000 interactions; medium, between 1,000 and 5,000 interactions; and high, over 5,000 interactions. This classification directly responds to the need to identify distinctive communication patterns and understand which content characteristics generate the most engagement, in line with RightNets' electoral transparency monitoring objectives.

The metrics for evaluating predictive models constitute the second group of KPIs. Five standard metrics in the field of machine learning, all micro-averaged, were used to measure the performance of classifiers in predicting the popularity of posts. Specifically, labeling TP as "true positives" (i.e., samples of a class classified correctly), TN as "true negatives" (i.e., samples of other classes classified correctly), and FP and FN as false positives and false negatives, accuracy (1) measures the percentage of correct classifications out of the total number of cases examined.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

Precision (2) quantifies the proportion of correct positive predictions out of the total number of positive predictions. Recall (3) or True Positive Rate (TPR) measures the model's ability to identify all positive cases actually present in the dataset.

$$Precision = \frac{TP}{TP + FP} \quad (2)$$

$$Recall = \frac{TP}{TP + FN} \quad (3)$$

The F1-score (4) represents the harmonic mean between precision and recall, balancing the ability to correctly identify positive cases while reducing false positives.

$$F1\text{-}score = 2\frac{Precision \times Recall}{Precision + Recall} \quad (4)$$

Furthermore, given the False Positive Rate (5), i.e., a measure of the rate of false positives relative to false negatives in the dataset, the Area Under the Curve (AUC) (6) provides an assessment of the model's discriminatory power considering the distribution of classes.

$$FPR = \frac{FP}{TN + FN} \quad (5)$$

$$\int_{-\infty}^{+\infty} \mathrm{TPR}(t) d\mathrm{FPR}(t) \quad (6)$$

These metrics were chosen for their complementarity and their ability to provide a multidimensional assessment of predictive performance, which is essential when dealing with multiclass classification problems.

The design of the experiments

The decision to conduct computational analysis experiments rather than develop a web-based dashboard was based on the need to empirically understand which characteristics of digital content most influence electoral engagement before implementing monitoring tools. A dashboard requires the 'a-priori' definition of consolidated metrics and behavioral patterns to be visualized, while the context of the 2024 European elections provided an opportunity to investigate precisely what these patterns are through machine learning approaches. The designed and conducted experiments therefore pursued the objective of identifying predictive relationships between visual characteristics of content and levels of popularity, providing an empirical basis for possible future developments of automated monitoring systems.

The experimental protocol adopted consists of three main phases. To use the proposed experimental methodology in the pilot case (i.e., the 2024 European Elections), the first phase of the experimental protocol involved preparing the Instagram image dataset, divided into training and test sets according to a 70-30 split, with popularity classes balanced to avoid predictive bias. The second phase involved the automatic extraction of visual features through four pre-trained convolutional neural network architectures: ResNet-50[2], a deep residual network effective in learning hierarchical representations, VGG-16[3], characterized by sequential architecture with small filters, Inception-v3[4], which uses multiple-sized filters to capture features at different scales, and ViT-B16, a Vision Transformer[5] that applies attention mechanisms to image patches. For each architecture, the final classification layer was removed to use the models exclusively as feature extractors, obtaining high-dimensional representation vectors that were subsequently flattened into one-dimensional arrays.

The third phase involved supervised classification of content using four distinct classifiers[6]. Logistic regression (LR) was selected as the linear baseline, particularly suitable for high-dimensional feature spaces and interpretable in terms of coefficients. Support Vector Machines (SVM) with nonlinear kernels were used for their ability to identify optimal separation hyperplanes in transformed spaces. The Gradient Boosting Classifier (GBT) is an ensemble approach that builds sequences of weak classifiers by iteratively correcting errors. The MultiLayer Perceptron (MLP) is a feedforward neural network with hidden layers trained by backpropagation, capable of learning complex nonlinear representations. The systematic combination of four extractors and four classifiers generated sixteen distinct experimental configurations, evaluated using the metrics described in the previous section.

The results obtained by applying the experimental protocol for each combination of selected classifiers with pre-trained networks for feature extraction, included in Table 1 and discussed in detail in the reference

[2] K. He, X. Zhang, S. Ren and J. Sun, "Deep Residual Learning for Image Recognition," 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 770-778, doi: 10.1109/CVPR.2016.90

[3] K. Simonyan e A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in Proc. Int. Conf. Learning Representations (ICLR), San Diego, CA, USA, 2015

[4] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, 2016, pp. 2818-2826, doi: 10.1109/CVPR.2016.308

[5] A. Dosovitskiy et al., "An Image is Worth 16×16 Words: Transformers for Image Recognition at Scale," in Proc. Int. Conf. Learning Representations (ICLR), 2021

[6] The hyperparameters of the classifiers and details regarding training are presented in the publication discussing the experimental results: P. Sernani, A. Cossiri, G. Di Cosimo, and E. Frontoni, "Analyzing digital political campaigning through machine learning: An exploratory study for the Italian campaign for European Union Parliament election in 2024," Computers, vol. 14, p. 126, 2025, doi: 10.3390/computers14040126.

publication[7], show that the combination of the ViT-B16 feature extractor and the SVM classifier achieves the best performance in terms of F1-score, while the VGG16 configuration with MLP achieves the highest accuracy.

Table 1 – Comparison of the performance of the models compared with the designed experiment.

Backbone	*Classifier*	*Accuracy*	*AUC*	*F1*	*Precision*	*Recall*
Inception V3	LR	0.595	0.682	0.541	0.563	0.595
	SVM	0.595	0.687	0.558	0.578	0.595
	GB	0.611	0.68	0.54	0.582	0.611
	MLP	0.603	0.669	0.549	0.567	0.603
VGG16	LR	0.627	0.756	0.566	0.634	0.627
	SVM	0.611	0.754	0.554	0.579	0.611
	GB	0.611	0.688	0.546	0.626	0.611
	MLP	0.643	0.786	0.574	0.745	0.643
ResNet50	LR	0.587	0.671	0.538	0.535	0.587
	SVM	0.556	0.641	0.522	0.513	0.556
	GB	0.627	0.71	0.578	0.623	0.627
	MLP	0.595	0.672	0.547	0.543	0.595
ViT-B16	LR	0.595	0.732	0.559	0.588	0.595
	SVM	0.603	0.739	0.577	0.592	0.603
	GB	0.611	0.691	0.551	0.704	0.611
	MLP	0.611	0.718	0.552	0.61	0.611

Conclusions

The activities documented in this report responded to the need to empirically identify which characteristics of digital content most influence voter engagement, a scientifically open question that is central to the transparency objectives of the RightNets project. The experiments conducted validated a systematic approach based on machine learning for the analysis of visual political communication, testing sixsteen experimental configurations resulting from the combination of four neural feature extractors and four supervised classifiers.

[7] All results are presented in detail in P. Sernani, A. Cossiri, G. Di Cosimo, and E. Frontoni, "Analyzing digital political campaigning through machine learning: An exploratory study for the Italian campaign for European Union Parliament election in 2024," Computers, vol. 14, p. 126, 2025, doi: 10.3390/computers14040126.

The definition of KPIs, divided into engagement metrics and predictive evaluation metrics, made it possible to quantify the communicative effectiveness of content and identify behavioral patterns relevant to monitoring electoral transparency.

The experimental methodology developed can be reused and generalized to different electoral contexts, both national and European, by adapting the data collection criteria and recalibrating the classification thresholds.

Looking ahead, the validated analytical tools provide a solid basis for the future development of automated digital campaign monitoring systems, supporting the definition of evidence-based policies aimed at ensuring fairness, transparency, and accountability in electoral competitions.