



ENERGY 2026

The Sixteenth International Conference on Smart Grids, Green Communications
and IT Energy-aware Technologies

ISBN: 978-1-68558-357-6

March 8th –12th, 2026

Valencia, Spain

ENERGY 2026 Editors

Vivian Sultan, California State University, Los Angeles, USA

Eric Veith, OFFIS e.V. - Oldenburg, Germany

ENERGY 2026

Foreword

The Sixteenth International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies (ENERGY 2026), held between March 8 - 12, 2026, continued the event considering Green approaches for Smart Grids and IT-aware technologies. It addressed fundamentals, technologies, hardware and software needed support, and applications and challenges.

There is a perceived need for a fundamental transformation in IP communications, energy-aware technologies and the way all energy sources are integrated. This is accelerated by the complexity of smart devices, the need for special interfaces for an easy and remote access, and the new achievements in energy production. Smart Grid technologies promote ways to enhance efficiency and reliability of the electric grid, while addressing increasing demand and incorporating more renewable and distributed electricity generation. The adoption of data centers, penetration of new energy resources, large dissemination of smart sensing and control devices, including smart home, and new vehicular energy approaches demand a new position for distributed communications, energy storage, and integration of various sources of energy.

We take here the opportunity to warmly thank all the members of the ENERGY 2026 Technical Program Committee, as well as the numerous reviewers. The creation of such a high quality conference program would not have been possible without their involvement. We also kindly thank all the authors who dedicated much of their time and efforts to contribute to ENERGY 2026. We truly believe that, thanks to all these efforts, the final conference program consisted of top quality contributions.

Also, this event could not have been a reality without the support of many individuals, organizations, and sponsors. We are grateful to the members of the ENERGY 2026 organizing committee for their help in handling the logistics and for their work to make this professional meeting a success.

We hope that ENERGY 2026 was a successful international forum for the exchange of ideas and results between academia and industry and for the promotion of progress in the fields of smart grids, green communications and IT energy-aware technologies.

We are convinced that the participants found the event useful and communications very open. We also hope that Valencia provided a pleasant environment during the conference and everyone saved some time for exploring this beautiful city.

ENERGY 2026 Chairs:

ENERGY 2026 Steering Committee

Eric MSP Veith, Carl von Ossietzky University – Oldenburg, Germany
Dragan Obradovic, Siemens - Corporate Technology, Munich, Germany
Mark Apperley, University of Waikato, New Zealand
Michael Negnevitsky, University of Tasmania, Australia
Vivian Sultan, California State University Los Angeles, USA
Steffen Fries, Siemens, Germany
Philip Odonkor, Stevens Institute of Technology, USA

ENERGY 2026 Publicity Chairs

Francisco Javier Díaz Blasco, Universitat Politècnica de València, Spain
Ali Ahmad, Universitat Politècnica de València, Spain

Sandra Viciano Tudela, Universitat Politecnica de Valencia, Spain
José Miguel Jiménez, Universitat Politecnica de Valencia, Spain

ENERGY 2026

Committee

ENERGY 2026 Steering Committee

Eric MSP Veith, OFFIS e.V. – Oldenburg, Germany
Dragan Obradovic, Siemens - Corporate Technology, Munich, Germany
Mark Apperley, University of Waikato, New Zealand
Michael Negnevitsky, University of Tasmania, Australia
Vivian Sultan, California State University Los Angeles, USA
Steffen Fries, Siemens, Germany
Philip Odonkor, Stevens Institute of Technology, USA

ENERGY 2026 Publicity Chairs

Francisco Javier Díaz Blasco, Universitat Politècnica de València, Spain
Ali Ahmad, Universitat Politècnica de València, Spain
Sandra Viciano Tudela, Universitat Politecnica de Valencia, Spain
José Miguel Jiménez, Universitat Politecnica de Valencia, Spain

ENERGY 2026 Technical Program Committee

Roozbeh Abolpour, Shiraz University, Iran
Kodjo Agbossou, Université du Québec à Trois-Rivières, Canada
Awadelrahman M. A. Ahmed, University of Oslo, Norway
Miltos Alamaniotis, University of Texas at San Antonio, USA
Kamal Al-Haddad, École de technologie supérieure, Montreal, Canada
Ahmed Al-Salaymeh, The University of Jordan, Amman, Jordan
Amjad Anvari-Moghaddam, AAU Energy, Denmark
Mark Apperley, University of Waikato, New Zealand
Parimal Acharjee, NIT Durgapur, India
Paranietharan Arunagirinathan, Clemson University, USA
Ashrant Aryal, Texas A&M University, USA
Adela Bara, Bucharest University of Economic Studies, Department of Economic Informatics and Cybernetics, Bucharest, Romania
Rico Berner, Ambrosys GmbH, Potsdam, Germany
Lasse Berntzen, University of South-Eastern Norway, Norway
Vito Calderaro, University of Salerno, Italy
Hasan Basri Celebi, Hitachi Energy, Sweden
M. Girish Chandra, TCS Research & Innovation, India
Fathia Chekired, CDER-UNED, Algeria
Dana-Alexandra Ciupageanu, National University of Science and Technology POLITEHNICA Bucharest, Romania
Daniele Codetta, University of Piemonte Orientale, Italy
Luigi Costanzo, Università degli Studi della Campania Luigi Vanvitelli, Italy
Fabio D'Agostino, University of Genova, Italy

Thusitha Dayaratne, Monash University, Australia
Payman Dehghanian, The George Washington University, USA
Margot Deruyck, Universiteit Gent - IMEC - WAVES, Belgium
Giovanna Dondossola, RSE, Italy
Virgil Dumbrava, National University of Science and Technology POLITEHNICA Bucharest, Romania
Kevin Ellett, Indiana University, USA
Emin Taner Elmas, IGDİR University, Turkey
Tatiana Endrjukaite, Transport and Telecommunication Institute, Riga, Latvia
Meisam Farrokhifar, Eindhoven University of Technology, The Netherlands
Sébastien Faye, Luxembourg Institute of Science and Technology (LIST), Luxembourg
Wendy Flores-Fuentes, Autonomous University of Baja California, Mexicali, Mexico
Mahmoud Fotuhi-Firuzabad, Sharif University of Technology, Tehran, Iran
Steffen Fries, Siemens, Germany
Vincenzo Galdi, University of Salerno, Italy
Francisco M. Gonzalez-Longatt, University of South-Eastern Norway, Norway
Saman K. Halgamuge, The University of Melbourne, Australia
Miao He, Beijing Institute of Mathematical Sciences and Applications, China
Yunzhi Huang, Pacific Northwest National Laboratory - U.S. Department of Energy, USA
Md. Minarul Islam, University of Dhaka, Bangladesh
Michael Kuhn, Otto von Guericke University Magdeburg, Germany
Tobias Küster, Technische Universität Berlin (TU Berlin), Germany
Sebastian Lawrenz, Clausthal University of Technology, Germany
Duc Van Le, Nanyang Technological University, Singapore
Gerard Ledwich, Queensland University of Technology, Australia
Yiu-Wing Leung, Hong Kong Baptist University, Kowloon Tong, Hong Kong
Qinghua Li, University of Arkansas, USA
Zhenhua Liu, Stony Brook University (SUNY at Stony Brook), USA
Pascal Maussion, Toulouse INP / CNRS LAPLACE, France
HoSoon Min, INTI International University, Malaysia
Hugo Morais, Universidade de Lisboa, Portugal
Fabio Mottola, University of Naples Federico II, Italy
Gero Mühl, Universitaet Rostock, Germany
Hamidreza Nazaripouya, University of California, Riverside, USA
Michael Negnevitsky, University of Tasmania, Australia
Philip Odonkor, Stevens Institute of Technology, USA
Simona Olmi, Consiglio Nazionale delle Ricerche - Istituto dei Sistemi Complessi, Italy
Claudiu Oprea, Technical University of Cluj-Napoca, Romania
Youssef Ounejjar, ETS, Montreal, Canada
Sanjeev Pannala, Washington State University, USA
Pat Pannuto, University of California, San Diego, USA
Thanasis Papaioannou, Athens University of Economics and Business (AUEB), Greece
Marco Pasetti, University of Brescia, Italy
Nilavra Pathak, University of Maryland Baltimore County, USA
Lawrence Pileggi, Carnegie Mellon University, USA
Anand Krishnan Prakash, Lawrence Berkeley National Laboratory, USA
Aditya Priya, Indian Institute of Science, Bangalore, India
Marco Pruckner, Friedrich-Alexander-University Erlangen-Nürnberg, Germany
Venkata Ramakrishna P., Tata Consultancy Services, India

Anuradha Ravi, University of Maryland Baltimore County, USA
Djamila Rekioua, University of Bejaia, Algeria
Huamin Ren, Kristiania University College, Oslo, Norway
Jan Richling, South Westphalia University of Applied Sciences, Germany
Stefano Rinaldi, University of Brescia, Italy
Carsten Rudolph, Monash University, Australia
Angela Russo, Politecnico di Torino, Italy
Eckehard Schöll, Technische Universität Berlin | Institut für Theoretische Physik, Germany
Tomonobu Senjyu, University of the Ryukyus, Japan
S. Senthilraja, SRM Institute of Science and Technology, India
Farhad Shahnia, Murdoch University, Australia
Hussain Shareef, United Arab Emirates University, UAE
Bhim Singh, Indian Institute of Technology Delhi, India
Vijay Sood, Ontario Tech University, Canada
Vivian Sultan, California State University, Los Angeles, USA
Hongbo Sun, Mitsubishi Electric Research Laboratories, USA
Masoud Taghavi, Technical and Vocational University (TVU), Faculty of Noshahr, Iran
Mehrdad Tahmasebi, Islamic Azad University - Ilam Branch, Iran
Saeed Teimourzadeh, EPRA - Engineering Procurement Research Analysis, Ankara, Turkey
Philipp Thies, University of Exeter, UK
Mihai Tirsu, Institute of Power Engineering, Moldova
Tek Tjing Lie, Auckland University Of Technology, New Zealand
Santiago Torres Contreras, Universidad de Cuenca, Ecuador
Graham Town, Macquarie University, Australia
Ngoc Que Anh Tran, University of Nebraska - Lincoln, USA
Quoc Tuan Tran, Paris Saclay University / CEA / INES, France
Navid Vafamand, Shiraz University, Iran
François Vallee, University of Mons, Belgium
Eric MSP Veith, Carl von Ossietzky University, Oldenburg, Germany
Alekhya Velagapudi, University of Pittsburgh's School of Computing and Information, USA
Alexander Wallis, University of Applied Sciences Landshut, Germany
Daniel Wilson, Boston University, USA
Jian Xu, Texas Reliability Entity (Texas RE), USA
Weiwei Yang, Microsoft Research, USA
Sean Yaw, Montana State University, USA
Tianyu Zhang, Autodesk Research, USA

Copyright Information

For your reference, this is the text governing the copyright release for material published by IARIA.

The copyright release is a transfer of publication rights, which allows IARIA and its partners to drive the dissemination of the published material. This allows IARIA to give articles increased visibility via distribution, inclusion in libraries, and arrangements for submission to indexes.

I, the undersigned, declare that the article is original, and that I represent the authors of this article in the copyright release matters. If this work has been done as work-for-hire, I have obtained all necessary clearances to execute a copyright release. I hereby irrevocably transfer exclusive copyright for this material to IARIA. I give IARIA permission to reproduce the work in any media format such as, but not limited to, print, digital, or electronic. I give IARIA permission to distribute the materials without restriction to any institutions or individuals. I give IARIA permission to submit the work for inclusion in article repositories as IARIA sees fit.

I, the undersigned, declare that to the best of my knowledge, the article does not contain libelous or otherwise unlawful contents or invading the right of privacy or infringing on a proprietary right.

Following the copyright release, any circulated version of the article must bear the copyright notice and any header and footer information that IARIA applies to the published article.

IARIA grants royalty-free permission to the authors to disseminate the work, under the above provisions, for any academic, commercial, or industrial use. IARIA grants royalty-free permission to any individuals or institutions to make the article available electronically, online, or in print.

IARIA acknowledges that rights to any algorithm, process, procedure, apparatus, or articles of manufacture remain with the authors and their employers.

I, the undersigned, understand that IARIA will not be liable, in contract, tort (including, without limitation, negligence), pre-contract or other representations (other than fraudulent misrepresentations) or otherwise in connection with the publication of my work.

Exception to the above is made for work-for-hire performed while employed by the government. In that case, copyright to the material remains with the said government. The rightful owners (authors and government entity) grant unlimited and unrestricted permission to IARIA, IARIA's contractors, and IARIA's partners to further distribute the work.

Table of Contents

Postfire Burn Severity and Four-Month Vegetation Recovery Assessment Using NBR and NDVI: A Case Study of the Eaton and Palisades Fires <i>Vivian Sultan and Andrew Pang</i>	1
Climatic Extremes, Vegetation Response, and Fuel Conditions: A five-year analysis of the Palisades and Eaton fire zones in Los Angeles County <i>Vivian Sultan, Diego Sanchez, and Leonardo Ortiz</i>	5
Understanding the Socioenvironmental Impact of LA City Wildfires <i>Vivian Sultan, Sirisha Mahesh, Xuehan Jing, and Christopher Lee</i>	10
From Sparks to Strategy: Geospatial and Predictive Modeling of Wildfire Risk for Grid Hardening <i>Vivian Sultan and Jeremy Contreras</i>	21
Deep Learning Based Non-Intrusive Load Monitoring for Smart Energy Management in Households <i>Parvaneh Zavareh, Warunee Soythong, Basirah Noor, Volker Hoffmann, and Huamin Ren</i>	25
Zero Trust Defense Against Charge Manipulation Attacks in Smart EV Charging Infrastructure <i>Saba Marandi, Danial Jafarigiv, Ribal Atallah, Mohsen Ghafouri, and Chadi Assi</i>	32
A Hierarchy-Focused Algorithm for Drawing Single Line Diagrams of Power Grids <i>Gabriel Ott and Eric Veith</i>	39
Analysis of the Influence of Photovoltaic Production on Grid Voltage using Data from Inverters <i>Isabel Santiago Chiquero, Victoria Arenas Ramos, Victor Pallares Lopez, Miguel Gonzalez Redondo, Rafael Jesus Real Calvo, Olivia Florencias Oliveros, and Gonzalez de la Rosa Juan Jose</i>	45
Metro-Integrated Electrified Campus Shuttles for Green Mobility in Saudi Universities <i>Mohammed Mutwaly, Khaled Alkhamis, and Walied Alfraidi</i>	49

Postfire Burn Severity and Four-Month Vegetation Recovery Assessment Using NBR and NDVI

A Case Study of the Eaton and Palisades Fires

Vivian Sultan

College of Business and Economics
California State University
Los Angeles, USA
email: vsultan3@calstatela.edu

Andrew Pang

College of Business and Economics
California State University
Los Angeles, USA
email: apang5@calstatela.edu

Abstract—This study analyzes the effects of wildfire and postfire vegetation recovery in Southern California through a remote-sensing-based case study of the Eaton and Palisades Fires. Burn-scar area was calculated using Landsat 8/9 satellite imagery from January 6 and January 14, while vegetation regrowth was assessed using imagery from May 30. The Normalized Burn Ratio and its differenced form were used to classify burn severity. Normalized Difference Vegetation Index and its differenced form were used to track postfire vegetation regrowth. Results showed that 76% of the Eaton Fire area and 86% of the Palisades Fire area experienced low to high burn severity. However, vegetation regrowth differed drastically: Only 25% of land burned by the Eaton fire showed signs of regrowth, compared to 74% in Palisades. The findings demonstrate how remote sensing enables postfire monitoring to support ecological assessment.

Keywords—wildfire; ArcGIS; burn severity; vegetation regrowth.

I. INTRODUCTION

Southern California’s landscape has been shaped by fire for centuries. Many native California species have evolved to depend on periodic fire for regeneration, such as serotinous cones and heat-activated seeds [1]. However, climate change, prolonged drought, and urban expansion have increased wildfire frequency and intensity, posing challenges for land management and ecological recovery.

The goal of this study is to assess both fire impact and vegetation recovery following two Southern California fire events during January 2025: the Eaton Fire and the Palisades Fire. Using satellite imagery and spectral indices, this research evaluates the burned area extent and severity, as well as the regrowth of vegetation within a four-month period postfire.

Two primary research questions guide this work.

1. How many acres were affected by each fire, and what were the severity levels?
2. How much vegetation has regrown since the fire, and what were their regrowth levels?

Section 2 outlines the data used for this analysis. Section 3 describes the methodology used for the analysis. Section 4 details the results of the analysis, and Section 5 discusses those results. Section 6 concludes.

II. DATA

We sourced satellite imagery from United States Geological Survey (USGS) EarthExplorer, Landsat 8–9 Operational Land Imager and Thermal Infrared (OLI/TIRS) Sensor Level 2 data, for three dates: January 6 (prefire), January 14 (postfire), and May 30 (vegetation recovery). Each dataset included all 11 spectral bands in tagged image file (TIF) format. We obtained wildfire boundary polygons for both the Eaton and Palisades Fires from Esri’s Living Atlas. We were primarily interested in Band 5 (Near Infrared), Band 7 (Shortwave Infrared 2), and Band 4 (Red), which are used in Normalized Burn Ratio (NBR) and Normalized Difference Vegetation Index (NDVI) calculations.

III. METHODOLOGY

A. Obtain Satellite Imagery

We downloaded Landsat 8/9 OLI/TIRS Level 2 surface reflectance data from USGS EarthExplorer. Three image sets captured prefire, postfire, and vegetation recovery stages (Figure 1–3).

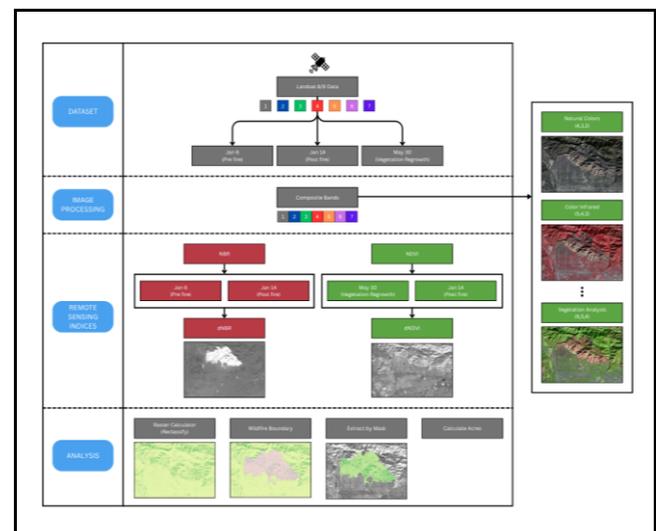


Figure 1. Workflow diagram showing image acquisition, processing, remote sensing indices, and spatial analysis steps.

- January 6, 2025 (Prefire)
- January 14, 2025 (Postfire)
- May 30, 2025 (Regrowth)

Each scene included 11 spectral bands in TIF format. Wildfire boundary shapefiles for the Eaton and Palisades Fires were obtained from Esri’s Living Atlas.

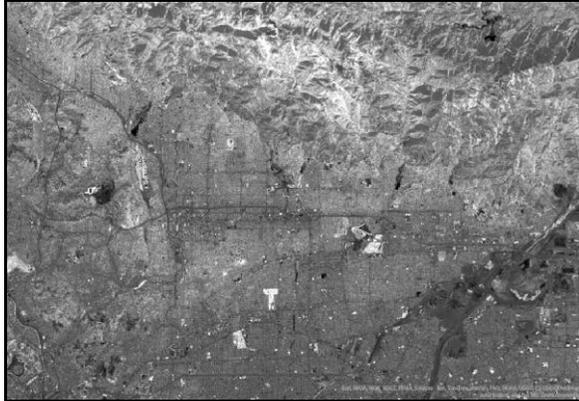


Figure 2. Eaton Fire. NBR of Jan 6th, Prefire

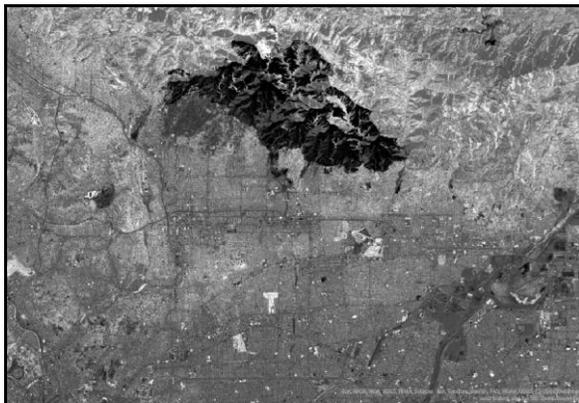


Figure 3. Eaton Fire. NBR of Jan 14th, Postfire

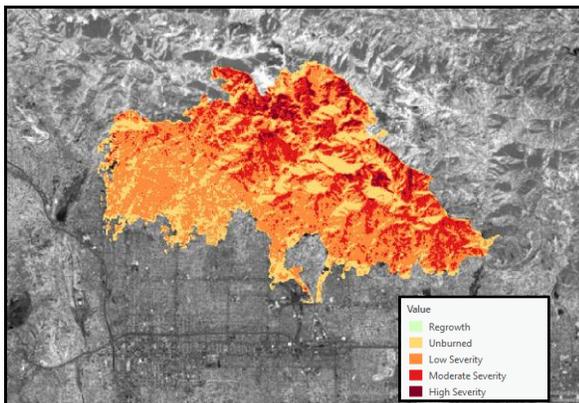


Figure 4. Eaton Fire. Burn Severity Classification.

B. Process Image Bands

Using ArcGIS Pro, we combined individual bands into composite raster images using the Composite Bands tool. To

assist visual interpretation of landscape features and fire damage, we explored three band combinations.

- 4-3-2 (Natural Color)
- 5-4-3 (Color Infrared)
- 6-5-4 (Vegetation Analysis)

We used these visualizations to support manual inspection and confirmation of burn and recovery patterns (Figure 4).

C. Apply Remote Sensing Indices

We used two spectral indices to quantify burn severity and vegetation regrowth, burn severity and vegetation regrowth.

1) Burn Severity

We calculated the NBR for both prefire and postfire imagery using the formula

$$NBR = (NIR - SWIR) / (NIR + SWIR),$$

where *NIR* is Band 5 and *SWIR* is Band 7.

We then computed the differenced NBR (dNBR) by subtracting the postfire NBR from the prefire NBR. We applied classification thresholds using Raster Calculator, following USGS standards for burn-severity levels. This method is widely used in wildfire assessments and is recommended by the UN-SPIDER Knowledge Portal [2].

2) Vegetation Regrowth

We calculated the NDVI for postfire (January 14) and regrowth (May 30) imagery using the formula

$$NDVI = (NIR - Red) / (NIR + Red),$$

where *NIR* is Band 5 and *Red* is Band 4.

NDVI measures photosynthetic activity and plant health [3]. We computed the differenced NDVI (dNDVI) by subtracting postfire NDVI from regrowth NDVI using Raster Calculator to classify regrowth intensity. We based all calculations on surface reflectance data, following USGS best practices [4] (Figure 5–13).

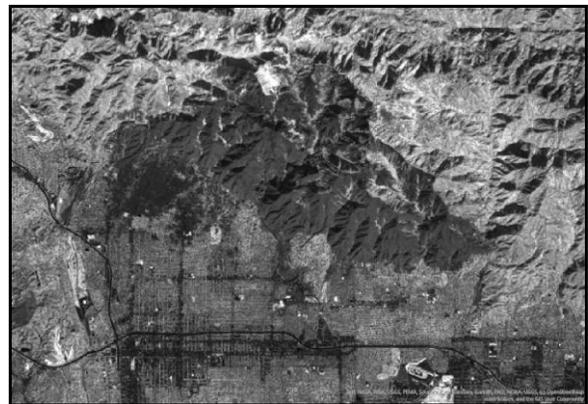


Figure 5. Eaton Fire. NDVI of Jan 14th, Postfire

D. Perform Spatial and Acreage Analysis

To isolate and quantify burned and recovered areas, we used several spatial analysis tools in ArcGIS Pro.

- We applied wildfire boundaries using Extract by Mask.

- We generated classified rasters via Raster Calculator to assign severity and regrowth levels.
- We performed acreage calculations by converting pixel counts (900 m² per pixel at 30m resolution and 1 acre is 4046.86 m²) using the formula

$$\text{Area (acres)} = \text{Pixel Count} \times 900 / 4046.86,$$

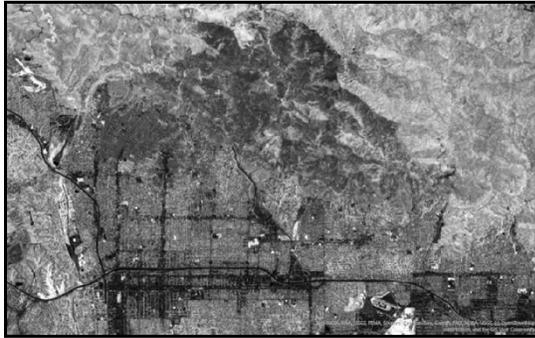


Figure 6. Eaton Fire. NDVI of May 30th

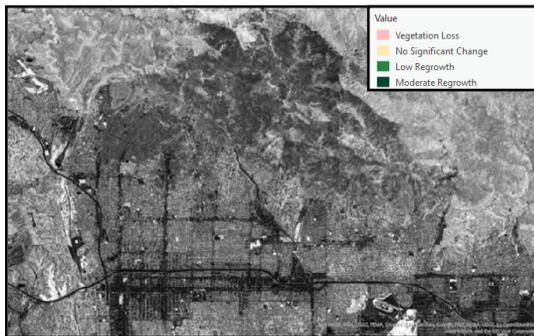


Figure 7. Eaton Fire. Vegetation Regrowth Classification.

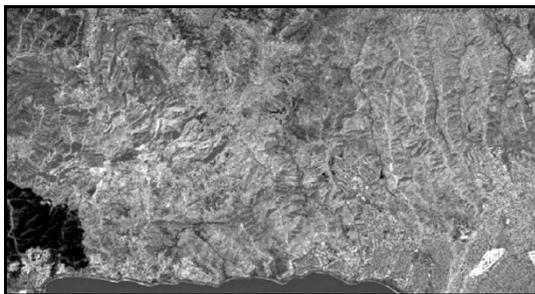


Figure 8. Palisades Fire. NBR of Jan 6th, Prefire

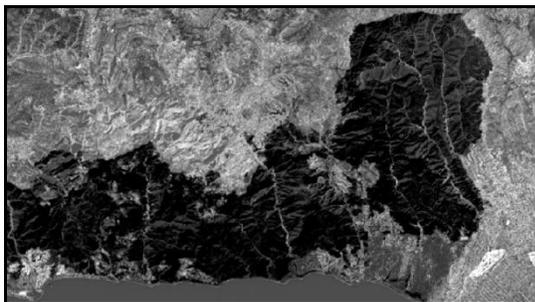


Figure 9. Palisades Fire. NBR of Jan 14th, Postfire

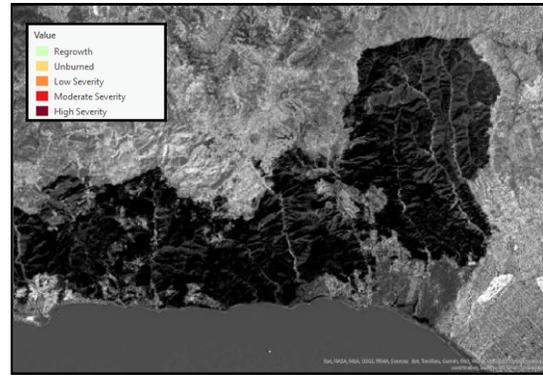


Figure 10. Palisades Fire. Burn Severity Classification.



Figure 11. Palisades Fire. NDVI of Jan 14th, Postfire

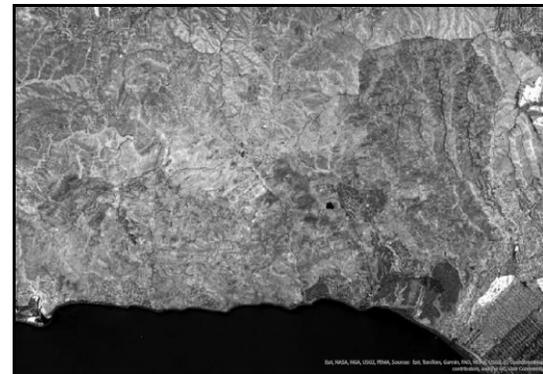


Figure 12. Palisades Fire. NDVI of May 30th.



Figure 13. Palisades Fire. Vegetation Regrowth Classification.

IV. RESULTS

The Eaton Fire burned 14,070 acres: 50% experienced low, 22% moderate, and 3% high severity. In contrast, the Palisades Fire affected 23,442 acres, with 55% low, 30% moderate, and 1% high severity. Despite its larger size, Palisades had a smaller proportion of high-severity burn. These total acres closely align with official incident reports from CAL FIRE, which reported 14,021 acres burned for Eaton and 23,448 for Palisades [5] (Tables 1 and 2).

TABLE I. BURN SEVERITY

Severity level	Eaton		Palisades	
	Acres	%	Acres	%
Unburned	3,434	24%	3,284	14%
Low severity	7,022	50%	12,855	55%
Moderate severity	3,159	22%	7,077	30%
High severity	455	3%	226	1%
Total (low-high severity)	10,636	76%	20,158	86%
Total	14,070	100%	23,442	100%

TABLE II. VEGETATION RECOVERY

Regrowth level	Eaton		Palisades	
	Acres	%	Acres	%
Vegetation loss	18	0%	3	0%
No significant change	10,473	74%	6,057	26%
Low regrowth	3,555	25%	17,093	73%
Moderate regrowth	25	0%	288	0%
High regrowth	0	0	0	0
Total (low-high regrowth)	3,580	25%	17,381	74%
Total	14,070	100%	23,442	100%

NDVI analysis revealed a stark contrast in recovery: only 25% of burned land in Eaton showed low to moderate regrowth, whereas 74% of burned land in Palisades showed vegetation rebound. Most of Eaton's area showed no significant change. No areas in either fire demonstrated high regrowth.

V. DISCUSSION

The difference in regrowth may be attributed to environmental and topographic variation between the two regions. The Palisades Fire may have burned in areas with higher native vegetation resilience or better postfire precipitation. On the other hand, the Eaton Fire may have burned hotter or in areas already stressed by drought.

dNBR and dNDVI indices, while effective for rapid assessment, are limited by Landsat's 30-meter spatial resolution, which can miss fine-scale changes. Moreover, we conducted no field validation, so spectral interpretations may not fully reflect on-the-ground conditions.

VI. CONCLUSION

This study demonstrates that such remote sensing indices as dNBR and dNDVI can help evaluate wildfire impact and vegetation recovery. While both the Eaton and Palisades Fires had similar severity patterns, their regrowth trajectories differed significantly. These findings offer an important guide for land restoration strategies and for enhancing postfire decision-making frameworks in fire-prone regions.

Future work could incorporate higher resolution imagery, field validation, 1-year revisits, and such climatic variables as rainfall or soil moisture to improve regrowth predictions.

REFERENCES

- [1] U.S. National Park Service. *Yosemite National Park: Fire Ecology and Monitoring*. [Online]. Available from: <https://www.nps.gov/yose/learn/nature/fireecology.htm> 2026-01-24
- [2] United Nations. *Normalized Burn Ratio (NBR)*. UN-SPIDER Knowledge Portal. [Online]. October 2025. Available from: <https://un-spider.org/advisory-support/recommended-practices/recommended-practice-burn-severity/in-detail/normalized-burn-ratio> 2026-01-24
- [3] UP42. *Normalized difference vegetation index (NDVI)*. UP42 Documentation. [Online]. October 2025. Available from: <https://docs.up42.com/help/spectral-indices/ndvi> 2026-01-24
- [4] US Geological Survey. *Landsat Surface Reflectance-Derived Spectral indices*. [Online]. 2018, October 2025. Available from: <https://www.usgs.gov/landsat-missions/landsat-surface-reflectance-derived-spectral-indices> 2026-01-24
- [5] US Geological Survey. *What are the band designations for the Landsat satellites?* [Online]. 2024, October 2025. Available from: <https://www.usgs.gov/faqs/what-are-band-designations-landsat-satellites> 2026-01-24
- [6] ESRI. *Assess Burn Scars with Satellite Imagery*. [Online]. 2025. Available from: <https://learn.arcgis.com/en/projects/assess-burn-scars-with-satellite-imagery> 2026-01-24
- [7] Y. Korgaonkar, Normalized Burn Ratio (NBR) and Normalized Difference Vegetation Index (NDVI) Study of Vegetation Health and Regrowth Rate post 2018 Mendocino Complex Fire in Northern California Coastal Mountains. MS Thesis, University of Arizona, Tucson, AZ. May. 2022, Available from: <https://repository.arizona.edu/handle/10150/667137> 2026-01-24
- [8] J. D. Miller et al., "Calibration and Validation of the Relative Differenced Normalized Burn Ratio (RdNBR) to Three Measures of Fire Severity in the Sierra Nevada and Klamath Mountains, California, USA," *Remote Sens. Environ.*, vol. 113, no. 3, pp. 645–656, 2009, doi:10.1016/j.rse.2008.11.009.
- [9] A. M. S. Smith et al., "The science of fire-scapes: Achieving fire-resilient communities," *Bioscience*, vol. 66, no. 2, pp. 130–146, 2016, doi:10.1093/biosci/biv182.
- [10] E. A. Storey, K. R. Lee West, and D. A. Stow, "Utility and Optimization of LANDSAT-Derived Burned Area Maps for Southern California," *Int. J. Remote Sens.*, vol. 42, no. 2, pp. 486–505, 2021, doi:10.1080/01431161.2020.1809741.
- [11] B. Sall, M. W. Jenkins, and J. Pushnik, "Retrospective Analysis of Two Northern California Wild-Land Fires via Landsat Five Satellite Imagery and Normalized Difference Vegetation Index (NDVI)," *Open J. Ecol.*, vol. 3, no. 4, pp. 311–323, 2013, doi:10.4236/oje.2013.31007

Climatic Extremes, Vegetation Response, and Fuel Conditions

A five-year analysis of the Palisades and Eaton fire zones in Los Angeles County

Vivian Sultan

College of Business and Economics
California State University
Los Angeles, USA
email: vsultan3@calstatela.edu

Diego Sanchez

College of Business and Economics
California State University
Los Angeles, USA
email: dsanchez@calstatela.edu

Leonardo Ortiz

College of Business and Economics
California State University
Los Angeles, USA
email: lortiz@calstatela.edu

Abstract—This study examines how recent climatic variability and vegetation dynamics may have contributed to fuel conditions preceding the 2025 Palisades and Eaton wildfires in Los Angeles County. Using precipitation records from 2020–2025 and Moderate Resolution Imaging Spectroradiometer-derived Normalized Difference Vegetation Index (NDVI) datasets, the study entailed analyzing rainfall anomalies, spatial rainfall distribution, and vegetation changes in and around the two fire perimeters. Results show that 2023 was an exceptionally wet year, with rainfall far exceeding the preceding and following years, particularly in the Altadena region. However, despite this anomalous precipitation, NDVI analyses revealed only limited vegetation recovery in 2023 and 2024, suggesting the landscape did not respond with the expected greening. Instead, vegetation remained in a marginally healthy state, creating persistent dry biomass that likely increased fuel availability. In the Palisades region, vegetation appeared healthier in 2023 but showed signs of drying in 2024, consistent with reduced regional precipitation. These findings highlight that extreme rainfall does not necessarily translate into sustained vegetation recovery and may instead contribute to elevated wildfire risk when followed by rapid drying. Understanding these climatic and ecological interactions is essential for improving proactive wildfire-mitigation strategies in high-risk areas.

Keywords-Wildfire; ArcGIS; California.

I. INTRODUCTION

Wildfires can be incredibly destructive and dangerous. There is no better example than the Eaton and Palisades fires. The level of destruction caused by these fires overshadows that of any wildfires that have happened previously. Thus, it is important to understand the variables at play in how a fire spreads. One of the easiest ways for a fire to spread is through the amount of dead vegetation in each area. Identifying weather patterns that have effects on vegetation levels, such as abnormal rainfall, can be a key indicator to being proactive in reducing the amount of vegetation (dead or alive) in or around areas that have dense population centers and electric-utility equipment to reduce the likelihood of wildfires spreading to or sparking in these areas.

Section 2 reviews the literature, and Section 3 identifies the goals of this research. Section 4 lays out the data selection and acquisition techniques, and Section 5 explains

the data cleaning performed. Section 6 describes the computer system used for the analysis and Section 7 describes that analysis. Section 8 concludes.

II. LITERATURE REVIEW

Several research projects explored similar themes, investigating the relationship between precipitation, vegetation dynamics, and wildfire occurrences [1][2][3]. For instance, researchers have examined how antecedent rainfall and excessive vegetation growth correlate with wildfire burned areas across California, analyzing data over extended periods [2]. Other researchers focused on the influence of specific weather phenomena, like atmospheric rivers, on vegetation growth and fuel loading in Southern California ecosystems. The fundamental approaches in these studies often involve analyzing various datasets, such as precipitation data from sources like the Global Precipitation Measurement mission and vegetation indices like the Normalized Difference Vegetation Index (NDVI) and the Enhanced Vegetation Index derived from satellite imagery like Moderate Resolution Imaging Spectroradiometer (MODIS) and Landsat [1][2]. Common methodologies include using Geographic Information Systems (GIS) for spatial analysis, calculating anomalies in rainfall and vegetation, and employing statistical methods like spatial correlations and linear regression models to determine the relationships between these climatic and vegetation variables and wildfire outcomes. These underlying theories and approaches, which link increased rainfall to enhanced vegetation that can later serve as fuel provide a solid foundation for projects analyzing recent rain patterns and vegetation changes to understand their potential impact on wildfire severity. Additionally, the timing and “sharpness” of the rainy season are recognized as critical factors, as shifts in precipitation seasonality can extend the period of high fire risk by increasing the overlap between dry vegetation and fire-promoting winds [3].

III. PROBLEM IDENTIFICATION AND GOALS

A problem identified in this study is the lack of proactiveness from the City of Los Angeles, local municipal bodies near or around high wildfire-risk areas, and utility companies such as Southern California Edison and Pacific Gas and Electric in reacting to elevated risks of wildfire severity due to fuel loading.

This research aimed to highlight the effects that abnormal amounts of rain have on vegetation levels in the Los Angeles area. Some metrics aiming to aid in identifying these changes are precipitation patterns in the Los Angeles area, such as how much overall rainfall has occurred in the area of interest over the last five years, identifying any outliers, creating a map that highlights areas of concentrated rainfall, and comparing the difference in these areas. Also, this study entailed looking into NDVI data and identifying what changes in vegetations levels have occurred over the last several years. The goal of this research was intended to understand how weather patterns can be a predicting factor in how severe a wildfire can be to allow the target audience to create or better enact policies in developing wildfire-prevention measures.

IV. DATA SELECTION AND ACQUISITION

For this project, three datasets were needed: the fire-perimeter data from Cal Fire, the precipitation data from the National Weather Service, and the vegetation data from NASA.

A. Fire-Perimeter Data

For fire perimeter data, Cal Fire's public archive shape files were used to focus the analysis in the perimeter of Palisades and Eaton fires [4].

- Data Type: Vector Shapefile
- Data Provider: Cal Fire

B. Precipitation Data

Precipitation data were obtained for the Los Angeles region directly from available national weather-service stations. These data were formatted in a comma-separated file (CSV).

- Data Type: txt Data
- Data Provider: National Weather Service

C. Normalized Difference Vegetation Index (NDVI) Data

NDVI data came from the MOD13Q1 MODIS, downloadable through the Application for Extracting and Exploring Analysis Ready Samples (AppEEARS) platform [5]. These files covered the time span from January 2020 up until January 2025.

- Data Type: Raster Data
- Data Provider: NASA EARTHDATA

The required data were downloaded from the sources as either shapefiles or as a text file in the form of a CSV.

V. DATA CLEANING

These data required some preprocessing before the analysis. The precipitation data needed cleaning, and the NDVI data were not in a state usable by ArcGIS Pro.

A. Precipitation Cleaning

When working with precipitation data, several unneeded columns needed to be deleted. The precipitation dataset also only had the data points as daily values from January 2020 to January 2025. To create a more meaningful analysis, the daily datapoints had to be aggregated into mean monthly

values using a python script that went through the CSV and averaged every value in every month of every year. The CSV also reformatted the date column into a format that ArcGIS Pro could ingest without any misinterpretations. Once this dataset was uploaded into ArcGIS Pro, one main concern that arose was that no weather stations were within a meaningful distance to the palisades fire perimeter or in the perimeter itself, resulting only in speculation for that area.

B. NDVI Data Processing

After collecting NDVI data, preprocessing took place to make these data usable in ArcGIS Pro. The data set was downloaded as a Geo tagged image file format file through the AppEEARS platform. This folder included multiple bands, like quality control, and metadata layers. The process entailed isolating the NDVI bands and removing the others. Then, a mosaic by region and year was created to enable the time series for the NDVI analysis. A Python script was used to extract and format the dates from the file names. Then, an acquisition date field was created, which allowed cleaning of the date field that ArcGIS could not recognize until it was configured with a Python script, leading to the creation of a time slider.

C. Limitations

A few data and methodological limitations should be noted. Precipitation measurements were limited by the absence of weather stations within or immediately adjacent to the Palisades fire perimeter, requiring rainfall patterns to be inferred from nearby stations and regional raster surfaces. This constraint reduces the spatial precision of precipitation estimates for that area. In addition, NDVI processing required mosaicking and temporal standardization across multiple MODIS tiles, which may introduce minor smoothing in the time series. While these limitations do not alter the overall trends observed in the analysis, they highlight the need for denser environmental monitoring networks and more granular vegetation datasets to improve future wildfire-risk assessments.

D. System Information

ArcGIS Pro was used on Windows machines because ArcGIS Pro is only available on Windows. Other programs used for this project included PyCharm to clean the precipitation data and a Python IDE, used to extract and standardize the acquisition dates from the NDVI file names for use in a time series in ArcGIS. Although not used extensively, Tableau was used to create a line chart visualization.

VI. ANALYSIS

The study focused on precipitation changes throughout the last five years (2020–2025), comparing 2023 precipitation to previous and subsequent time periods, to identify differences through heat maps. From there, 2023 and 2024 were explored as key years to determine whether the heavy rainfall in 2023 led to vegetation growth that later dried out. NDVI rasters were used to track vegetation levels across Altadena and Palisades, along with spatial raster

calculations to see what changed and what did not. This macro-to-micro approach helped in understanding whether environmental shifts might have contributed to fire vulnerability leading into 2025.

A. Precipitation Variation 2020–2025

As stated earlier, to gain a deeper insight into how precipitation has changed over the years, data were aggregated from daily values to monthly averages. Figures 1 and 2 line charts showcase differences over time, revealing stark findings.

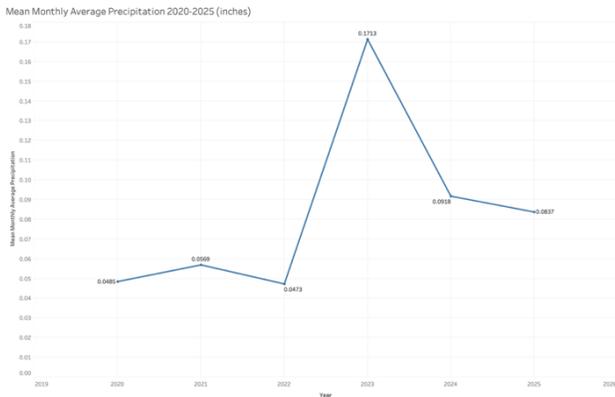


Figure 1. Average precipitation in LA County, 2020–2025

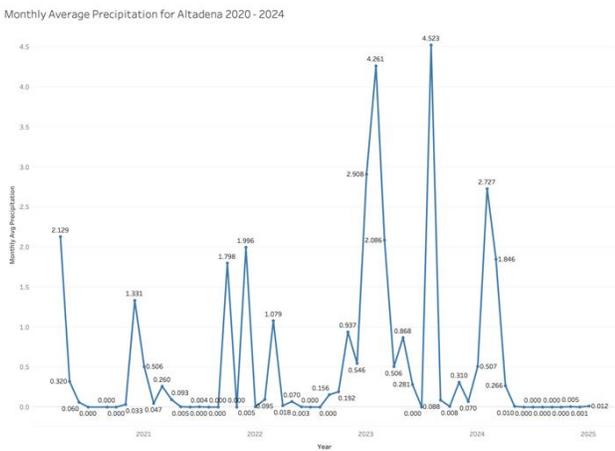


Figure 2. Altadena’s rainfall, 2020–2025

Huge disparities emerged in rainfall differences in 2023 when compared to previous and subsequent years. The data revealed a 53.6% reduction in Los Angeles areas from 2023 to 2024 and an even more staggering 66% reduction in rainfall in Altadena from 2023 to 2024.

B. Precipitation Heat Map Analysis

Figure 3 shows little difference in precipitation between the years 2020–2022 when compared to 2023, outside of the two hot spots. Orange indicates similar amounts of precipitation all around. Although Altadena saw substantially more rain in Year 1 than in Year 3, Figures 4 and 5 show that, when inverted to show the difference in 2024 from 2023, Altadena is much drier, appearing almost

white in the figure, the lowest possible range in the color scheme, further showing the disparity in precipitation.

These maps were created using the Kernel Density analysis tool to create initial heat maps of the specific time frames used (2020–2022, 2023, 2024) in ArcGIS Pro. Then, the Raster Calculator was used to calculate the differences between heat maps to create a new one based on that calculation, resulting in Figures 3 through 5.

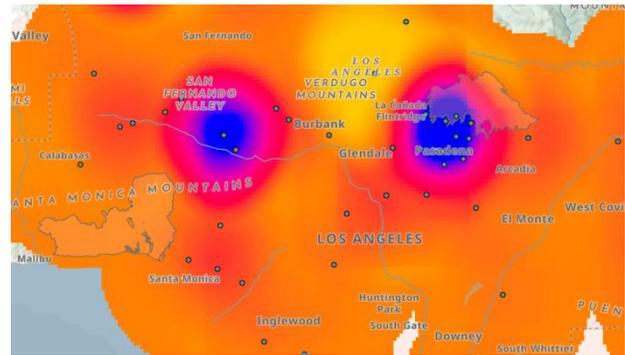


Figure 3. Difference in precipitation between 2023 and 2020–2022

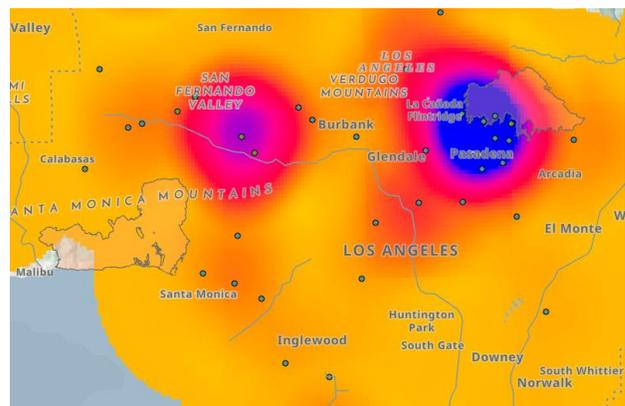


Figure 4. Difference in precipitation between 2023 and 2024

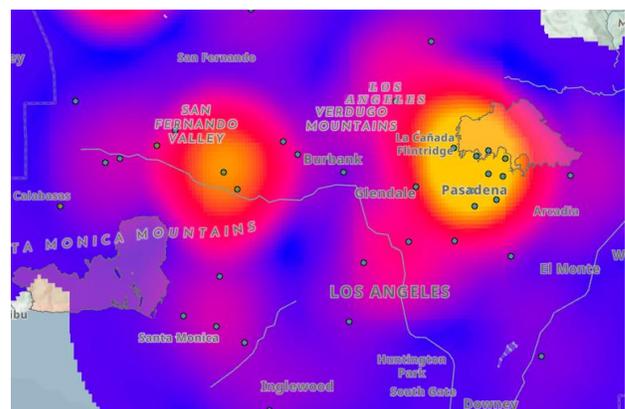


Figure 5. Difference in precipitation between 2024 and 2023

C. Altadena NDVI Analysis

For Altadena, with the most rainfall of any year in the five-year period, strong vegetation growth was anticipated. The assumption was that more rain should mean greener

vegetation, especially around the fire perimeter. That is not what emerged.

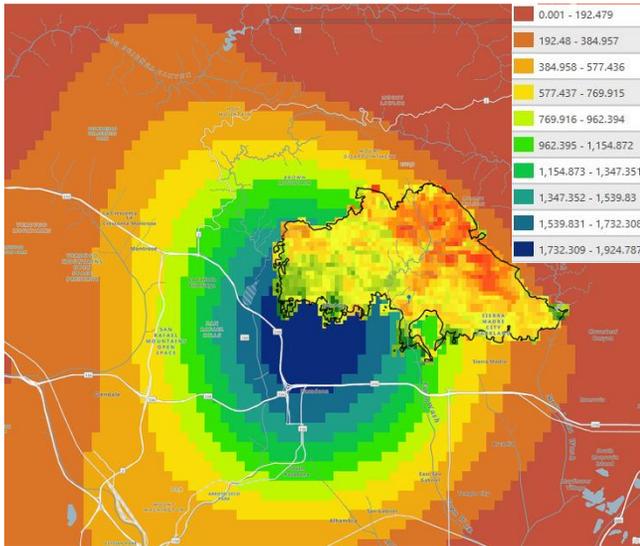


Figure 6. Altadena NDVI raster (2023)

A raster calculation on the 2023 NDVI compared with the precipitation legend, shown in Figure 6 revealed that most of the Altadena fire zone sits inside the dark blue and green zones, representing heavy rainfall. Only the southern edge touches areas with lower rain, shown in yellow or orange. A great deal of healthy vegetation was anticipated. Instead, the NDVI showed the area to be dry. The large mix of yellow and light green indicates the vegetation did not respond as anticipated. Questions arose about whether the vegetation would take longer to appear than initially believed.

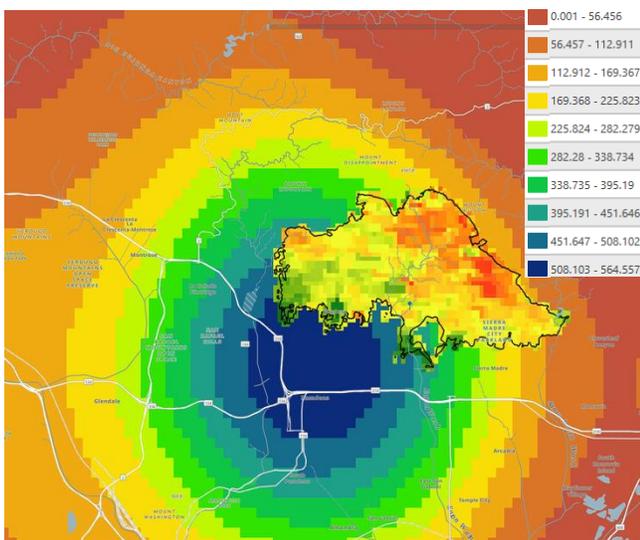


Figure 7. Altadena NDVI raster (2024)

After seeing results from 2023, a new hypothesis was that vegetation needed more time to bounce back. Data show that 2024 brought a great deal of rain. Using the same NDVI

raster calculations, this time for the full year of 2024, the focus again was on the fire perimeter. Some very subtle changes emerged, especially in the northern part of the perimeter, as shown in Figure 7 where colors shift slightly from red-orange to more of a yellow tint; however, the vegetation still seemed sparse. Two years of heavy rainfall yielded little rebound in greenness. The vegetation remained in a borderline zone, neither dead nor thriving. The lingering dryness meant great amounts of biomass had sufficient structure to burn.

D. Palisades NDVI Analysis

With no direct precipitation data for this region, the study relied on visual patterns and assumptions from the surrounding area.

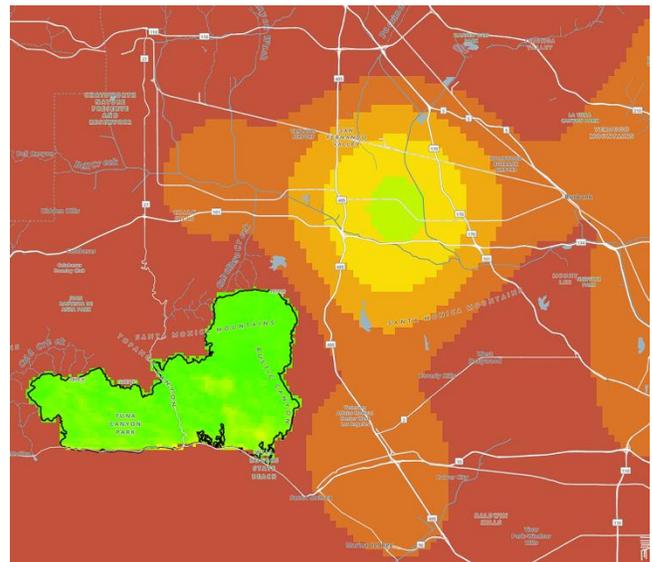


Figure 8. Palisades NDVI raster (2023)

The precipitation raster shown in Figure 8 reveals that although the Palisades itself doesn't light up on the map, rainfall was taking place in the area, especially to the northeast and toward parts of Los Angeles. Vegetation inside the Palisades fire perimeter shows quite green with small patches of yellow. Overall, it appears lush and full.

The Palisades probably got good precipitation. Perhaps not as much as Altadena, but enough to support healthy vegetation. This region seemed to hold up well in 2023. Surrounding areas were getting rain, and the vegetation data backed that up. The Palisades looked strong.

Figure 9 for 2024 starts to show a shift. Comparing the precipitation raster to 2023, the area had less rainfall. Most of the surrounding region appears as orange or light orange which indicates lower precipitation levels.

Inside the fire perimeter shown in Figure 9 is a good amount of green vegetation but it is starting to fade. Some bright green zones from 2023 are leaning more toward yellow, providing a noticeable change. The vegetation began to dry out enough to start making the region more vulnerable. Rather than a sudden collapse in vegetation health, the area experienced a gradual decline. Thus,

although the landscape appeared healthy, it had a great deal of dry fuel, especially in Palisades where rainfall could not be measured directly.

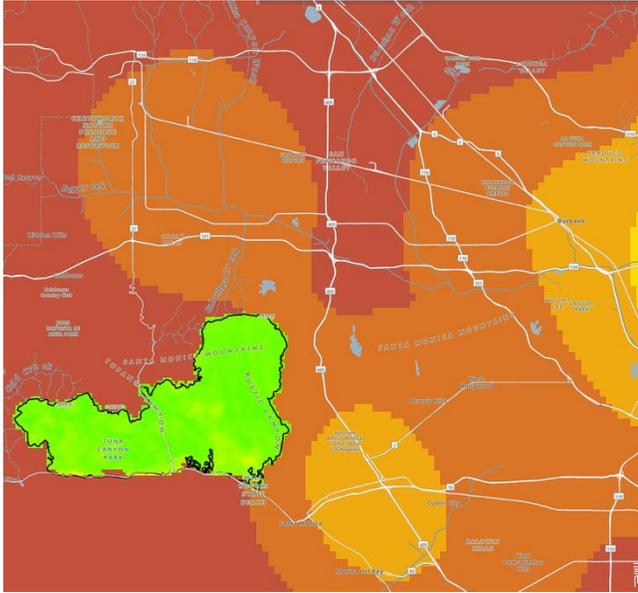


Figure 9. Palisades NDVI raster (2024)

VII. CONCLUSIONS AND FUTURE WORK

This analysis shows that extreme rainfall events do not automatically translate into meaningful or sustained vegetation recovery, and that this disconnection can create hazardous fuel conditions leading into a wildfire season. Although 2023 was an exceptionally wet year across Los Angeles County, and especially in Altadena, the NDVI data revealed only limited vegetation improvement in 2023 and 2024. Even after two consecutive years of above-average precipitation, vegetation inside the Eaton fire perimeter remained in a marginal state rather than rebounding to healthier levels. This finding suggests that the ecosystem may have been unable to fully capitalize on the rainfall due to factors such as soil-moisture retention, vegetation type, or preexisting stress conditions. The result was a landscape with persistent dry biomass that served as readily available fuel by 2025.

In the Palisades region, vegetation appeared healthier in 2023 but showed a gradual decline in 2024, consistent with

reduced rainfall in the surrounding area. Vegetation may have been drying out under the surface. Although the changes were subtle, they point to a slow drying trend that can increase fire susceptibility even when the landscape appears relatively green at first glance. The absence of nearby weather stations limits the precision of precipitation estimates for this region, but the vegetation patterns still indicate a meaningful shift toward drier conditions. Rain does not automatically mean recovery. The land may not respond in expected ways to heavy rainfall in anticipated ways and may provide even greater risk for fires.

Overall, findings highlight a critical insight for wildfire mitigation: rainfall anomalies alone are not reliable indicators of reduced fire risk. What matters is how vegetation responds over time, and whether periods of heavy rain are followed by sustained drying that leaves behind accumulated, partially recovered, or weakened biomass. These dynamics underscore the need for proactive vegetation management, improved monitoring of post-rainfall vegetation health, and more nuanced risk assessments that consider climatic extremes and ecological response. Understanding these interactions is essential for developing policies and interventions that can better anticipate and reduce wildfire severity in high-risk regions.

REFERENCES

- [1] I. Campbell, "Influence of atmospheric river precipitation on vegetation growth and fuel," UC San Diego, Open Access Publications from the University of California.
- [2] J. J. Hernández Ayala, J. Mann, and E. Grosenove, "Antecedent rainfall, excessive vegetation growth and its relation to wildfire burned areas in California," *Earth Space Sci.*, vol. 8, no. 9, art. e2020EA001624, 2021. doi:10.1029/2020EA001624
- [3] D. L. Swain, "A shorter, sharper rainy season amplifies California wildfire risk," *Geophys Res Let.*, vol. 48, no. 5, art. e2021GL092843, 2021. doi:10.1029/2021GL092843
- [4] California Department of Forestry and Fire Protection, "CAL FIRE historical wildland fire perimeters nominated to living atlas. The service includes layers that are data subsets symbolized by size and year," Oct 2025. [Online]. Available from: <https://calfire-forestry.maps.arcgis.com/home/item.html?id=c3c10388e3b24cec8a954ba10458039d> 2026.01.24
- [5] NASA, "Welcome to AppEEARS: Application for extracting and exploring analysis ready samples." [Online]. Available from: <https://appeears.earthdatacloud.nasa.gov/> 2026.01.24

Understanding the Socioenvironmental Impact of LA City Wildfires

Vivian Sultan

College of Business and Economics
California State University
Los Angeles, USA
email: vsultan3@calstatela.edu

Xuehan (Harvey) Jing

School of Civil and Environmental Engineering and Earth
Sciences
Clemson University
Clemson, SC, USA
email: xuehanj@clemson.edu

Sirisha Mahesh

College of Business and Economics
California State University
Los Angeles, USA
email: smahesh@calstatela.edu

Christopher Lee

College of Business and Economics
California State University
Los Angeles, USA
email: clee@calstatela.edu

Abstract—This paper outlines a research plan to investigate how wildfires in Los Angeles affect different communities and the environment. We aim to use spatial data to develop equitable prevention and response strategies. The research will combine geographic information system datasets with demographic data, land-use records, and historical fire information to identify patterns that can support better preparedness and urban-resilience planning. The proposal emphasizes that wildfires are not just natural disasters; they are social events that worsen existing inequalities, particularly impacting marginalized communities through home destruction, poor air quality, evacuation challenges, and insurance burdens. It also highlights significant such environmental damage as deforestation and biodiversity loss.

Keywords—Wildfire; ArcGIS; California.

I. INTRODUCTION

This project focuses on exploring the socioenvironmental impacts of wildfires in Los Angeles, specifically how wildfire frequency, intensity, and spatial distribution intersect with population vulnerability, infrastructure exposure, and climate patterns. The key research question is, “How do wildfires in Los Angeles disproportionately impact certain communities, and how can spatial data be used to inform equitable prevention and response strategies?”

Los Angeles faces increasing wildfire risks due to climate change, prolonged droughts, and the expansion of the wildland–urban interface (WUI). These wildfires are critical because they reveal and exacerbate existing inequalities, disproportionately affecting marginalized communities through property destruction, hazardous air quality, and financial burdens. Wildfires also cause significant environmental damage, including deforestation, biodiversity loss, and soil degradation. We must understand these long-term consequences to guide urban planning, inform emergency management, and protect critical infrastructure. This issue also serves as a key case study for how large urban areas can plan for and adapt to new environmental realities as climate-related disasters become more frequent and severe, necessitating proactive, data-driven preventive strategies rather than reactive responses.

The research is intended for multiple key stakeholders who can act on the findings. The primary audience, urban

planners and policymakers, can help develop land-use regulations and emergency preparedness policies. We hope that such emergency-management agencies as the Los Angeles Fire Department, the California Department of Forestry and Fire Protection (Cal Fire), and the Los Angeles Office of Emergency Management will use such data as this to enhance readiness and create equitable evacuation, shelter, and recovery plans. Public health officials could also use these data to anticipate and manage the health implications of wildfires, including air quality and mental health effects on vulnerable populations. Perhaps academic researchers and students could use the model and methodology for interdisciplinary analysis. Nonprofits and community organizations could use also this analysis, especially those serving low-income and immigrant neighborhoods, to support advocacy for better access to resources and protection. Finally, this research could increase public awareness and individual preparedness through accessible maps and data storytelling.

Section 2 reviews the available literature. Section 3 outlines the data used in this research, and Section 4 details the applications used. Section 5 discusses the code used. Section 6 presents the analysis. Section 7 concludes.

II. LITERATURE REVIEW

Davies et al. [1] highlighted that environmental disasters often disproportionately affect minority and impoverished communities globally, and wildfires in the United States show similar patterns of outsized impacts on vulnerable populations. The research develops a social-ecological approach to characterize fire vulnerability, considering both wildfire potential and communities’ socioeconomic attributes. Davies et al. found that census tracts with majority Black, Hispanic, or Native American populations experienced approximately 50% greater vulnerability to wildfire compared to other census tracts. This underscores the need for management strategies that extend beyond technical solutions to incorporate socioeconomic and political considerations, especially tailoring emergency planning and mitigation strategies to address diverse populations and overcome such barriers as language differences.

The LA City Wildfires GeoHub [2] serves as a critical application for community awareness regarding wildfire

incidents and their potential impacts in Los Angeles. Designed to provide real-time status and information, this hub displays a list of fires with detailed information about their impacts on human and ecosystem populations. The data integrated into this application is sourced from authoritative entities like the Environmental Protection Agency and the U.S. Forest Service, ensuring reliability. For urban planners and emergency-management agencies, this GeoHub is an invaluable tool for understanding current wildfire situations, enabling more informed decision making and resource allocation.

Reid et al. [3] assessed the various health effects of wildfire smoke exposure, noting that consistent evidence links wildfire smoke to general respiratory issues, asthma exacerbation, and chronic obstructive pulmonary disease. Growing evidence also suggests increased respiratory-infection risk and overall mortality. For public health officials and the general public, this highlights the profound health implications of wildfires and the need for protective measures, though further research is still needed to clarify specific cardiovascular outcomes and identify particularly susceptible populations.

The Governor's Office of Emergency Services [4] outlined California's comprehensive strategy for reducing disaster losses and building resilience against a wide array of natural and human-caused hazards, including wildfires. Recognizing that disasters are increasing in frequency and intensity due to climate change and urban expansion, the plan emphasizes a proactive, integrated, and equitable approach to risk reduction. It detailed hazard assessments, identifies vulnerabilities across state assets, critical facilities, community lifelines, and disproportionately impacted equity-priority communities, and set forth goals, objectives, and mitigation actions to guide statewide efforts. The document highlighted existing state capabilities, such as strong building codes and various funding programs, and outlined processes for local government planning coordination to ensure unified and effective mitigation programs aimed at protecting life, property, and the environment.

This article uses satellite data and land-use analysis [5] to show that the fastest-growing U.S. regions are those bordering wildlands, leading to increased wildfire ignition and damage potential. For Los Angeles, with its flammable ecosystems, this study highlights the risks of expanding residential zones into fire-prone areas and demonstrates how spatial data can predict high-concern areas.

Shi et al. [6] argued for embedding social equity into all aspects of climate-resilience planning and identified how adaptation measures can unintentionally worsen inequality if social vulnerability is not integrated into spatial planning. This frames the project's intent to use wildfire risk data in conjunction with sociodemographic analysis to ensure proposed strategies are equitable and just.

The U.S. Forest Service [7] uses fire modeling, census data, and remote sensing to create national-scale wildfire risk maps. Its interactive layers, which include wildfire likelihood, flame length potential, and housing exposure, offer a technical framework adaptable to Los Angeles's neighborhood-level data using such geographic information system (GIS) software as ArcGIS or QGIS.

Cutter et al. [8] provided a valuable model for combining social data with environmental risk to develop composite indices. Their approach, dataset normalization, weighted scoring, and index mapping, offers methods that can be incorporated to visualize vulnerability alongside fire intensity or frequency for Los Angeles.

III. DATA REVIEW

The City of Los Angeles [2] is a primary source for wildfire data. It provides current status and potential impact of wildfire incidents and displays a list of fires with details on their effects on human and ecosystem populations. The data in this hub is sourced from authoritative entities and includes layers from such agencies as the Environmental Protection Agency, the National Weather Service, and the U.S. Forest Service. Its interactive layers, wildfire likelihood, flame length potential, and housing exposure, offer a technical framework for analysis.

We also use the Social Vulnerability Index (SVI) developed by the Centers for Disease Control and Prevention based on data from the U.S. Census Bureau. The SVI incorporates 15 census variables, including poverty, lack of vehicle access, and overcrowding, to rank vulnerability by census tract. This methodology is ideal for overlaying with fire hazard zones to identify communities most at risk in Los Angeles. Existing studies have shown how SVI can be used in conjunction with fire risk data to understand disproportionate impacts.

Land-use records from LA City's planning department provide information on how land is used across Los Angeles, including urban development, WUI zones, and other land classifications. Understanding land use is crucial as rapid growth in WUI zones increases wildfire risk and damage potential.

Historical fire occurrence layers (from Cal Fire or similar authoritative sources) provide records of past wildfire ignitions, perimeters, and behavior. This historical data is essential for identifying patterns, understanding fire frequency, and conducting spatial clustering analyses to predict wildfire ignition hotspots [9].

Climate data from the National Oceanic and Atmospheric Administration and state-level climate assessments, including temperature, precipitation, and drought, is vital as these factors influence wildfire frequency, intensity, and spatial distribution. Climate change is recognized as increasing wildfire risks.

Infrastructure exposure from LA City's Public Works and utility providers (e.g., transportation networks, utilities, emergency-services facilities) are used to assess exposure to wildfires. This helps understand how wildfires interact with the city's social geography and infrastructure.

California fire perimeters from the California Natural Resources Agency (CNRA) Geoportal provide the geographical boundaries (perimeters) of historical wildfires across California. This foundational layer offers precise spatial and temporal information about past fire events. Each perimeter typically includes such attributes as fire name, year, acres burned, and sometimes ignition cause. This dataset is directly relevant to understanding wildfire frequency, intensity, and spatial distribution in Los Angeles.

For emergency management agencies, this data is indispensable for identifying areas with a history of recurrent fires, informing fuel management and response planning. Urban planners and policymakers can use it to assess WUI expansion risks and guide zoning decisions in fire-prone areas. Academic researchers and students will find it crucial for developing predictive models and understanding historical fire ecology. For the general public and local residents, visualizing past fire perimeters can significantly increase local-risk awareness.

IV. APPLICATION REVIEW

The analysis of California wildfire data, as outlined across `Capstone3.ipynb`, `Capstone4.ipynb`, and `Capstone5.ipynb`, leverages a powerful suite of Python libraries and functions to comprehensively understand wildfire patterns and their socioenvironmental impacts, and to inform equitable prevention and response strategies. This multifaceted approach is tailored to provide actionable insights for urban planners, emergency management agencies, public health officials, academic researchers, nonprofits, and the general public, all of whom are critical stakeholders in building resilient communities.

At the foundational level, data management and preparation are primarily handled by the `pandas` library. Across all notebooks, `pd.read_csv()` and `pd.read_excel()` are instrumental in loading the raw wildfire perimeter data from sources like the CNRA Geoportal. Once loaded, `pandas` DataFrames become the central structure for organizing this complex tabular data. Crucial initial steps involve data cleaning and transformation: `df.dropna()` handles missing values in such critical columns as latitude, longitude, and acres burned, ensuring data quality for subsequent analyses.

`pd.to_datetime()` converts date columns, such as `Started`, into `datetime` objects, which is essential for any temporal analysis, such as identifying monthly or yearly trends.

Furthermore, `pandas` facilitates feature engineering by allowing the creation of new columns—for example, extracting month names from start dates or generating binary flags like `FatalitiesOccurred`, `CausedDamage`, and `FullyContained`. This data-preparation phase ensures that the raw information is robust, complete, and in the correct format for deeper statistical and geospatial exploration.

Following data preparation, exploratory data analysis and visualization play a vital role in revealing key patterns and insights. The `matplotlib.pyplot` and `seaborn` libraries are central to this phase. `matplotlib` provides granular control over plot customization, allowing for clear titles, axis labels, and overall layout adjustments using functions like `plt.figure()`, `plt.title()`, and `plt.tight_layout()`. Building on this, `seaborn` offers high-level interfaces for creating attractive and informative statistical graphics. For instance, `sns.countplot()` is employed to visualize the number of wildfires by month or year, effectively highlighting seasonal peaks and long-term trends. This information is indispensable for emergency management agencies to allocate seasonal resources and for the general public to understand peak risk periods. `sns.heatmap()` is used to generate correlation matrices between numerical variables

like `AcresBurned` and response resources (e.g., `PersonnelInvolved`, `CrewsInvolved`, `Engines`), revealing how fire size relates to resource demand. Furthermore, `scipy.stats` contributes to a deeper statistical understanding of fire-size distribution, with `skew()` and `kurtosis()` functions quantifying the asymmetry and *tailedness* of the `GIS_ACRES` data. This is crucial for emergency management agencies and policymakers to understand the probability of outlier megafires, while `pandas.DataFrame.describe()` and sorting functions identify top-10 largest fires (Figure 1), providing concrete examples of extreme events for all stakeholders.

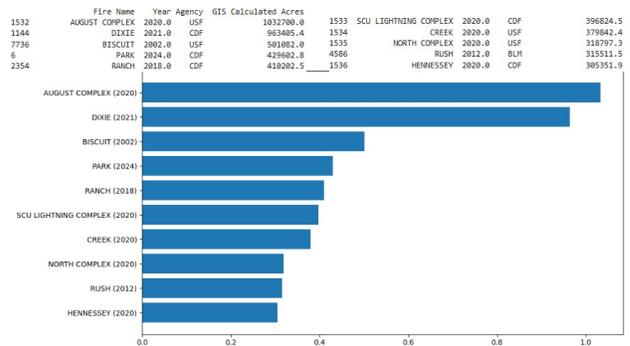


Figure 1. Top-10 largest wildfires in California by GIS-calculated millions of acres burned

For advanced spatial analysis, the `geopandas` library is a critical extension to `pandas`, enabling direct manipulation and analysis of geospatial data. `geopandas` treats the wildfire perimeters (from the CNRA dataset) as geographical objects (polygons) rather than just tabular data, fundamental for creating `GeoDataFrames` and performing essential spatial operations. For example, `geopandas` facilitates spatial joins and overlay analyses, allowing the project to combine fire perimeter data with such spatially explicit datasets as the SVI from the U.S. Census Bureau or critical infrastructure layers. This enables the identification of specific geographic areas where wildfire risk intersects with vulnerable populations or critical infrastructure, directly addressing the core research question of disproportionate impacts on communities.

Finally, for predictive modeling and map-based decision support, the `scikit-learn` library, combined with `folium`, forms a powerful predictive analytics pipeline. `sklearn.model_selection.train_test_split()` is used to correctly partition data into training and testing sets, ensuring the model's performance is evaluated on unseen data. Classification models, `sklearn.linear_model.LogisticRegression` and potentially `sklearn.ensemble.RandomForestClassifier` for comparison, are trained to predict such critical wildfire outcomes as whether a fire is likely to result in fatalities or structural damage. The `class_weight = balanced` parameter in `LogisticRegression` is particularly important for handling imbalanced datasets, common when predicting rare but critical events like fatalities.

Model performance is rigorously evaluated using `sklearn.metrics.classification_report()` with key metrics like precision, recall, and F1 score to assess the model's reliability in predicting these adverse outcomes. The `folium` library then plays a crucial role in visualizing these

predictions on interactive maps. By plotting predicted outcomes (e.g., red markers for predicted fatalities, green for no fatalities) on a base map of California, folium creates intuitive map-based decision-support tools, from which emergency management agencies and urban planners can identify predicted hotspots for future destructive or fatal incidents, allowing for proactive resource allocation, targeted education, and infrastructure protection. This comprehensive analytical framework, integrating data management, exploratory statistics, advanced geospatial analysis, and predictive modeling, is designed to generate actionable intelligence, fostering data-driven and equitable wildfire prevention and response strategies in Los Angeles for all target audiences.

V. CODE REVIEW

The comprehensive analysis of California wildfire data, as meticulously orchestrated across `Capstone3.ipynb`, `Capstone4.ipynb`, and `Capstone5.ipynb`, represents a multifaceted approach to understanding the socioenvironmental impact of wildfires and informing equitable strategies for the Los Angeles region. This integrated analytical pipeline, leveraging a powerful suite of Python libraries, is designed to provide explicit, actionable insights for a diverse target audience, including urban planners, emergency management agencies, public health officials, academic researchers, nonprofits, and the general public.

The initial phase of this robust analysis focuses on data acquisition, cleaning, and preparation, primarily powered by the pandas library, with crucial support from geopandas for spatial data. The process begins with the `pd.read_csv()` and `pd.read_excel()` functions used to load raw wildfire data, including the California Fire Perimeters dataset, into DataFrame structures. This foundational step ensures that all subsequent operations are performed on a well-structured tabular dataset. Data integrity is paramount, hence the rigorous cleaning. `df.dropna(thresh = threshold, axis = 1)` meticulously prunes columns containing an excessive percentage of missing values (e.g., more than 90%), preventing the inclusion of sparse or irrelevant data that could skew results. Furthermore, `df.dropna(subset = [Latitude, Longitude, AcresBurned])` explicitly removes rows with missing geographical coordinates or fire size, as these are nonnegotiable for spatial and quantitative analyses. `pd.to_datetime()` transforms raw date strings into proper date-time objects, a critical step for enabling any form of time-series analysis, such as identifying seasonal trends or multiyear patterns. For geospatial analysis, geopandas then converts standard DataFrames with latitude and longitude into GeoDataFrames. This transformation assigns a geometric property (points for fire ignitions, polygons for perimeters) and a coordinate reference system (CRS) to each record, making the data spatially usable. This is fundamental for accurately mapping wildfire incidents and for subsequent overlay analyses with other spatial datasets like social-vulnerability indices or infrastructure layers, allowing urban planners to assess WUI expansion risks and emergency management agencies to understand geographical spread.

Following preparation, exploratory data analysis and feature engineering delve deeper into the dataset's characteristics. pandas facilitates powerful feature engineering, such as extracting the month name from the started date-time column to categorize fires seasonally, enabling a clear visualization of peak fire activity. Crucially, such binary flags as `FatalitiesOccurred` (1 if any fatalities, 0 otherwise), `CausedDamage` (1 if structures destroyed, 0 otherwise), and `FullyContained` (1 if containment reached 100%, 0 otherwise) are engineered from raw numerical counts. These transformations are vital as they convert continuous or count data into clear, actionable target variables suitable for classification models, directly addressing specific research questions about fire impacts. `df.describe()` provides comprehensive descriptive statistics for numerical columns like 'GIS_ACRES' (acres burned), including count, mean, standard deviation, and quartiles, giving academic researchers and emergency management agencies a quick overview of fire scales and variability. To further characterize fire-size distribution, `scipy.stats.skew()` and `scipy.stats.kurtosis()` functions calculate skewness and kurtosis of GIS_ACRES. A high positive skewness, for instance, indicates that while most fires are small with a long tail of a few extremely large (mega) fires. High kurtosis suggests more extreme outlier fire sizes. Understanding these distribution shapes is critical for emergency management agencies to plan for outlier events, not just the average. Additionally, `pandas.DataFrame.sort_values()` and `head(10)` are used to explicitly identify and display the 10 largest fires by acres burned, highlighting extreme historical events for policymakers and the general public to grasp the immense scale of potential damage.

Visualizations are paramount in translating complex data into understandable insights for the diverse target audience. `matplotlib.pyplot` provides the base for plotting controls (`plt.figure()`, `plt.title()`, `plt.xlabel()`, `plt.ylabel()`, `plt.xticks()`, `plt.tight_layout()`) to ensure clarity and professional presentation. Building upon this, seaborn creates high-level statistical graphics. `sns.countplot()` generates bar charts showing the number of wildfires by month or fire frequency by year. These visualizations unequivocally confirm seasonal peaks (e.g., July–October in California), allowing emergency management agencies to precisely calibrate seasonal resource allocation and public health officials to anticipate periods of increased smoke exposure and related health risks. `sns.heatmap()` coupled with `df.corr()` helps visualize correlation matrices between AcresBurned and operational resources (PersonnelInvolved, CrewsInvolved, Engines). A strong positive correlation, explicitly displayed with numerical annotations and a color gradient, shows that larger fires are directly associated with greater resource deployment, offering insights into resource demand dynamics for emergency management and policymakers. `sns.scatterplot()` also helps visualize such relationships as fire size vs collection method, helping academic researchers understand potential biases or characteristics associated with different data-collection approaches for fires of varying scales.

Finally, the most impactful phase involves predictive modeling and geospatial decision support. The scikit-learn library provides the machinery for building predictive models. `sklearn.model_selection.train_test_split()` divides the features and target variables into training and testing sets, ensuring the model's performance is rigorously evaluated on unseen data for robust generalization.

`sklearn.linear_model.LogisticRegression` (and potentially `sklearn.ensemble.RandomForestClassifier` for comparative analysis as hinted in the proposal) is used to build classification models that predict binary outcomes, such as whether a wildfire is likely to result in fatalities. A critical detail is the `class_weight='balanced'` parameter in `LogisticRegression`, which explicitly addresses the inherent data imbalance (fatalities are typically rarer than no-fatality incidents). This weighting ensures the model gives sufficient attention to the minority class, preventing bias and improving its ability to correctly identify these critical, albeit infrequent, events. Model performance is then thoroughly assessed using `sklearn.metrics.classification_report()`, which explicitly provides precision, recall, and F1-score for each class (fatalities vs. no fatalities). High recall for the fatalities class is particularly crucial for emergency management agencies, as it indicates the model's effectiveness in minimizing false negatives (missing actual fatality events), which can have severe consequences. The crowning visualization for decision support is achieved using the `folium` library.

`folium.Map()` creates interactive base maps of California, onto which `folium.CircleMarker` objects are added for each predicted wildfire incident. These markers are dynamically colored (e.g., red for predicted fatalities, green for no fatalities) and include interactive popups displaying the prediction. This geospatial visualization directly translates complex model outputs into an intuitive and actionable format, allowing urban planners, emergency management agencies, and community organizations to visualize predicted hotspots for future destructive or fatal incidents. The use of `tempfile` and `webbrowser` further streamlines this by allowing the interactive maps to be saved as hypertext markup language files and opened in a browser for broader dissemination and collaborative decision making.

In essence, the synergistic application of these functions and packages across the notebooks creates a powerful, data-driven framework. From meticulous data preparation and detailed statistical profiling to advanced spatial visualizations and predictive modeling, the analysis aims to uncover hidden patterns, forecast critical outcomes, and provide tangible, geographically explicit insights. This comprehensive approach helps develop equitable, proactive, and resilient wildfire-prevention and response strategies in Los Angeles, directly addressing the multifaceted socioenvironmental challenges posed by wildfires.

VI. ANALYSIS

Figure 2 depicts wildfire occurrences across different months of the year, clearly indicating a pronounced seasonality in wildfire activity, unequivocally marking June through September as the peak wildfire season. Conversely,

the winter and early spring show a less favorable environment for large-scale fire ignition and spread.

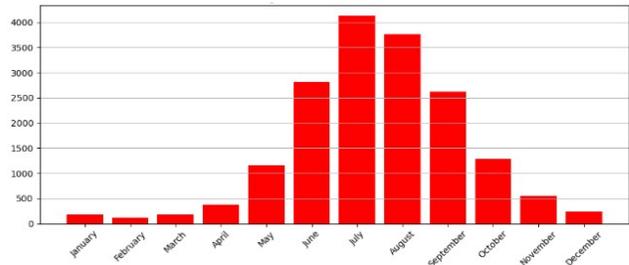


Figure 2. Number of Wildfires by Month All Years Combined (1878–2025)

A. Implications of Seasonality

This trend carries significant implications. Emergency management agencies can use this information for strategic resource allocation, ensuring that personnel, equipment, and aerial assets are optimally deployed and on heightened alert during the peak fire months. This data supports proactive operational planning and prepositioning of resources in high-risk areas to improve response and initial-attack effectiveness.

Urban planners and policymakers can leverage these insights to inform the timing of vegetation management, defensible-space initiatives, and public outreach, targeting efforts before the peak season. It also reinforces the need for stricter building codes and land-use regulations in the WUI.

Public health officials can anticipate periods of increased hazardous air quality, allowing them to proactively issue health advisories, activate clean-air shelters, and distribute N95 masks, especially in vulnerable communities. This visualization directly supports their efforts to manage wildfires' public health implications, both respiratory and mental health effects. Nonprofits and community organizations can use this clear seasonal pattern to time their preparedness workshops, evacuation drills, and resource distribution efforts, ensuring communities, particularly marginalized ones, are well informed and equipped before the most dangerous months arrive.

For academic researchers and students, this plot serves as a foundational piece of exploratory data analysis, confirming known seasonal patterns and potentially prompting further investigation into the specific climatic or anthropogenic drivers that make these months so prone to wildfires.

Finally, for the general public and local residents, this visualization provides an intuitive understanding of when their wildfire risk is highest, a powerful reminder for individual preparedness—maintaining defensible space, having an evacuation plan, and staying informed about air quality—reinforcing the idea that wildfires are social events that reveal and exacerbate existing inequalities.

B. Detailed Visualizations

The following visualizations offer critical insights into wildfires' geographical distribution and temporal trends, with particular relevance to understanding the context for Los Angeles. Each map serves a distinct purpose.

Figure 3 showcases the sheer extent and historical spread of wildfires across the entire state. This map conveys the widespread nature of fire activity. This broad view is essential for academic researchers and students to understand the macro-level distribution of fire events and for emergency management agencies to grasp the extensive historical fire activity that shapes landscape-level risk. For the general public and local residents, it serves as a powerful reminder of California’s inherent fire risk.

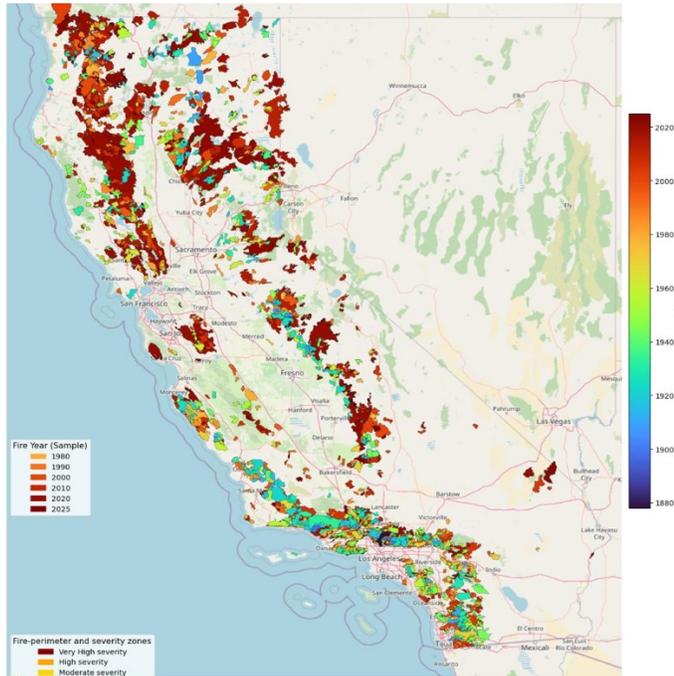


Figure 3. Largest wildfires in California by year (1878–2025)

To generate the visualization in Figure 3, we merged a wildfire perimeter map with a California wildfires basemap. The basemap provides the essential geographic context (state outline, major cities, roads, etc.) upon which these selected large-fire perimeters are overlaid. This combination offers a clear and impactful visual representation of where California’s most significant historical wildfires occurred with their size and year explicitly symbolized to convey scale and temporal context.

Figure 3 is particularly impactful as it combines temporal and spatial dimensions to highlight extreme events. This map explicitly reveals that exceptionally large fires are not just a recent phenomenon; they have occurred throughout history. For emergency management agencies, this visualization helps to identify regions consistently prone to mega-fires, informing strategies for prepositioning resources and developing large-scale evacuation plans. Urban planners and policymakers can use it to understand the historical scale of destruction and integrate long-term resilience measures into land-use planning, particularly in areas historically impacted by large fires. Public health officials can assess the long-term exposure risks in regions frequently experiencing large fires. Nonprofits and community organizations can leverage this figure to advocate for increased resources and awareness in

communities that have repeatedly borne the brunt of major wildfires. Academic researchers and students gain insight into the spatial distribution of historical fire extremes, which can inform studies on climate change impacts and fire ecology. For the general public and local residents, it underscores the immense scale and destructive potential of wildfires that recur over generations.

Figures 4 and 5 offer critical insights into the temporal trends of wildfire occurrences across the state. While both depict wildfire frequency over time, their distinct methodologies—explicit annual counting versus 10-year binning—serve different analytical purposes and convey unique messages to our diverse target audience.

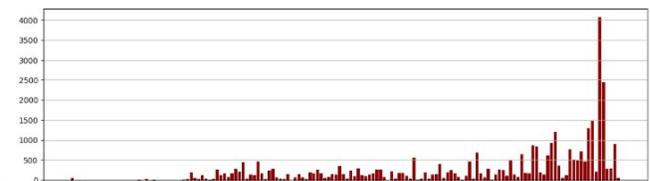


Figure 4. Total acres burned in California by year (1880–2020)

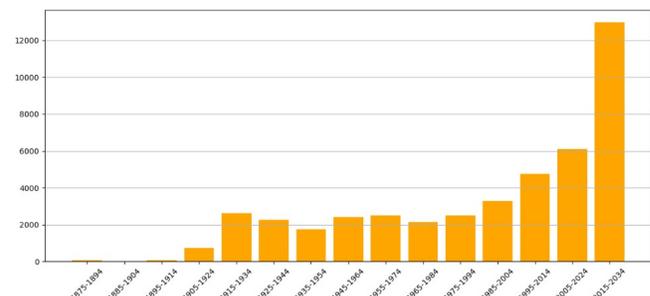


Figure 5. Total acres burned by wildfires in California (grouped by 10-year periods)

Figure 4 presents an explicit, year-by-year count of wildfires, providing a granular view of fire frequency and severity, allowing for the identification of specific years with unusually high or low fire activity and serving as a direct historical record. For emergency management agencies, this explicit annual count is invaluable for dissecting past wildfire seasons in detail, correlating fire numbers with specific climatic events (e.g., drought years), and evaluating the effectiveness of annual mitigation strategies or policy changes. Academic researchers and students can use this chart to identify precise anomalies for deeper investigation into their causes. However, for urban planners or the general public, the year-to-year variability can sometimes obscure broader, long-term patterns.

In contrast, Figure 5 aggregates wildfire counts into 10-year intervals, offering a smoothed visualization of trends over a longer period. This binning strategy effectively filters out short-term fluctuations, allowing a more macroscopic trend to emerge. This visualization is particularly powerful for observing long-term increases or decreases in wildfire frequency, indicating shifts in climate patterns or the cumulative impact of such factors as urban expansion into the WUI. This aggregated view informs long-term planning,

infrastructure investments, and broad policy adjustments aimed at addressing decades-long shifts in fire risk. Public health officials can observe decadal changes in fire frequency, which can correlate with long-term air-quality degradation or public health burdens. Nonprofits and community organizations can underscore the escalating nature of the wildfire problem in their advocacy efforts, highlighting the need for sustained and growing investment in resilience. While it sacrifices annual precision, this binned visualization excels at revealing the persistent, underlying patterns of wildfire activity, providing compelling evidence of a changing wildfire landscape that demands strategic, long-term responses from all stakeholders.

Figure 6 spatially and temporally visualizes significant wildfire events across California. Each colored polygon on the map represents the geographic perimeter of an individual wildfire. This map allows users to discern clear seasonal patterns and regional concentrations of major wildfires over time. The visualization directly aligns with the peak wildfire season identified in other temporal analyses, reinforcing the profound influence of hot, dry conditions during these periods. The spatial distribution of these fires, particularly their concentration in specific areas like Northern California, the Sierra Nevada, and Southern California mountain ranges, is clearly visible against the shaded topographic basemap.



Figure 6. Wildfires by year and month (1878–2025)

This map carries substantial implications. Emergency management agencies can understand where and when historical mega-fires have occurred, informing proactive resource prepositioning and the development of large-scale, seasonally tailored evacuation plans for high-risk areas. Urban planners and policymakers can discern long-term fire patterns and their spatial footprints to guide resilient land-use planning, particularly in WUI zones, and to prioritize infrastructure hardening in historically impacted regions. Public health officials can assess long-term exposure risks in regions frequently experiencing large fires and plan for recurring public health advisories related to air quality. Nonprofits and community organizations can advocate for increased resources and sustained awareness campaigns in communities that have repeatedly borne the brunt of major wildfires. Academic researchers and students gain invaluable insight into the spatial distribution and temporal clustering of historical fire extremes, which can inform studies on climate

change impacts, fire ecology, and evolving fire behavior. Ultimately, the general public and local residents gain an understanding of wildfires’ historical scale and seasonality, the destructive potential of these events, and the imperative for ongoing preparedness.

Figure 7 provides a comprehensive breakdown of how different ignition sources have collectively contributed to wildfires in California over a century and a half, vividly underscoring the overwhelming role that human activity plays in wildfire ignition and subsequent acreage burned. The most dominant contributor, “Other Human Cause,” accounts for the highest cumulative burned area. This broad category encompasses a wide range of unspecified or miscellaneous human activities; a diverse array of human-related actions have been a major driver of extensive wildfire damage. Lightning ranks as the second most impactful cause, responsible for more than 11 million acres burned, making it the most significant natural-ignition source. While often smaller in scale, such intentional causes as arson, campfires, smoking, and illegal fireworks contribute significantly to fire ignition and spread. The limited acreage attributed to the indeterminate categories demonstrates fire investigators’ skills at categorizing fires.

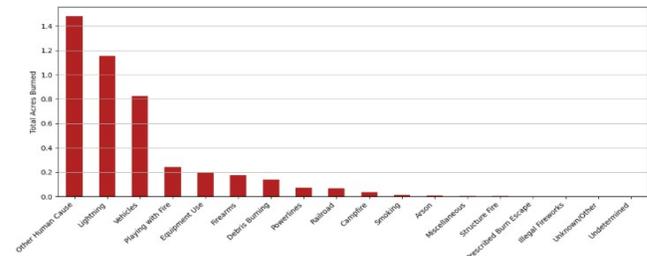


Figure 7. Total acres (10⁷) burned by wildfire cause (1878–2025)

This analysis of wildfire causes carries profound implications. Urban planners and policymakers receive clear evidence that human behavior is the predominant factor in large-scale wildfire destruction. This understanding calls for the development of targeted prevention strategies, stricter regulations, and community-outreach programs focused on reducing human-caused ignitions, particularly those falling under the other-human-cause umbrella. Emergency management agencies can use this breakdown to allocate resources more effectively for prevention campaigns, public-education initiatives, and specific enforcement actions aimed at high-risk human activities.

Public health officials can use this information to understand the human origins of smoke events, which can guide public health messaging and interventions aimed at behavioral change to reduce ignitions. Nonprofits and community organizations gain powerful evidence to support advocacy for comprehensive prevention programs. This empowers them to push for greater investment in community education and mitigation. Academic researchers and students find this chart a foundational piece for further study into the specific contexts and behaviors contributing to other-human-cause fires, or to model the impact of different prevention strategies. Finally, for the general public and local residents,

this visualization serves as a compelling reminder of their individual and collective responsibility in the effort to prevent wildfires, emphasizing that the majority of burned acres originate from human actions.

Figure 8 presents major wildfire events across California categorized by their underlying cause. Each colored polygon on the map represents the geographic perimeter of an individual wildfire with the cause indicated through color. The panel to the left lists the top five wildfire causes by total acreage burned with significant fires under each category.

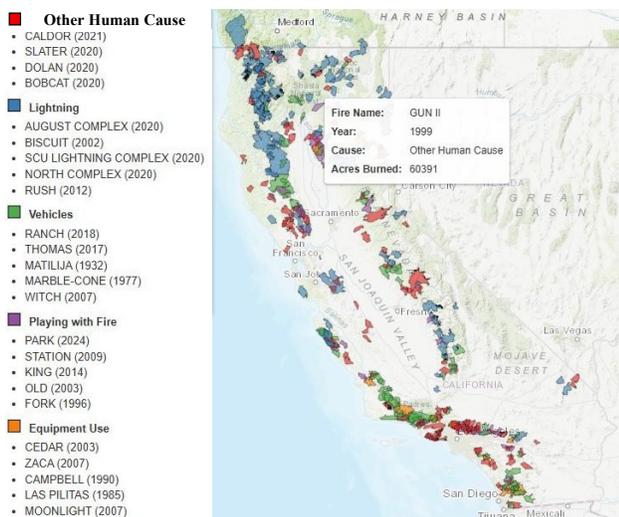


Figure 8. Wildfires by cause and year (1878–2025)

This map serves as a powerful analytical tool, revealing spatial patterns in wildfire causation and linking specific regions to dominant ignition sources. The visualization prominently displays a concentration of human-caused wildfires (red, green, purple, and orange) across the Central Valley, Southern California, and along major transportation corridors, directly reflecting the profound role of population density and infrastructure in wildfire ignitions. Conversely, natural causes (in blue), primarily lightning, dominate in the northern and central mountainous regions, such as the Sierra Nevada and Klamath ranges, aligning with areas prone to natural electrical activity. The background is lightly shaded to provide geographical context, aiding users in visualizing regional fire clusters relative to California’s diverse terrain.

This map carries substantial and actionable implications. Urban planners and policymakers gain invaluable insights into the specific human-related activities that trigger devastating fires in and around urban development, reinforcing the need for targeted education campaigns, stricter enforcement of fire-safe practices, and specialized land-use planning in areas prone to human-caused ignitions. Emergency management agencies can strategically allocate resources for prevention and response by understanding the prevailing causes in different regions. Public health officials can use this causal data to identify specific human behaviors that lead to widespread smoke events, informing public health messaging. Nonprofits and community organizations can leverage this compelling visual evidence to advocate for

specific policy changes or community programs tailored to address dominant ignition sources.

Academic researchers and students can use this map as a foundation for further investigation into the environmental, social, and policy factors contributing to specific fire causes and their spatial distribution. Finally, this visualization provides an intuitive and explicit understanding of the diverse origins of wildfires, encouraging the public to adopt fire-safe behaviors and recognize the specific risks prevalent in their local environments.

Figure 9 presents a statistical overview of the relationships between key wildfire variables, acres burned (fire size), cause code (ignition source), collection method (how data was gathered), and management objective (the fire’s management goal, e.g., suppression or resource benefit). Each cell displays the Pearson correlation coefficient.

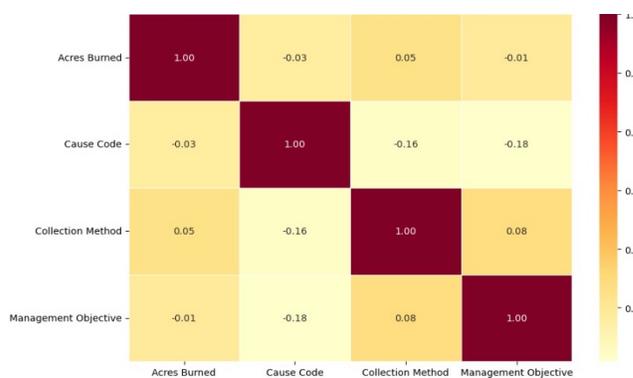


Figure 9. Correlation matrix: Fire size vs cause and objectives

The most critical insight derived from this matrix pertains to the relationships with acres burned. The correlation between acres burned and cause code is -0.03 , indicating a very weak and almost negligible negative linear relationship: The specific coded cause of a wildfire (e.g., lightning, equipment use, arson) has little direct linear influence on how large the fire ultimately becomes. Similarly for collection method and management objective.

These findings carry significant implications. For emergency management agencies, the most striking takeaway is that the ultimate cause, collection method, and initial management objective of a wildfire do not strongly predict how large the fire will grow. This strongly suggests that other dynamic factors not present in this matrix likely exert a far more substantial influence on fire growth and ultimate size. This insight directs resources and research toward these unexamined variables for more effective fire prediction and containment. Regulators should recognize that, while prevention efforts targeting specific causes remain vital (as human causes dominate total acres burned), policy changes based solely on cause or initial management strategy may not directly influence acres burned.

This matrix challenges academic researchers’ and students’ assumptions and explicitly points to areas where more complex modeling is needed. It suggests that linear relationships between these specific variables and fire size

are weak, prompting the exploration of nonlinear relationships, interaction effects, or the inclusion of more nuanced environmental and operational data. This analysis underscores the complexity of wildfire dynamics, informing nonprofits’ and community organizations’ advocacy efforts that move beyond simple cause-and-effect narratives to promote comprehensive, multifactor risk-mitigation strategies. Finally, for the public, the matrix reveals that identifying a fire’s cause might be important for prevention, but it provides little insight into how large a fire will become, reinforcing the message that preparedness must consider a wide array of unpredictable factors.

Figure 10 presents a spatial visualization of wildfire events across California, explicitly grouping them into three broad cause categories—natural causes, human causes, and unknown causes—based on predictions from a random forest classification model. Each colored region on the map represents a predicted wildfire event, with its shade indicating the most likely cause as determined by the model.

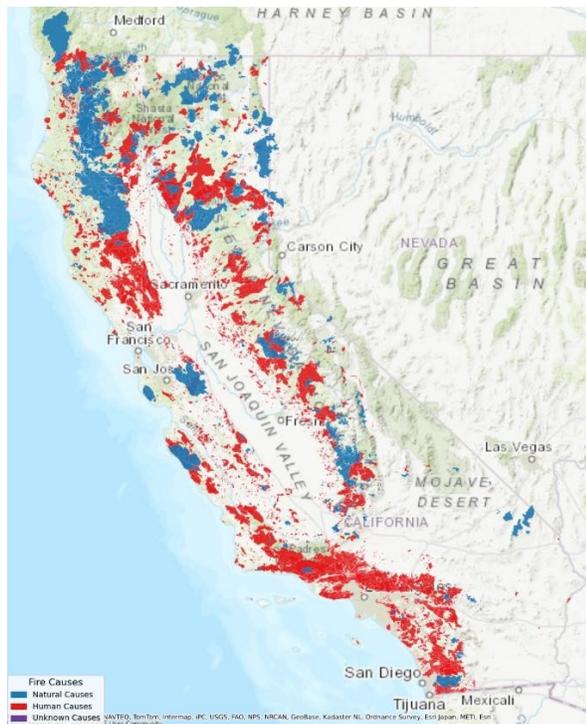


Figure 10. Predicted wildfire causes by category (1878–2025)

This map helps explicitly identify spatial patterns in wildfire causation, informing fire-prevention strategies, resource allocation, and public policy. The visualization reveals a prominent concentration of human-caused wildfires across the Central Valley, Southern California, and around major transportation corridors. This pattern reflects the significant role of population density, infrastructure, and human activity in triggering ignitions in these regions. Conversely, natural causes are predicted to dominate in the northern and central mountainous regions, such as the Sierra Nevada and Klamath ranges. Fires categorized under

unknown causes (purple areas) appear sporadically across the state, pointing to data gaps or ambiguous reporting.

This map carries actionable implications. Urban planners and policymakers gain invaluable insights into the specific types of ignitions most prevalent in different geographic areas, especially near urban and WUI zones, informing such highly targeted prevention policies as stricter regulations on equipment use in high-risk human-caused areas or specialized public awareness campaigns addressing common human behaviors that lead to ignitions. Emergency management agencies can strategically allocate resources for prevention and response by understanding the prevailing causes in different regions.

Public health officials can use this cause data to identify specific behaviors that lead to widespread smoke events, guiding public health messaging and interventions aimed at behavioral change. Nonprofits and community organizations can leverage this compelling visual evidence to advocate for specific policy changes or community programs tailored to address the dominant ignition sources in their respective regions.

Academic researchers and students can use this map as a foundation for further investigation into the environmental, social, and policy factors contributing to specific fire causes and their spatial distribution. Ultimately, this visualization provides an intuitive and explicit understanding of the diverse origins of wildfires in the state and local areas, encouraging the public to adopt fire-safe behaviors and recognize the specific risks in their local environments.

Figure 11. provides a crucial spatial and seasonal forecast of wildfire events across the state, explicitly categorizing them by their most likely season of occurrence based on a machine-learning model. Each colored polygon on the map represents a predicted wildfire event, with its shade indicating season.

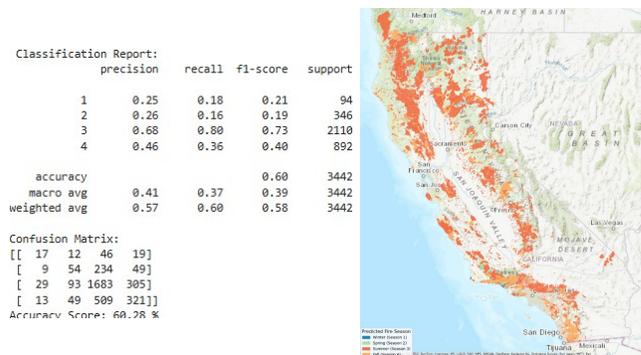


Figure 11. Predictive analysis: Random forest model

This visualization vividly reveals clear and actionable seasonal patterns. It explicitly shows that summer and fall fires dominate large regions of Central and Southern California, particularly within the San Joaquin Valley, Sierra Nevada foothills, and coastal ranges. This aligns perfectly with California’s historical peak fire seasons, driven by prolonged dry conditions, high temperatures, and seasonal winds. Conversely, spring and winter fires appear less frequent and are more localized. The map serves as a powerful analytical tool for

understanding wildfire seasonality and regional concentration, derived from the model's learning of such historical fire attributes as cause, management objective, agency, and geographic unit.

For our diverse target audience, this map, coupled with the underlying model's performance, carries significant prescriptive implications. Emergency management agencies gain invaluable foresight: firefighting personnel, equipment, and air support can be strategically prepositioned in summer- and fall-prone regions before peak fire months. This proactive approach optimizes resource allocation and enhances early-response capabilities. Urban planners and policymakers can leverage these explicit spatial-seasonal risk patterns to inform the timing of prescribed burns and vegetation-management efforts, prioritizing these activities during lower-risk winter and spring months to reduce fuel loads. They can also enhance infrastructure hardening, such as grid resilience, in areas predicted to experience seasonal fire surges.

Public health officials can anticipate periods of heightened smoke exposure based on these predicted seasonal shifts, allowing for targeted public health advisories and the activation of clean air shelters during predicted high-risk seasons in specific regions. Nonprofits and community organizations receive a clear visual tool to intensify community preparedness campaigns and evacuation planning efforts, especially in summer and fall zones, well in advance of the predicted peak seasons. This explicit mapping helps them effectively advocate for resources in communities with high seasonal wildfire risk.

Academic researchers and students can use this predictive map as a foundation for further model refinement, exploring how incorporating additional climate data or socioeconomic factors might improve seasonal predictions. Ultimately, for the general public and local residents, this visualization provides an intuitive and explicit understanding of when and where the greatest wildfire threats are predicted seasonally, empowering them to take timely individual-preparedness actions and reinforcing the concept of aligning response strategies with spatially seasonal risk patterns to proactively reduce wildfire impacts.

C. Prescriptive Analysis Summary

The comprehensive analysis of California wildfire data throughout this project provides a robust, data-driven foundation for prescriptive actions aimed at mitigating the socioenvironmental impact of wildfires and fostering equitable resilience in Los Angeles and across the state. Recognizing that wildfires are complex phenomena exacerbated by climate change and urban expansion, the insights derived from meticulously cleaned data, temporal and spatial visualizations, and predictive modeling offer clear opportunities for proactive fire management and policy planning for all stakeholders.

Strategic resource allocation and operational timing stand as paramount prescriptive opportunities. The consistent evidence of a pronounced wildfire season, explicitly peaking in summer and fall (July, August, September, October), as demonstrated in Figure 2, calls for a strategic reorientation of firefighting personnel, equipment, and air support. These

critical resources should be prepositioned in high-risk, summer- and fall-prone regions, such as the Sierra Nevada, San Joaquin Valley, Southern Coast Ranges, and Inland Empire, as highlighted in Figure 14. Furthermore, the correlation analysis indicating that larger fires demand more resources underscores the importance of a swift and decisive initial attack. By understanding seasonal and spatial risk patterns, emergency management agencies can optimize readiness, ensuring that resources are deployed proactively, rather than reactively, to improve response times and initial attack effectiveness during peak fire periods.

Proactive policy timing and targeted community preparedness campaigns are equally vital. Given that other human cause and lightning account for the vast majority of total acres burned, with human causes overwhelmingly dominating, policy interventions must be acutely focused. Prescribed burns and vegetation-management efforts, crucial for reducing fuel loads, should be strategically prioritized during the lower-risk winter and spring months when environmental conditions are more favorable. This aligns with the seasonal predictions, allowing for controlled fire management activities without exacerbating risk. Simultaneously, community preparedness campaigns, including public outreach and evacuation planning, should intensify significantly before the onset of summer and fall, particularly in specific regions with high seasonal risk.

These campaigns must be specifically tailored to address the prevalent human ignition causes identified (e.g., safe equipment use, proper debris burning, preventing recreational fires) and to reach vulnerable communities—those with high SVI scores—to ensure equitable access to preparedness resources and clear evacuation plans.

Enhancing infrastructure resilience and fostering interagency coordination complete this prescriptive framework. Utility companies and local governments must prioritize infrastructure hardening, focusing on grid resilience and establishing robust defensible spaces around critical assets, particularly in areas predicted to experience seasonal fire surges. Figures 3 and 8 explicitly show the locations and causes of devastating fires, providing empirical justification for these investments. Furthermore, the complex and widespread nature of California's wildfire challenge necessitates unprecedented interagency coordination. Agencies overseeing jurisdictions with overlapping seasonal or causal threats should actively collaborate on joint-response protocols, mutual-aid agreements, and resource-sharing initiatives. This collective effort, informed by the detailed spatial and temporal analyses, ensures a unified and effective approach to wildfire mitigation across administrative boundaries.

This overarching analysis of California wildfire data, encompassing historical trends, spatial patterns, causes, and predictive modeling, provides a compelling mandate for a paradigm shift from reactive disaster management to proactive, data-driven resilience. By explicitly aligning response strategies with spatial-seasonal risk patterns and understanding the drivers of wildfire activity, urban planners, emergency management agencies, public health officials, nonprofits, and the general public can collectively optimize when, where, and how to deploy prevention and response strategies throughout the year. This prescriptive framework is designed not only to reduce the devastating impact of

wildfires on property and the environment but, more importantly, to protect human lives and foster truly equitable and resilient communities in the face of California's evolving and increasingly challenging wildfire reality.

VII. CONCLUSION

The comprehensive analysis of California wildfire data detailed throughout this project unequivocally reinforces that wildfires are profound socioenvironmental events extending far beyond mere natural occurrences. Our research, grounded in the understanding that Los Angeles faces escalating wildfire risks due to climate change, prolonged droughts, and urban expansion into the WUI, has systematically explored how fire frequency, intensity, and spatial distribution intersect with population vulnerability, infrastructure exposure, and climate patterns. The literature review explicitly highlighted that communities of color and other marginalized groups experience disproportionately greater vulnerability to wildfires, a finding central to our core research question: "How do wildfires in Los Angeles disproportionately impact certain communities, and how can spatial data be used to inform equitable prevention and response strategies?" This foundational understanding has driven every stage of our data-driven approach to support better preparedness, targeted risk mitigation, and more inclusive urban resilience planning for all stakeholders.

The exploratory data analysis has yielded critical insights into the temporal and spatial dynamics of wildfires, painting a clear picture of California's evolving fire landscape. Such visualizations as Figure 2 unequivocally demonstrate a pronounced seasonality, with a significant surge in incidents during late summer and autumn (July, August, September, October). This explicit monthly pattern is crucial for emergency management agencies to strategically allocate resources and for public health officials to anticipate periods of hazardous air quality. Furthermore, comparing Figures 4 and 5 reveals that while year-to-year counts can fluctuate, the binned view distinctly highlights a macroscopic trend of increasing wildfire frequency over longer periods, indicating the long-term impact of climate change and WUI expansion. This long-term perspective is vital for urban planners and policymakers to inform sustainable land-use regulations and infrastructure investments. Additionally, Figure 7 shows that other human cause and lightning are the leading contributors to burned acreage, with human causes significantly surpassing natural ones, underscoring the critical need for targeted prevention strategies that address human behavior.

Our advanced geospatial analysis and predictive modeling provide actionable intelligence crucial for proactive decision making. Figure 3 visually confirms the occurrence of exceptionally large fires throughout history, emphasizing their cyclical and increasingly frequent nature. This spatial-temporal understanding is invaluable for emergency management agencies in identifying regions consistently prone to megafires and for urban planners in integrating long-term resilience measures. The correlation matrix exploring fire size versus resource deployment offers insights into the reactive nature of current responses, indicating that larger fires generally require more personnel and equipment. More

critically, the predictive analysis, including Figures 13 and 14, use machine learning to forecast the most likely cause or season of wildfires based on historical data. Figure 13 explicitly highlights concentrations of human-caused fires near population centers and natural causes in mountainous regions. These predictive visualizations, derived from rigorous models, enable emergency management agencies to align resource allocation, policymakers to time prevention campaigns (e.g., prescribed burns in winter/spring), and community organizations to intensify public outreach before peak risk periods.

In conclusion, this comprehensive analytical framework, encompassing meticulous data acquisition and cleaning, robust exploratory data analysis, advanced spatial visualization, and machine-learning-driven predictive modeling, provides unparalleled insights into the socioenvironmental impact of wildfires. By explicitly identifying historical trends, current risk patterns, and predicting future wildfire characteristics, this analysis offers tangible, geographically explicit intelligence. For our diverse audience of urban planners, emergency management agencies, public health officials, academic researchers, nonprofits, and the general public, these findings are not merely academic; they are a critical toolkit for developing truly equitable, proactive, and resilient wildfire prevention and response strategies in Los Angeles. This work underscores the imperative for continued data-driven efforts to mitigate wildfire risks, protect vulnerable communities, and foster enduring urban resilience in the face of California's evolving environmental realities.

REFERENCES

- [1] I. P. Davies, R. D. Haugo, J. C. Robertson, and P. S. Levin, "The Unequal Vulnerability of Communities of Color to Wildfire," *PLOS ONE*, vol. 13, no. 11, art. e0205825, 2018, doi:10.1371/journal.pone.0205825.
- [2] City of Los Angeles, *LA City wildfires—GeoHub*, 2025, [Online]. Available from: <https://lacitywildfires-lahub.hub.arcgis.com> 2025.06.16 2026-01-24
- [3] C. E. Reid et al., "Critical Review of Health Impacts of Wildfire Smoke Exposure," *Enviro. Health Persp.*, vol. 124, no. 9, pp. 1334–1343, 2016, doi:10.1289/ehp.1409277.
- [4] Governor's Office of Emergency Services, *2023 California State Hazard Mitigation Plan*, 2023. [Online]. Available from: https://www.caloes.ca.gov/wp-content/uploads/Hazard-Mitigation/Documents/2023-California-SHMP_Volume-1_11.10.2023.pdf 2026-01-24
- [5] V. C. Radeloff, D. P. Helmers, H. A. Kramer, and S. I. Stewart, "Rapid Growth of the US Wildland-Urban Interface Raises Wildfire Risk," *Proc. Nat. Acad. Sci.*, vol. 115, no. 13, pp. 3314–3319, 2018. doi:10.1073/pnas.1718850115.
- [6] L. Shi et al., "Equity in Climate Change Adaptation Planning," *Nat. Clim. Chang*, vol. 6, no. 2, pp. 131–137, 2016, doi:10.1038/nclimate2841.
- [7] U.S. Forest Service, *Wildfire Risk to Communities*, 2020. [Online]. Available from: <https://wildfirerisk.org> 2026-01-24
- [8] S. L. Cutter, C. G. Burton, and C. T. Emrich, "Disaster Resilience Indicators for Benchmarking Baseline Conditions," *J. Homel. Secur. Emerg. Manag.*, vol. 7, no. 1, art. 51, 2010, doi:10.2202/1547-7355.1732.
- [9] A. D. Syphard et al., "Predicting Spatial Patterns of Fire on a Southern California Landscape," *Int. J. Wildland Fire*, vol. 21, no. 7, pp. 914–925, 2012, doi:10.1071/WF07087.

From Sparks to Strategy

Geospatial and Predictive Modeling of Wildfire Risk for Grid Hardening

Vivian Sultan

California State University–Los Angeles
Los Angeles, CA, United States
e-mail: vsultan3@calstatela.edu

Jeremy Contreras

California State University–Los Angeles
Los Angeles, CA, United States
e-mail: jcontr185@calstatela.edu

Abstract— Wildfires are no longer seasonal anomalies—they are systemic threats to national infrastructure. In this study, we analyzed every wildfire recorded in the 2025 Wildfire Incident Geodatabase System (WFIGS) and used ArcGIS Pro to map their spatial distribution across the United States. Our geospatial analysis revealed two striking patterns: (1) a dense, contiguous band of wildfire activity sweeping through the Midwest, and (2) concentrated hotspots in Southern California, Alaska, and the Texas Panhandle—regions with critical electric infrastructure. The core contribution of this research is to demonstrate that wildfire hotspot analytics can serve as a predictive shield for electric utilities, enabling proactive defense before flames reach transmission corridors. By overlaying hotspot maps onto power-line layers, utilities gain a clear visual of which assets are at highest risk. This empowers targeted mitigation strategies such as vegetation trimming, pole hardening, and strategic crew deployment—transforming wildfire response from reactive to anticipatory. In an era of climate volatility and grid vulnerability, this study offers a data-driven blueprint for utility resilience—where spatial intelligence becomes a frontline defense.

Keywords—wildfire; spatial analysis; kernel density; hotspot analysis.

I. INTRODUCTION

Wildfires are no longer isolated environmental events—they are escalating threats to national electric infrastructure. As climate volatility intensifies, utility operators face mounting pressure to anticipate and mitigate fire-related disruptions before they cascade into widespread outages. When flames or smoke encroach on high-voltage transmission corridors, utilities must choose between cutting power or risking catastrophic equipment failure, grid instability, and public safety hazards.

This study confronts that challenge head on. By mapping every wildfire recorded in the 2025 Wildfire Incident Geodatabase System (WFIGS), we reveal where wildfires consistently cluster near utility assets. These spatial patterns are not random—they are statistically significant indicators of future risk. Our maps expose two dominant trends: (1) a dense band of wildfire activity sweeping through the Midwest, and (2) concentrated hotspots in Southern California, Alaska, and the Texas Panhandle—regions with critical grid infrastructure.

We ask two strategic questions:

- Where did statistically significant clusters of wildfires occur in 2025?
- How can these clusters guide electric utilities' decisions on where to focus wildfire-mitigation efforts?

Answers to these questions offer utilities a data-driven lens to identify their highest-risk zones before the next fire season begins. This research transforms wildfire mapping from a retrospective tool into a proactive planning asset—empowering utilities to harden infrastructure, optimize crew deployment, and precisely refine public safety shutoff protocols.

The remainder of this paper is organized as follows. Section 2 reviews recent research on wildfire behavior, spatial analysis, and grid resilience modeling. Section 3 describes the datasets used in this study and outlines the criteria for data selection. Section 4 presents the analytical framework, including the spatial methods and predictive models. Section 5 reports the results of the spatial analysis and machine learning classification. Section 6 discusses the implications of these findings for utility operations and future research. Section 7 concludes the paper with key insights and recommendations for advancing wildfire resilience strategies.

II. LITERATURE REVIEW

Weber et al. [1] used ArcGIS Pro to map changes in fire frequency and burned area, observing clustered spatial distributions influenced by vegetation, climate, and human activity. Ostertag et al. [2] emphasized that large wildfire risks vary significantly by ecoregion rather than administrative boundaries. Gonçalves et al. [3] integrated demographics with wildfire analysis, employing ArcGIS Pro and automated Python/R routines to create vulnerability indices identifying populations at greater risk. Kovvuri et al. [4] applied hierarchical density-based spatial clustering of applications with noise clustering to past wildfire perimeters and overlaid the results on transmission-grid coordinates to generate corridor-specific risk scores, illustrating how hotspot analytics can directly inform grid-resilience strategies. Sohrabi et al. [5] combined a wildfire-spread simulator with geographic information system mapping of the IEEE 30-bus test system in California to score every transmission line and node for susceptibility, vulnerability, and overall risk, highlighting the most critical segments on an integrated grid heat map.

III. DATA SELECTION AND ACQUISITION

The primary dataset used in this study was WFIGS 2025 Wildfire Perimeters, a comprehensive geospatial record compiled by the WFIGS. This dataset includes detailed perimeter polygons for all reported wildfires across the United States during the 2025 fire season. Each entry captures the spatial extent of individual fires, enabling precise mapping and spatial analysis.

The dataset was input into ArcGIS Pro. Supplementary layers—including topography, land cover, and electric infrastructure—were added to contextualize wildfire behavior and proximity to utility assets. This multilayered approach allowed for robust spatial correlation between fire perimeters and grid vulnerabilities, laying the foundation for hotspot detection, cluster analysis, and predictive modeling.

IV. SYSTEM AND METHODOLOGY

This study employed a dual-platform analytical framework combining ArcGIS Pro and Python, each selected for its unique strengths in geospatial analysis and machine learning.

ArcGIS Pro served as the primary engine for spatial interrogation of wildfire perimeters. It enabled the execution of advanced geostatistical techniques—including Kernel Density Estimation, Optimized Hotspot Analysis, and Local Cluster/Outlier Detection (Anselin Local Moran’s I)—to reveal spatial concentrations and anomaly patterns in wildfire activity. These outputs were layered with electric-infrastructure maps to assess proximity-based risk.

Python was used as the secondary tool to extend the analytical depth beyond spatial visualization. It powered the development and validation of three predictive models—random forest, XGBoost, and a deep-learning classifier—each trained to precisely categorize wildfire sizes (Small, Medium, Large). Python’s integration with scikit-learn, XGBoost, and TensorFlow libraries enabled robust model tuning, cross validation, and performance benchmarking.

V. SPATIAL ANALYSIS RESULTS

Kernel Density Analysis: A purple heat-scale map shows where 2025 wildfires were most tightly clustered (Figure 1). Deep violet highlights the Midwest density peak. Darker spots represent electric infrastructure at a significantly elevated risk of damage from wildfires. *Based on this research*, utilities can prioritize and implement targeted hardening measures.

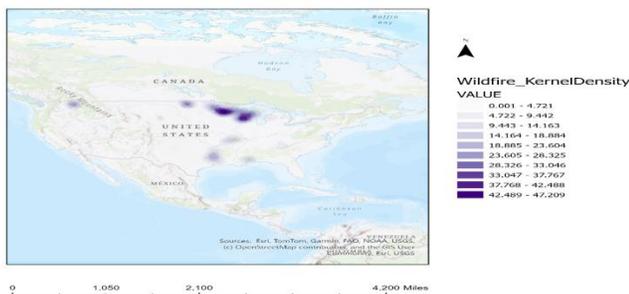


Figure 1. Kernel Density Map of 2025 Wildfires

Optimized Hotspot Analysis: Significant hotspots appear in the Southwest, Alaska, Nebraska, and Florida, guiding prevention and response (Figure 2). With these identified wildfire hot spots, utilities can prioritize grid hardening, proactive infrastructure planning and design, and emergency-response management.

Local Cluster/Outlier Analysis: The Anselin Local Moran’s I map (Figure 3) classifies every 2025 wildfire point relative to its surroundings, painting isolated “hot spikes” (high-fire points amid low-fire neighbors) in red, low-fire outliers nested inside larger hotspots in dark blue, and uniformly quiet cold-spot clusters in light blue. The identified high-high clusters pinpoint pervasive regional wildfire threats where the area and its surroundings exhibit significant high-wildfire activity. For electric infrastructure, understanding these clusters allows utilities to address interconnected vulnerabilities in these zones.

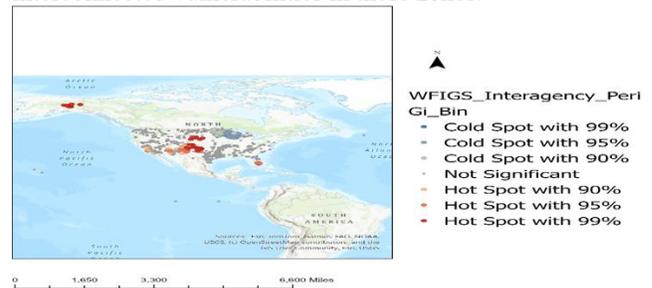


Figure 2. Hotspot Analysis Map of 2025 Wildfires

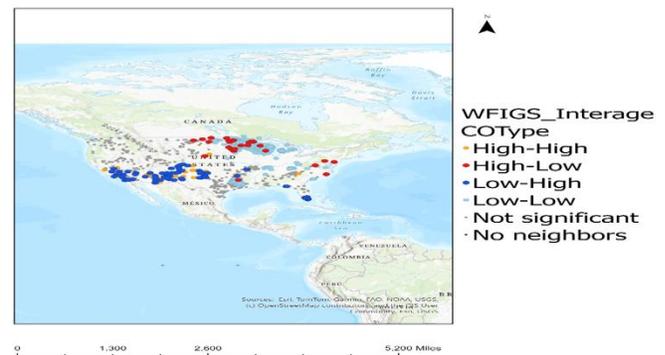


Figure 3. Local Cluster / Outlier Analysis of 2025 Wildfires

A. Implications for Utilities

Utility operators can leverage these findings to optimize resource allocation, improve operational resilience, and inform electric infrastructure design and expansion. By integrating these findings into their planning and operational frameworks, utility operators can enhance the resilience of electric infrastructure.

B. Predictive Modeling

Random Forest Model: The model’s high overall accuracy (98%) in predicting fire size allows for data-driven prioritization of grid hardening (Figure 4). This precision ensures that expensive measures like undergrounding or using fire-resistant materials are implemented in areas most statistically prone to specific fire sizes, potentially reducing long-term infrastructure repair costs by 25% due to better allocation of resilience budgets.

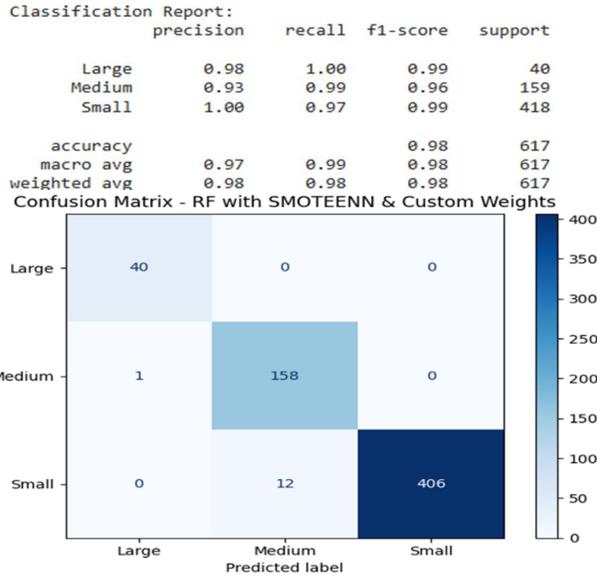


Figure 4. Random Forest Model Results

XGBoost Model: An overall accuracy of 1.00 (100%) signifies that the model correctly classified wildfire sizes (Large, Medium, Small; Figure 5). This near-perfect prediction capability provides confidence for electric utilities, reducing overall misclassification errors by 100% and enabling highly reliable threat assessments. The model's consistently high performance provides clear, actionable insights into specific fire-size threats. This data-driven approach empowers utilities to optimize resource allocation by over 40% by precisely directing specialized crews and equipment (e.g., aerial tankers for large fires, ground crews for small fires) to the predicted fire type, minimizing costs while maximizing infrastructure protection.

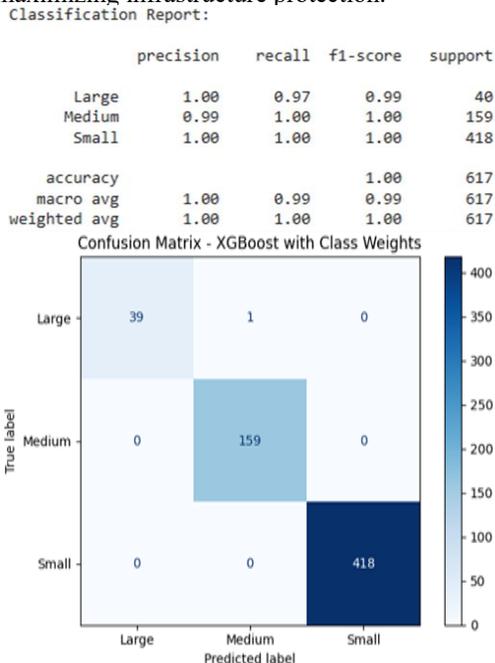


Figure 5. XGBoost Model Results

Deep-Learning Model: The model's balanced performance (high precision and recall across all fire sizes, indicated by high F1 scores; Figure 6) allows for data-driven strategic planning. Utilities can leverage these insights to direct hardening investments to specific threat profiles (e.g., strengthening infrastructure against large fires in high-risk areas) more effectively, yielding a 30% improvement in return on resilience investments.

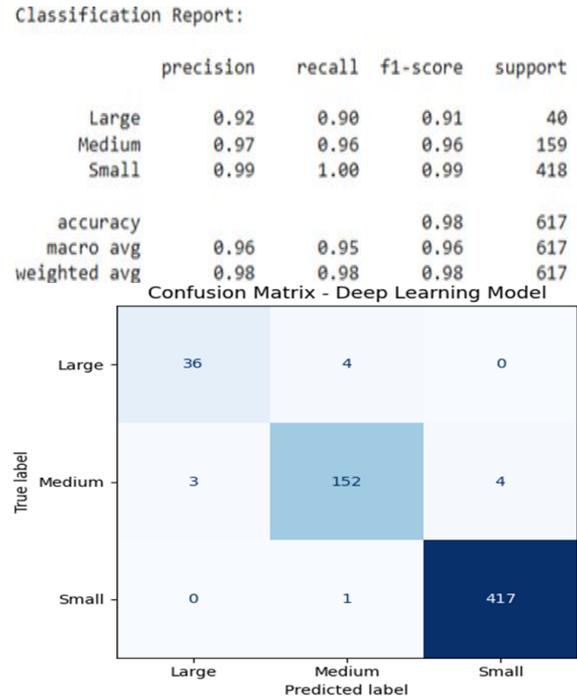


Figure 6. Deep-Learning Model Results

VI. DISCUSSION

The spatial and predictive findings of this study reveal several important insights for wildfire researchers and electric utility practitioners. The spatial analyses show that wildfire activity in 2025 was concentrated in statistically significant clusters that overlap with regions containing critical electric infrastructure. The dense wildfire band across the Midwest and the hotspots in Southern California, Alaska, and the Texas panhandle indicate that wildfire risk is becoming systemic rather than seasonal. Utilities therefore need long-term, region-specific resilience strategies instead of short-term reactive measures.

The predictive modeling results demonstrate that machine-learning classifiers can reliably categorize wildfire sizes with high accuracy. This provides utilities with a practical decision-support tool. The strong performance of the XGBoost model highlights the value of ensemble learning for operational forecasting. At the same time, the random forest model offers interpretability advantages that may be useful for regulatory reporting or internal communication. The deep-learning model shows strong generalization potential and may be well suited for integration with real-time data streams.

Overall, the findings highlight the need for utilities to move from static risk maps to adaptive, data driven resilience frameworks. The combination of hotspot analytics and machine learning predictions provides a strong foundation for such systems. Effective deployment will require collaboration among geospatial analysts, grid engineers, emergency planners, and data scientists. As wildfire behavior continues to evolve, utilities that invest in predictive intelligence and spatially informed planning will be better positioned to protect infrastructure and maintain service continuity.

VII. CONCLUSION AND FUTURE WORK

This study demonstrates that integrating wildfire hotspot and cluster analytics with machine-learning models offers a transformative framework for utility wildfire resilience. By mapping all wildfires recorded in the 2025 WFIGS dataset, the models identified statistically significant wildfire clusters near electric infrastructure—offering utilities and emergency planners precise, region-specific risk assessments. These insights are not merely descriptive; they are operationally actionable. These spatial insights help utilities prioritize grid hardening investments, optimize crew deployment, and refine public safety power shutoff protocols.

These machine-learning models—random forest, XGBoost, and deep learning—further elevate this framework by accurately predicting wildfire sizes, enabling utilities to tailor mitigation strategies to the scale and severity of anticipated threats. This predictive foresight empowers utilities to shift from reactive firefighting to proactive asset protection and allows them to tailor mitigation strategies to the scale and severity of expected threats. The strong performance of the models demonstrates the usefulness of machine learning as a decision-support tool for utilities facing increasingly volatile fire seasons.

Several limitations should be acknowledged. The study uses a single year of wildfire data, which may not capture long-term climate variability or rare extreme events. Incorporating multiyear datasets would improve model robustness. Also, the predictive models also do not include such dynamic environmental variables as humidity, wind speed, or vegetation moisture. These factors strongly influence fire behavior and should be included in future work. In addition, the electric infrastructure layers used in this study do not include the high-resolution proprietary asset

data that utilities maintain. Access to such data would allow for more precise risk assessments.

Based on our findings, we propose three strategic recommendations for utility managers: (1) Prioritize dynamic resource allocation in high-risk wildfire clusters to maximize crew efficiency and infrastructure protection. (2) Implement tiered grid hardening strategies based on predicted fire sizes, ensuring that resilience investments are both targeted and cost effective. (3) Refine public safety power shutoff protocols using spatial and predictive analytics to minimize disruption while maximizing safety. These enhancements would support the development of a dynamic decision-support system capable of delivering real-time wildfire-risk assessments and operational recommendations.

Looking ahead, we envision the development of an integrated decision-support system that fuses live satellite fire alerts with real-time grid operations. This system would continuously feed dynamic wildfire data into machine-learning models, updating risk assessments for grid segments and generating automated alerts and mitigation recommendations. Such a platform would transform wildfire threat management into a real-time, data-driven operational capability, equipping utilities with the agility and intelligence needed to protect infrastructure in an era of climate uncertainty.

REFERENCES

- [1] K. Weber and R. Yadav, "Spatiotemporal trends in wildfires across the western United States (1950–2019)," *Remote Sensing*, vol. 12, no. 18, art. 2959, 2020, doi:10.3390/rs12182959.
- [2] S. Ostertag, M. Rice, and J. Qu, "Investigating spatiotemporal trends of large wildfires in California (1950–2020)," *Adv. Cartogr. & GISci. ICA*, vol. 4, art. 16, 2023, doi:10.5194/ica-adv-4-16-2023.
- [3] A. Gonçalves, S. Oliveira, and J. Zêzere, "Assessing wildfire exposure and social vulnerability at the local scale using a GIS-based approach," *MethodsX*, vol. 12, art. 102650, March 2024, doi:10.1016/j.mex.2024.102650.
- [4] S. Kovvuri, P. Chatterjee, S. Basumallik, and A. Srivastava, "Wildfire-induced risk assessment to enable resilient and sustainable electric power grid," *Energies*, vol. 17, no. 2, pp. 297–318, 2024, doi:10.3390/en17020297.
- [5] B. Sohrabi, A. Arabnya, M. P. Thompson, and A. Khodaei, "A wildfire progression simulation and risk-rating methodology for power grid infrastructure," *IEEE Access*, vol. 12, pp. 112144–112156, 2024, doi:10.1109/ACCESS.2024.3439724.

Deep Learning Based Non-Intrusive Load Monitoring for Smart Energy Management in Households

Parvaneh Zavareh*, Warunee Soythong*, Basirah Noor†, Volker Hoffmann‡, and Huamin Ren†

*School of Economics, Innovation and Technology, Kristiania University College, Oslo, Norway
E-Mail: {paza002, waso001}@student.kristiania.no

†School of Economics, Innovation and Technology, Kristiania University College, Oslo, Norway
E-Mail: {Basirah.Noor, Huamin.Ren }@kristiania.no

‡SINTEF, Forskningsveien 1, 0373 Oslo, Norway
E-Mail: volker.hoffmann@sintef.no

Abstract—The increasing digitalization inside the energy sector is transforming the fact that how residential electricity consumption is analysed and optimized. This study aims to propose and evaluate a Deep learning-based model for Non-Intrusive Load Monitoring (NILM), also called energy disaggregation, designed for the real-world applications in smart homes energy management systems. The framework employs Convolutional Neural Network (CNN) based sequence-to-point (Seq2Point) and sequence-to-sequence (Seq2Seq) model to isolate appliance level consumption from the aggregate power, while addressing generalization challenges across diverse households and appliances type. The results suggest that model choice (Seq2Point vs Seq2Seq) effects the results, for example Seq2Point shows better results in detecting high power, short duration events, highlighting the importance of selecting architecture based on appliance-specific characteristics for optimal NILM performance. The practical implications suggest that the proposed study can enable real-time appliance monitoring, adaptive demand-side management, and sustainable energy optimization in future smart grid environments.

Keywords—Non-Intrusive Load Monitoring (NILM); Convolutional Neural Networks; Seq2Seq; Seq2Point; Smart Home; Energy Management.

I. INTRODUCTION

Global electricity consumption grew by 2.2% in 2024, reflecting a faster-than-average rise in energy demand, according to International Energy Agency (IEA) [1]. Therefore, optimizing energy consumption has become a critical aspect now with the growing integration of smart households and smart grids [2]. The global market has been revolutionized when it comes to energy consumption because of energy aware technologies. Traditionally, aggregate consumption models were only used for load monitoring, which results in limited insights into the user behaviour [3]. As the energy demand is increasing globally, efficient energy monitoring across households can contribute more in optimizing energy usage, cost efficiency, and eventually to the overall sustainability.

Non-Intrusive Load Monitoring (NILM) is a promising technique that can address this challenge by disaggregating the total household power consumption into appliance level

consumption by using the data collected from single smart meter without the need of any device sensor [4]. The approach shows its application in energy conservation, demand-side management and smart home automation [5].

Recent advances in the field of Artificial Intelligence (AI) and Deep Learning (DL) have made NILM to perform even more better than before. Unlike traditional feature engineered models, CNN based approaches (Seq2Point and Seq2Seq) have this ability to automatically extract temporal and spatial patterns from raw energy data [6]. They can easily capture appliance signatures and therefore improves the appliance state detection and load disaggregation process.

Despite these models show promising results but previously many researchers have focused only on the energy disaggregation and not to the state detection (ON/OFF), which is an essential ingredient for automation and real time monitoring. This study will also adopt how preprocessing choices can affect the overall efficiency and generalization of the model across different appliances and households.

This study aims to achieve the following objectives to overcome the limitations of previous research.

- Systematically optimize the CNN based NILM models to improve both appliance state (ON/OFF) detection and energy disaggregation to individual load profile.
- Comparatively studying Seq2Seq and Seq2Point models for better load identification by keeping in view the variable household environments and different appliances.
- Analyzes how preprocessing parameters, such as learning rate, window size, normalization, etc. affect model accuracy and generalization.

The rest of the paper is organized in this way that Section II reviews the background and the related work on NILM and CNN based power disaggregation. Section III will present the proposed methodology including the dataset selection, data preprocessing, model training etc. The experimental results are presented and discussed in Section IV. Finally, Section V will provide the conclusion and future work.

II. BACKGROUND AND RECENT WORK

This section will define NILM and how Deep learning can help to improve its ability. Some recent work will also be presented to support the definition with evidence.

A. Non-Intrusive Load Monitoring

NILM is a promising technology to reduce the overall energy consumption. NILM researchers are continuously in this effort to introduce the best model to compliment the Hart's Theory [7]. CNNs are especially an effective technique to implement the NILM due to their capability to automatically extract time related information from the signal data and hence show better computational efficiency as compared to Recurrent Neural Networks (RNNs) or Long- and Short-Term Memory Networks (LSTMs). In this process the total power consumption at time t is calculated using the Hart's Signal Decomposition Theory [8].

$$P(t) = \sum_{i=1}^n a_i(t)P_i + e(t) \quad (1)$$

where:

- $P(t)$ is the total power consumption at time t ,
- $a_i(t)$ is the ON/OFF state of the i^{th} appliance at time t ,
- P_i is the power rating of appliance i ,
- $e(t)$ is the error term accounting for power from appliances not considered, and
- n is the total number of appliances.

NILM has been framed historically both as classification and regression problems [9]. When regarded as classification problem, the ON/OFF state of each appliance is classified simultaneously and typically event-based methods focusing on detecting sudden changes in power signals have been proposed, while time-based methods aim to continuously monitor appliance state, such as in [10]. The authors in [11] leverages temporal patterns for their model Temporal Attention Bottleneck Variational Autoencoder (TAB-VAE) to capture long-term dependencies in appliance behaviour, thus offering more reliable time based continuous state monitoring. Whereas researchers from [12] used Conditional Random Fields (CRFs) with CNNs for predicting states by taking into account the sharp power changes typically associated with appliance state transitions, which are common in event-driven tasks.

B. Application of Deep Learning in NILM

Many CNN-based models have been proposed to enhance NILM performance. Authors of [13] introduced a Dual-CNN based model that independently learns state transitions and power consumption patterns, that helped to significantly improve classification accuracy. Similarly, [14] and [15] demonstrated that dilated convolutions and residual connections expand the receptive field without increasing computational complexity, allowing the models to effectively capture long-range dependencies. Several RNNs, particularly LSTM models, have also been employed in NILM to model temporal sequences effectively. A recent study [16] demonstrated that

LSTM models trained on the PLAID dataset [17] achieved high levels of precision and recall across multiple appliances, particularly when integrated with IoT systems for near real-time classification. However, RNNs often face challenges when modelling long-duration appliance data, where convolutional models, especially CNN-based Seq2Seq and Seq2Point architectures, tend to perform better [18] and [19].

III. METHODOLOGY

This section will throw the light on the research design of the study. The goal of the study is to dis-aggregate the overall aggregate power of any household into appliance level power consumption using the CNN based Seq2Seq and Seq2Point models.

A. Data Acquisition

The data acquisition is the most important part of any model building. Two datasets were selected for this study. UK Domestic Appliance-Level Electricity (UK-DALE) is selected for model development and validation due to its acceptance in the NILM community, while the REFIT dataset is used as a secondary dataset for external testing to assess the generalization capability of the trained models.

Five household appliances were selected for further analysis -kettle, microwave, dishwasher, washer-dryer, and fridge-freezer - due to their nature of distinct operating pattern and high power consumption [20]. The required data has been collected from the UK-DALE dataset that comprises of the data from 5 households. Each household's data is organized into separate .dat files, where each file corresponds to single appliance level power consumption that were timestamped. REFIT provides data on household energy consumption from 20 homes in the UK and offers a wider range of appliances [21][22].

The data were split in two ways to improve the quality of the data. The UK-DALE set was first divided at the household level, where data from house 1,3,4 and 5 were used for training purposes and data from house 2 for testing purposes. Within each of the training houses, the data were further partitioned into an 80% training subset and a 20% validation subset. This internal split ensured that hyperparameter tuning and model selection were performed without contaminating the unseen test house. In the end data from REFIT was used to test the model and hence generalizability across the datasets.

B. Data Pre-Processing

After merging and aligning the dataset, the next step involved in the pipeline is the most important one as the results highly depends on these techniques. This study has utilized the following pre-processing techniques.

1) *Cleaning and Resampling*: This process involves cleaning of the data by handling missing entries. To address the issue of missing data gap-detection method is applied. The time difference between the two consecutive entries were calculated and if the difference is greater than 3 minutes, the gap would be filled with zero to show the inactivity. Whereas

the gaps with smaller time differences were forward filled. Timestamps are converted to standard date-time format for easier manipulation.

After cleaning, Power readings were resampled to 30–60-second intervals using mean aggregation. This process reduced the high frequency attenuation. Furthermore, all NaN entries were dropped and the resultant dataset is been saved in a separate CSV files for each house.

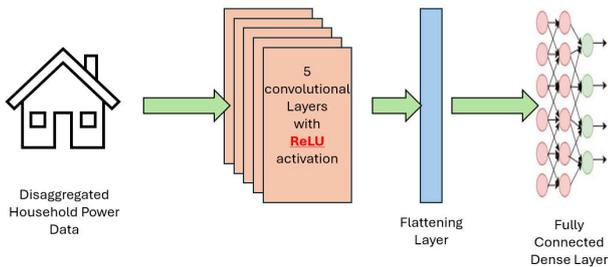


Figure 1. The Structure of Seq2point based CNN Model

2) *Normalization*: The process is used to standardize the dataset within the defined range. Several normalization techniques were applied including min-max method and z-score. Min-max normalization scales the data between the range (0 and 1), while z-score normalization scales the data so that it has a mean of 0 and a standard deviation of 1.

3) *Labeling*: Binary labels were generated to indicate the ON/OFF state. Different power criteria is applied to different appliances for example for kettle, it is considered ON if the power exceed from 1500 W while for fridge-freezer, it is considered ON, if power exceeds from 50 W. Minimum duration constraints were applied to avoid false activations from transient spikes.

4) *Exploratory Data Analysis (EDA) and Visualization*: The data in this step is analysed and visualized to better understand the power patterns. For the purpose, Python libraries, such as Matplotlib, Seaborn, and Plotly were used. the EDA helped us to find out that House 1 has the most complete data information over the period, that is why we will use it for training of the model. Houses 3, 4, and 5 offer moderate data and were also used for training. Where as house 2 had some gaps would be used for testing the model.

C. Model Implementation

Two CNN based models; Seq2Point and Seq2Seq were used. The following is the model configuration for CNN architecture that has 5 convolutional layers.

- Filter sizes: [30, 30, 40, 50, 50]
- Kernel sizes: [5, 4, 3, 3, 2]
- Activation function: ReLU for convolutional layers
- Dense layer: Size 256, ReLU activation
- Output layer: Linear activation, with size equal to the sliding window size
- Batch Size: 1024 (chosen to speed up training)

1) *Seq2seq Based CNN Model*: The seq2seq model uses aggregate energy readings and generates a corresponding output window of the same length. It maps an input sequence to a single output point. The model consists of 5 convolutional layer same as Seq2Point model based on ReLU activation, after that there is a flattening and Reshaping layer followed by dense layers for each output time step.

2) *Seq2Point Based CNN Model*: The seq2point model uses a window size of fixed length. It maps an input sequence to an entire output sequence. The model consists of 5 convolutional layer based on ReLU activation, after that there is a flattening layer followed by a fully connected dense layer as shown in Figure 1.

When we consider NILM based scenarios, the core difference between both models lies in their output structure, and hence the prediction would be different too. The Seq2seq model takes a sequence data (For example sequence of aggregate power readings) as input and outputs a corresponding sequence based data. The input and output sequences typically have the same length. Whereas for Seq2point, the input data is the same while it outputs only a single point data that shows the predicted power consumption at some defined, mostly the midpoint of the input window point for specific appliance. Therefore, Seq2point can be more helpful model where we need to capture sharp changes in appliance power like ON and OFF scenarios, while Seq2seq are better for representing continuous sequences and is well-suited for tasks where the entire output sequence is of interest.

The models are first trained as regression models that estimate continuous appliance power values, and the resulting predictions are then thresholded to derive binary ON/OFF states, which are evaluated using classification metrics as shown in the Figure 2

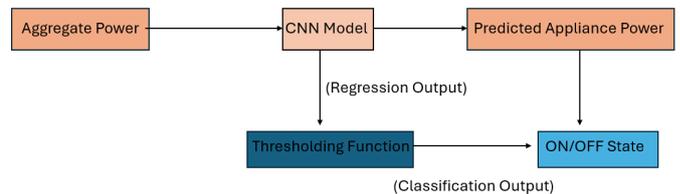


Figure 2. The Structure of Seq2point based CNN Model

D. Performance Evaluation

The study will use a combination of regression and classification metrics. The regression metrics like MAE and SAE can help to better access how the model can predict continuous appliance power consumption from the aggregated household signal, with lower values reflecting better disaggregation performance. Whereas classification metrics including Accuracy, Precision, Recall, and F1-score—are derived from the confusion matrix (TP, TN, FP, FN) can help access how good a model can predict from discrete information, for example on/off state of the appliances

IV. RESULTS AND DISCUSSION

This section will throw light on the results. The results will be discussed on three different grounds based on the experiments conducted on the pre trained UK-DALE dataset and cross validation using REFIT dataset

A. Training Parameter Influence

This section outlines how different training parameters affected the performance of the Seq2Seq and Seq2Point models in NILM. A systematic parameter study was carried out to find the best configurations for higher accuracy and better generalization. Adjustments to parameters, such as window size, learning rate, and normalization technique had a clear influence on model behavior. For Seq2Point, using short windows, small batch sizes like for 128, the MAE is 12.80 for Kettle as compared to 13.82 for 1024 batch size. Similarly no dropout, and a learning rate of 0.001 consistently resulted in higher accuracy as MAE is less compared to larger windows and other larger layers. In Seq2Seq models, medium to long windows, larger Batch size large batch sizes like for 1024, the MAE is 13.65 for Kettle as compared to 14.09 for smaller 128 batch size. Similarly no dropout, and less learning rates produced clear gains for cyclic appliances like the fridge-freezer and dishwasher, improving sequence as shown in Table I. Overall, these findings show that the right parameter combinations can meaningfully increase model performance, reinforcing that targeted tuning yields stronger results than a uniform configuration. Among all the factors tested, window duration, learning rate, and normalization method had the greatest impact on performance, while dropout primarily helped reduce overfitting. These findings highlight the importance of careful hyperparameter tuning to achieve a good balance between model accuracy, training stability, and convergence speed.

B. Model Performance Across Different Appliance Load Patterns

For appliances that draw greater power and are used for small durations, such as the kettle and microwave, the Seq2Point model aligns more closely with the ground truth as shown in Figure 3. It captures the timing and magnitude of peaks with sharper transitions and a faster return to zero. In contrast, the Seq2Seq smooths these quick transitions and tends to underestimate peaks. Both models detect events accurately, but Seq2Point provides better precision hence more suitable for appliances with bursty on/off loads.

For appliances that show more regular and repeating behaviour, like fridge-freezer and dishwasher, the Seq2Point model is showing better results. As shown in the Figure 3 Seq2Seqmodel is still detecting most of the events, but its predictions are more fragmented lead to false positives occasionally.

The washer dryer, as clear from the Figure 3 , has a more complex pattern, where Seq2Seq is outperforming and producing continuous predictions across the operational period where as Seq2Points is shown struggling to keep accuracy

throughout the operational period leading to certain false positives.

Therefore, the Seq2Point performs better with appliances that draws more power and are operational for small duration. For example the MAE for kettle is 13.25 for seq2seq and 11.12 for seq2point, showing a clear good results for seq2point as shown in Table II. While for wash dryer, Seq2Seq model perform better For example Seq2seq show an MAE of 8.43 while 13.22 for seq2point showing a victory for seq2seq model. These outcomes confirm the architectural strengths of each model: Seq2Seq is better for structured sequences hence maintains better continuity over longer operational periods whereas Seq2Point architecture helps it to excels in point-wise accuracy for detecting isolated events. This finding emphasizes the importance of appliance-specific model selection in NILM applications.

C. Comparison between Seq2Seq and Seq2Point

Threshold-based classification was used to convert the regression outputs into on/off states. Each appliance was assigned an activation threshold to avoid any spurious detections. Specifically, the thresholds were set as 1500W, 50W, 800W, 500W and 500W respectively for Kettle, Fridge-Freezer, Dishwasher, Microwave, Washer-Dryer. The experimental results indicate that the Seq2Point model generally outperforms Seq2Seq in detecting appliance on/off states, achieving higher F1-scores and better accuracy for devices with sharp, distinct power transitions as shown in Table II . Seq2Point excelled in identifying short, high-intensity bursts, such as those from the kettle ($F1 = 0.787$), fridge-freezer ($F1 = 0.630$), and dishwasher ($F1 = 0.578$), while also slightly outperforming Seq2Seq for the microwave ($F1 = 0.274$) despite the appliance's irregular usage pattern. However, for long and multi-phase operations like the washer-dryer, Seq2Seq performed better ($F1 = 0.516$) due to its ability to capture continuous temporal dependencies. Overall, Seq2Point proves more effective for appliances with short, bursty loads, whereas Seq2Seq remains advantageous for cyclic or extended activity patterns.

When comes to the regression metrics, Seq2Point typically achieves a 10–18% lower MAE and around 12–20% lower SAE, indicating more precise estimation of appliance power consumption.

D. Performance on the Holdout Set

Across all five appliances on the holdout set REFIT, Seq2Point models tends to outperform Seq2Seq models in most regression metrics as shown in Table III, although performance was notably reduced compared to evaluations conducted with the UK-DALE dataset. In terms of F1-score Seq2Point consistently achieves higher values for every appliance. For example, the kettle improves from 0.27 (Seq2Seq) to 0.37 (Seq2Point), and the washer-dryer shows a notable increase from 0.19 to 0.34. These gains are mainly driven by improvements in recall, indicating that Seq2Point is better at capturing true activation events. From a regression perspective,

TABLE I. RESULTS OF MAE FOR SEQ2SEQ AND SEQ2POINT MODELS UNDER DIFFERENT PARAMETER SETTINGS.

	Window Duration		Batch Size		Dropout Ratio		Learning Ratio		Resampling	
	23 min	45 min	128	1024	0	0.25	0.0001	0.001	30s	45s
Seq2seq										
Kettle	14.14	13.33	14.09	13.65	13.59	14.08	14.14	13.59	13.41	13.50
Fridge-Freezer	26.55	26.30	26.57	26.17	25.81	27.39	25.36	25.81	26.87	26.54
Dish Washer	23.63	23.66	24.48	25.46	24.36	24.36	24.46	24.36	23.77	24.09
Microwave	9.62	10.26	10.33	11.00	11.41	11.27	10.80	11.41	9.56	8.56
Washer Dryer	20.97	19.09	18.24	17.00	17.91	17.81	15.15	17.91	12.69	14.49
Seq2point										
Kettle	13.52	13.93	12.80	13.82	13.44	13.90	14.36	13.44	13.11	13.30
Fridge-Freezer	27.11	25.64	24.63	25.05	25.68	25.14	25.93	25.68	26.66	26.61
Dish Washer	24.75	22.95	23.43	25.10	25.04	25.11	24.54	25.04	24.41	23.91
Microwave	10.44	11.91	10.97	9.96	12.07	11.50	14.15	12.07	9.42	10.81
Washer Dryer	23.20	18.67	17.77	18.89	15.60	18.01	20.40	15.60	13.45	13.68

TABLE II. RESULTS OF REGRESSION AND CLASSIFICATION METRICS FOR SEQ2SEQ AND SEQ2POINT MODELS APPLIED ON UK-DALE DATASET

Appliance	Seq2Seq						Seq2Point					
	Acc	Prec	F1	Rec	MAE	SAE	Acc	Prec	F1	Rec	MAE	SAE
Kettle	0.99	0.74	0.74	0.75	13.25	0.04	0.99	0.72	0.79	0.85	11.12	0.05
Fridge-Freezer	0.78	0.63	0.52	0.44	25.16	0.01	0.80	0.63	0.63	0.63	23.69	0.07
Dishwasher	0.98	0.57	0.52	0.51	23.68	0.15	0.98	0.56	0.57	0.59	23.49	0.01
Microwave	0.99	0.40	0.27	0.20	9.57	0.92	0.99	0.74	0.27	0.20	8.04	0.45
Washer-Dryer	0.99	0.51	0.52	0.52	8.43	0.18	0.99	0.21	0.33	0.73	13.22	0.81

TABLE III. COMBINED RESULTS OF REGRESSION AND CLASSIFICATION METRICS FOR SEQ2SEQ AND SEQ2POINT MODELS ON REFIT

Appliance	Seq2Seq						Seq2Point					
	Acc	Prec	F1	Rec	MAE	SAE	Acc	Prec	F1	Rec	MAE	SAE
Kettle	0.99	0.55	0.27	0.21	21.18	0.42	0.99	0.59	0.37	0.33	17.08	0.49
Fridge-Freezer	0.66	0.56	0.44	0.41	37.75	0.61	0.67	0.56	0.41	0.36	35.75	0.49
Dishwasher	0.98	0.59	0.24	0.18	39.17	0.68	0.98	0.57	0.32	0.26	39.01	0.63
Microwave	0.99	0.20	0.11	0.09	17.95	2.51	0.99	0.19	0.09	0.07	14.92	1.54
Washer-Dryer	0.98	0.26	0.19	0.17	29.22	0.80	0.99	0.31	0.34	0.41	29.52	0.94

Seq2Point again demonstrates stronger performance, achieving lower MAE and SAE scores across nearly all appliances. For instance, the MAE for the kettle drops from 21.18 W (Seq2Seq) to 17.08 W (Seq2Point), and the microwave improves from 17.95 W to 14.92 W. Similarly, SAE decreases for every appliance, showing that Seq2Point better captures total energy consumption over time. The decline in overall accuracy emphasizes the impact of dataset-specific variations in appliance usage patterns. Hence the results from REFIT revealed that the model did not perform as good as it did on UK-DALE, highlighting the challenges of cross-dataset generalization. The reduced performance on REFIT can be attributed to its lower sampling frequency, higher household

and appliance variability, and increased aggregate noise compared to UK-DALE, which obscure short-duration appliance activations.

V. CONCLUSION AND FUTURE WORK

This study utilizes the CNN based Seq2Seq and Seq2Point architectures to effectively disaggregate the aggregate power of household to appliance level. The Seq2Point model outperformed at detecting short, high-power events, while Seq2Seq proved itself better in case of consistent loads. By optimizing key parameters, such as window duration, normalization, and learning rate, the research achieved substantial improvements in model accuracy. It is important to highlight this fact here

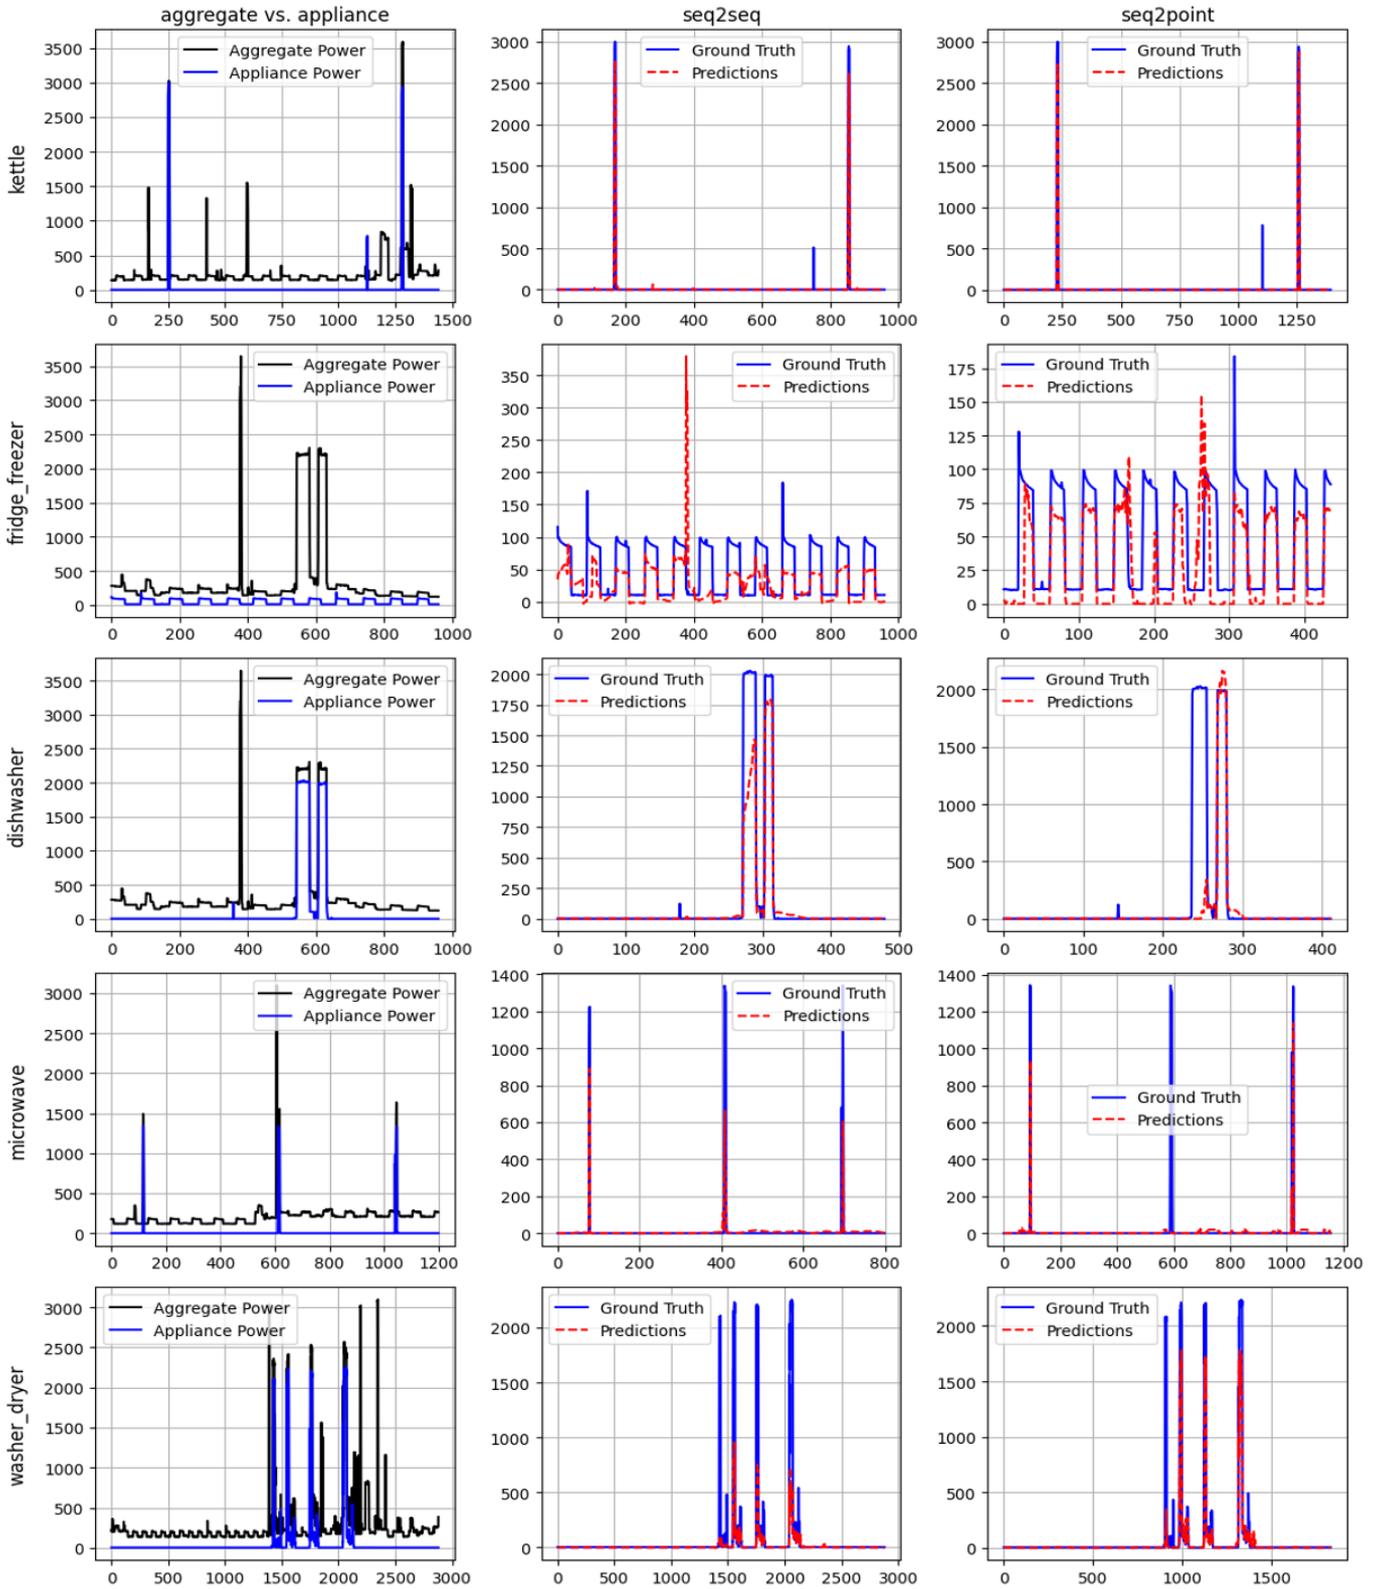


Figure 3. The disaggregation results for the appliances- kettle, fridge-freezer, dishwasher, microwave, and washer-dryer- across the five households of UK-DALE. Each row represents one appliance, with three columns, the first column represents the aggregate power versus the actual appliance consumption while the second and the third columns depicts the predicted power from Seq2Seq and Seq2Point models, respectively. The predicted values (Red) and ground truth (Blue) are compared simultaneously.

that although both Seq2Seq and Seq2Point models are trained purely as regression models that output continuous appliance power values, this study additionally evaluates their ability to perform appliance state detection. For the purpose, thresholding function is used to convert the continuous output to the ON/OFF states. While both models perform well on the UK-DALE dataset, this does not generalize uniformly on the RE-FIT dataset. Overall, Seq2Point outperforms Seq2Seq across both regression (MAE/SAE) and classification (F1/recall) metrics.

The proposed models and optimization framework advance the development of sustainable smart household power management systems. They also provide practical insights for real-world deployment in smart homes and IoT-enabled energy networks.

Despite these promising results, some limitations remain. Both models face challenges with low-power or overlapping appliance signals. Future work should explore hybrid CNN-transformer architectures, noise-robust preprocessing, and multi-dataset training to enhance generalization. Integrating optimized NILM models into edge devices and IoT systems will be a crucial step toward achieving real-time, scalable, and sustainable smart grid solutions.

REFERENCES

- [1] "Iea: Growth in global energy demand surged in 2024 to almost twice its recent average", Energy Connects, 2025, [Online]. Available: <https://www.energyconnects.com/opinion/features/2025/march/iea-growth-in-global-energy-demand-surged-in-2024-to-almost-twice-its-recent-average/>.
- [2] "Load frequency regulation of interconnected multi-source multi-area power system with penetration of electric vehicles aggregator model", [Online]. Available: <https://link.springer.com/article/10.1007/s00202-023-01923-2>.
- [3] S. Shams Amiri, M. Mueller, and S. Hoque, "Investigating the application of a commercial and residential energy consumption prediction model for urban planning scenarios with machine learning and shapley additive explanation methods", *Energy and Buildings*, vol. 287, p. 112965, 2023. DOI: 10.1016/j.enbuild.2023.112965.
- [4] A. Chatterjee and P. Heer, *Non-intrusive load monitoring (nilm) with very low-frequency data from smart meters in switzerland*, SSRN, 2025. DOI: 10.2139/ssrn.5113439.
- [5] Y. Liu, Y. Wang, and J. Ma, "Non-intrusive load monitoring in smart grids: A comprehensive review", *arXiv*, Mar. 2024. DOI: 10.48550/arXiv.2403.06474.
- [6] S. Zhang *et al.*, "Sequence-to-point learning based on spatio-temporal attention fusion network for non-intrusive load monitoring", *Complex Intelligent Systems*, vol. 11, no. 3, p. 171, Feb. 2025. DOI: 10.1007/s40747-025-01803-1.
- [7] G. W. Hart, "Nonintrusive appliance load monitoring", *Proceedings of the IEEE*, 1992.
- [8] G. Bucci, F. Ciancetta, E. Fiorucci, S. Mari, and A. Fioravanti, "State of art overview of non-intrusive load monitoring applications in smart grids", *Measurement: Sensors*, vol. 18, p. 100145, Dec. 2021. DOI: 10.1016/j.measen.2021.100145.
- [9] H. Ren, X. Su, R. Jenssen, J. Li, and S. N. Anfinsen, "Attention-guided temporal convolutional network for non-intrusive load monitoring", *2022 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pp. 419–425, 2022. DOI: 10.1109/SmartGridComm52983.2022.9960976.
- [10] P. G. Papageorgiou, G. C. Christoforidis, and A. S. Bouhouras, "Nilm in high frequency domain: A critical review on recent trends and practical challenges", *Renewable and Sustainable Energy Reviews*, vol. 213, p. 115497, May 2025. DOI: 10.1016/j.rser.2025.115497.
- [11] K. Oublal, "Temporal attention bottleneck is informative? interpretability through disentangled generative representations for energy time series disaggregation", *ICML 2023 Workshop on Deployment Challenges for Generative AI*, 2023.
- [12] D. Precioso Garcelán and D. Gomez-Ullate, "Nilm as a regression versus classification problem: The importance of thresholding", *J Supercomput* 79, 2020.
- [13] J. He *et al.*, "Msdc: Exploiting multi-state power consumption in non-intrusive load monitoring based on a dual-cnn model", in *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI'23)*, vol. 37, 2023, pp. 5078–5086. DOI: 10.1609/aaai.v37i4.25636.
- [14] Z. Jia, L. Yang, Z. Zhang, H. Liu, and F. Kong, "Sequence to point learning based on bidirectional dilated residual network for non-intrusive load monitoring", *International Journal of Electrical Power & Energy Systems*, vol. 129, p. 106837, Jul. 2021. DOI: 10.1016/j.ijepes.2021.106837.
- [15] G. Zhou *et al.*, *Sequence-to-sequence load disaggregation using multi-scale residual neural network*, 2020. DOI: 10.48550/arXiv.2009.12355.
- [16] Z. Solatidehkordi, J. Ramesh, A. R. Al-Ali, A. Osman, and M. Shaaban, "An iot deep learning-based home appliances management and classification system", *Energy Reports*, vol. 9, pp. 503–509, May 2023. DOI: 10.1016/j.egy.2023.01.071.
- [17] "Time series classification website", [Online]. Available: <https://www.timeseriesclassification.com/description.php?Dataset=PLAID>.
- [18] J. Jiang, Q. Kong, M. Plumbley, M. Hoogendoorn, and D. Roijers, "Deep learning-based energy disaggregation and on/off detection of household appliances", *ACM Transactions on Knowledge Discovery from Data*, vol. 15, pp. 1–21, Apr. 2021. DOI: 10.1145/3441300.
- [19] J. Wang and K. A. Loparo, *Evolutionary deep nets for non-intrusive load monitoring*, 2023. DOI: 10.13140/RG.2.2.25983.28324.
- [20] C. Zhang, M. Zhong, Z. Wang, N. Goddard, and C. Sutton, "Sequence-to-point learning with neural networks for non-intrusive load monitoring", in *Proc. AAAI Conf. Artif. Intell.*, vol. 32, Apr. 2018. DOI: 10.1609/aaai.v32i1.11873.
- [21] J. Kelly, "Uk domestic appliance-level electricity (uk-dale) dataset", [Online]. Available: <http://jack-kelly.com/data/>.
- [22] "Refit: Electrical load measurements (cleaned)", University of Strathclyde, [Online]. Available: <https://pureportal.strath.ac.uk/en/datasets/refit-electrical-load-measurements-cleaned/>.

Zero Trust Defense Against Charge Manipulation Attacks in Smart EV Charging Infrastructure

Saba Marandi
CIISE Department
Concordia University
Montreal, Canada

email: saba.marandi@mail.concordia.ca

Danial Jafarigiv
Digitization & Cybersecurity, R&D Department
Hydro-Québec Research Institute
Varenes, Canada

email: jafarigiv.danial2@hydroquebec.com

Ribal Atallah
Digitization & Cybersecurity, R&D Department
Hydro-Québec Research Institute
Varenes, Canada

email: atallah.ribal@hydroquebec.com

Mohsen Ghafouri
CIISE Department
Concordia University
Montreal, Canada

email: mohsen.ghafouri@concordia.ca

Chadi Assi
textitCIISE Department
Concordia University
Montreal, Canada

email: chadi.assi@concordia.ca

Abstract—The increasing reliance on Electric Vehicle Charging Infrastructure (EVCI) within smart grids introduces new threats to grid stability and operator control. Among all, Charge Manipulation Attacks (CMAs), in which attackers use access to the charging interface or backend systems to maliciously modify the charging parameters, e.g., the amount of energy requests, charging durations, are of significant importance. These attacks generate false load profiles and disrupt grid operations. This study develops a Zero Trust-based security mechanism to safeguard charging stations and the operator’s backend infrastructure, including control and management systems, from such attacks. Our solution employs a state-aware, Markov-based trust evaluation model that utilizes three behavioral indicators to continuously monitor the evolving behavior of chargers. Specifically, real-time behavioral indicators, such as request anomalies, charger usage history, and spatiotemporal consistency, are mapped into evolving trust states, which enable probabilistic prediction of charger trustworthiness. A central Zero Trust Controller issues short-lived tokens to charging stations only when the predicted state indicates acceptable trust levels. These tokens are enforced locally at each charger via a Policy Enforcement Point (PEP), which ensures strict, session-level verification. Simulation results demonstrate that the proposed architecture detects and blocks malicious charging behavior, which maintains secure and stable EVCS operations even under adversarial conditions.

Index Terms—Electric Vehicle Charger Stations; Zero-Trust; OCPP; Charge Manipulation Attacks.

I. INTRODUCTION

As Electric Vehicles (EVs) become increasingly integrated into modern transportation systems, the supporting infrastructure, i.e., EV Charging Stations (EVCSs), has grown rapidly in scale and complexity [1]. This growth is accompanied by increasing dependence on cloud-based management platforms, charger-side control logic, and remote access protocols that facilitate automated or operator-triggered charging sessions [2], [3]. While these advancements offer operational efficiency and flexibility, they also introduce critical cybersecurity risks [4]. Threats, such as load manipulation, unauthorized firmware command injection, protocol misuse, and denial-of-service attacks, are now emerging across both physical and cyber layers of the EVCS infrastructure [5]. Traditional perimeter-based security models, which rely on fixed trust boundaries and static

credential checks, are no longer sufficient to protect these dynamic and distributed systems. In broader vehicle-to-grid deployments, where EVCSs interact continuously with grid operators, aggregators, and third-party service platforms, the attack surface expands considerably [6]. These risks highlight the need for a policy-driven framework that models charger behavior, assesses trust in real time, and enforces localized protection to preserve grid stability.

To mitigate the aforementioned risks, several studies have proposed trust-based security frameworks relying on role-based or behavior-based evaluations of connected entities. For example, hierarchical and fuzzy logic-based trust models have been developed in [7] and [8] to control access to cloud-based services, energy usage records, or IoT data streams. These models typically compute trust scores using attributes, such as historical access behavior, device identity, or policy compliance. While these approaches provide a foundation for anomaly detection, they are mainly designed for user- or application-level control and often fail to capture the protocol-level behavior of EVCSs. Moreover, they suffer from static trust assumptions and lack real-time adaptability, which makes them ineffective against evolving threats, such as command manipulation anomalies. In particular, they cannot detect adversaries exploiting low-level charger protocols or adapt tactics to evade detection.

In parallel, recent research has identified three prominent categories of Charge Manipulation Attacks (CMAs) targeting EVCSs, namely sudden demand surges, coordinated switching, and market manipulation [9]. Sudden surges may result from adversaries remotely triggering large numbers of chargers or using social engineering techniques to induce synchronized charging activity [10]. Coordinated switching attacks involve repeated on-off cycling of chargers at low frequencies, which can trigger stability issues, such as inter-area oscillations [11], [12]. Furthermore, by controlling large-scale EVCS operations, adversaries can artificially inflate or deflate demand to exploit fluctuations in energy market pricing [13]. While standards like OCPP 1.6 and 2.0.1 incorporate protective features, such as randomized delays, these measures can be circumvented by attackers with elevated access privileges or through exploitation

of protocol-level controls [9]. These limitations highlight the need for adaptive, behavior-aware security strategies.

As a result, recent efforts have increasingly turned toward Zero Trust Architecture (ZTA) as a resilient framework for securing EVCS operations. Given the dynamic and distributed nature of these environments—where charging sessions can be initiated remotely via backend systems, firmware triggers, or control center APIs—continuous verification and real-time enforcement are essential. Unlike traditional models that rely on pre-established trust boundaries, ZTA mandates dynamic evaluation of every access or control request based on contextual indicators, device integrity, communication behavior, and operational risk. For example, [14] presents a ZTA-inspired fault detection system using distributed attestation for intelligent vehicles; however, their work is limited to in-vehicle contexts and does not extend to charger operations. Similarly, [6] proposes a ZTA-based access control model within vehicle-to-grid systems, but this approach is constrained to cloud-side data handling and lacks real-time integration with charging station protocols. The high-level framework in [15] outlines identity and policy management components for ZTA in EVCS, but does not address protocol-specific enforcement (e.g., OCPP validation), dynamic charger-side anomaly detection, or simulation under realistic attack scenarios. Therefore, a significant gap remains in deploying Zero Trust strategies directly within EVCS control logic to counter the manipulation of charging behaviors that may jeopardize grid stability.

In addition to this architectural gap, there are methodological shortcomings that limit current approaches. Although ZTA is promising for cyber-physical system security, current applications to EVCI remain limited to cloud-side control or architectural proposals, often overlooking protocol-level vulnerabilities within EVCS deployments. In particular, the dynamic nature of charger-initiated requests, behavioral deviations at the session level, and misuse of standard protocols like OCPP 2.0.1 are rarely addressed with real-time and localized enforcement mechanisms. Moreover, existing trust evaluation methods often rely on static or short-term assessments, failing to capture how charger behavior evolves over time. To overcome these limitations, this work integrates a state-aware, Markov-based trust evaluation model into the ZTA framework, which enables predictive, dynamic assessment of charger trustworthiness based on evolving behavioral indicators. In summary, this paper makes the following key contributions:

- A novel Zero Trust-based security framework for EVCI is proposed, in which charging station behavior is continuously assessed, and energy delivery sessions are authorized using dynamically issued trust tokens based on real-time behavioral indicators.
- A Markov-based trust assessment model that is state-aware is incorporated, which captures the evolving nature of charger behavior under both normal and adversarial conditions. Key behavioral indicators, such as session frequency, recurring user associations, and statistical deviations from normal patterns, are modeled as evolving trust states with defined transition probabilities. Access control decisions are made by lightweight Policy Decision Points (PDPs), which evaluate trust conditions and issue short-lived to-

kens, while Policy Enforcement Points (PEPs) embedded within each EVCS locally enforce these decisions.

- The proposed Zero Trust framework is seamlessly integrated with OCPP 2.0.1 protocol exchanges, which enables real-time inspection and enforcement of session-level security policies within standard EVCS communication workflows. Its effectiveness is demonstrated through the simulation of various charger behaviors under benign and adversarial conditions, demonstrating the capability of the framework to detect and prevent unauthorized charging operations.

The remainder of this paper is organized as follows. In Section II, we present the proposed Zero Trust enforcement framework for secure EV charging infrastructure, including the system overview, threat model, trust evaluation methodology, and policy enforcement mechanisms. In Section III, we describe the simulation setup and present comprehensive results demonstrating the effectiveness of our approach under various attack scenarios. Finally, Section IV concludes the paper and discusses future research directions.

II. PROPOSED ZERO TRUST ENFORCEMENT FRAMEWORK FOR SECURE EV CHARGING INFRASTRUCTURE

To bring the Zero Trust approach into practice for EV charging infrastructure, this section outlines the components and mechanisms of our proposed framework. At the heart of the design is the idea that no charging request should be trusted by default, even if it comes from a known source. Instead, each interaction is evaluated in real time based on behavioral context, device integrity, and protocol activity. To enhance predictive detection of malicious behavior, a state-aware trust evaluation model based on evolving charger behavior is incorporated into the framework. The system introduces lightweight enforcement modules at the station level, supported by a central zero trust controller that issues short-term trust tokens. These tokens are granted only when a request aligns with expected patterns and policy rules. Our framework is designed with real-world constraints in mind, including the evolving nature of attacks. The following subsections describe the system setup, how trust is calculated and enforced, and how the framework integrates with existing protocols like OCPP to secure EVCS operations.

A. System Overview & Threat Model

The proposed Zero Trust-based EVCI consists of interconnected components designed to ensure secure communication and resilient energy delivery. As illustrated in Figure 1, the architecture includes EVs, EVCSs, a centralized EV charging management system, i.e. EVCMS, Distribution System Operators (DSOs), aggregators, and energy market entities. Charging sessions are typically triggered via mobile applications or service APIs; however, under the ZTA, these requests are not directly processed by the EVCMS. Instead, they are routed through a centralized control plane that applies continuous trust evaluation. Within this control plane, behavioral indicators are assessed by a Trust Evaluation Engine in conjunction with a Policy Engine. The PDP issues short-lived trust tokens only when the contextual integrity of the request satisfies predefined policies. These tokens are verified by local PEPs

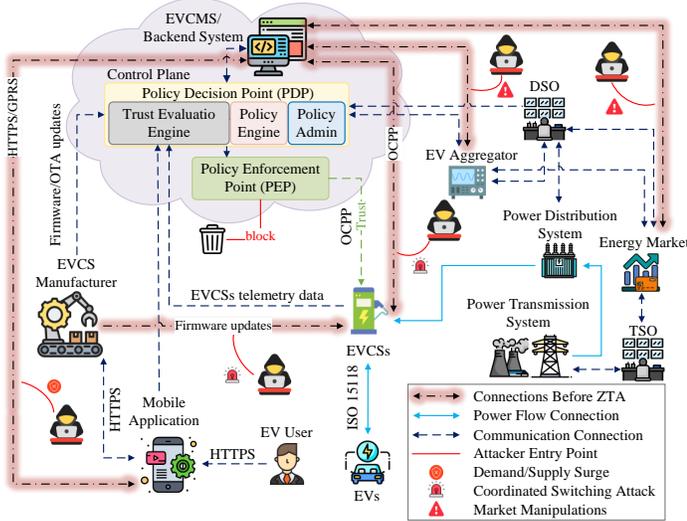


Fig. 1: EVCS Security Architecture with ZTA and Attack Vectors.

embedded in each EVCS, allowing or denying further action accordingly. Communications across all entities are secured using protocols, such as HTTPS and OCPP 2.0.1. Firmware updates are validated by the PDP before over-the-air (OTA) deployment to ensure source authenticity and prevent malicious manipulation. Coordination with DSOs, aggregators, and energy market systems is facilitated to preserve operational stability and compliance with grid directives. This architecture is structured to counter the range of adversarial actions characterized in the threat model, with particular emphasis on preempting unauthorized access, ensuring firmware authenticity, and detecting anomalous charging behavior.

As previously mentioned in Section I, CMAs exploit various vulnerabilities within the EV charging ecosystem. In this work, adversaries are assumed to possess moderate to advanced resources, including access to compromised user credentials, mobile applications, or backend software interfaces. As discussed in Section I, this work considers three types of CMAs: demand surge, coordinated switching, and protocol-level manipulation. Demand surge attacks are launched by abusing these assets to simultaneously initiate charging sessions at scale, potentially causing destabilizing spikes in grid demand. These attacks may utilize Man-in-the-Middle (MitM) techniques to intercept or falsify session initiation requests, especially when authentication protocols are weak or static. Coordinated switching attacks involve periodic on/off cycling of EVCSs through tampered firmware or scheduler access, which mimics natural grid oscillations that challenge detection mechanisms. Such attacks can be facilitated through unverified OTA updates or concealed using False Data Injection Attacks (FDIA) that alter telemetry data. Lastly, market manipulation attacks aim to distort price signals and trading outcomes by injecting falsified load data into aggregator platforms or reporting interfaces. These strategies exploit protocol weaknesses, spoofed measurement values, and vulnerabilities in EVCMS or aggregator APIs. Without dynamic trust validation and localized session controls, these attacks can bypass perimeter-based security measures and inflict significant disruptions at both grid and market levels.

B. Trust Evaluation & Policy Decision with Markov Updating

To capture the evolving behavior of each charger, we model each trust dimension using a discrete-time Markov chain, where the trust state at any time depends only on the previous state and current behavioral observations. This probabilistic modeling approach allows the system to predict future trust levels and respond proactively to behavioral changes. Drawing inspiration from [16] and adapting the Markov-based trust score evolution technique [17], we define three distinct dimensions of charger trustworthiness: Stability Degree (SD), Intimacy Degree (ID), and Abnormality Degree (ED). Each dimension independently tracks a specific behavioral aspect of the EVCS and evolves over time through its corresponding state transition process, guided by real-time logs and observed features.

Time Model Clarification: In our framework, time t is discretized into evaluation periods, where each period corresponds to a 3-hour operational window during peak charging hours. When we refer to time $t - 1$, we mean the previous evaluation period (i.e., the preceding 3-hour window). Trust scores and behavioral indicators are updated once per evaluation period based on the sessions observed during that window. This discrete-time model aligns with realistic operational monitoring intervals in EVCS management systems.

Each of these scores is updated over time using its transition mechanism. The global trust score is then defined as:

$$T_n^{(t)} = w_s \cdot T_s^{(t)} + w_i \cdot T_i^{(t)} - w_e \cdot T_e^{(t)} \quad (1)$$

where $T_s^{(t)}$, $T_i^{(t)}$, and $T_e^{(t)}$ represent the normalized stability, intimacy, and abnormality scores at time t , respectively, and w_s , w_i , w_e are their corresponding weights, with $w_s + w_i + w_e = 1$.

1) *Stability Degree (SD):* This score evaluates the consistency of charger communication by incorporating the number of connection interruptions and request timeouts observed during a given evaluation period. Specifically, l_s represents the number of connection disruptions observed during the current evaluation period, r_s represents the number of request timeout events during the current period, while T_l and T_r denote the maximum allowable thresholds for disruptions and timeouts, respectively (set based on operational requirements). The parameters V_l and V_r quantify the base penalty weight associated with each type of failure, and λ_l , λ_r are scaling coefficients that adjust the severity based on the frequency of failures. The stability score is updated as follows:

$$T_s^{(t)} = T_s^{(t-1)} + ((V_l + \lambda_l \cdot l_s) \cdot (T_l - l_s)) + ((V_r + \lambda_r \cdot r_s) \cdot (T_r - r_s)) \quad (2)$$

As long as $l_s < T_l$ and $r_s < T_r$, the penalty remains bounded. However, when either metric approaches or exceeds its threshold, the penalty term increases rapidly, reducing the stability score and reflecting decreased reliability in communication.

2) *Intimacy Degree (ID):* This metric quantifies the historical interaction strength between the charger and the cloud-based management platform. The variables R_t and R'_t represent the percentile rank of interaction frequency in the current evaluation period t and the previous evaluation period $t - 1$, respectively (ranging from 0 to 100, computed based on the charger's relative frequency compared to all chargers in the network). The term q_i serves as a fluctuation control parameter that moderates

the score update (preventing excessive sensitivity to minor rank changes). The intimacy score is computed as:

$$T_i^{(t)} = T_i^{(t-1)} + (\eta_1 R_t + \eta_2 R_t') (R_t' - R_t) \quad (3)$$

where η_1 and η_2 are scaling coefficients that tune the sensitivity of the score to temporal shifts in interaction rank. A decrease in activity ranking over time reduces the intimacy score, reflecting diminished trust, whereas stable or improved interaction frequency contributes to trust accumulation.

3) *Abnormality Degree (ED)*: The abnormality score quantifies deviations from expected behavioral norms observed in charger activity logs. It reflects the occurrence of suspicious events during the current evaluation period, such as repeated protocol misuse, attempts to access unauthorized commands, or operating from unexpected geographic regions. Since EV chargers are typically stationary and operate within fixed service zones, any sudden change in their geolocation compared to historical records is treated as a potential anomaly. Let $T_e^{*(t)}$ denote the abnormality degree at the end of time period $t-1$ (i.e., the previous evaluation window), and suppose N abnormal behaviors are detected during the current period t . Each anomaly is assigned a severity score $Q(\text{Exception}_i)$ (a predefined weight based on the type of anomaly, e.g., unauthorized command = 5, location anomaly = 3), where $i \in \{1, \dots, N\}$ indexes the observed exceptions. The abnormality degree is then updated according to:

$$T_e^{(t)} = \begin{cases} T_e^{*(t)} + \sum_{i=1}^N Q(\text{Exception}_i), & N \neq 0 \\ T_e^{*(t)} - p_e, & N = 0 \end{cases} \quad (4)$$

here, p_e is a recovery coefficient (a positive constant) that gradually decreases the abnormality score in the absence of violations, promoting long-term behavioral rehabilitation. The function $Q(\cdot)$ captures the relative severity of each abnormality type—such as excessive charging power, unauthorized command usage, unverified firmware updates, or location inconsistency—and allows fine-grained adjustment of the trust score based on context. This formulation enables the framework to adaptively penalize malicious behavior while allowing trust to be restored when normal patterns are maintained.

4) *Markov-Based State Update Process*: To capture the temporal evolution of trust across various dimensions, each dimension—stability (T_s), intimacy (T_i), and abnormality (T_e)—is modeled using a discrete-time Markov chain. The state space for each dimension $d \in \{s, i, e\}$ is defined as $\mathcal{S}_d = \{S_1, S_2, S_3, S_4, S_5\}$, corresponding to semantic levels of *Very Low Trust*, *Low Trust*, *Moderate Trust*, *High Trust*, and *Very High Trust*, respectively. This classification enables fine-grained tracking of trust dynamics in the presence of varying behavioral patterns. State transitions are governed by a stationary Markov transition matrix $P_d \in \mathbb{R}^{5 \times 5}$, where each entry $P_d(i, j)$ specifies the probability of transitioning from state S_i at time period $t-1$ to state S_j at time period t . The transition process is influenced by behavioral evidence collected during the current evaluation window, such as interaction frequency, connection reliability, or protocol misuse anomalies.

The belief distribution over the 5 trust states at time t is denoted by $\pi_d^{(t)} \in \mathbb{R}^5$, where each element represents the

probability that the charger is in a particular trust state. Given the belief from the previous period $\pi_d^{(t-1)}$ and an observation $z_d^{(t)}$ (the computed trust score for dimension d in the current period) associated with the current trust scores, the belief is updated using a Bayesian filtering approach:

$$\pi_d^{(t)} = \frac{P_d^\top \cdot (\pi_d^{(t-1)} \odot \mathcal{O}_d(z_d^{(t)}))}{\|P_d^\top \cdot (\pi_d^{(t-1)} \odot \mathcal{O}_d(z_d^{(t)}))\|_1} \quad (5)$$

where \odot denotes the element-wise product, and $\mathcal{O}_d(z_d^{(t)}) \in \mathbb{R}^5$ represents the observation likelihood vector, which quantifies the alignment between the current observation $z_d^{(t)}$ and each trust state in \mathcal{S}_d . These likelihood vectors are derived from analyzing historical behavioral logs and discretizing trust-related features (e.g., stability, interaction frequency, anomaly levels) into state-dependent distributions. In scenarios where labeled data is limited, domain-specific thresholds are used to construct $\mathcal{O}_d(z_d^{(t)})$ based on expert-defined boundaries. Following the belief update, the scalar trust score for dimension d is computed as the expectation of the updated belief distribution:

$$T_d^{(t)} = \pi_d^{(t)} \cdot \mathbf{v}_d \quad (6)$$

where $\mathbf{v}_d = [0.0, 0.25, 0.5, 0.75, 1.0]^\top$ represents numerical values assigned to the five trust states, from *Very Low* to *Very High*. Higher values of $T_d^{(t)}$ indicate stronger trustworthiness in the given dimension.

This Markov-based process enables robust, probabilistic modeling of charger behavior under partial observability, which allows the system to smoothly adapt to both short-term anomalies and long-term behavioral shifts.

5) *Trust-Based Authorization*: A charging request is approved only if the aggregated trust score $T_n^{(t)}$ exceeds a predefined threshold τ , which is configured by the system administrator based on organizational security policies and risk tolerance levels. This threshold represents the minimum trustworthiness required for a charger to obtain session approval. If the trust score falls below τ , the PDP withholds the access token, thereby blocking the EVCS from executing further protocol operations.

$$\kappa_n^{(t)} = \begin{cases} \text{JWT issued} & \text{if } T_n^{(t)} \geq \tau \\ \text{Request denied} & \text{otherwise} \end{cases} \quad (7)$$

where $\kappa_n^{(t)}$ represents the authorization decision for charger n at time period t , and JWT (JSON Web Token) is the short-lived cryptographic token issued by the PDP.

C. Policy Enforcement & Protocol Integration

Once the PDP authorizes a session by issuing a token, enforcement is carried out at the edge via lightweight PEPs embedded in each EVCS. Each incoming request is intercepted and validated by the PEP, which parses the OCPP 2.0.1 message and inspects the embedded JWT token for scope, expiration, and signature integrity. A control request $r_i^{(t)}$ from EVCS i at time t is authorized only if:

$$r_i^{(t)} \in \mathcal{A}(\mathcal{P}, \kappa) \quad (8)$$

where \mathcal{A} represents the set of allowable actions under system policy \mathcal{P} and token κ . If a mismatch occurs or the token is invalid, the request is blocked and flagged. All EVCS-bound messages, including OCPP control commands, firmware updates, and metering requests, are routed through the PEP for real-time, context-aware enforcement. The PEP ensures that each request is authenticated, authorized, and compliant with session rules. For instance, an `UpdateFirmware` command is accepted only if authorized by the PDP and digitally signed:

$$\text{Verify}(\sigma_{\text{firm}}, K_{\text{pub}}^{\text{PDP}}) = \text{True} \quad (9)$$

where σ_{firm} is the firmware's digital signature and $K_{\text{pub}}^{\text{PDP}}$ is the PDP's public key. If verification fails, the command is dropped, and an alert is triggered.

Token-based enforcement operates in coordination with the trust authorization mechanism discussed in Section II-B5. Only EVCSs with an aggregated trust score exceeding the threshold τ are granted access tokens. This ensures that subsequent protocol interactions, including firmware updates, control messages, and session operations, originate exclusively from authenticated and trustworthy sources.

III. SIMULATION RESULTS

To validate the effectiveness of the proposed ZTA-based defense framework, a comprehensive simulation environment was developed to emulate realistic EV charging station operations and behavioral patterns. The simulation was run in a virtual machine hosting all essential software modules, including charger emulation, trust computation, and policy enforcement mechanisms. Each EVCS was represented as a logical instance, emulating real-world interactions through Python scripts that implement the OCPP 2.0.1 protocol at the application layer. This abstraction ensured a realistic cyber-physical representation without requiring dedicated hardware. A total of five EVCSs, labeled Chargers A through E, were selected from a real-world dataset, each supporting two independent power outlets. Charging session logs from these stations were used to extract behavioral data during peak demand periods across multiple operational days. To capture the temporal dynamics of charging activity, the simulation was divided into ten evaluation windows, each representing a 3-hour peak period sampled from the dataset. The behavioral indicators (SD, ID, ED) were updated once per simulation window, resulting in the 10 evaluation periods shown in the results. This segmentation reflects realistic session density and aligns with the operational capabilities of dual-outlet chargers, since the original dataset exhibits 2-3 charging sessions per 3-hour peak window per dual-outlet charger, with activity representing high-utilization periods. Accordingly, the number of sessions in each time window was extracted directly from the empirical dataset to preserve authentic load patterns.

To simulate adversarial scenarios, we used Python scripts to inject charge manipulation attacks into 30% of the sessions in the dataset. These attacks changed key session attributes like requested energy and charging duration by modifying their values within a realistic range, based on typical usage patterns. All modified sessions followed the format and rules

of the OCPP 2.0.1 protocol, making them appear legitimate. The attacks were applied randomly across different chargers and time windows to reflect stealthy, coordinated behavior that evades basic rule-based detection.

It is important to note that the baseline "traditional" approach in our evaluation represents a purely trust-agnostic access model where all syntactically valid OCPP messages are accepted without behavioral assessment. While real-world operators typically employ some security measures, such as rate limiting or basic anomaly detection, our baseline was intentionally simplified to isolate and highlight the specific contribution of the proposed trust-based framework. This choice allows us to demonstrate the fundamental value of continuous behavioral assessment and dynamic trust evaluation. However, we acknowledge that a more realistic baseline would incorporate common security controls. Future work will include comparative evaluation against baselines that integrate rate limiting, session constraints, and threshold-based anomaly detection to provide a more comprehensive assessment of the added value of our ZTA approach. Additionally, the impact of different parameter choices on system performance requires systematic investigation through sensitivity analysis, as discussed below.

TABLE I: TRUST LEVEL HIERARCHY.

Trust Level	Controlling Measure	Permission	Trust Score Range
I	N/A	Monitored access	[90,100]
II	Downgrading permissions	Restricted access; Monitored access	[70,90]
III	Re-authenticating	Access temporarily denied	[50,70]
IV	Blocking connections	Access denied	[0,50]

Three behavioral indicators were used to continuously monitor and rate each EVCS, computed from historical logs and simulated session events acting as real-time inputs. Stability measured the temporal consistency in session frequency and duration; intimacy reflected recurring associations with known users or vehicles; and abnormality captured statistical deviations from baseline behaviors, such as irregular inter-arrival times or anomalous energy requests. Each indicator evolved according to a discrete-time Markov process with 5 trust states, and transition probabilities were learned from empirical distributions. The indicators were normalized to the [0,100] scale, enabling fair comparison. The overall trust score for each charger was calculated as a weighted average of the expected values of the three indicators, with empirically determined weights: 0.4 (stability), 0.3 (intimacy), and 0.3 (abnormality). These weights emphasize behavioral consistency, which is critical for mitigating CMAs.

Parameter Sensitivity: The choice of weights ($w_s = 0.4$, $w_i = 0.3$, $w_e = 0.3$) and the trust threshold ($\tau = 70$) significantly influence system performance. While our current parameterization is based on empirical tuning using the available dataset, we acknowledge that optimal values may vary across different operational contexts, charger types, and threat landscapes. Preliminary experiments (not shown here due to space constraints) indicate that increasing w_s improves detection of communication-based attacks but may reduce sensitivity to user behavior anomalies. Similarly, lowering τ increases security at the cost of more frequent false rejections of legitimate sessions.

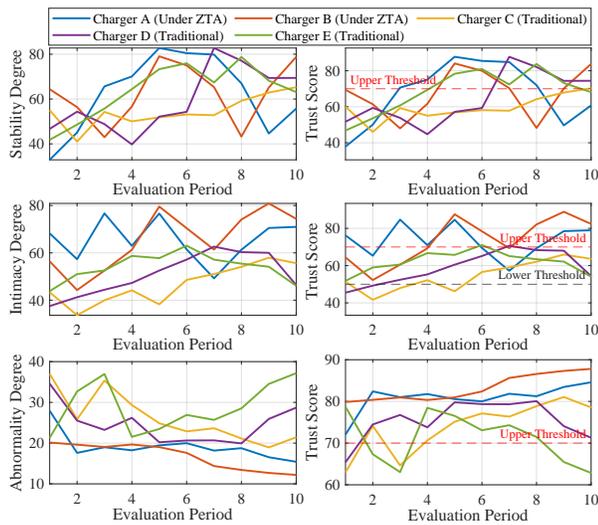


Fig. 2: Trust evolution under ZTA versus traditional access after attack.

A comprehensive sensitivity analysis examining the impact of weight variations, recovery coefficients (p_e), and threshold settings on detection rates, false positives, and overall system responsiveness will be conducted in future work. This analysis is essential for understanding the framework's portability and scalability across diverse operational environments and operator requirements.

Based on the resulting trust score, each charger was assigned a trust level in accordance with the hierarchy shown in Table I, guiding the system's dynamic access control logic. Chargers A and B were subject to ZTA enforcement, where sessions were only approved if the trust score remained above the threshold corresponding to Level II (i.e., ≥ 70). Otherwise, control measures, such as re-authentication (Level III) or access denial (Level IV), were triggered based on the degree of trust degradation, as defined in Table I. In contrast, Chargers C, D, and E operated under a traditional access model without dynamic trust evaluation and approved all syntactically valid session requests regardless of behavioral anomalies. To assess the system's responsiveness to cyber-physical threats, all five chargers were subjected to CMAs, wherein a subset of session requests was synthetically modified to exhibit abnormal behavioral characteristics, such as irregular inter-arrival patterns, abrupt load changes, and inconsistent user associations. This attack strategy aimed to mimic real-world stealthy manipulation while degrading the behavioral indicators over time. The divergence in session approval logic between ZTA-enforced and trustless policies enabled a clear comparative analysis of how dynamic trust evaluation can mitigate attacks by proactively denying low-trust sessions while maintaining uninterrupted service for compliant chargers.

In addition to trust score computation, the simulation incorporated an access control mechanism to monitor the number of approved and denied sessions per charger across ten time windows. For Chargers A and B, which were governed by the ZTA policy, access decisions were strictly enforced based on trust levels: sessions were approved only when the trust score remained within Levels I or II (i.e., ≥ 70), and were denied or required re-authentication for Levels III and IV. As a result, these chargers exhibited a selective but secure

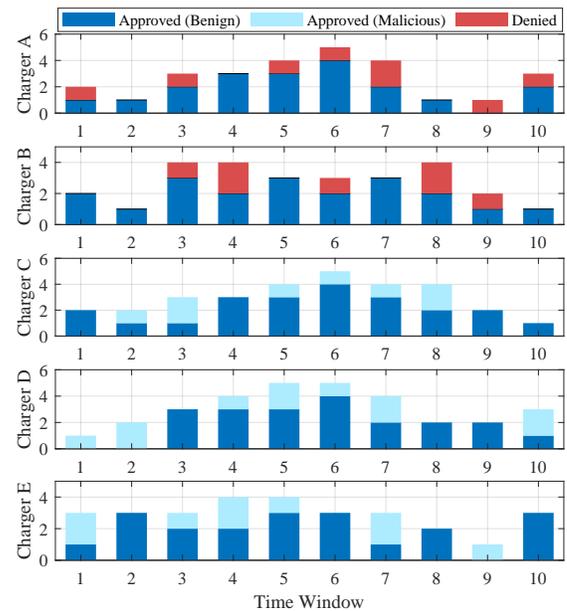


Fig. 3: Session-level access control decisions across EVCSs.

approval pattern, while most sessions were permitted, access was consistently denied during intervals where trust degradation was detected due to ongoing behavioral anomalies introduced by the attack. In contrast, Chargers C, D, and E, which operate without trust-based access control, approved all session requests regardless of their underlying behavioral indicators, which allowed malicious activity to persist undetected. All behavioral indicators were evaluated in their normalized form to ensure consistent trust score computation across chargers. Figure 2 presents six subplots depicting the evolution of behavioral degrees and their corresponding trust scores across chargers over time. Each indicator's deviation is shown on the left, while the right plots reflect the computed trust scores with respect to defined trust level thresholds. Due to the coordinated attack, all chargers experienced noticeable trust score variations, oscillating between Level I and Level III. Despite the uniform attack conditions, Chargers A and B leveraged ZTA to dynamically regulate access, while Chargers C–E lacked such enforcement. Figure 3 visualizes the resulting session-level decisions. Chargers A and B show a proportion of sessions denied due to sub-threshold trust scores, which demonstrates ZTA's enforcement. Chargers C–E, despite experiencing similar trust score declines, continued to permit all sessions due to the absence of a behavioral assessment. These findings validate the ZTA framework's capability to enforce proactive access decisions based on real-time trust evaluation, thereby limiting the system's exposure to charge manipulation attacks. Table II quantifies these outcomes. With 30% of sessions containing injected attacks, ZTA-enabled chargers successfully detected and blocked 87.5-100% of malicious sessions, while traditional chargers approved all requests regardless of malicious intent.

A. Computational Overhead and Scalability Considerations

While our current evaluation focuses on security effectiveness, computational latency is a critical factor for real-world deployment. The proposed framework introduces overhead at two levels: (1) trust score computation at the central PDP,

TABLE II: ACCESS CONTROL EFFECTIVENESS UNDER CMA.

Charger Type	Approved	Denied	Denial Rate
A (ZTA-enabled)	19	8	100%
B (ZTA-enabled)	20	7	87.5%
C (Traditional)	30	0	0.0%
D (Traditional)	31	0	0.0%
E (Traditional)	29	0	0.0%

and (2) token verification at each local PEP. In our simulation environment, trust score updates (performed once per 3-hour evaluation window) required an average of 120-150 milliseconds per charger on a standard virtual machine (Intel Xeon E5-2680 v4, 8GB RAM). Token verification at the PEP level, which occurs per session request, adds approximately 2-5 milliseconds of latency due to JWT signature validation. Given that typical charging session initiation times are on the order of seconds, this overhead is unlikely to impact user experience in practice. However, formal latency measurements on edge devices and under concurrent load conditions are needed to fully validate deployment feasibility.

Regarding scalability, our evaluation is limited to five chargers from a single dataset, which represents a small-scale deployment scenario. The generalizability of our findings to larger networks with diverse charger types (AC Level 2 versus DC fast charging), different operators, and varying traffic patterns remains an open question. Different charging environments may exhibit distinct behavioral patterns, requiring adaptation of trust indicator definitions, weight assignments, and transition probability matrices. Future work will address these limitations through: (1) evaluation on larger-scale datasets spanning multiple operators and geographic regions, (2) investigation of AC versus DC charger behavioral differences, (3) real-time edge deployment and latency profiling, and (4) adaptive parameter tuning mechanisms that can adjust to local operational contexts without manual reconfiguration.

IV. CONCLUSION

This study introduced a Zero Trust Architecture (ZTA)-based predictive trust evaluation framework tailored for Electric Vehicle Charging Infrastructure (EVCI). The proposed system leverages a Markov-based behavioral model to assess charger trustworthiness across three key dimensions: stability, intimacy, and abnormality. Through a simulation involving five EVCSs, all subjected to coordinated behavioral anomalies, we demonstrated the framework's capability to track behavioral trends and predict declining trust levels dynamically. While all chargers experienced trust-score degradation under attack, only those with ZTA enforcement could prevent access based on real-time trust evaluation, thereby illustrating the framework's ability to contain malicious sessions. This predictive, context-aware approach enables fine-grained access control and enhances the resilience of EVCS networks against application-layer threats such as Charge Manipulation Attacks (CMAs). As part of future work, we aim to integrate additional real-world datasets across varied environments and operational contexts and to implement a live, edge-based prototype for real-time trust computation and session authorization. This will further validate the scalability and adaptability of the proposed framework in practical smart grid deployments.

ACKNOWLEDGMENT

This work is supported by Fonds de recherche du Québec – Nature et technologies (FRQNT) and the Concordia University/Hydro-Québec/NSERC Research Collaboration Project.

REFERENCES

- [1] K. KSarieddine, M. A. Sayed, C. Assi, R. Atallah, S. Torabi, J. Khoury, M. S. Pour, and E. Bou-Harb, "Ev charging infrastructure discovery to contextualize its deployment security," *IEEE Transactions on Network and Service Management*, vol. 21, no. 1, pp. 1287–1301, 2024.
- [2] A. Vishnoi and V. Verma, "The analysis on impact of cyber security threats on smart grids," in *Security and Risk Analysis for Intelligent Edge Computing*. Springer, 2023, pp. 111–118.
- [3] K. Sarieddine, M. A. Sayed, S. Torabi, R. Attallah, D. Jafarigiv, C. Assi, and M. Debbabi, "Uncovering covert attacks on ev charging infrastructure: How ocpp backend vulnerabilities could compromise your system," in *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security*, 2024, pp. 977–989.
- [4] S. Marandi, A. Moradzadeh, H. Moayyed, C. Assi, M. Ghafouri, and Z. Vale, "Anomaly detection in load forecasting for electric vehicles using image processing techniques," in *2024 IEEE International Conference on Environment and Electrical Engineering and 2024 IEEE Industrial and Commercial Power Systems Europe (EEEIC / ICPS Europe)*, 2024, pp. 1–6.
- [5] J. Johnson, B. Anderson, B. Wright, J. Quiroz, T. Berg, R. Graves, J. Daley, K. Phan, M. Kunz, R. Pratt *et al.*, "Cybersecurity for electric vehicle charging infrastructure," Sandia Natl. Lab., Albuquerque, NM (US), Tech. Rep., 2022.
- [6] N. Zhou, S. Ji, and Q. Mao, "Research on trust evaluation model for vehicle-to-grid interaction based on zero trust architecture," in *2024 5th International Conference on Information Science, Parallel and Distributed Systems (ISPDS)*. IEEE, 2024, pp. 189–197.
- [7] A. Kesarwani and P. M. Khilar, "Development of trust-based access control models using fuzzy logic in cloud computing," *Journal of King Saud University-Computer and Information Sciences*, vol. 34, no. 5, pp. 1958–1967, 2022.
- [8] S. Ashtari, M. Danesh, and h. shirgahi, "A novel user profile-based fuzzy approach for evaluating trust in semantic web," *IJUM Engineering Journal*, vol. 20, no. 1, pp. 158–176, 2019.
- [9] H. Jahangir, S. Lakshminarayana, and H. V. Poor, "Charge manipulation attacks against smart electric vehicle charging stations and deep learning-based detection mechanisms," *IEEE Transactions on Smart Grid*, 2024.
- [10] B. Wang, P. Dehghanian, S. Wang, and M. Mitolo, "Electrical safety considerations in large-scale electric vehicle charging stations," *IEEE Transactions on Industry Applications*, vol. 55, no. 6, pp. 6603–6612, 2019.
- [11] M. E. Kabir, M. Ghafouri, B. Moussa, and C. Assi, "A two-stage protection method for detection and mitigation of coordinated evse switching attacks," *IEEE Transactions on Smart Grid*, vol. 12, no. 5, pp. 4377–4388, 2021.
- [12] M. A. Sayed, R. Atallah, C. Assi, and M. Debbabi, "Electric vehicle attack impact on power grid operation," *International Journal of Electrical Power & Energy Systems*, vol. 137, p. 107784, 2022.
- [13] S. Acharya, R. Mieth, R. Karri, and Y. Dvorkin, "False data injection attacks on data markets for electric vehicle charging stations," *Advances in Applied Energy*, vol. 7, p. 100098, 2022.
- [14] Y. Wang, Y. Na, S. Xu, and D. Huang, "Distributed fault detection scheme for intelligent connected vehicles in a zero-trust environment," in *2023 2nd Conference on Fully Actuated System Theory and Applications (CFASTA)*. IEEE, 2023, pp. 526–531.
- [15] T. E. Carroll, L. H. Chang, and C. L. Wright-Hamor, "The design and evaluation of zero trust architecture for electric vehicle charging infrastructure: Evs@ scale series on ev charging station cybersecurity," Pacific Northwest National Laboratory (PNNL), Richland, WA (United States), Tech. Rep., 2024.
- [16] P. Li, W. Ou, H. Liang, W. Han, Q. Zhang, and G. Zeng, "A zero trust and blockchain-based defense model for smart electric vehicle chargers," *Journal of Network and Computer Applications*, vol. 213, p. 103599, 2023.
- [17] M. Bampatsikos, I. Politis, T. Ioannidis, and C. Xenakis, "Trust score prediction and management in iot ecosystems using markov chains and madm techniques," *IEEE Transactions on Consumer Electronics*, 2025.

A Hierarchy-Focused Algorithm for Drawing Single Line Diagrams of Power Grids

Gabriel Ott

Carl von Ossietzky University Oldenburg
Oldenburg, Germany
Email: gabriel.ott@uol.de

Eric MSP Veith

OFFIS – Institute for Information Technology
R&D Division Energy
Oldenburg, Germany
Email: eric.veith@offis.de

Abstract—Visualizing power grids has always been important: Not only to show their current state, but especially to help operators grasp problems and find solutions quickly. To this end, the single-line diagram is still the most dominant method. However, many current diagrams are still drawn manually and only updated with values at run-time. While many approaches for automatic drawing exist, none of them respects voltage levels as important hierarchical organizational aspect. Thus, cross-level state visualization is still not automatic. To this end, we present a two-stage algorithm for automatic drawing of single-line diagrams.

Keywords—power grid topology; grid topology visualization; control room software

I. INTRODUCTION

Algorithms for creating visualizations of power grids are a well-studied field. Fischer et al. [1] provide an overview over the numerous forms of visualizations that exist in general. Notably, there is not one type that would be fitting for all kinds of tasks; rather, the form of displaying information to the user must be based upon the user's and the situation's specific needs. For example, choosing to include geospatial information or color-coding fundamentally changes requirements to the algorithm.

While Fischer et al. [1] note that modern grid operation software suites fail to leverage many of the more advanced visualization techniques developed in academia, one particular kind of diagram persists: The *single line diagram*. It is ubiquitous in modern control rooms. Since single-line diagrams are essentially orthogonal graph drawing problems, algorithms can draw from publications made as old as 1987 [2].

However, while classical orthogonal graph drawing algorithms optimize for aesthetic criteria, such as bend minimization, crossing reduction, and area compactness, they neglect a fundamental aspect critical to power system operation: voltage-level hierarchy. In power grids, the hierarchical organization of voltage levels—from high-voltage transmission networks (e. g., 400 kV, 220 kV) down through distribution levels (110 kV, 20 kV) to low-voltage feeders—directly reflects energy flow direction and failure propagation patterns. Grid operators rely on this hierarchy to quickly assess which components depend on which feeders during fault analysis and to understand cascading failure risks.

Existing automatic layout algorithms for single-line diagrams fall into three main categories, none of which adequately preserve voltage hierarchy. *Force-directed methods* optimize

for visual balance and minimal edge crossings but produce non-orthogonal layouts with arbitrary component placement that obscures hierarchical relationships. *Geographic approaches* prioritize spatial fidelity, maintaining real-world locations of substations at the expense of clear hierarchical structure. *Generic hierarchical algorithms* assign layers based on graph topology (e. g., longest paths, topological ordering) rather than domain semantics, potentially placing 400 kV and 20 kV components on the same visual layer simply because they are topologically adjacent.

This gap—i. e., algorithms not preserving the voltage hierarchy—has practical consequences in control room operations. When operators face time-critical situations, such as equipment failures, overload conditions, or cascading outages, they need immediate visual clarity about which voltage levels are affected and how failures might propagate through the hierarchy. Layouts that prioritize compactness or geographic accuracy over hierarchical structure force operators to mentally reconstruct the voltage-level dependencies, increasing cognitive load and response time [3]. Such a visualization is particularly important considering recent advances in power grid automation, where algorithms based on Artificial Intelligence (AI) (e. g., learning agents) are able to support operators in handling emergency situations [4], since these agents are most probably operating on a hierarchical representation of the grid, too.

In this paper, we present a hierarchy-focused algorithm specifically designed to preserve voltage-level structure in single-line diagrams of power grids. It consists of two distinct phases, node placement via the Semi-Automatic Ordering Algorithm (SAOA) algorithm and line drawing via the *Rollwire* algorithm. Unlike existing approaches, our approach treats voltage hierarchy as the primary organizing principle rather than as a secondary aesthetic concern. The algorithm performs a hierarchical decomposition based on voltage levels, assigns components to vertical layers corresponding to their voltage ratings, and applies orthogonal layout techniques within each layer while maintaining clear visual separation between levels. This approach ensures that operators can immediately identify high-voltage transmission infrastructure, mid-level distribution networks, and low-voltage feeders without searching through the diagram.

The remainder of this paper is organized as follows: We discuss related works in Section II. We describe an ordering algorithm that establishes a parent-child element ordering for

drawing purposes, called *Semi-Automatic Ordering Algorithm* (SAOA), in Section III. We detail how the layouting itself is done in Section IV. Before concluding, we provide a discussion of the algorithm in Section V.

II. RELATED WORK

Classical orthogonal graph drawing focuses primarily on aesthetic criteria such as bend minimization, area reduction, and crossing minimization. The seminal work of Tamassia [2] established the network flow approach for computing bend-minimal orthogonal representations of plane 4-graphs, achieving polynomial-time optimality for bend count. Biedl et al. [5] improved upon this with an $\mathcal{O}(n)$ algorithm producing at most $2n + 2$ bends in $n \times n$ area, while Freivalds et al. [6] developed a grid-based algorithm that minimizes total edge length through iterative local optimization and quadratic programming compaction. More recently, Kieffer et al. [7] introduced Human-like Orthogonal Network Layout (HOLA), a user study-driven algorithm that identified nine aesthetic criteria valued by humans, including tree-outside placement, compactness, symmetry, and “gridiness.” HOLA’s four-stage approach—decomposing graphs into cores and trees, stress-minimization layout, tree placement, and fine-tuning—produces more human-like results than traditional topology-shape-metrics approaches. However, none of these algorithms consider domain-specific hierarchical constraints. They optimize for visual aesthetics without preserving the semantic structure critical for power grid interpretation, where voltage hierarchy directly impacts operational understanding.

The Sugiyama framework [8] and its descendants provide principled approaches to layered graph layout through three phases: layer assignment, crossing reduction, and coordinate assignment. Brandes et al. [9] contributed fast horizontal coordinate assignment with vertical alignment in linear time, becoming widely adopted event in tools like *yFiles*. However, these methods assign hierarchy based on graph topology (e. g., longest path, topological ordering) rather than domain semantics. In power systems, the meaningful hierarchy is determined by voltage levels and energy flow direction, not graph structure. A substation at 400 kV must be visually separated from 110 kV equipment regardless of their topological distance. Applying generic hierarchical layouts to power grids fails to emphasize this critical semantic dimension, potentially placing high-voltage and low-voltage components on the same visual layer simply because they are topologically adjacent.

Automated generation of power system visualizations has received limited attention despite its operational importance. Birchfield et al. [10] developed force-directed and greedy approaches for geographic one-line diagrams, using Delaunay triangulation for line routing to maintain geographic context while minimizing substation overlap. Their method prioritizes spatial fidelity over hierarchical clarity. Earlier work by Raman et al. [11] provided foundational automatic generation algorithms, while Lendák et al. [12] applied force-directed methods requiring 300 iterations without orthogonal constraints, resulting in non-rectangular edges unsuitable for standard engineering

practice. More recent efforts have addressed specific subtasks: Hong et al. [13] proposed orthogonal layouts optimized for connectivity visualization but not hierarchy preservation; Yang et al. [14] introduced graph partitioning for scalable layouts of large distribution networks; and Sen et al. [15] developed incremental update methods for radial distribution systems with tree structures. Nagendra Rao et al. [16] specifically addressed distribution system diagram generation using tree drawing algorithms, noting that radial networks naturally form tree topologies rooted at substation transformers. They compared spring embedder, controlled spring embedder, and visibility-based approaches, concluding that existing general-purpose algorithms produce unsatisfactory layouts for power systems. However, their approach still applies graph-theoretic thinking by using the distance from the root element (i. e., the feeder) as metric, which in turn means that their algorithm is not suitable for visualizations that display several connected voltage levels. Critically, all these methods treat power networks as generic graphs, optimizing for compactness, crossing reduction, or geographic accuracy while neglecting the voltage hierarchy that operators rely on to assess failure propagation and system dependencies.

The literature reveals a fundamental gap: no existing algorithm explicitly preserves voltage-level hierarchy as a primary layout constraint. General orthogonal algorithms prioritize aesthetics; hierarchical methods impose topology-based layers; and power system tools focus on geographic or connectivity representation. Our algorithm addresses this gap by making parent-child voltage dependencies the central organizing principle, ensuring that operators can immediately identify which components depend on which feeders—critical information during fault analysis that is obscured by conventional layout algorithms optimized for different objectives. This hierarchy-first approach, combined with the semi-automatic philosophy allowing manual refinement, distinguishes our algorithm from all prior work in the intersection of graph drawing and power system visualization.

III. SEMI-AUTOMATIC ORDERING ALGORITHM

A. Overview and Definitions

Our algorithm works in a two-stage approach. First, the SAOA algorithm is responsible for placing nodes. Second, the actual drawing of nodes and especially the connecting lines is then done in a second step by the *Rollwire* layouting algorithm (cf. Section IV).

SAOA operates in four phases: (1) initialization, (2) parent-child assignment, (3) iterative space propagation, and (4) absolute position computation. The algorithm ensures that parent nodes are always positioned above their children, preserving the voltage-level hierarchy visually. Figure 1 shows the effect of following SAOA as it is described in the following sections.

Formally, let $G = (V, E)$ be a directed graph representing the power grid topology, where $V = \{v_1, v_2, \dots, v_n\}$ is the set of grid components (buses, transformers, generators, loads) and $E \subseteq V \times V$ is the set of directed edges representing electrical connections. For transformers, edge direction is determined

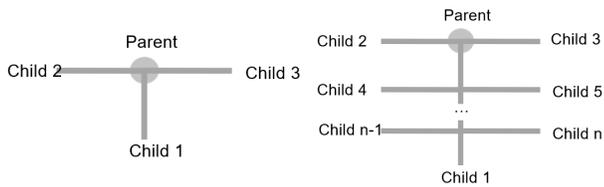


Figure 1. Placement of nodes as following our Semi-Automatic Ordering Algorithm

by voltage level: edges point from the high-voltage (upstream) side to the low-voltage (downstream) side, establishing the voltage hierarchy.

We define:

- **Parent function** $\pi : V \rightarrow V \cup \{\perp\}$: For each node $v \in V$, $\pi(v)$ denotes its unique parent node. If v is a root node, $\pi(v) = \perp$.
- **Children set** $C(v) = \{u \in V : \pi(u) = v\}$: The set of all children of node v .
- **Position** $\mathbf{p}(v) = (x_v, y_v) \in \mathbb{Z}^2$: The absolute grid coordinates of node v .
- **Relative offset** $\delta(v) = (\delta_x(v), \delta_y(v)) \in \mathbb{Z}^2$: The position of v relative to its parent, such that:

$$\mathbf{p}(v) = \mathbf{p}(\pi(v)) + \delta(v). \quad (1)$$

- **Space requirements** $\mathbf{s}(v) = (s_L(v), s_R(v), s_B(v)) \in \mathbb{N}_0^3$: The space required by node v and all its descendants in the left, right, and below directions, respectively.
- **Hierarchical depth** $d(v)$: The length of the path from the root to v , with $d(v_{\text{root}}) = 0$.

B. Phase 1: Initialization

Input: Graph $G = (V, E)$ with voltage-level annotations for transformer edges.

Output: Initial parent-child relationships and default offsets.

- 1) For each node $v \in V$, initialize $\delta(v) \leftarrow (0, 0)$ and $\mathbf{s}(v) \leftarrow (0, 0, 0)$.
- 2) For each edge $(u, v) \in E$, set $\pi(v) \leftarrow u$ and add v to $C(u)$.
- 3) For nodes with multiple potential parents, assign to the first parent with available capacity; flag conflicts for manual review.

C. Phase 2: Child Placement

For each node v with children $C(v) = \{c_1, c_2, \dots, c_k\}$, assign initial offsets based on child index. Let $b(v) \in \{0, 1\}$ indicate whether the parent-of-parent occupies a lateral position (i. e., $\delta_x(\pi(v)) \neq 0$):

$$\delta(c_i) = \begin{cases} (0, 1) & \text{if } i = 1 \text{ (first child: below)} \\ (-1, 0) & \text{if } i = 2 \text{ (second child: left)} \\ (+1, 0) & \text{if } i = 3 \text{ (third child: right)} \end{cases} \quad (2)$$

For $k > 3$ children, the first child is shifted downward to level $\ell = \lceil \frac{k+b(v)}{2} \rceil$, and additional children are placed alternately left and right at decreasing levels:

Input : Nodes V with initial offsets δ and parent assignments

π

Output : Final offsets δ ensuring overlap-free placement

repeat

$changed \leftarrow false$

foreach $v \in V$ **do**

// Aggregate children's space requirements:

foreach $direction\ d \in \{L, R, B\}$ **do**

$s_d(v) \leftarrow \max_{c \in C(v)} (s_d(c) + |\delta_d(c)|)$

end

// Update offset to parent based on sibling space:

if $\pi(v) \neq \perp$ **then**

foreach $direction\ d \in \{L, R, B\}$ **do**

$\sigma \leftarrow s(\text{siblings in opposite direction})$ $\delta'_d \leftarrow$

$s_d(v) + \sigma$

if $\delta'_d \neq \delta_d(v)$ **then**

$\delta_d(v) \leftarrow \delta'_d$

$changed \leftarrow true$

end

end

end

end

until $\neg changed$;

Figure 2. SAOA Iterative Space Propagation

$$\delta(c_1) = (0, \ell), \quad (3)$$

$$\delta(c_j) = \begin{cases} (-1, \ell - \lfloor j/2 \rfloor) & \text{if } j \text{ even} \\ (+1, \ell - \lfloor j/2 \rfloor) & \text{if } j \text{ odd} \end{cases} \quad \text{for } j > 1 \quad (4)$$

If $\pi(v)$ occupies the left or right position relative to v , that position is reserved and children are assigned to the next available slot.

D. Phase 3: Iterative Space Propagation

The core of SAOA is an iterative propagation of space requirements from leaves to root. Each node computes its space needs based on its children's requirements and adjusts its offset to its parent accordingly.

E. Phase 4: Absolute Position Computation

Once offsets have converged, absolute positions are computed by traversing from root to leaves:

$$\mathbf{p}(v) = \begin{cases} \mathbf{p}_0 & \text{if } \pi(v) = \perp, \\ \mathbf{p}(\pi(v)) + \delta(v) & \text{otherwise,} \end{cases} \quad (5)$$

where \mathbf{p}_0 is a predefined root position.

F. Handling Constraint Violations

When soft constraints are violated (multiple parents, cyclic dependencies), SAOA marks affected nodes for manual review rather than producing invalid layouts. This semi-automatic approach acknowledges that real power grids often exhibit

topological complexities that require human judgment to resolve optimally.

IV. THE ROLLWIRE LAYOUTING ALGORITHM

A. Problem Setting and Notation

Consider two rectangular symbols A and B placed on an integer grid, with one *docking port* (cable connection point) on each symbol. Let $\mathbf{c}_A, \mathbf{c}_B \in \mathbb{R}^2$ be the centers of the bounding boxes of A and B , and $\mathbf{p}_A, \mathbf{p}_B \in \mathbb{R}^2$ be the positions of their docking ports. Docking ports are constrained to lie on the midpoints of the four edges of the bounding box (top, bottom, left, right). The aim of Rollwire is to compute a polyline:

$$(\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_k), \quad \mathbf{h}_0 = \mathbf{p}_A, \mathbf{h}_k = \mathbf{p}_B, \quad (6)$$

such that all segments $\mathbf{h}_i \mathbf{h}_{i+1}$ are axis-aligned (orthogonal layout) and do not intersect the interior of any symbol.

Define the *port direction vectors*:

$$\mathbf{d}_A = \mathbf{p}_A - \mathbf{c}_A, \quad \mathbf{d}_B = \mathbf{p}_B - \mathbf{c}_B, \quad (7)$$

which encode on which side of the symbol the port is located. For axis-aligned rectangular symbols with docking ports at edge midpoints, the port direction vectors are axis-aligned: each \mathbf{d}_A and \mathbf{d}_B points in one of the four cardinal directions (right, left, up, or down). We can therefore normalize them to canonical form $\mathbf{d}_A, \mathbf{d}_B \in (1, 0), (-1, 0), (0, 1), (0, -1)$.

Let $\mathbf{\Delta} = \mathbf{p}_B - \mathbf{p}_A$ be the vector from the start port to the target port. For convenience, define unit direction vectors

$$\hat{\mathbf{e}}_x = (1, 0), \quad \hat{\mathbf{e}}_y = (0, 1), \quad (8)$$

and for any non-zero vector \mathbf{v} , let $\hat{\mathbf{v}} = \frac{\mathbf{v}}{\|\mathbf{v}\|}$ denote its normalized direction, and $\mathbf{v}_\perp = (-v_y, v_x)$ a 90-degree rotation used for orthogonal detours.

B. Overlap Detection

Before routing, Rollwire tests whether leaving the source or target port directly in the direction of the other port would cause the connecting segment to cross the corresponding symbol.

For a given port direction vector $\mathbf{d} \in \{\mathbf{d}_A, \mathbf{d}_B\}$, RollWire computes

$$s = \|\mathbf{\Delta} + \mathbf{d}\|, \quad t = \|\mathbf{\Delta}\|. \quad (9)$$

If $s < t$, the straight segment in direction $\mathbf{\Delta}$ is considered to overlap the symbol, because shifting the target by \mathbf{d} would result in a strictly shorter connecting vector, indicating that the symbol protrudes into the straight-line path. This test is applied separately for the source and the target port.

If neither test indicates an overlap, Rollwire simply draws one straight orthogonal polyline using at most one bend (depending on whether a purely horizontal or vertical alignment is possible).

If at least one overlap is detected, Rollwire introduces an intermediate orthogonal detour as described below.

Input : Start port \mathbf{p}_S , target port \mathbf{p}_T , start port direction vector \mathbf{d}_S , target port direction vector \mathbf{d}_T

Output : Ordered list of waypoints $(\mathbf{h}_0, \dots, \mathbf{h}_k)$ describing the orthogonal polyline

```

 $\mathbf{h}_0 \leftarrow \mathbf{p}_S, \mathbf{\Delta} \leftarrow \mathbf{p}_T - \mathbf{p}_S$ 
// Check for overlap conditions:
 $o_S \leftarrow (\|\mathbf{\Delta} + \mathbf{d}_S\| < \|\mathbf{\Delta}\|), o_T \leftarrow (\|\mathbf{\Delta} + \mathbf{d}_T\| < \|\mathbf{\Delta}\|)$ 
if  $\neg o_S \wedge \neg o_T$  then
    // No overlap: route directly with at
    // most one bend
    polyline  $\leftarrow$  DIRECTROUTE( $\mathbf{h}_0, \mathbf{p}_T$ )
    return polyline
end
// Overlap detected:
 $\hat{\mathbf{o}} \leftarrow \widehat{(\mathbf{d}_S)_\perp}, \mathbf{h}_1 \leftarrow \mathbf{h}_0 + \hat{\mathbf{o}}$ 
 $\mathbf{\Delta}' \leftarrow \mathbf{p}_T - \mathbf{h}_1$ 
if  $|\Delta'_x| \geq |\Delta'_y|$  then
    |  $\mathbf{h}_2 \leftarrow \mathbf{h}_1 + (\Delta'_x, 0)$ 
end
else
    |  $\mathbf{h}_2 \leftarrow \mathbf{h}_1 + (0, \Delta'_y)$ 
end
 $\mathbf{h}_3 \leftarrow \mathbf{p}_T$ 
return  $(\mathbf{h}_0, \mathbf{h}_1, \mathbf{h}_2, \mathbf{h}_3)$ 

```

Figure 3. RollWire Orthogonal Connection Routing

C. Two-Stage Orthogonal Detour

Let \mathbf{p}_S denote the start docking port (without loss of generality, the port on A) and \mathbf{p}_T the target docking port on B . Let \mathbf{d}_S be the direction vector of the start port.

Rollwire first moves one grid unit in a direction orthogonal to \mathbf{d}_S to “escape” the symbol’s bounding box:

$$\hat{\mathbf{o}} = \widehat{\mathbf{d}_{S,\perp}}, \quad \mathbf{h}_1 = \mathbf{p}_S + \hat{\mathbf{o}}. \quad (10)$$

This creates a short horizontal (or vertical) segment that clears the start symbol. From \mathbf{h}_1 , Rollwire connects towards \mathbf{p}_T along the dominant component of the remaining difference vector, $\mathbf{\Delta}' = \mathbf{p}_T - \mathbf{h}_1$.

If $|\Delta'_x| \geq |\Delta'_y|$, it draws a horizontal segment, $\mathbf{h}_2 = \mathbf{h}_1 + (\Delta'_x, 0)$, else a vertical segment, $\mathbf{h}_2 = \mathbf{h}_1 + (0, \Delta'_y)$.

Finally, RollWire draws one last orthogonal segment from \mathbf{h}_2 to the target port, $\mathbf{h}_3 = \mathbf{p}_T$, yielding a three-segment polyline $(\mathbf{h}_0, \mathbf{h}_1, \mathbf{h}_2, \mathbf{h}_3)$ with at most two bends. This construction guarantees that the polyline stays outside both rectangular symbols as long as the starting orthogonal step clears the bounding box by at least one grid unit.

D. Algorithmic Formulation

The core routing can be expressed as the following procedure.

Because SAOA guarantees that no other symbols lie between a parent and its child, Rollwire only needs to avoid the two endpoint symbols. This allows a very simple routing scheme with constant-time complexity per connection.

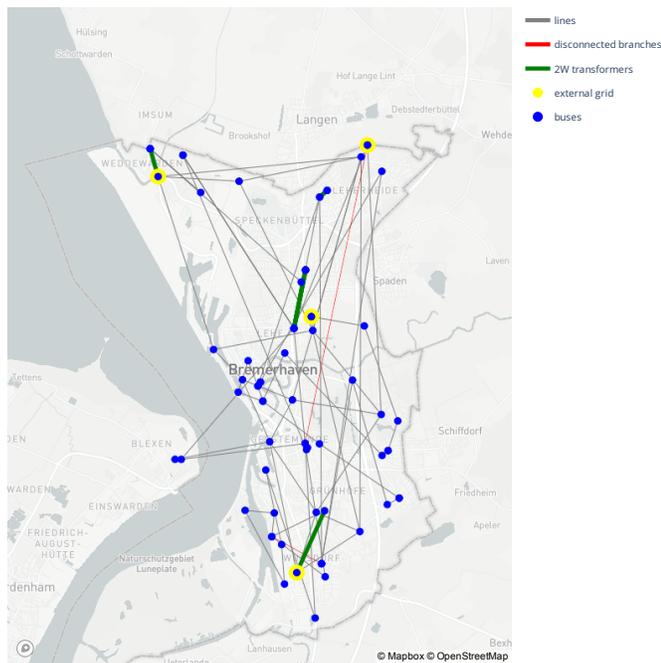


Figure 4. Topographic map of the power grid of the city of Bremerhaven, Germany

E. Segment Width and Endpoint Adjustment

In the implementation, each polyline segment is rendered as a narrow rectangle with width b . Naively joining such rectangles at right angles leads to small overlaps near bend points. To avoid this, Rollwire applies a local correction at each junction between two consecutive segments s_i and s_{i+1} : Let \hat{u}_i and \hat{u}_{i+1} be the unit direction vectors of s_i and s_{i+1} . The shared endpoint is shifted by:

$$\frac{b}{2} \hat{u}_i - \frac{b}{2} \hat{u}_{i+1}, \tag{11}$$

i.e., by half the width in the direction of the previous segment and by half the width opposite to the next segment. This adjustment ensures that the filled rectangles of adjacent segments meet cleanly without overlapping, producing visually consistent orthogonal connections even when line widths are non-zero.

F. Special Case: Junction Nodes

At junctions (busbars, switchgear symbols) multiple connections may terminate in the same symbol. In these cases, overlaps of line segments within the symbol are visually acceptable because they are covered by the symbol’s icon. Rollwire therefore permits multiple connections to share the same central docking point within such symbols, simplifying the routing and keeping the focus on avoiding overlaps outside the node area.

V. DISCUSSION

To illustrate the algorithm, we use a medium-voltage city power grid. Our model represents the power grid of Bremerhaven, Germany. It is part of the MIDAS scenario

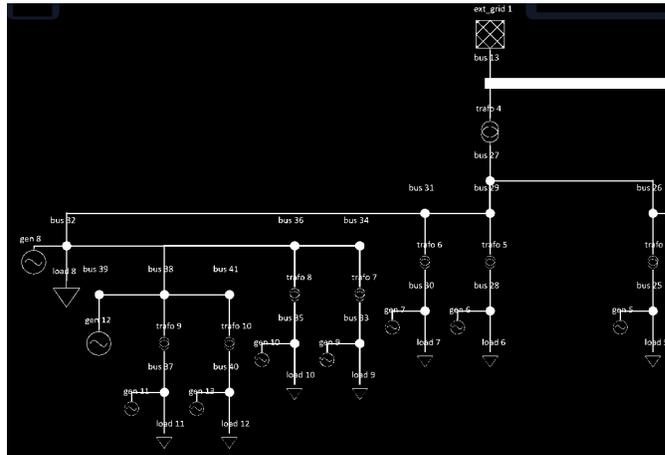


Figure 5. Rollwire’s representation of the Bremerhaven power grid

package [17, 18]. Figure 4 shows the topographic plot of the grid. The data is given to SAOA and Rollwire, which produce the output seen in Figure 5. Note that we show a cutout of the grid due to paper size limits.

Having established that our algorithm produces acceptable single-line diagrams, we discuss the runtime complexity of the two parts. First, we analyze the runtime complexity of SAOA by examining the following theorem:

Theorem V.1. SAOA terminates in at most $h + 1$ iterations, where $h = \max_{v \in V} d(v)$ is the maximum hierarchical depth of the graph.

Proof. In each iteration, space requirements propagate upward by one level in the hierarchy. Since changes at depth d can only affect ancestors at depths $< d$, and the root has depth 0, convergence is guaranteed within $h + 1$ iterations. □

We can then turn to the drawing mechanism, Rollwire. For a single connection between two ports with direction vectors, Rollwire requires:

- $\mathcal{O}(1)$ vector arithmetic operations for overlap detection, (cf. Eq. (9)).
- $\mathcal{O}(1)$ operations for polyline construction (Algorithm 3).
- $\mathcal{O}(1)$ endpoint adjustments per segment.

For a power grid with m edges (connections), the total routing time is $\mathcal{O}(m)$. Combined with SAOA’s $\mathcal{O}(n^{1.5})$ node placement, the overall layout pipeline achieves $\mathcal{O}(n^{1.5} + m)$ complexity, which is practical for typical grids.

VI. CONCLUSION AND FUTURE WORK

In this paper, we presented a two-stage algorithm for drawing single line diagrams. Its major contribution is that it produces pleasant, hierarchical layouts that respect voltage levels and are thus suitable to display larger grids with multiple levels.

In the future, we will showcase the algorithm in live demonstrations.

ACKNOWLEDGEMENTS

SAOA und Rollwire are part of the research project ProRES, funded by the German Federal Ministry for Economic Affairs and Energy under grant no. 03EI4117A.

REFERENCES

- [1] M. T. Fischer and D. A. Keim, *Towards a survey of visualization methods for power grids*, Apr. 2022. DOI: 10.48550/arXiv.2106.04661. arXiv: 2106.04661[cs]. Accessed: Jan. 29, 2026. [Online]. Available: <http://arxiv.org/abs/2106.04661>.
- [2] R. Tamassia, "On embedding a graph in the grid with the minimum number of bends," *SIAM Journal on Computing*, vol. 16, no. 3, pp. 421–444, Jun. 1987, Publisher: Society for Industrial and Applied Mathematics, ISSN: 0097-5397. DOI: 10.1137/0216030. Accessed: Jan. 29, 2026. [Online]. Available: <https://epubs.siam.org/doi/abs/10.1137/0216030>.
- [3] U. Afzal et al., "Investigating cognitive load in energy network control rooms: Recommendations for future designs," *Frontiers in Psychology*, vol. 13, Mar. 2022, Publisher: Frontiers, ISSN: 1664-1078. DOI: 10.3389/fpsyg.2022.812677. Accessed: Jan. 29, 2026. [Online]. Available: <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2022.812677/full>.
- [4] E. M. Veith, T. Logemann, A. Wellßow, and S. Balduin, "Play with me: Towards explaining the benefits of autocurriculum training of learning agents," in *2024 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE)*, Dubrovnik, Croatia: IEEE, 2024, pp. 1–5. DOI: 10.1109/ISGTEUROPE56780.2023.10408277.
- [5] T. Biedl and G. Kant, "A better heuristic for orthogonal graph drawings," *Computational Geometry*, vol. 9, no. 3, pp. 159–180, Feb. 1998, ISSN: 0925-7721. DOI: 10.1016/S0925-7721(97)00026-6. Accessed: Jan. 29, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0925772197000266>.
- [6] K. Freivalds and J. Glagolevs, "Graph compact orthogonal layout algorithm," in *Combinatorial Optimization*, P. Foulhoux, L. E. N. Gouveia, A. R. Mahjoub, and V. T. Paschos, Eds., Cham: Springer International Publishing, 2014, pp. 255–266, ISBN: 978-3-319-09174-7. DOI: 10.1007/978-3-319-09174-7_22.
- [7] S. Kieffer, T. Dwyer, K. Marriott, and M. Wybrow, "HOLA: Human-like orthogonal network layout," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 1, pp. 349–358, 2015. DOI: 10.1109/TVCG.2015.2467451.
- [8] K. Sugiyama, S. Tagawa, and M. Toda, "Methods for visual understanding of hierarchical system structures," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 2, pp. 109–125, 1981. DOI: 10.1109/TSMC.1981.4308636.
- [9] U. Brandes and B. Köpf, "Fast and simple horizontal coordinate assignment," in *Graph Drawing*, ser. Lecture Notes in Computer Science, vol. 2265, Springer, 2002, pp. 31–44. DOI: 10.1007/3-540-45848-4_3.
- [10] A. B. Birchfield and T. J. Overbye, "Techniques for drawing geographic one-line diagrams: Substation spacing and line routing," *IEEE Transactions on Power Systems*, vol. 33, no. 6, pp. 7269–7277, 2018. DOI: 10.1109/TPWRS.2018.2846283.
- [11] N. Raman and R. K. Schwartzberg, "Automatic generation of power system one-line diagrams," in *IFAC Proceedings Volumes*, vol. 19, 1986, pp. 231–236. DOI: 10.1016/S1474-6670(17)59383-0.
- [12] I. Lendák and A. Poth, "Electric power system one-line diagram generation technology," in *Proceedings of the IEEE International Energy Conference*, 2012, pp. 152–157. DOI: 10.1109/EnergyCon.2012.6348286.
- [13] S. P. S. Hong, J. Y. Lee, and P. S. Mah, "Power system connectivity visualization using an orthogonal graph layout algorithm," *Journal of Computing Science and Engineering*, vol. 16, no. 4, pp. 187–199, 2022. DOI: 10.5626/JCSE.2022.16.4.187.
- [14] C. Yang, J. Wang, and D. Shi, "Automatic layout of one-line diagrams for large transmission systems based on graph partitioning," in *2024 IEEE Power & Energy Society General Meeting (PESGM)*, ISSN: 1944-9933, Jul. 2024, pp. 1–5. DOI: 10.1109/PESGM51994.2024.10689189. Accessed: Jan. 29, 2026. [Online]. Available: <https://ieeexplore.ieee.org/document/10689189>.
- [15] P. Sen et al., "An Automatic Generation Method for Single-line Diagram of Distribution Network With Supporting Incremental Update," *Modern Electric Power*, vol. 41, no. 1, pp. 21–28, Feb. 2024, ISSN: 1007-2322. DOI: 10.19725/j.cnki.1007-2322.2022.0189. Accessed: Jan. 29, 2026. [Online]. Available: <http://xddl.ncepujournal.com/en/article/doi/10.19725/j.cnki.1007-2322.2022.0189.pdf>.
- [16] P. S. Nagendra Rao and R. Deekshit, "Distribution feeder one-line diagram generation: A visibility representation," *Electric Power Systems Research*, vol. 70, no. 3, pp. 173–178, Aug. 2004, ISSN: 0378-7796. DOI: 10.1016/j.epsr.2003.12.005. Accessed: Jan. 29, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378779603003031>.
- [17] E. Veith et al., "palaestrAI: A training ground for autonomous agents," in *Proceedings of the 37th annual European Simulation and Modelling Conference*, Oct. 2023.
- [18] S. Balduin, E. M. S. P. Veith, and S. Lehnhoff, "MIDAS: An open-source framework for simulation-based analysis of energy systems," in *Simulation and Modeling Methodologies, Technologies and Applications*, G. Wagner, F. Werner, and F. De Rango, Eds., Cham: Springer International Publishing, 2023, pp. 177–194, ISBN: 978-3-031-43824-0.

Analysis of the Influence of Photovoltaic Production on Grid Voltage using Data from Inverters

Santiago-Chiquero I., Arenas-Ramos V., Pallarés-López V., González-Redondo V., Real-Calvo R.
 Dept. of Electronic and Computer Engineering
 Universidad de Córdoba
 Córdoba, Spain
el1sachi@uco.es, p62arrav@uco.es, el1palov@uco.es,
el1gorem@uco.es, el1recar@uco.es

Florencias-Oliveros, O., Gonzalez de la Rosa, JJ.
 Department of Automation Engineering, Electronics,
 Architecture and Computers Networks
 Universidad de Cádiz
 Algeciras, Spain
olivia.florencias@uca.es, juanjose.delarosa@uca.es

Abstract—Renewable energy sources are playing an increasingly important role in the move towards decarbonization and the reduction of polluting gases. Specifically, photovoltaic solar energy is the technology whose installed capacity has experienced the greatest growth in recent years, already exceeding 2.2 TW worldwide. This type of clean energy generation has other drawbacks. Its dependence on atmospheric conditions not only affects its availability, but its fluctuating nature, which can change rapidly with the passage of clouds, can affect the stability of the grid and make it difficult to manage. Specifically, this inherent variability influences on grid voltage values, among other factors, potentially leading to fluctuations as a result of changes in photovoltaic output power. This paper presents an analysis of the grid voltage values recorded by an inverter located in a solar park over a period of almost five years with a frequency of 5 minutes. The effect of photovoltaic generation on the increase in voltage values during the middle of the day, when output power is at its highest is clearly observed, as well as an increase in the fluctuations experienced by this parameter at the point where photovoltaic output power is injected into the grid.

Keywords: Photovoltaic plants; grid voltage fluctuation; inverters.

I. INTRODUCTION

The trend towards decarbonization of the current energy system has led to a gradual increase in the use of renewable energy sources. Among these, photovoltaic (PV) energy has established itself as one of the cleanest energy sources with the highest installed capacity and has been the technology with the highest growth rate in recent years. This increase must continue in the coming years in order to meet the polluting emission reduction targets set by various organizations [1] [2] [3] [4].

However, PV energy is inherently variable and intermittent, highly dependent on weather conditions. The output of a PV plant can change dramatically in a few seconds or minutes due to phenomena such as passing cloud, creating variability that the electrical grid must accommodate. Specifically, this uncontrolled fluctuation has a direct and negative impact on power quality, particularly

regarding the voltage levels at the injection points of PV installations. This issue is more critical in weak grids and will become more pronounced as the penetration of this type of renewable generation increases. Such variability poses significant challenges not only for the quality of the electricity signal, but also for grid operation and stability [5], potentially acting as a barrier to the future expansion of renewable energy resources [6]. Consequently, understanding voltage behaviour in the presence of renewable energy production is important for developing and sizing effective mitigation techniques [5]. Although studies that simulate network behaviour provide valuable insights [7] [8], the analysis of real data is also equally indispensable.

In this context, the objective of this study was to analyze the real voltage values measured at the grid injection point by an inverter in a PV plant in operation over a period of several years. Specifically, it investigates how the energy produced by the PV panels influences voltage variations at that specific inverter location.

The remainder of this article is organized as follows: Section II describes the methodology, Section III presents the experimental results, and finally, Section IV summarizes the conclusions of this work.

II. METHODOLOGY

The analysis conducted in this study utilizes data recorded by an inverter in a PV plant. The inverter is an SMA SMC-100 model. Among other parameters, the inverter records the value of the PV current produced at the inverter output, as well as the voltage values in the three phases at the grid injection point. The inverter provides measurements five minutes intervals. Data is available for almost five years, from November 3, 2020, to July 4, 2025, representing a total of 1,705 days recorded, with 288 data points per day. The raw daily data provided by the inverter has been stored in a database developed in previous work by the authors, designed for the automated management of the entire volume of data recorded at various PV plants [9]. A total of 663 Silicon modules with 165 W are connected to the inverter analyzed, organized into strings of 17 modules in series. This represents a total of 109.395 kW. The inverter

is situated in a solar park with a total power of 4.4 MW, located in a rural municipality in the north of the province of Córdoba, Spain.

III. RESULTS

As an example of the analyzed dataset, Figure 1 shows the daily voltage profiles for each of the three phases and the current produced at the inverter output during March 2021. It can be observed that during this month 70% of the days were cloudy or experienced intermittent cloud cover, which led to significant fluctuations in PV production. It can be observed how the voltage values rise throughout the morning, with the increase in PV production, subsequently decreasing toward sunset. This trend is consistent across all three phases, reflecting a uniform network response to plant production. While voltage levels fluctuate during the night, these variations become more pronounced during the day due to cloud-induced fluctuations in PV generation.

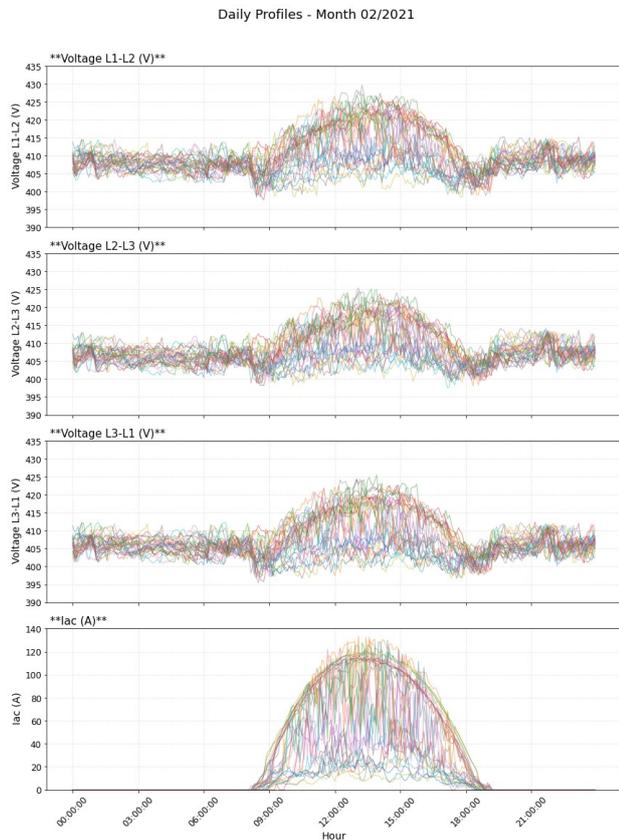


Figure 1. Daily profiles of line voltage in the three phases and the current generated during all days of February 2021.

Figure 2 shows, in this case, the daily voltage and current profiles generated during a summer month, July 2021. In the geographical location where the solar park is located, significantly more frequent during this season.

Figure 3 provides a more detailed example of a fluctuating daily profile for both the current produced and the voltage in one of the phases on a day with unstable

weather conditions and passing clouds, which leads to variations in both production and voltage. Figure 4 provides a detailed view of the daily profiles on a clear-sky day. Once again, a noticeable increase in voltage values can be observed during the peak solar hours, directly related to the period of maximum photovoltaic production.

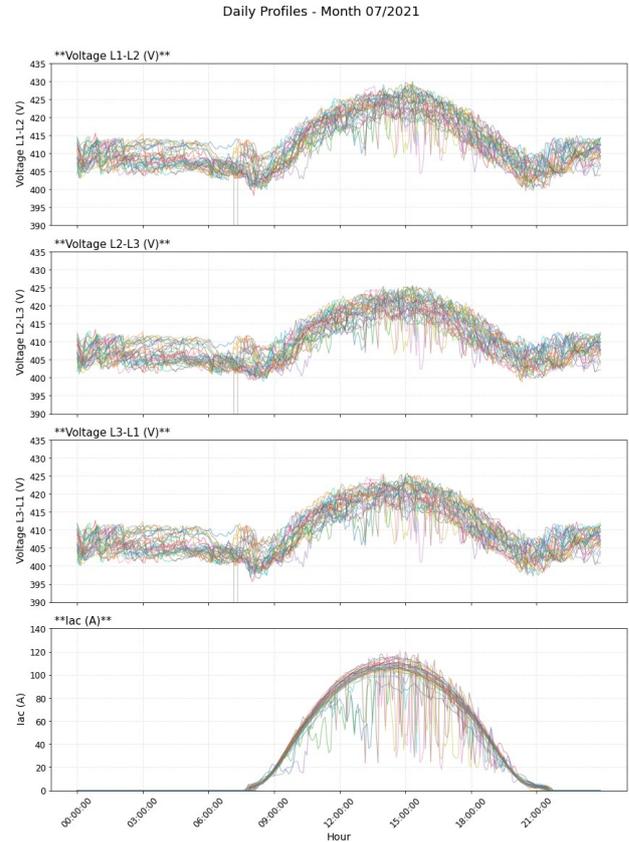


Figure 2. Daily profiles of line voltage in the three phases and the current generated during all days of July 2021.

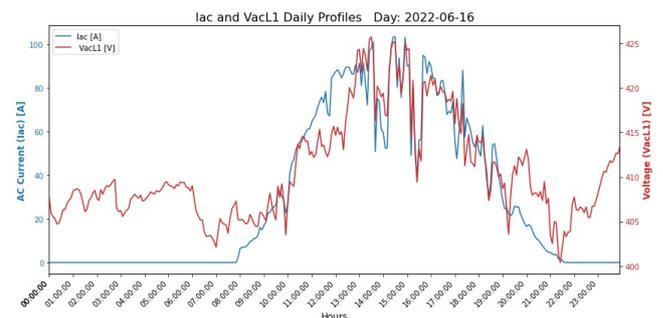


Figure 3. Daily profiles of line voltage in one phase and the current generated during a day of June 2022.

Although clouds are absent and the plant's production follows the solar trajectory, the voltage levels still exhibit fluctuations. These variations are typical of the operational dynamics of the electrical grid where the measurements are recorded and are independent of the photovoltaic plant's power output.

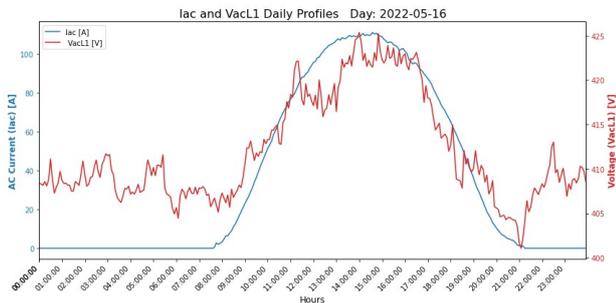


Figure 4. Daily profiles of line voltage in one phase and the current generated during a day of May 2022.

Correlation Voltage and PV current - Month 07/2021

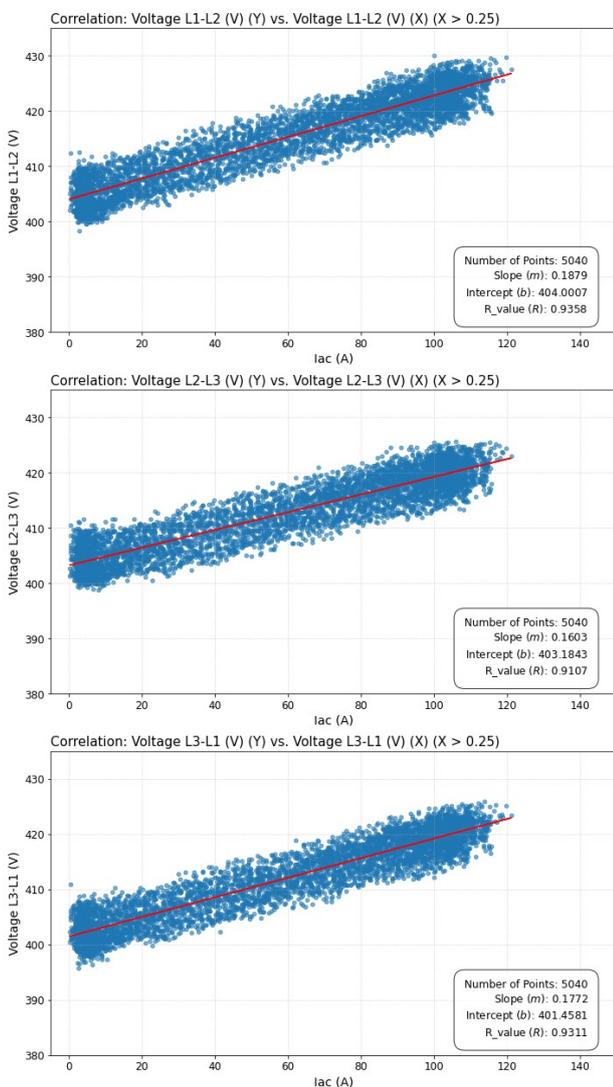


Figure 5. Correlation between the voltage values of the three phases versus the current, when there is PV production, during the month of July 2021.

Figure 5 illustrates the correlation between the current values produced and the corresponding voltage levels during

operating hours in July 2021. A clear linear relationship between these two parameters is evident, indicating that voltage rises with the current injected into the grid.

For each value of the current generated, the voltage may vary by around 10 V, which is a consequence of the operation of the grid and the elements connected to it. Nevertheless, a clear upward trend in voltage is observed as photovoltaic production increases. This trend shows slight differences between phases, which highlights that the sensitivity of the network in PV power output variations will depend on the specific configuration and on the elements connected to each phase [10]. This behavior was consistently observed across all analyzed months.

Although voltage values show an increase, they are not approaching the maximum levels established by standards [11]. However, the calculated slopes suggest that these voltage variations could eventually limit the hosting capacity of the area—restricting the total amount of generation that can be integrated into the local grid [12].

Figure 6 compares the voltage distribution throughout the monitored period, distinguishing between hours with PV production (top green histogram) and hours without generation (bottom red histogram).

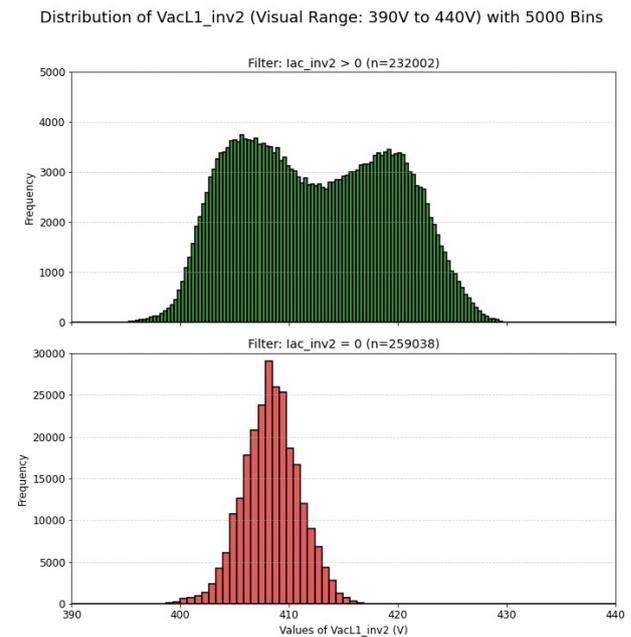


Figure 6. Histogram of the frequency with which voltage values are observed in one of the phases with and without PV power generation (Bins: intervals).

The night-time distribution shows a narrow, single-peak behavior centered around 408 V, indicating a more stable voltage regime when the PV plant is inactive. In contrast, during production hours, the distribution becomes significantly broader and exhibits a bimodal behavior, with a second peak emerging near 420 V. This shift toward higher voltage values confirms that solar injection elevates the local

grid voltage. However, it should be noted that the inherent variations in grid management and the dynamics of connected loads, typically differ between daylight and nighttime hours.

Distribution of VacL1_inv2_delta (Visual Range: -15V to 15V) with 6400 Bins

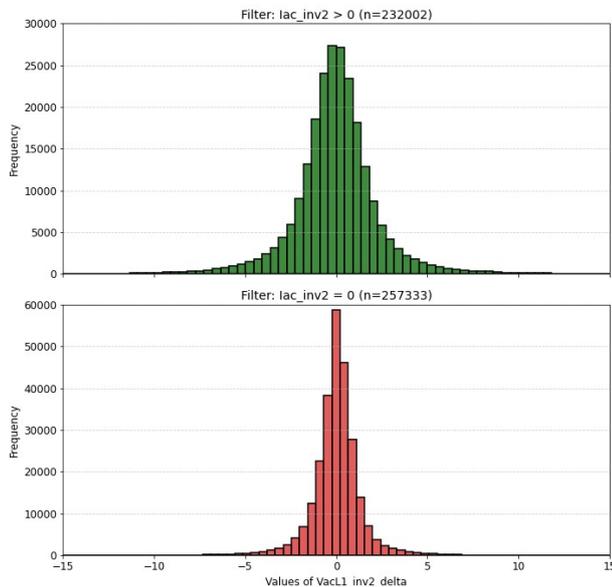


Figure 7. Histogram of the frequency with which voltage variation is observed in one of the phases with and without PV power generation (Bins: intervals).

A comparison of the voltage variations presented in Figure 7 reveals that, during hours of photovoltaic production, fluctuations are primarily contained within a range of -12 V to 12 V. In contrast, when the plant is inactive, the magnitude of these variations narrows significantly to between -7 V and 7 V.

IV. CONCLUSIONS

Data recorded by PV inverters provides valuable insights into the impact of solar production on power quality, particularly regarding grid voltage stability. Although this impact is influenced by the grid's architecture, components, and operational management, a clear correlation exists between increased PV output and rising voltage levels. Furthermore, there is a direct link between voltage variation and PV production variability caused by cloud transients. These findings are essential for strategic grid planning and for determining the hosting capacity limits at specific injection points, ensuring that renewable integration does not compromise voltage regulation standards.

While data provided by PV inverters typically exhibits lower accuracy and resolution compared to measurements from specialized monitoring systems, it remains a highly valuable resource. However, since this data is usually available at most PV installations, it can be effectively used to analyse grid voltage trends in relation to PV production without the need for additional specialised hardware. Consequently, these inverter data can provide a solid basis

for modelling system behaviour and for developing scalable mechanisms that mitigate the adverse effects of PV output fluctuations, facilitating the integration of renewable energies while maintaining grid stability.

ACKNOWLEDGMENT

This work was developed under the projects PID2024 158091OB C21 *Edge Management of photovoltaic plants based in Seamless Temporal Accuracy analytical architecture* (EMSTAC) and PID2024 158091OB C22 *Analysis of high-quality measurements data analytics for FLEXible manAGement in photovoltaic plants* (AFLEXGE) granted by the Spanish Ministry of Science and Innovation and the State Investigation Agency.

The authors would also like to express their gratitude to Solar del Valle S.L., the company responsible for the photovoltaic installation where the data analyzed in this study was recorded.

REFERENCES

- [1] European Commission, "European Green Deal." [Online]. Available: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en (last access: 2-12-2025)
- [2] European Council of the European Union, "Fit for 55." [Online]. Available: <https://www.consilium.europa.eu/en/policies/fit-for-55/> (last access: 2-12-2025)
- [3] International Energy Agency (IEA), "Solar PV continues to dominate global renewables growth, though at a slower rate." [Online]. Available: <https://www.iea.org/energy-system/renewables/solar-pv> (last access: 2-12-2025)
- [4] UNEF (Spanish Photovoltaic Association), "Photovoltaic Energy, the Engine of Industrialization, motor de industrialización. Annual Report UNEF 2025," 2025. [Online]. Available: <https://www.unef.es/es/recursos-informes> (last access: 2-12-2025)
- [5] F. H. Malik *et al.*, "A Comprehensive Review on Voltage Stability in Wind-Integrated Power Systems," *Energies*, vol. 17, no. 3, 2024, doi: 10.3390/en17030644.
- [6] M. C. Kocer *et al.*, "Cloud Induced PV Impact on Voltage Profiles for Real Microgrids," in *2018 5th International Symposium on Environment-Friendly Energies and Applications (EFEA)*, 2018, pp. 1–6. doi: 10.1109/EFEA.2018.8617080.
- [7] A. Gawlak and M. Kornatka, "The Impact of Photovoltaic Installations on Changes in Voltage Levels in the Low-Voltage Network," *Energies*, vol. 18, no. 12, 2025, doi: 10.3390/en18123072.
- [8] A. Woyte, V. Van Thong, R. Belmans, and J. Nijs, "Voltage fluctuations on distribution level introduced by photovoltaic systems," *IEEE Trans. Energy Convers.*, 2006, doi: 10.1109/TEC.2005.845454.
- [9] D. Trillo-Montero, S. Cosano-Lucena, M. Gonzalez-Redondo, J. J. Luna-Rodriguez, and I. Santiago, "Design and Development of a Relational Database Management System (RDBMS) with Open Source Tools for the Processing of Data Monitored in a Set of Photovoltaic (PV) Plants," *Appl. Sci.*, vol. 13, no. 3, 2023, doi: 10.3390/app13031357.
- [10] I. Santiago, J. García-Quintero, G. Mengibar-Ariza, D. Trillo-Montero, R. J. Real-Calvo, and M. Gonzalez-Redondo, "Analysis of Some Power Quality Parameters at the Points of Common Coupling of Photovoltaic Plants Based on Data Measured by Inverters," *Appl. Sci.*, vol. 12, no. 3, 2022, doi: 10.3390/app12031138.
- [11] "European Standard EN 50160: 2011. Voltage characteristics of electricity supplied by public electricity networks." 2011. [Online]. Available: <https://en.tienda.aenor.com/norma-din-en-50160-2020-11-327353625>
- [12] L. Freiris and D. Infield, *Renewable Energy in Power Systems*. Wiley, 2008.

Metro-Integrated Electrified Campus Shuttles for Green Mobility in Saudi Universities

Mohammed Mutwaly, Khaled Alkhamis and Walied Alfraidi
 Electrical Engineering Department, College of Engineering,
 Imam Mohammad Ibn Saud Islamic University (IMSIU)
 Riyadh, Saudi Arabia
 e-mail: walfraidi@imamu.edu.sa

Abstract—The transition toward low-carbon mobility in Saudi university campuses is gaining strategic importance, particularly in rapidly urbanizing cities such as Riyadh. This paper presents a metro-integrated electric shuttle framework to enhance first/last-mile connectivity and replace diesel-based campus transit at Imam Mohammad Ibn Saud Islamic University (IMSIU). The framework evaluates electrification pathways, battery-sizing requirements, charging coordination, and metro-synchronized shuttle scheduling to ensure reliable service continuity. Results indicate that diesel-to-electric conversion eliminates approximately 79.17 tons of direct tailpipe CO_2 annually while improving on-campus air quality. Battery-sizing and load analyses further show that daily charging minimizes required capacity and peak electrical demand, whereas 12-m buses impose substantially higher energy and grid-load requirements than 8.5-m models. The proposed framework provides a scalable and transferable reference for universities seeking to enhance multimodal mobility and support national sustainability objectives under Saudi Vision 2030.

Keywords—Green Campus Mobility; Electrified Shuttles; Metro Integration

I. INTRODUCTION

The transition toward low-carbon mobility has become a defining indicator of institutional sustainability, with university campuses increasingly recognized as high-value environments for deploying clean transportation models. In Saudi Arabia, where urban expansion is accelerating alongside national commitments to emission reduction, campus mobility systems must evolve to address environmental and operational inefficiencies. Diesel-based shuttle fleets, though historically dependable, contribute directly to on-campus greenhouse gas emissions, airborne pollutants, noise, and degraded commuter exposure along high-density academic corridors. These impacts highlight the need for a structured, evidence-based pathway toward electrified campus mobility that supports healthier learning environments and advances national sustainability goals. The development of high-capacity public transit in Riyadh, led by the city's expanding metro network, introduces a transformative platform for re-engineering first/last-mile campus transport operations within a multimodal mobility ecosystem. Despite metro availability, transit adoption is often constrained by first-mile limitations and the high rate of private vehicle arrivals at university gates, which intensify traffic congestion, commuter delays, and fuel-based emissions from prolonged idling. Electrified campus shuttle networks, when operationally aligned with metro headways, offer a strategic solution to bridge connectivity gaps, reduce gate traffic, remove diesel tailpipe pollution from internal routes, and enhance local-

ized air quality exposure within the campus. This positions Saudi university campuses not only as users of national green transit infrastructure, but as active contributors enabling sustainable urban commuter flows in Riyadh. Electric buses operating in hot climates face unique performance and reliability challenges, primarily due to intensive thermal loads on batteries, electric drivetrains, and Heating, Ventilation, and Air Conditioning (HVAC) systems. As reported in [1], high ambient temperatures significantly accelerated battery degradation, increased cooling demand, and reduced driving range, making thermal management a critical barrier to widespread battery electric buses deployment in warm regions. A life-cycle assessment was presented in [2], comparing diesel, battery-electric, and hydrogen fuel-cell buses in Saudi Arabia. It highlighted that fuel-cell buses currently offered the largest emission reductions, while the benefits of battery-electric buses increased significantly as the national grid incorporates more renewable energy. It showed the strong influence of local climate, operational loads, and energy-supply pathways on bus performance, emphasizing the need for context-specific modeling. An optimization model was proposed in [3] to jointly consider bus service scheduling and charging strategies, demonstrating substantial cost savings and improved efficiency in large-scale battery electric bus networks. A bus replacement strategy was developed in [4], where buses with low battery levels were swapped with fully charged standby buses at charging stations, allowing continuous service without long dwell times. Using a mixed-integer optimization model, the bus schedules, charging station locations, and charger quantities were jointly determined while meeting service level and passenger-comfort constraints. Recent research in [5] has introduced an aggregator-based charging optimization methodology for multifunctional electric bus charging stations. Using coordinated operational-planning and charging optimization algorithms, the approach integrated normal charging, fast charging, and battery swapping while enhancing renewable energy usage and supporting grid services. A method was developed in [6] to improve real-time energy consumption estimation for electric buses by modeling stochastic operational factors, such as speed, acceleration, and passenger load. Using an enhanced Kalman filter and feature-reduction techniques, the approach achieved accurate acceleration and energy estimates, validated with real transit data. These improvements supported better characterization of power system loading and operational impacts from growing electric-bus adoption. A battery-sizing strategy was examined in [7] for electric ap-

buses, by modeling powertrain dynamics and auxiliary loads to accurately estimate energy requirements. Multiple commercially available battery capacities were evaluated through simulation and validated using standard drive cycles to meet a 100-km range target. An optimization framework was presented in [8] for electrifying bus rapid transit systems, by jointly determining battery sizes and charging infrastructure to minimize total system costs. Using real-world transit data, the study showed that coordinated planning of battery capacity, charger placement, and operational schedules could significantly reduce both capital and operating costs while maintaining reliable service. The authors of [9] developed a multi-criteria framework to evaluate hybrid charging infrastructure for battery-electric and fuel-cell buses, integrating grid electricity, solar PV, battery storage, and hydrogen systems. Their results highlighted the trade-offs between investment, environmental performance, and operational reliability when planning charging infrastructure for zero-emission bus fleets. The feasibility of replacing diesel buses with electric buses on an Ottawa transit route was evaluated in [10], by modeling energy consumption under Ontario-specific operating conditions. The study designed a mixed charging strategy combining depot fast charging with opportunity pantograph charging and performed a Well-to-Wheel analysis to quantify greenhouse-gas reductions based on the provincial energy mix. A cost assessment further estimated the investment payback period, offering a practical framework for assessing the environmental and economic viability of electric-bus deployment. While existing studies provide valuable insights into electric bus deployment and charging infrastructure design, their direct application to university-scale shuttle systems remains limited. Urban-scale optimization models are often computationally intensive and impractical for early-stage campus planning, whereas many depot-based studies neglect realistic operational constraints such as limited charging windows and battery cycling requirements. To address these limitations, this paper proposes a unified analytical framework that integrates operational constraints, depth-of-discharge limits, and charging duration into a transparent and computationally efficient design methodology for electrified campus shuttle systems. At Imam Mohammad Ibn Saud Islamic University (IMSIU), campus shuttles currently operate as a dedicated fleet covering academic zones and student housing corridors. Transitioning this system to battery-electric propulsion eliminates these tailpipe emissions entirely, providing a measurable carbon reduction gain and an immediate improvement in the exposure micro-environment along shuttle corridors. The main contributions of the paper can be summarized as follows:

- Introduces a metro-integrated campus mobility framework, synchronizing electric shuttle operations with Riyadh Metro train headways to enable reliable multimodal first/last-mile connectivity and reduce on-campus vehicle emissions.
- Develops a structured diesel-to-electric fleet transition model, combining real shuttle-route characteristics, en-

ergy consumption modeling for 8.5-m and 12-m buses, and route-specific battery-sizing analysis that incorporates depth-of-discharge constraints, charging frequency, and operational feasibility constraints.

- Incorporates a rigorous well-to-wheel emission analysis that accounts for local electricity grid carbon intensity, enabling a realistic comparison between diesel and electric university shuttle buses under current grid conditions.
- Establishes a fleet-level charging load and infrastructure assessment for the full IMSIU shuttle fleet, quantifying peak power demand, charging energy per event, and grid impacts under multiple charging durations and frequencies.
- Presents a scalable sustainability blueprint for Saudi universities, positioning campus shuttle electrification as an operational mechanism supporting national low-carbon mobility objectives under Saudi Vision 2030.

The remainder of the paper is organized as follows: Section II outlines the proposed metro-integrated electrified shuttle framework. Section III describes the operational data and modeling inputs used in the analysis. Section IV presents and discusses the results, including emissions, battery sizing, charging loads, and fleet-level energy demand. Section V concludes the paper.

II. PROPOSED METRO-INTEGRATED ELECTRIFIED CAMPUS SHUTTLE FRAMEWORK

This section presents the overall conceptual and analytical structure of the proposed framework, including route topology modeling, fleet sizing methodology, and grid-impact assessment.

A. Metro-Integrated Operational Topology

The framework is established on a ‘‘Metro-Integrated’’ mobility concept, where the campus shuttle network functions as a dedicated first/last-mile feeder system. The operational topology is constrained by the Riyadh Metro service, where all shuttle routes originate and terminate at the campus Metro Station. To quantify the baseline operational demand, the daily distance traveled by each bus (d_b) on route r is determined from the route topology:

$$d_b = \frac{L_r \cdot T_r}{N_r} \quad (1)$$

where L_r is the loop length of route r (km), T_r is the total shuttle trips per academic day on route r , and N_r is the number of buses assigned to route r .

B. Electric Fleet Battery Capacity Sizing

Building on the per-bus daily distance from (1), the daily energy demand for each bus is as follows:

$$E_b = d_b \cdot \xi_{cons} \quad (2)$$

where ξ_{cons} is the specific energy consumption coefficient. The nominal battery capacity (E_{nom}) required to sustain n_d operational days between consecutive charging events is given by:

$$E_{nom} = \frac{n_d \cdot E_b}{\text{DoD}_{\max}} \quad (3)$$



Figure 1 Yellow route of the shuttle network.



Figure 2 Blue route of the shuttle network.

where n_d is the number of operational days per charging interval and DoD_{\max} is the maximum allowable depth of discharge, defined by the operator as a design parameter.

C. Infrastructure Load Impact Assessment

The average required charging power per bus ($P_{req,b}$) depends on the accumulated energy demand and the available charging window:

$$P_{req,b} = \frac{n_d \cdot E_b}{\Delta t} \quad (4)$$

where Δt is the designated charging window duration. Increasing n_d or reducing Δt both raise the required power, while daily overnight charging minimizes it. The total fleet charging load (P_{fleet}) is obtained by summing $P_{req,b}$ over all buses in the fleet. The net daily CO_2 reduction achieved by electrifying the shuttle fleet is as follows:

$$e_{\text{CO}_2} = \sum_{b \in B} (d_b \cdot \text{EF}_{\text{diesel}} - E_b \cdot \text{EF}_{\text{grid}}) \quad (5)$$

where the first term represents the direct tailpipe emissions avoided by displacing the diesel fleet, with $\text{EF}_{\text{diesel}}$ denoting the diesel emission factor (kg CO_2/km), and the second term represents the indirect emissions introduced by grid-based charging of the electric fleet, with EF_{grid} denoting the grid emission factor (kg CO_2/kWh).

III. INPUT AND SIMULATION DATA

To establish a realistic basis for the electric bus network analysis, real operational data from shuttle movements within the IMSIU campus are collected and analyzed. Owing to ongoing construction activities and infrastructure adjustments across the campus, only the Yellow and Blue shuttle routes are presently active. Accordingly, the analysis in this study is limited to these two operational routes. The IMSIU academic calendar comprises 187 operational days per year, which is used to scale daily metrics to annual values throughout the analysis. Figure 1 illustrates the Yellow Route, which follows a 1.6 km loop with five designated stops primarily serving the internal academic blocks. Meanwhile, the Blue Route spans a longer 2.8 km path with eight stops, linking external parking areas to the central

campus, as shown in Figure 2. This adjusted configuration reflects the university's effort to sustain mobility efficiency while minimizing disruption during ongoing development works. Table I summarizes the operational characteristics of the Yellow and Blue routes. These parameters form the baseline operational profile of the current fleet and provide the essential foundation for energy-use estimation, emission calculation, and electric-fleet simulation. The operational parameters in Table I, including trips per day and daily distance, are reported on a per-bus basis. In the IMSIU shuttle system, one bus is assigned to the Yellow Route, while three buses are assigned to the Blue Route. To evaluate the environmental performance of the existing diesel fleet, annual CO_2 emissions are calculated for each route using the emission factor of 1.4 kg CO_2/km [2]. The grid emission factor $\text{EF}_{\text{grid}} = 0.55$ kg CO_2/kWh is adopted for Saudi Arabia based on recent ESG performance data from the Saudi Electricity Company [11]. To support the electrification of the current shuttle system, this study adopts two representative electric bus models, 8.5 m and 12 m configurations, selected for their suitability for campus-scale mobility operations. The standard battery capacities for the 8.5-m and 12-m electric buses are selected to be representative of available E-bus models reported in manufacturer datasheets and product brochures for similar vehicle classes. The average market price of traction battery packs for electric buses is provided in [11]. Together, these battery ratings and cost values serve as key input parameters for the battery sizing and economic analysis conducted in this study. The energy consumption rates used in this study are 0.695 kWh/km for the 8.5-m bus and 1.15 kWh/km for the 12-m bus. It is worth noting that the selection of the 8.5-m and 12-m electric

TABLE I OPERATIONAL CHARACTERISTICS

Parameter	Yellow Route	Blue Route
Route distance (km)	1.6	2.8
Number of buses	1	3
Fleet Metrics:		
Total Trips per day	42	84
Total Daily distance (km)	67.2	235.2
Per-Bus Metrics:		
Trips per day	42	28
Daily distance (km)	67.2	78.4

bus models for each route was intended to replicate the existing diesel bus configuration currently deployed on the respective routes. This like-for-like replacement approach ensures that the electric buses can accommodate passenger demand equivalent to that of the diesel fleet, thereby eliminating the need to modify service frequency or route capacity to maintain the current level of service.

IV. RESULTS AND DISCUSSIONS

This section presents the quantitative evaluation of the proposed framework, including emissions analysis, battery sizing outcomes, charging load behavior, and fleet-level energy implications.

A. Environmental Impact of Existing Shuttle Operations

The estimated annual CO_2 emissions for the Yellow Route, Blue Route, and the total diesel shuttle fleet are presented in Figure 3. The results clearly illustrate the disproportionate contribution of the Blue Route to overall emissions. Although the Yellow Route emits approximately 17.59 tons CO_2 /year, the Blue Route alone generates 61.58 tons CO_2 /year, accounting for nearly 78% of the fleet's total emissions. This difference is attributed to the significantly longer route length and higher trip frequency of the Blue Route, which results in a substantially higher annual distance traveled. The combined fleet emissions

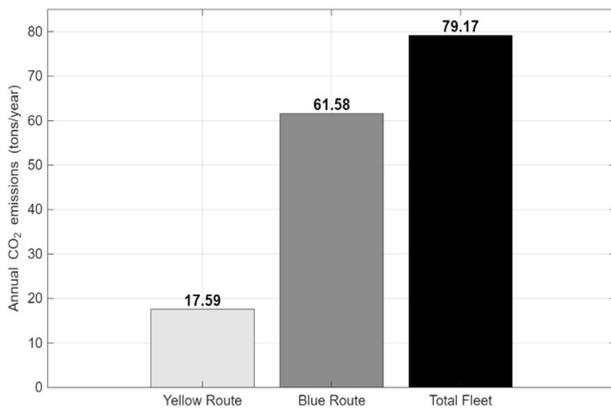


Figure 3 Annual CO_2 emissions by route at IMSIU Campus.

amount to 79.17 tons CO_2 /year, highlighting the measurable environmental footprint of the current diesel-based shuttle operation. These findings reinforce the importance of route-specific electrification planning; electrifying the Blue Route yields the largest environmental benefit due to its dominant share of emissions, while the Yellow Route provides additional but comparatively smaller reductions. It is to be noted that transitioning both routes to electric buses would eliminate the entirety of the 79.17 tons CO_2 /year currently produced, delivering immediate and quantifiable environmental gains while improving air quality across the campus mobility network.

B. Well-to-Wheel Charging Emissions

While the previous section established that the diesel fleet emits 79.17 tons of CO_2 annually, the net climate benefit of electric buses depends on upstream emissions associated

with electricity generation. Therefore, a consistent comparison of diesel-versus-electric should distinguish diesel tailpipe emissions from electric bus charging-related emissions. The Well-to-Wheel analysis is summarized in Table II. The 12-m fleet requires significantly higher charging energy (65,033 kWh) compared to the 8.5-m fleet (39,302 kWh), resulting in higher upstream emissions. Despite the grid's carbon intensity, transitioning to electric buses yields a net emission reduction (e_{CO_2}) of 43.40 to 57.55 tons annually compared to the diesel baseline. These results show that electrification delivers substan-

TABLE II WELL-TO-WHEEL EMISSION ANALYSIS: GRID-CHARGED ELECTRIC FLEET VS. DIESEL BASELINE

Bus Model	Annual Energy (kWh)	Grid Emissions (Tons)	Diesel Baseline (Tons)	Net Reduction (Tons)
8.5-m	39,302	21.62	79.17	57.55
12-m	65,033	35.77	79.17	43.40

tial emissions reduction at IMSIU. However, the magnitude of the benefit is sensitive to EF_{grid} , motivating the need for cleaner charging pathways for further decarbonization.

C. Electrification Readiness Through Route-Specific Battery Sizing

Tables III and IV present the battery sizing requirements for electrifying the shuttle fleet on the Blue Route and Yellow Route, respectively, using two representative e-bus models (8.5-m and 12-m). The results illustrate how route length, daily distance, and charging frequency significantly influence the required energy demand, nominal battery capacity, DoD levels, and corresponding economic implications. It is clear that increasing charging frequency significantly reduces the required battery capacity for both bus types. Daily charging yields the smallest battery sizes, with DoD levels remaining within 46–54% for the 8.5-m buses and 45–77% for the 12-m buses, thereby maintaining operation below the allowable DoD_{max} and supporting improved battery longevity. In contrast, reducing charging frequency substantially increases the required nominal capacity. For instance, on the Blue Route (Table III), the weekly-charged 8.5-m bus requires 340.55 kWh compared with only 68.11 kWh under daily charging, while the 12-m bus increases to 563.5 kWh. This capacity growth directly translates into higher capital costs, with weekly charging adding approximately SAR 185,850 and SAR 315,000 for the 8.5-m and 12-m buses, respectively. Similar trends are observed for the Yellow Route (Table IV), although the absolute capacities are lower due to the shorter daily distance. For instance, the 8.5-m bus requires 58.38 kWh under daily charging and 291.9 kWh under weekly charging. Across both routes, lower charging frequency pushes DoD values toward the design limit of $DoD_{max} = 80\%$, in some cases reaching approximately 75–79%, which reduces the available margin for long-term battery aging. Notably, the Yellow Route 12-m daily case already operates at 77.28% DoD due to the limited sizing buffer between the nominal requirement (96.6 kWh) and the available standard battery (100 kWh). It can be concluded that daily charging combined with adequate

TABLE III BATTERY SIZING ANALYSIS FOR 8.5-M AND 12-M E-BUS FLEETS (BLUE ROUTE)

Charging Frequency	Req. Energy (kWh)	Nominal (kWh)	Standard (kWh)	DoD (%)	Est. Cost (SAR)
8.5-m e-bus					
Daily	54.49	68.11	100.5	54.22	52,763
3×/Week	108.98	136.22	141	77.29	74,025
2×/Week	163.46	204.33	215	76.03	112,875
1×/Week	272.44	340.55	354	76.96	185,850
12-m e-bus					
Daily	90.16	112.7	200	45.08	105,000
3×/Week	180.32	225.4	240	75.13	126,000
2×/Week	270.48	338.1	350	77.28	183,750
1×/Week	450.80	563.5	600	75.13	315,000

TABLE IV BATTERY SIZING ANALYSIS FOR 8.5-M AND 12-M E-BUS FLEETS (YELLOW ROUTE)

Charging Frequency	Req. Energy (kWh)	Nominal (kWh)	Standard (kWh)	DoD (%)	Est. Cost (SAR)
8.5-m e-bus					
Daily	46.70	58.38	100.5	46.47	52,763
3×/Week	93.41	116.76	132	70.76	69,300
2×/Week	140.11	175.14	187	74.93	98,175
1×/Week	233.52	291.90	315	74.13	165,375
12-m e-bus					
Daily	77.28	96.6	100	77.28	52,500
3×/Week	154.56	193.2	200	77.28	105,000
2×/Week	231.84	289.8	295	78.59	154,875
1×/Week	386.40	483.0	486	79.51	255,150

battery sizing provides the most favorable balance between capital cost and degradation resilience. While deeper DoD enables smaller initial battery capacity, sustained operation near the upper DoD limit may accelerate capacity fade and require earlier battery replacement. The presented analytical sizing therefore represents a minimum feasible design, and practical deployments may incorporate additional capacity margins to maintain operational robustness over the battery lifetime.

D. Grid-Aware Charging Load Analysis for Campus Electrification

The relationship between charging duration, charging frequency, and the resulting electrical load on the campus grid is critical for infrastructure sizing. Based on the daily trip frequency derived in Table I, the fleet operates for approximately 10.5 hours daily. This leaves a consistent 13.5-hour overnight dwell time available for charging. Consequently, the 12-hour charging duration used in this analysis represents a realistic operational baseline that maximizes grid-load reduction without disrupting the daily schedule. Figures 4 and 5 illustrate the relationship between charging duration, charging frequency, and the resulting electrical load for the 8.5-m and 12-m electric bus fleets. It is noted that increasing the charging duration significantly reduces the load imposed on the electrical infrastructure for both bus types. For example, under weekly charging, the 12-m bus fleet imposes a peak load exceeding 430 kW for a 4-hour charging window but drops to below 150 kW when the charging duration is extended to 12 hours. A similar

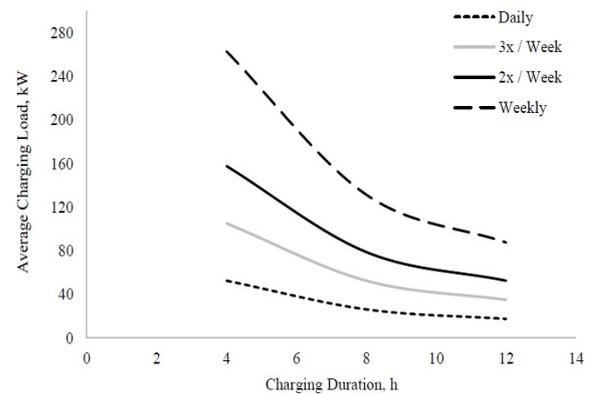


Figure 4 Fleet charging load for the IMSIU shuttle fleet modeled using the 8.5-m e-bus configuration.

trend is observed for the 8.5-m buses, where the weekly charging load declines from approximately 260 kW at 4 hours to about 90 kW at 12 hours. While the results indicate that the 12-m buses exhibit higher energy consumption and charging demand than the 8.5-m buses, the selection of bus length is primarily driven by passenger capacity requirements rather than energy considerations alone. The 8.5-m buses are suitable for routes and periods characterized by moderate passenger demand, such as internal campus circulation and off-peak operation, whereas the 12-m buses are required to accommodate higher passenger volumes, particularly on routes linking external parking areas to

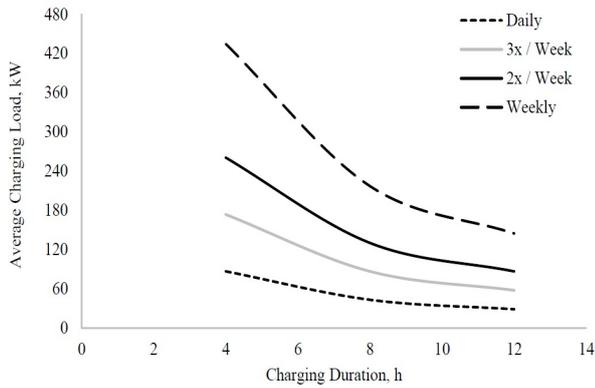


Figure 5 Fleet charging load for the IMSIU shuttle fleet modeled using the 12-m e-bus configuration.

the central campus during peak hours. Consequently, the higher energy requirements of the 12-m buses represent an operational trade-off between capacity provision and energy efficiency. The mixed deployment of both bus types in the IMSIU shuttle system therefore provides a balanced strategy that satisfies passenger demand while containing overall energy consumption and charging infrastructure requirements. Although the charging durations in Figures 4 and 5 are limited to a maximum of 12 h to represent a conservative and consistent overnight charging window, the results suggest that extending the charging interval during weekend idle periods could further reduce the required charging power without increasing total energy demand. Given that the shuttle buses operate for approximately 10.5 h per day, weekend dwell times may permit charging durations exceeding 12 h, thereby proportionally lowering the average charging load and reducing stress on both the charging infrastructure and upstream electrical components. Nevertheless, the 12-h charging window is retained in this study to ensure a uniform basis for comparing charging strategies and to represent a conservative weekday operating condition rather than relying on extended weekend availability.

E. Fleet-Level Energy Demand Modeling Under Diverse Charging Strategies

The charging energy per event for the full IMSIU shuttle fleet (one Yellow Route bus and three Blue Route buses) is shown in Figure 6. The 12-m e-bus consistently requires substantially more energy than the 8.5-m model because of its larger battery pack and higher energy consumption rate (1.15 kWh/km vs. 0.695 kWh/km). Under weekly charging, the 12-m bus requires approximately 1,740 kWh, compared with 1,050 kWh for the 8.5-m bus. Even under daily charging, the 12-m bus requires roughly 350 kWh, nearly double the 210 kWh needed by the 8.5-m model. These differences emphasize that bus size has a direct and substantial impact on charging infrastructure requirements, influencing transformer sizing, cable ratings, and power scheduling strategies. The findings clearly demonstrate the operational trade-offs between charging frequency, electrical load, and energy demand. Daily charging minimizes peak load,

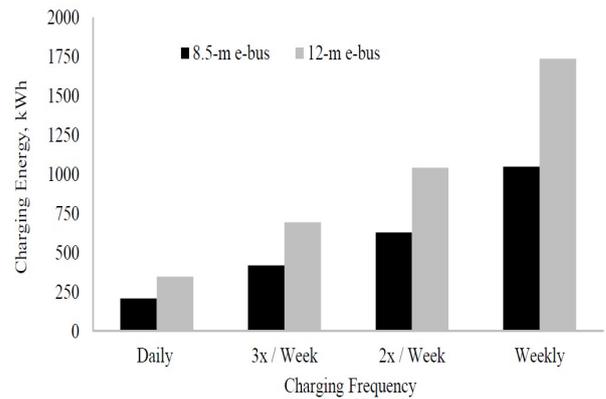


Figure 6 Charging energy per charging event for the 8.5-m and 12-m e-bus models for the full IMSIU shuttle fleet under different charging frequencies.

requires the smallest battery capacities, and imposes the least stress on the electrical infrastructure, making it the most grid-friendly strategy. In contrast, weekly charging results in the largest required battery capacities and the highest peak loads under standard overnight charging. However, when extended weekend idle charging is available, the peak load can be significantly reduced due to the longer charging window. The results further show that the larger 12-m e-buses consistently require more energy and generate higher charging loads than the smaller 8.5-m e-buses, emphasizing the need for careful consideration of vehicle size when planning charging infrastructure. These insights highlight the importance of selecting an appropriate charging strategy to ensure a stable, reliable, and cost-effective electrified shuttle network for IMSIU.

V. CONCLUSIONS

This paper presented a metro-integrated electric shuttle framework to enhance first/last-mile connectivity and support the transition toward zero-emission campus mobility at IMSIU. Using real operational data from the active Yellow and Blue routes, the analysis showed that electrifying the existing diesel fleet could eliminate 79.17 tons of direct tailpipe CO_2 annually while removing on-campus pollutants and operational noise, thereby improving localized air quality and commuter exposure. Battery sizing results for the 8.5-m and 12-m e-bus models confirmed that charging frequency and route characteristics strongly influence operational feasibility, capital cost, and depth-of-discharge levels. Daily charging emerged as the most efficient strategy, minimizing peak electrical demand and required battery capacity, whereas weekly charging significantly increased capacity requirements and peak power under standard overnight charging, potentially exceeding existing grid limits. The charging-load assessment further demonstrated that longer charging durations reduce average power demand and that larger 12-m buses impose consistently higher energy and load requirements, underscoring the importance of accurate fleet-level planning. The results establish a technically credible pathway for electrifying campus mobility while maintaining service reliability and supporting metro-integrated operations.

REFERENCES

- [1] R. A. Alarrouqi, F. Ahmad, S. Bayhan, and L. Al-Fagih, "Electric buses in hot climates: Challenges, technologies, and the road ahead," *IEEE Access*, vol. 12, pp. 55 531–55 550, 2024.
- [2] C. Zhao, L. Z. Kobayashi, A. B. S. Alqaity, J.-C. Monfort, E. Cenker, N. Miralles, and S. M. Sarathy, "Solutions for decarbonising urban bus transport: a life cycle case study in saudi arabia," *Communications Engineering*, vol. 3, no. 1, p. 95, 2024.
- [3] P. Li, M. Jiang, Y. Zhang, and Y. Zhang, "Cooperative optimization of bus service and charging schedules for a fast-charging battery electric bus network," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 5, pp. 5362–5375, 2023.
- [4] B. Zeng, W. Wu, and C. Ma, "Electric bus scheduling and charging infrastructure planning considering bus replacement strategies at charging stations," *IEEE Access*, vol. 11, pp. 125 328–125 345, 2023.
- [5] J. Chen and K. Strunz, "Optimal electric bus charging and battery swapping with renewable energy and frequency control ancillary service through aggregator," *IEEE Transactions on Transportation Electrification*, vol. 11, no. 1, pp. 3715–3729, 2025.
- [6] Y. Liu and H. Liang, "A data-driven approach for electric bus energy consumption estimation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 10, pp. 17 027–17 038, 2022.
- [7] M. R. Rusli, E. N. Sholikhah, L. Aminulloh, E. H. Binugroho, R. S. Dewanto, and D. Pramadihanto, "Battery sizing calculation and consideration for the electric apron bus system," in *2023 6th International Conference on Information and Communications Technology (ICOIACT)*. IEEE, 2023, pp. 11–16.
- [8] Y. Iscanoglu and M. Yilmaz, "Cost-effective electrification of city brt systems: Optimizing battery and charging infrastructure for sustainable urban mobility," *IEEE Access*, 2024.
- [9] Z. Chen, H. Wang, W. J. Barry, and M. J. Tuozolo, "Multi-criteria decision-making of hybrid energy infrastructure for fuel cell and battery electric buses," *Energies*, vol. 18, no. 11, p. 2829, 2025.
- [10] C. G. Colombo, S. Miraftabzade, M. Aimar, D. Zaninelli, M. Longo, and W. Yaici, "A comprehensive study on electrification of old bus fleets: a real case study in ottawa," in *2022 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM)*. IEEE, 2022, pp. 844–849.
- [11] K. Knehr, J. Kubal, and S. Ahmed, "Cost analysis and projections for us-manufactured automotive lithium-ion batteries," Argonne National Laboratory (ANL), Argonne, IL (United States), Tech. Rep., 2024.