



ADVCOMP 2025

The Nineteenth International Conference on Advanced Engineering Computing
and Applications in Sciences

ISBN: 978-1-68558-289-0

September 28th - October 2nd, 2025

Lisbon, Portugal

ADVCOMP 2025 Editors

Andreas Rausch, TU Clausthal, Clausthal-Zellerfeld, Germany

ADVCOMP 2025

Forward

The Nineteenth International Conference on Advanced Engineering Computing and Applications in Sciences (ADVCOMP 2025), held between September 28th, 2025, and October 2nd, 2025, in Lisbon, Portugal, continued a series of international events meant to bring together researchers from the academia and practitioners from the industry in order to address fundamentals of advanced scientific computing and specific mechanisms and algorithms for particular sciences.

With the advent of high-performance computing environments, virtualization, distributed and parallel computing, as well as the increasing memory, storage and computational power, processing particularly complex scientific applications and voluminous data is more affordable. With the current computing software, hardware and distributed platforms, effective use of advanced computing techniques is more achievable.

The conference provided a forum where researchers were able to present recent research results and new research problems and directions related to them. The conference sought contributions presenting novel research in all aspects of new scientific methods for computing and hybrid methods for computing optimization, as well as advanced algorithms and computational procedures, software and hardware solutions dealing with specific domains of science.

We take here the opportunity to warmly thank all the members of the ADVCOMP 2025 technical program committee, as well as all the 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 effort to contribute to ADVCOMP 2025. We truly believe that, thanks to all these efforts, the final conference program consisted of top-quality contributions. We also thank the members of the ADVCOMP 2025 organizing committee for their help in handling the logistics of this event.

We hope that ADVCOMP 2025 was a successful international forum for the exchange of ideas and results between academia and industry for the promotion of progress related to advanced engineering computing and applications in sciences.

ADVCOMP 2025 Chairs

ADVCOMP 2025 Steering Committee

Dean Vucinic, Vrije Universiteit Brussel (VUB), Belgium, FERIT, Croatia

Juha Röning, University of Oulu, Finland

Marcin Hojny, AGH University of Science and Technology, Poland

Andreas Rausch, TU Clausthal, Clausthal-Zellerfeld, Germany

Alice E. Koniges, University of Hawai'i at Mānoa, USA

ADVCOMP 2025 Publicity Chairs

Laura Garcia, Universidad Politécnica de Cartagena, Spain

Lorena Parra Boronat, Universidad Politécnica de Madrid, Spain

ADVCOMP 2025 Committee

ADVCOMP 2025 Steering Committee

Dean Vucinic, Vrije Universiteit Brussel (VUB), Belgium, FERIT, Croatia
Juha Rönning, University of Oulu, Finland
Marcin Hojny, AGH University of Science and Technology, Poland
Andreas Rausch, TU Clausthal, Clausthal-Zellerfeld, Germany
Alice E. Koniges, University of Hawai'i at Mānoa, USA

ADVCOMP 2025 Publicity Chairs

Laura Garcia, Universidad Politécnica de Cartagena, Spain
Lorena Parra Boronat, Universidad Politécnica de Madrid, Spain

ADVCOMP 2025 Technical Program Committee

Waleed H. Abdulla, University of Auckland, New Zealand
José Abellán, Catholic University of Murcia, Spain
Mohamed Riduan Abid, Alakhawayn University, Morocco
Rashmi Agrawal, Manav Rachna International Institute of Research and Studies, India
Francisco Airton Silva, Federal University of Piauí, Brazil
M. Azeem Akbar, Nanjing University of Aeronautics and Astronautics, China
Haifa Alharthi, Saudi Electronic University, Saudi Arabia
Sónia Maria Almeida da Luz, Polytechnic Institute of Leiria - School of Technology and Management, Portugal
Madyan Alsenwi, Kyung Hee University, Global Campus, South Korea
Mohamed E. Aly, California State Polytechnic University, Pomona, USA
Daniel Andresen, Kansas State University, USA
Anindya Das Antar, University of Michigan, USA
Abhinav Arora, Meta Platforms, USA
Ehsan Atoofian, Lakehead University, Canada
Vadim Azhmyakov, Universidad Central, Bogota, Republic of Colombia
Carlos Becker Westphall, University of Santa Catarina, Brazil
Raoudha Ben Djemaa, ISITCOM | University of Sousse, Tunisia
Peter Bentley, University College London, UK
Alessandro Borri, CNR-IASI Biomathematics Laboratory, Rome, Italy
David Bouck-Standen, Kingsbridge Research Center, UK
Sofiane Bououden, University Abbes Laghrour Khenchela, Algeria
Dominique Fabio Briechle, ISSE - Institute for Software and Systems Engineering of the Clausthal University of Technology, Germany
Xiao-Chuan Cai, University of Colorado Boulder, USA
Patricia Camacho Magriñán, Universidad de Cádiz, Spain
Jadson Castro Gertrudes, Federal University of Ouro Preto, Brazil
Graziana Cavone, Polytechnic of Bari, Italy

Sébastien Cayrols, University of Tennessee Knoxville, USA
Mete Celik, Erciyes University, Turkey
Jieyang Chen, Oak Ridge National Laboratory, USA
Jinyuan Chen, Louisiana Tech University, USA
Luca Davoli, University of Parma, Italy
Vassilios V. Dimakopoulos, University of Ioannina, Greece
Inês Domingues, IPO Porto Research Centre (CI-IPOP), Portugal
Maha Elarbi, University of Tunis, Tunisia
Javier Fabra, Universidad de Zaragoza, Spain
Akemi Galvez, University of Cantabria, Spain / Toho University, Japan
Tong Geng, Boston University, USA
Jing Gong, KTH Royal Institute of Technology, Sweden
Teofilo Gonzalez, UC Santa Barbara, USA
Maki Habib, American University in Cairo, Egypt
Yang He, University of Technology Sydney, Australia
Mohd Helmy Abd Wahab, Universiti Tun Hussein Onn Malaysia, Malaysia
Marcin Hojny, AGH University of Science and Technology, Poland
Wladyslaw Homenda, Warsaw University of Technology, Poland
Tzung-Pei Hong, National University of Kaohsiung, Taiwan
Mehdi Hosseinzadeh, Washington University in St. Louis, USA
Paul Humphreys, Ulster University | Ulster University Business School, UK
Andres Iglesias, University of Cantabria, Spain / Toho University, Japan
Joanna Isabelle Olszewska, University of West Scotland, UK
Hiroshi Ishikawa, Tokyo Metropolitan University, Japan
Félix J. García Clemente, University of Murcia, Spain
Rishabh Joshi, Google Research - Brain Team, USA
Zaheer Khan, University of the West of England, UK
Alice E. Koniges, University of Hawai'i at Mānoa, USA
Sonia Lajmi, University of Sfax, Tunisia / Al Baha University, Saudi Arabia
Yahia Lebbah, University of Oran, Algeria
Seyong Lee, Oak Ridge National Laboratory, USA
Maurizio Leotta, University of Genova, Italy
Clement Leung, Chinese University of Hong Kong, Shenzhen, China
Yiu-Wing Leung, Hong Kong Baptist University, Hong Kong
Jianwen Li, East China Normal University, Shanghai, China
Yiheng Liang, Bridgewater State University, USA
Stephane Maag, Telecom Sud Paris, France
Elbert E. N. Macau, Federal University of Sao Paulo - UNIFESP at Sao Jose dos Campos, Brazil
Rafael Magdalena Benedicto, University of Valencia, Spain
Marcin Markowski, Wroclaw University of Science and Technology, Poland
Mirko Marras, University of Cagliari, Italy
René Meier, Hochschule Luzern, Switzerland
Yuan Meng, University of Southern California, USA
Mohamed Wiem Mkaouer, Rochester Institute of Technology, USA
Zewei Mo, University of Pittsburgh, USA
Sébastien Monnet, Savoie Mont Blanc University (USMB), France
Shana Moothedath, University of Washington, Seattle, USA
Kiran Nalla, Microsoft, USA

Laurent Nana, University of Brest, France
Ehsan Nekouei, City University of Hong Kong, Hong Kong
Kaiming Ouyang, Nvidia, USA
Marcin Paprzycki, Systems Research Institute | Polish Academy of Sciences, Poland
Prantosh Kumar Paul, Raiganj University, India
Biagio Peccerillo, University of Siena, Italy
Damien Pellier, Université Grenoble Alpes, France
Sonia Pérez-Díaz, University of Alcalá, Spain
Antonio Petitti, Institute of Intelligent Industrial Systems and Technologies for Advanced Manufacturing (STIIMA) - National Research Council of Italy (CNR) , Italy
Tamas Pflanzner, University of Szeged, Hungary
Agostino Poggi, Università degli Studi di Parma, Italy
Evgeny Pyshkin, University of Aizu, Japan
Andreas Rausch, Technische Universität Clausthal, Germany
Michele Roccotelli, Politecnico di Bari, Italy
Ivan Rodero, Rutgers University, USA
Juha Röning, University of Oulu, Finland
Diego P. Ruiz, University of Granada, Spain
Bibhudatta Sahoo, National Institute of Technology, Rourkela, India
Julio Sahuquillo, Universitat Politècnica de València, Spain
Subhash Saini, NASA, USA
Shailaja Sampat, Arizona State University, USA
Hamed Sarvari, George Mason University, USA
Alireza Shahrabi, Glasgow Caledonian University, Scotland, UK
Justin Shi, Temple University, USA
Piotr Sowiński, Systems Research Institute, Polish Academy of Sciences, Poland
Sudarshan Srinivasan, Oak Ridge National Laboratory, USA
Mohammed Tanash, Kansas State University, USA
Costas Vassilakis, University of the Peloponnese, Greece
Bhavan Vasu, Oregon State University, USA
Flavien Vernier, LISTIC – Savoie University, France
Juan Vicente Capella Hernández, Universitat Politècnica de València, Spain
Dean Vucinic, Vrije Universiteit Brussel (VUB), Belgium / FERIT, Croatia
Guangjing Wang, Michigan State University, USA
Hanrui Wang, Massachusetts Institute of Technology, USA
Lei Wang, University of Connecticut, USA
Adriano V. Werhli, Universidade Federal do Rio Grande - FURG, Brazil
Gabriel Wittum, Goethe University Frankfurt, Germany
Zongshen Wu, University of Wisconsin, Madison, USA
Mudasser F. Wyne, National University, USA
Chenhao Xie, Beihang University, Beijing, China
Cong-Cong Xing, Nicholls State University, USA
Feng Yan, University of Nevada, Reno, USA
Limin Yang, University of Illinois at Urbana-Champaign, USA
Jinsongdi Yu, Fuzhou University, China
Carolina Yukari Veludo Watanabe, Federal University of Rondônia, Brazil
Michael Zapf, Technische Hochschule Nürnberg Georg Simon Ohm (University of Applied Sciences Nuremberg), Germany

Vesna Zeljkovic, Lincoln University, USA
Ruochen Zeng, NXP Semiconductors, USA
Penghui Zhang, Arizona State University, USA
Pengmiao Zhang, University of Southern California, USA
Qian Zhang, Liverpool John Moores University, UK

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

Orientation Prediction in Robotics: A Study of Trigonometric Decomposition Methods Across Synthetic and Real-World Datasets <i>Antonio Gambale, Sonya Coleman, Dermot Kerr, Philip Vance, Emmett Kerr, Cornelia Fermuller, and Yiannis Aloimonos</i>	1
Investigating the Adaptability of ALE-AMR Hydrocode for Darcy Flow and Geothermal Simulations <i>Alice Koniges, David Eder, Jonghyun Lee, Jiawei Shen, Aaron Fisher, and Tzanio Kolev</i>	8
Modeling the Interaction of Laser-Produced Proton Beams with Matter <i>Jack McKee, David Eder, Aaron Fisher, Alice Koniges, Claudia Parisua, Maxence Gauthier, Elizabeth McBride, Frank Seiboth, Chandra Breanne Curry, Mungo Frost, Eric Galtier, and Siegfried Glenzer</i>	15
Artificial Intelligence-Based Local Weather Forecasting for Agricultural Digital Twins <i>Miguel Zaragoza-Esquerdo, Alberto Ivars-Palomares, Irene Eiros-Fonseca, Sandra Sendra, and Jaime Lloret</i>	20

Orientation Prediction in Robotics: A Study of Trigonometric Decomposition Methods Across Synthetic and Real-World Datasets

Antonio Gambale , Sonya Coleman , Dermot Kerr , Philip Vance 

School of Computing, Engineering & Intelligent Systems,

Ulster University, Londonderry, Northern Ireland

e-mail: {gambale-a | sa.coleman | d.kerr | p.vance}@ulster.ac.uk

Emmett Kerr 

Dept. of Electronic & Mechanical Engineering,
Atlantic Technological University, Letterkenny, Ireland

e-mail: emmett.kerr@atu.ie

Cornelia Fermüller , Yiannis Aloimonos 

Institute for Advanced Computer Studies,
University of Maryland, USA

e-mail: {fermulcm | yiannis}@umd.edu

Abstract—Orientation prediction is a critical task for robotics as it enables robots to understand and interact with their environment more effectively. By accurately determining an object's position and orientation, robots can perform a range of complex tasks. This in turn will advance smart manufacturing facilities to achieve higher levels of automation, increase efficiency, and enable more flexible production systems. Hence, we present a comparative study of shallow regression models, integration strategies, and trigonometric encoding schemes for planar orientation prediction in robotics, using synthetic and real-world datasets. Results demonstrate that XGBoost 1.7, combined with vector integration and quadrant encoding, achieves the best balance of accuracy, robustness to angular boundary discontinuities, and computational efficiency, significantly outperforming alternative approaches in real-world scenarios.

Keywords—Computer Vision; Robotics; Manipulation; Machine Learning; Smart Manufacturing.

I. INTRODUCTION

Autonomous grasping serves as the foundation for numerous robotic operations, yet current approaches frequently generate arbitrary grasp poses without considering the object's pose for subsequent manipulation requirements [1] [2]. Contemporary robotic grasping research typically focuses on two-finger parallel grippers [1] [2], using complex seven or five dimensional grasp representations [3] [4], often referred to as grasping rectangles [5]. Some works have explored approaches to simplify grasp parameters whilst maintaining effectiveness. Notably, the work in [6] demonstrated success in representing object orientation using a single planer angle, predicted through ImageNet Convolutional Neural Network (CNN) feature extraction with a Support Vector Regressor (SVR) for angle prediction. However, the work demonstrated in [6] relies on RGB-Depth (RGB-D) data and complex hierarchical regression techniques to handle angle discontinuity at the $0^\circ/360^\circ$ boundary. This is accomplished using a two-step process: first, the angle is classified into one of four 90° intervals; then, a separate SVR model is trained for each interval. While effective, this method introduces additional complexity compared with single stage direct angle regression approaches that utilise RGB data. Additionally, alternative

existing strategies present other notable drawbacks: methods that rely on high-dimensional grasp representations or hierarchical angle classification typically involve significant computational overhead, increased susceptibility to annotation errors, and limited generalisation when object pose varies beyond training data distributions. In contrast, approaches that require depth sensing or RGB-D inputs restrict applicability on platforms equipped only with RGB cameras. These constraints hinder practical deployment in dynamic or resource-constrained environments.

This research, therefore, aims to enhance post-grasp task execution through the use of a novel methodology which represents object orientation as a single angle in a 360° coordinate system. This approach provides orientation information that not only facilitates the intended use of the object but also enables derivation of the appropriate grasp angle. Additionally, this representation enables learning architectures to focus specifically on orientation prediction of one value (θ), circumventing the complexity of multiple parameter estimations and facilitating grasping with manipulators of all types.

To address the scarcity of annotated grasping datasets suitable for single angle object pose estimation, this work draws direct inspiration from [7], which utilised traditional computer vision with object geometric data to generate object poses. Here, a modified version of the methodology outlined in [7] is employed to augment the MetaGraspNet dataset: rather than relying on edge maps, ground-truth object segmentation masks are used to extract geometric features such as centroid, handle-tip displacement, and axis of symmetry, enabling the calculation of precise orientation annotations for each object instance. This automated annotation process was validated through manual inspection to ensure geometric fidelity and orientation accuracy, resulting in a high-quality synthetic dataset well-suited for model training. Recognising that synthetic datasets alone cannot capture the full complexity of real-world scenarios, a complementary real-world dataset was created and rigorously annotated. This secondary dataset was curated to reflect a range of object configurations, occlusions, and lighting conditions encountered in practical applications. The combination of

these two datasets (a large, systematically annotated synthetic set and a smaller, carefully validated real-world set) enables robust evaluation of model generalisation and domain transfer, providing a more comprehensive assessment of each method's practical viability. This research systematically evaluates the influence of regression architectures, integration strategies, and encoding schemes on orientation prediction performance. Both branched and vector-based integration approaches are assessed alongside various trigonometric encoding schemes, using synthetic and real-world data. Cross-validation, computational benchmarking, and error analysis are integrated to determine practical suitability, advancing the current state of planar orientation prediction for robotic manipulation. The study clarifies the interplay between encoding, integration, and model architecture, while also providing actionable insights for deploying robust, efficient orientation predictors in real-world robotic pipelines.

The remainder of this paper is organised as follows. Section II details the datasets used for training and evaluation. Section III describes the methodology, outlining the regression models, target encoding strategies, and integration approaches assessed in this work. Section IV presents the experimental results, including performance metrics and computational benchmarking. Section V provides a comprehensive discussion and evaluation of the findings, with a focus on model robustness, error distributions, and practical implications. Finally, Section VI concludes the paper and outlines future work.

II. DATASETS

Training Dataset MetaGraspNet: The following work employs two distinct datasets: the first is used for training and validating the proposed methods, while the second serves to evaluate performance under conditions more closely aligned with real-world applications. The initial training of the proposed methods utilised a carefully curated subset of the MetaGraspNet dataset [8] reduced to focus exclusively on Phillips and flat-head screwdrivers. The subset selected for this work corresponds to the single-class, multiple-instance configuration within the MetaGraspNet framework. The choice of screwdrivers as the focal object class is deliberate as these objects possess geometric properties, which satisfy the requirements outlined in [7]: screwdrivers possess an axis of symmetry and a mass distribution that is biased to one side perpendicular to this axis. These characteristics permit the application of the approach described in [7] to compute the centroid (derived from the segmentation mask), the distances from the centroid to the screwdriver's handle and tip, and hence the identification of the direction based on the displacement of the centre of mass relative to the tip. This directional information is then recorded as an angular annotation, ranging from 0° to 359° and provides the angular information required to use this dataset to train machine learning models such as those employed in this paper.

The resulting curated dataset provides 7,932 annotations across 2,691 images of size 1200×1200 pixels [9]. Initial evaluation of the dataset revealed that orientation calculations for occluded objects with areas less than 10,000 pixels were

unreliable. Consequently, such instances were removed, reducing the dataset to 5,709 annotations. A 10% sample of these annotations was manually verified, revealing 55 screwdrivers with incorrectly detected angles (errors ranging from 1° to 9°), yielding a Mean Absolute Angular Error (MAAE) of 0.313° across the sample set. This inspection demonstrated that the majority of angular errors were minor and would not significantly compromise robotic grasping performance [10]. The final dataset thus comprises 4,567 training samples and 1,142 testing samples with associated images and annotations (see examples in Figure 1).



Figure 1. Example images from the MetaGraspNet training dataset.

Secondary Evaluation Dataset: A secondary real-world dataset, containing 81 annotations, was constructed to evaluate the trained models using real-world RGB images of screwdrivers. This dataset mirrors the structure of the curated MetaGraspNet subset employed for training, closely adhering to the difficulty levels defined in [9] but uses real screwdrivers rather than synthetic data. The MetaGraspNet difficulty levels are designed to progressively increase in complexity: Levels 1 and 2 represent single objects with no occlusion, while higher levels introduce multiple objects and increasing degrees of occlusion and clutter. In line with this, the custom real-world evaluation dataset contains 8 images at difficulty level 1 or 2 (single screwdrivers, no occlusion), 5 images at difficulty level 4 (moderate clutter and occlusion), and 14 images at difficulty level 5 (high clutter and occlusion). Example images from the evaluation set are shown in Figure 2. No images from difficulty level 3 were included, consistent with the original subset selection. Table I summarises the dataset details.

The secondary evaluation dataset was acquired using a camera equipped with a Samsung ISOCELL GN9 sensor, capturing images of size 3072×4080 pixels, with a lens aperture of $f/1.9$, an exposure time of $1/100$ second, and an

International Organisation for Standardisation (ISO) sensitivity of 386, with images being collected from both overhead and 45° angled perspectives. The dataset features three screwdriver variants, red/black Torx, blue/black flathead/starhead, and solid black starhead, randomly arranged on a white backdrop under uncontrolled ambient lighting conditions. To ensure natural illumination variability, 27 images were taken across daylight hours, resulting in 81 annotations.



Figure 2. Example images from the real-world evaluation dataset.

Ground truth segmentation masks and orientation angles were annotated by three domain experts [11]. Inter-annotator agreement on segmentation masks was assessed using mean Intersection-over-Union (IoU), where IoU is calculated as the area of overlap between two masks divided by the area encompassed by either mask, providing a measure of agreement for object segmentation; this analysis yielded a high consistency of 0.95. Discrepancies were resolved via consensus review to refine the segmentation boundaries. For orientation, annotators measured shaft angles against the right horizontal axis, with a mean absolute difference of $\pm 1.8^\circ$ between predictions. Disagreements exceeding 5° were resolved by taking the circular mean for the affected objects; otherwise, the final orientation annotations were selected from a random sample

TABLE I. SUMMARY OF IMAGES AND ANNOTATIONS BY DIFFICULTY LEVEL

Level	No. of Images	No. of Annotations
1-2	8	8
4	5	17
5	14	56

of the independent measurements, except where consensus averaging was required as described.

For inference in the custom real-world dataset, segmentation masks were generated using the Segment Anything Model (SAM) [12]. The accuracy of these masks was validated against the manual annotations, achieving a mean IoU of 0.95. This strong agreement supported the use of SAM-generated masks for subsequent stages in the pipeline.

III. METHODOLOGY

This work systematically evaluates planar orientation prediction using three regression architectures, two integration strategies, and four target encoding schemes. The methodology extends trigonometric decomposition to address circular data challenges and combines encoding scheme analysis with shallow learning models. Each combination's performance is assessed on both a MetaGraspNet subset and a custom real-world dataset.

A. Target Encoding Schemes

Four encoding configurations were designed, in which the data are derived from ground truth angles to form the study:

- **Base:** Fundamental trigonometric components:

$$\mathbf{y}_{\text{base}} = [\sin(\theta), \cos(\theta)]$$

- **Quadrant:** Base + one-hot quadrant encoding (Q1-Q4):

$$\mathbf{y}_{\text{quad}} = [\sin(\theta), \cos(\theta), \mathbf{1}_{Q1}, \mathbf{1}_{Q2}, \mathbf{1}_{Q3}, \mathbf{1}_{Q4}]$$

- **Polar:** Base + radian displacement:

$$\mathbf{y}_{\text{polar}} = [\sin(\theta), \cos(\theta), \theta_{\text{rad}}]$$

- **Full:** Comprehensive representation:

$$\mathbf{y}_{\text{full}} = [\sin(\theta), \cos(\theta), \mathbf{1}_{Q1}, \mathbf{1}_{Q2}, \mathbf{1}_{Q3}, \mathbf{1}_{Q4}, \theta_{\text{rad}}]$$

These encodings represent target variables used to train various shallow learning models, as outlined in Table II. The selection of these models is motivated by specific technical requirements of the orientation prediction task and practical deployment constraints. Random Forest (RF) was chosen for its native multi-output capability and robustness to noisy features, particularly relevant given the variable quality of geometric features extracted from segmentation masks. SVR was selected for its strong generalisation via kernel methods, enabling effective handling of non-linear relationships between ResNet50 features and trigonometric targets; Multiple-Output Support Vector Regression (M-SVR) extends this capability to joint optimisation across all target variables. XGBoost 1.7 implementations were included due to their established performance in regression tasks involving high-dimensional feature spaces and their gradient boosting approach's ability to iteratively correct prediction errors, a property especially valuable for circular data where small angular errors can compound. The XGBoost 2.0 variant specifically addresses multi-output limitations present in the earlier versions (1.7). Deep learning approaches were deliberately excluded due to the relatively modest dataset size (5,709 training samples),

which favours shallow models that avoid overfitting, as well as the need for rapid inference in robotic applications where computational efficiency is paramount.

TABLE II. SHALLOW LEARNING MODELS TESTED

Model	Description
RF	Random Forest
SVR	Support Vector Regression
M-SVR	Multiple-Output Support Vector Regression
XGBoost 1.7	Standard XGBoost implementation
XGBoost 2.0	Multi-output enabled XGBoost variant

Along with the multiple target encoding schemes and different model architectures, two proposed model integration strategies were introduced. This refers to how each of the models handles the target variables required to perform the analysis of encoding schemes. The two proposed integration strategies are as follows:

B. Branched integration strategy

This integration strategy uses multi-output models split into branches which are trained in parallel, one predicting the sine component and the other predicting the cosine component. In each of these branches, the additional target variables required for the encoding scheme analysis remain, meaning the only change is that the branch simply does not receive its alternative sin/cosine pair, but it does receive all other target variables required for the encoding scheme testing. During testing and subsequent inference, the predicted sine and cosine values are then combined using the inverse tangent function (\tan^{-1}) to provide the predicted angle.

C. Vector integration strategy

The second proposed integration strategy leverages a vector-based approach, however, depending on the model's architecture, the meaning of this differs slightly. For tree-based approaches, such as decision trees and RF, multiple outputs are natively supported, allowing a single model to predict all target variables directly, allowing for a joint relationship between targets sharing their influence on the scorers, loss function and predictions [13].

However, for algorithms lacking native multi-output support (e.g., SVR or XGBoost 1.7), a wrapper-based framework is required [13]. This wrapper fits one independent regressor per target variable, however it differentiates itself from the branched approaches by using shared hyperparameters tuned globally across all targets. This allows for a more balanced performance optimisation across all target variables through the use of a custom scorer that can be changed but does not alter individual model training objectives [14].

As noted previously, wrapper-based approaches face limitations when applied to models lacking native multi-output regression capabilities. To address this constraint and enable comprehensive comparative analysis, two additional implementations were evaluated: XGBoost 2.0 [15] and a M-SVR [16]. These vector-based approaches introduce critical architectural enhancements over standard wrapper methods. XGBoost 2.0 implements multi-output trees, where leaf nodes contain vector

outputs spanning all targets simultaneously. This enables feature splits during tree construction to directly consider cross-target relationships, a capability absent when isolated models produced by Multi-Output Regressor wrappers. Concurrently, M-SVR extends traditional SVR by jointly optimising all targets through a unified ε -insensitive loss function, eliminating the hyperparameter compromises inherent in wrapper approaches. These additional models, when used with this integration strategy, should better enable implicit enforcement of trigonometric relationships (e.g., $\sin^2 \theta + \cos^2 \theta = 1$) more akin to the use of RF models.

D. Feature extraction

The pipeline (Figure 3) processes object segmentation masks (provided in MetaGraspNet annotations) to extract screwdriver patches. Each patch is placed on a white background, resized with aspect ratio preserved, and padded to 224×224 pixels. A ResNet50 CNN pre-trained on ImageNet, with classification head removed and global average pooling applied, extracts 2048 dimensional feature vectors [17]. These features serve as inputs to regression models trained on synthetic data from MetaGraspNet subset. For testing on real-world data, the same pipeline is applied using the custom dataset, with segmentation provided by SAM [12].

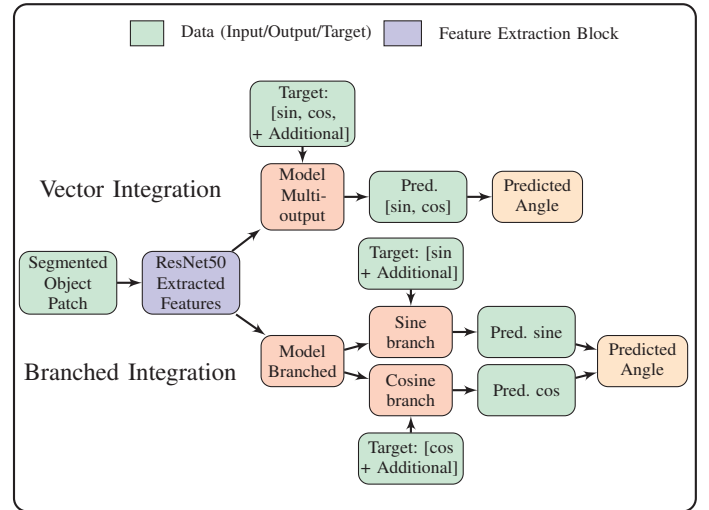


Figure 3. Object orientation prediction pipeline.

IV. RESULTS

To assess the predictive performance of the tested models, integration strategies, and encoding approaches, the MAAE was employed as the primary metric. This measure was computed independently for each encoding method and dataset, ensuring a granular evaluation of angular accuracy. Robustness was further enhanced through the application of 5-fold cross validation across all models and encoding scheme, providing a reliable estimate of performance.

A comprehensive computational benchmarking analysis was also conducted to evaluate inference speed. For each model, inference was performed on identical image features, with each test comprising 1000 repetitions per run and five aggregate

runs to account for variability. To ensure fairness, the sequence of model evaluation was randomised for each aggregate run, and system memory was cleared between measurements to minimise interference. All experiments were executed on a Linux 6.8.0 system equipped with Python 3.8.3, an Intel Xeon w7-3445 CPU (20 cores, 40 threads), and 62.3 GB of RAM. This rigorous protocol facilitates a thorough and equitable comparison of both predictive accuracy and computational efficiency across all models and integration strategies.

The following results begin with the MAAE measurements across the MetaGraspNet subset. The results compare different integration strategies (vector vs branched encoding) and encoding configurations (full target set vs reduced sets). All values represent degrees of angular error, with lower values indicating better performance. The initial results on the MetaGraspNet subset are displayed in Table III.

TABLE III. MAAE FOR METAGRASPNET DATASET (DEGREES)

Model	Base	Quadrant	Polar	Full
XGBoost 1.7 (Vector)	5.15	5.15	5.15	5.15
XGBoost 1.7 (Branched)	5.15	5.15	5.15	5.15
XGBoost 2 (Vector)	4.92	5.01	5.46	5.17
M-SVR (Vector)	8.04	8.04	8.04	8.04
SVR (Vector)	5.01	5.12	5.01	5.01
SVR (Branched)	5.01	5.12	5.01	5.01
RF (Vector)	5.48	5.08	5.87	5.67
RF (Branched)	7.43	5.44	5.93	5.88

TABLE IV. MAAE RESULTS FOR REAL-WORLD DATASET (DEGREES)

Model	Base	Quadrant	Polar	Full
XGBoost 1.7 (Vector)	9.61	8.15	8.96	9.61
XGBoost 1.7 (Branched)	9.61	11.66	9.15	11.14
XGBoost 2 (Vector)	17.09	13.91	7.18	10.46
M-SVR (Vector)	28.86	28.86	28.82	28.82
SVR (Vector)	23.13	23.54	23.14	28.03
SVR (Branched)	23.15	23.54	23.14	23.28
RF (Vector)	14.81	14.49	11.84	17.43
RF (Branched)	16.45	11.62	12.90	17.50

TABLE V. INFERENCE TIME PER IMAGE PATCH (MILLISECONDS)

Model Type	Base	Quadrant	Polar	Full
XGBoost 1.7 (Vector)	0.76	1.86	0.91	1.73
XGBoost 1.7 (Branched)	0.50	0.83	1.20	3.61
XGBoost 2 (Vector)	0.76	1.86	0.91	0.29
M-SVR (Vector)	17.76	17.78	17.80	17.76
SVR (Vector)	13.48	41.94	20.78	51.25
SVR (Branched)	13.55	70.23	27.68	75.49
RF (Vector)	59.27	56.50	57.90	45.42
RF (Branched)	119.15	117.08	117.83	91.04

Table IV summarises model performance on the real-world dataset using the same MAAE metric to enable direct comparison with the MetaGraspNet results. In addition to angular accuracy, Table V reports the inference time per

image (in milliseconds) for each model and condition. This enables an explicit comparison of computational efficiency across integration strategies and model types, with lower values reflecting faster processing.

V. DISCUSSION | EVALUATION

While the MAAE metrics presented in Section IV form the core of this discussion, it is essential to first address the primary challenge that this work seeks to overcome: the issue of boundary discontinuity. To thoroughly assess each method's performance, we conduct a granular analysis of the inference data on the custom dataset, as visualised in Figure 4. Figure 4 displays the prediction error as a function of the ground truth angle for each model and encoding configuration. Specifically, the x -axis represents the ground truth angle of the object, while the y -axis shows the signed angular error, calculated as the shortest difference between the predicted and actual angles, wrapped to the interval $[-180^\circ, 180^\circ]$. Each point corresponds to a single prediction, positioned horizontally by its ground truth angle and vertically by the deviation from the true value.

A detailed inspection of these plots reveals substantial differences in error distributions between models, which are not always reflected in the aggregate MAAE values. For instance, the M-SVR model exhibits pronounced and frequent error spikes at the 0° and 359° boundaries, indicating a persistent struggle with boundary discontinuity and a lack of robustness in these critical regions. These errors are not isolated; the M-SVR model demonstrates erratic behaviour across much of the angular range, with large, abrupt deviations that suggest poor generalisation and reliability. In contrast, XGBoost 1.7 (both vector and branched variants) stands out for its consistent and stable error profile. Across nearly the entire angular range, prediction errors remain tightly clustered around zero, with only occasional moderate spikes, mostly at angular boundaries. This stability is indicative of a model that not only achieves a low mean error, as seen in Table V, but also avoids catastrophic failures, making it more suitable for real-world deployment where reliability is paramount.

XGBoost 2, while achieving competitive MAAE values in some configurations, displays a more volatile error pattern. Notably, it exhibits significant errors not only at the 0° and 359° boundaries but also around 180° , suggesting that its generalisation may be compromised at multiple critical angles. This behaviour underscores the importance of evaluating models beyond mean metrics, as a low MAAE can mask underlying instability. The SVR models, both vector and branched, provide mixed results. While their errors are generally moderate, both models are prone to sporadic, large prediction failures at various angles, particularly near the boundaries and occasionally in the mid-range. These large spikes indicate that, although SVR may perform adequately on average, it is susceptible to unpredictable outliers that could undermine its practical utility. Random Forest models show moderate stability, with the vector variant generally outperforming the branched version in terms of error consistency. While occasional spikes are present, these are less frequent and less severe than those observed in M-SVR or

SVR, positioning Random Forest as a reasonable compromise between stability and accuracy, something reflected by its relatively low MAAE when paired with quadrant encoding and a branched integration strategy as seen in Table V.

It is important to note that all models were trained exclusively on a synthetic dataset which is inherently simpler and more controlled than the real-world test set. As a result, the MAAE values achieved using the synthetic MetaGraspNet subset (Table III) are substantially lower across all models, reflecting the training-domain familiarity. For example, XGBoost 1.7 achieves a MAAE of 5.15° using the synthetic data, compared with $8.15\text{--}9.61^\circ$ using the real-world dataset, while M-SVR and SVR also show marked increases in error when transitioning to real-world evaluation. This domain gap highlights the challenge of generalising from synthetic to real data, and underscores the value of robust error analysis using the real-world test set.

In addition to predictive accuracy, computational efficiency was also evaluated. Table V quantifies inference times per image patch, demonstrating XGBoost 1.7's superior performance: configurations predominantly achieve sub-2ms inference times (one outlier). In contrast, M-SVR/SVR models exhibit 10-100 \times slower performance (13.48-75.49ms), while RF demonstrates the poorest efficiency, consistently exceeding 45ms and frequently surpassing 100ms. This efficiency advantage positions XGBoost 1.7 as optimal for applications requiring both rapid inference and angular reliability.

Taken together, these results demonstrate that XGBoost 1.7 provides the most robust and reliable predictions across the full angular range using real-world data, effectively managing boundary discontinuities and avoiding large, erratic errors, while also offering leading computational efficiency. In contrast, models such as M-SVR and SVR are hindered by frequent and severe outliers, particularly at critical boundaries, and Random Forest occupies a middle ground, providing reasonable stability but not matching the overall reliability or speed of XGBoost 1.7. This comprehensive evaluation highlights the necessity of considering both aggregate metrics and detailed error distributions, as well as computational efficiency, when selecting models for applications where consistent and timely performance across all angles is essential.

A key focus of this work is the interplay between encoding strategies, model architectures, and integration approaches. Therefore, these results also reveal that model responsiveness to encoding varies considerably:

- **XGBoost 1.7** demonstrates strong robustness across all encodings, but achieves its best real-world performance with quadrant encoding in the vector integration configuration, yielding the lowest MAAE (8.15°). Notably, the addition of quadrant or polar information generally improves performance over the base (sin, cos) encoding, suggesting that XGBoost 1.7 is able to leverage richer target representations to better manage boundary effects and reduce systematic errors. The full encoding does not consistently outperform quadrant or polar, indicating that the full encoding may introduce redundancy or noise for this architecture.

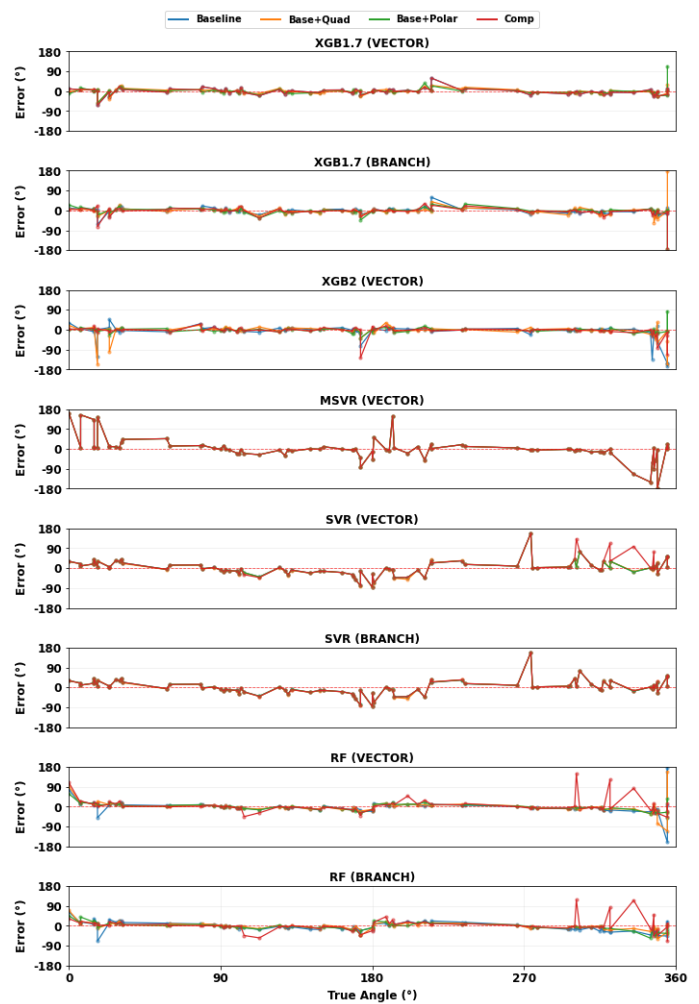


Figure 4. Model prediction error vs ground truth angle on real-world test set (missing data in 355° and 359° range).

- **XGBoost 2.0** is more sensitive to the encoding choice, achieving its best real-world MAAE (7.18°) with the polar encoding. However, its error distribution is less stable, with notable spikes at both boundary and mid-range angles, suggesting that, while certain encodings can lower mean error, they may not guarantee robust predictions.
- **SVR and M-SVR** models show limited benefit from more complex encodings. Both models exhibit high MAAE and frequent large errors regardless of encoding, indicating that their architectures are less capable of exploiting additional target information to improve generalisation, especially in the presence of boundary discontinuities.
- **Random Forest** benefits moderately from quadrant encoding, particularly in the branched configuration, but does not reach the accuracy or stability of XGBoost models. Its performance is more consistent than SVR/M-SVR but less robust to encoding changes than XGBoost 1.7.

The choice between branched and vector integration strategies also influences model performance:

- For **XGBoost 1.7**, the vector strategy paired with quadrant

encoding is optimal, while the branched approach is less effective, especially when combined with more complex encodings.

- **Random Forest** shows a preference for the branched strategy when used with quadrant encoding, achieving its lowest MAAE (11.62°), but remains slower and less accurate than XGBoost.
- **SVR** and **M-SVR** do not benefit significantly from either integration strategy, with both approaches yielding high error and erratic predictions.

In summary, these results demonstrate that XGBoost 1.7, with vector integration and quadrant encoding, offers the best trade-off between accuracy, robustness to boundary discontinuities, and computational efficiency using real-world data. While more complex encodings marginally improve mean error for certain models, the quadrant encoding strikes an effective balance between informativeness and generalisation. Models such as M-SVR and SVR are hindered by frequent and severe outliers and do not benefit meaningfully from richer encodings or alternative integration strategies. Random Forest provides moderate stability but cannot match the overall reliability or speed of XGBoost 1.7. These findings underscore the need to consider both aggregate metrics and error distributions, alongside computational efficiency, when choosing models for consistent, timely performance.

VI. CONCLUSION AND FUTURE WORK

This work systematically evaluates a range of shallow regression models, integration strategies, and target encoding schemes for the challenging task of planar orientation prediction using both synthetic and real-world datasets. The results demonstrate that among the tested configurations, **XGBoost 1.7**, when paired with the **vector integration strategy** and **quadrant encoding**, consistently delivers the best trade-off between predictive accuracy, robustness to angular boundary discontinuities, and computational efficiency on real-world data. While more complex encodings, such as the full or polar representations, can marginally improve mean error for some models, the quadrant encoding achieves a superior balance between informativeness and generalisation, avoiding the pitfalls of overfitting or redundancy. Building on the insights gained from this study, future work will prioritise scaling to larger and more diverse datasets, extending evaluation to complex object categories, and exploring deep learning and hybrid approaches to benchmark gains over shallow architectures. Finally, efforts will focus on full pipeline integration and real-world robotic deployment, enabling end-to-end assessment in closed-loop manipulation scenarios.

REFERENCES

- [1] X. Gao and Z. Wen, "Task-oriented robotic grasping for intelligent manufacturing", in *2023 3rd International Symposium on Artificial Intelligence and Intelligent Manufacturing (AIIM)*, IEEE, 2023, pp. 101–104.
- [2] H. Sekkat, O. Moutik, L. Ourabah, B. ElKari, Y. Chaibi, and T. A. Tchakoucht, "Review of reinforcement learning for robotic grasping: Analysis and recommendations", *Statistics, Optimization & Information Computing*, vol. 12, no. 2, pp. 571–601, 2024.
- [3] L. Chen, P. Huang, Y. Li, and Z. Meng, "Detecting graspable rectangles of objects in robotic grasping", *International Journal of Control, Automation and Systems*, vol. 18, no. 5, pp. 1343–1352, 2020.
- [4] J. Redmon and A. Angelova, "Real-time grasp detection using convolutional neural networks", in *2015 IEEE international conference on robotics and automation (ICRA)*, IEEE, 2015, pp. 1316–1322.
- [5] S. Kumra and C. Kanan, "Robotic grasp detection using deep convolutional neural networks", in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2017, pp. 769–776.
- [6] M. Schwarz, H. Schulz, and S. Behnke, "RGB-D object recognition and pose estimation based on pre-trained convolutional neural network features", in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 1329–1335. DOI: 10.1109/ICRA.2015.7139363.
- [7] A. Gambale *et al.*, "A comparative study of hough transform and pca for bolt orientation detection", in *2024 IEEE 22nd International Conference on Industrial Informatics (INDIN)*, IEEE, 2024, pp. 1–6.
- [8] Y. Chen, E. Z. Zeng, M. Gilles, and A. Wong, "Metagraspnet_v0: A large-scale benchmark dataset for vision-driven robotic grasping via physics-based metaverse synthesis", *arXiv preprint arXiv:2112.14663*, 2021.
- [9] M. Gilles, Y. Chen, T. R. Winter, E. Z. Zeng, and A. Wong, "Metagraspnet: A large-scale benchmark dataset for scene-aware ambidextrous bin picking via physics-based metaverse synthesis", in *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, IEEE, 2022, pp. 220–227.
- [10] A. Gambale *et al.*, "Computing the orientation of hardware components from images using traditional computer vision methods", *Engineering Proceedings*, vol. 65, no. 1, p. 8, 2024.
- [11] A. Gambale, "Screwdrivers dataset", <https://app.roboflow.com/mvtec-iihve/screwdrivers-kmij6/2> (accessed: 2025-06-13), 2025.
- [12] A. Kirillov *et al.*, "Segment anything", in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 4015–4026.
- [13] J. Brownlee, *Ensemble Learning Algorithms With Python: Make Better Predictions with Bagging, Boosting, and Stacking*. Machine Learning Mastery, 2021.
- [14] H. Borchani, G. Varando, C. Bielza, and P. Larranaga, "A survey on multi-output regression", *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 5, no. 5, pp. 216–233, 2015.
- [15] T. Chen and C. Guestrin, *XGBoost 2.0: Scalable tree boosting system*, <https://xgboosting.com/xgboost-2.0/> (accessed: 2025-06-13), 2024.
- [16] Y. Bao, T. Xiong, and Z. Hu, "Multi-step-ahead time series prediction using multiple-output support vector regression", *Neurocomputing*, vol. 129, pp. 482–493, 2014, ISSN: 0925-2312. DOI: <https://doi.org/10.1016/j.neucom.2013.09.010>.
- [17] A. Depierre, E. Dellandréa, and L. Chen, "Scoring graspability based on grasp regression for better grasp prediction", in *2021 IEEE international conference on robotics and automation (ICRA)*, IEEE, 2021, pp. 4370–4376.

Investigating the Adaptability of ALE-AMR Hydrocode for Darcy Flow and Geothermal Simulations

Alice Koniges

Information and Computer Sciences
University of Hawai'i at Mānoa
Honolulu, HI, USA
email: koniges@hawaii.edu

David Eder

Physics and Astronomy
University of Hawai'i at Mānoa
Honolulu, HI, USA
email: dceder@hawaii.edu

Jonghyun Lee

Civil & Env Eng and Water Res Center
University of Hawai'i at Mānoa
Honolulu, HI, USA
email: jonghyun.harry.lee@hawaii.edu

Jiawei Shen

Civil & Env Eng and Water Res Center
University of Hawai'i at Mānoa
Honolulu, HI, USA
email: jiaweish@hawaii.edu

Aaron Fisher

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
Livermore, CA, USA
email: fisher47@llnl.gov

Tzanio Kolev

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
Livermore, CA, USA
email: kolev1@llnl.gov

Abstract—The PISALE ALE-AMR hydrocode suite is an advanced computational tool that combines the Arbitrary Lagrangian-Eulerian (ALE) method with Adaptive Mesh Refinement (AMR) to simulate complex multi-physics problems involving substantial material deformation. The suite currently includes physics modules for heat conduction and radiation transport, which are handled by a finite element diffusion solver operating on a structured, adaptive mesh infrastructure provided by the SAMRAI library. This paper investigates the feasibility of extending this framework to simulate fluid flow in porous media as described by Darcy's law, a critical component for subsurface applications like geothermal energy extraction and hydrogeology. We analyze the mathematical parallels between diffusion and Darcy flow, assess the suitability of the existing solver, and consider the integration of the more general MFEM finite element library. The primary objective is to evaluate the potential of the ALE-AMR methodology for Darcy flow simulations and to outline the necessary modifications and implementation steps, including addressing challenges related to integrating different AMR and grid formulations.

Keywords—ALE-AMR; Darcy flow; geothermal simulation; porous media; finite element method; MFEM; hydrocode; computational fluid dynamics.

I. INTRODUCTION

The PISALE ALE-AMR hydrocode suite represents an advanced computational tool that combines the Arbitrary Lagrangian-Eulerian (ALE) method with Adaptive Mesh Refinement (AMR) [1]. This combination enables the simulation of physical phenomena characterized by substantial material deformation, via the Lagrangian approach, while simultaneously addressing challenges associated with mesh distortion and optimizing computational efficiency through localized mesh refinement, which are key features of AMR methodologies. Initially published as ALE-AMR, the code is more recently referred to as PISALE (Pacific Island Structured-AMR with ALE), with various specialized versions developed to model a diverse range of applications [2].

Currently, the hydrocode suite incorporates physics modules dedicated to simulating heat conduction and radiation

transport, both of which are modeled using a finite element diffusion solver [3]. This solver is specifically engineered to operate on the dynamically adapting mesh structures generated by the AMR technique. The foundational support for adaptive mesh refinement within the hydrocode is provided by the SAMRAI (Structured AMR Application Infrastructure) library. Given its established infrastructure for managing complex mesh geometries and integrating various physics modules, we discuss here the potential for further expansion to simulate additional physical phenomena, including Darcy flow in subsurface applications.

Darcy's law governs the movement of fluid through a given material, and for this application, we consider flow through a porous medium [4]. This law states that the rate at which a fluid flows through a permeable material is directly proportional to the pressure gradient driving the flow and the intrinsic permeability of the material, while being inversely proportional to the viscosity of the fluid. In hydrogeology, it serves as a basis for analyzing groundwater flow; in petroleum engineering, it is used for multiphase flow modeling (e.g., the behavior of oil and gas reservoirs); and in geothermal energy, it is used for understanding the transport of both heat and fluids within the Earth's subsurface [5].

This paper presents an investigation into the feasibility of using the existing ALE-AMR methodology, in conjunction with its current finite element diffusion solver and/or the MFEM solver, for simulating fluid flow in porous media governed by Darcy's law. A detailed analysis examines the mathematical similarities and differences between diffusion and Darcy flow, assess the suitability of the current solver and/or replacing it with a more general finite element package, MFEM, for modeling Darcy flow. The primary objective is to evaluate the potential of ALE-AMR for Darcy flow simulations and outline the requisite steps for implementation.

The MFEM solver is a general purpose high-performance, open-source finite element library developed for solving partial differential equations. It provides a flexible and scalable frame-

work for discretizing and solving a wide range of problems. MFEM's capabilities include support for various finite element formulations, efficient linear and nonlinear solvers, and advanced mesh management techniques. By leveraging MFEM's strengths, this investigation aims to explore its potential as a viable alternative or complement to the existing diffusion solver for Darcy flow simulations within the ALE-AMR framework. One problem with direct application of MFEM is that it has its own AMR and grid formulation which needs adaptation for PISALE. We discuss solutions to this obstacle later in the paper.

II. THE ALE-AMR HYDROCODE AND ITS FINITE ELEMENT DIFFUSION SOLVER

The ALE method, as implemented in ALE-AMR, provides flexibility in handling various geometries and scenarios involving large-scale deformations and multiple objects within a domain. The mathematics follows a Lagrangian description, where the mesh follows the material, and an optional and/or modified remap to an Eulerian description, where the mesh appears via the remap to remain fixed in space. The AMR capabilities integrated within ALE-AMR serve to enhance the efficiency of computations by selectively refining the mesh in regions identified as requiring higher accuracy, such as areas with sharp gradients or complex flow patterns. This localized refinement ensures that computational resources are focused where they are most needed, without incurring the overall cost of a uniformly fine mesh. This ALE-AMR framework, which was developed initially for pure gas dynamics simulations, has matured into a comprehensive multiphysics framework capable of addressing a broad spectrum of applications, including phenomena in high-energy-density physics, material impacts, and laser target modeling [2]. Some very preliminary results on the application of PISALE to groundwater flow are given elsewhere [6].

The physics modules responsible for simulating heat conduction and radiation transport within ALE-AMR are enabled by a finite element diffusion solver specifically engineered to function on the composite meshes generated by the AMR. This solver employs a nodal-based approach, where the primary variables of interest are defined at the nodes of the computational mesh. A critical feature of this solver is its utilization of transition elements to effectively manage the interfaces that arise between regions of the mesh with different levels of refinement, a direct consequence of the AMR technique. These specialized elements are designed to ensure the continuity of the solution across these coarse-fine boundaries by appropriately handling the hanging nodes, edges, and faces that are characteristic of such interfaces. For the necessary 3:1 refinement ratio (or multiples, thereof) employed by the ALE-AMR framework for consistency, a variety of transition element types are used, depending on which of the element's sides are subject to refinement.

The finite element method necessitates the use of numerical integration techniques, which are implemented through

quadrature rules within the solver. For the transition elements, compound Gauss-Legendre quadrature[3] is employed to maintain a level of integration accuracy comparable to that achieved on standard elements. Additionally, mass lumping quadrature rules, which strategically place integration points at the element nodes, are utilized to produce diagonal mass matrices, a property that can be advantageous for certain time-stepping schemes. To overcome the challenge of undefined derivatives at the transition faces, which complicates the computation of the stiffness matrix, "blurred" quadrature rules are implemented. These rules work by averaging the evaluations of derivatives taken from different regions within the element, thereby ensuring the accurate assembly of the stiffness matrix.

The discretization of the diffusion equation is achieved using the standard Galerkin approach, a method where the equation is multiplied by a test function, integrated over the computational domain, and then subjected to integration by parts to derive the weak form. Both the solution being sought and the test functions used in the formulation are approximated using a basis set composed of shape functions defined on both standard and transition elements. This process culminates in a system of linear algebraic equations, typically represented in matrix form as $Au = f$, where A is the system/stiffness matrix, u is the vector containing the unknown nodal values of the solution, and f is the vector representing the source terms. This linear system is then solved using the HYPRE GMRES solver, an iterative algorithm particularly well-suited for handling large, sparse systems of equations, often enhanced by the use of a preconditioner to accelerate the convergence of the solution. The current finite element framework within PISALE/ALE-AMR is based on first-order H1 quadrilateral elements in two dimensions and hexahedral elements are required in three dimensions. These element types are recognized as being well-suited for the diffusion equation solvers that underpin the heat conduction and radiation diffusion modules.

The PISALE diffusion solver serves as the foundation for modeling heat conduction through the dynamic diffusion equation, which accounts for the temporal evolution of temperature and the flow of heat within the material. This equation incorporates parameters such as specific heat, thermal conductivity, and the absorptivity of the medium. Similarly, radiation transport can be modeled using the diffusion approximation, a simplification of the more complex radiative transfer equations that is applicable under certain conditions, such as in optically thick media. This approach involves formulating equations for both the energy density of radiation and the temperature of the material, with coupling terms that describe the absorption and emission of radiation. A significant hurdle in integrating these physics modules with the ALE-AMR framework arises from the inherent difference in how physical variables like temperature and energy are represented within the code. Specifically, the finite element method uses nodal representations, while PISALE uses cell-centered values for certain variables. To bridge this gap, we employ projection integrals as a means of mapping variables between the nodes and the cell centers,

a technique that ensures the conservation of energy during the transfer process. This mapping involves calculating the differences in cell temperatures after the hydrodynamic step, using the specific heat capacity to determine the corresponding energy and specific heat differences at the nodes, updating the nodal temperatures based on these differences, and then transferring these changes back to the cells to update their internal energy.

III. DARCY'S LAW AND FLOW IN POROUS MEDIA

In its most fundamental form, Darcy's law describes the rate of fluid flow (Q) through a porous medium as being directly proportional to the cross-sectional area (A) of the flow path and the pressure difference (ΔP) over a given length (L), and inversely proportional to the viscosity (μ) of the fluid. Mathematically, this relationship is expressed as $Q = -(kA/\mu)(\Delta P/L)$, where k represents the permeability of the porous medium. The negative sign in the equation signifies that the direction of flow is from regions of higher pressure to regions of lower pressure. Permeability (k) is an intrinsic property of the porous medium that quantifies its capacity to transmit fluids. This property is influenced by the grain size, shape, and interconnectedness of the pores within the material. Permeability can be uniform in all directions, in which case it is termed isotropic, or it can vary with direction, in which case it is termed anisotropic and is mathematically represented as a tensor. The differential form of Darcy's law relates the Darcy velocity (v), which is the volumetric flow rate per unit cross-sectional area, to the gradient of the pressure (∇p) and is given by $v = -(k/\mu)\nabla p$. The ratio k/μ multiplied by the specific weight (ρg) is often referred to as the hydraulic conductivity (K), particularly when considering the flow of a specific fluid with a known viscosity. Hydraulic conductivity can also incorporate the effect of gravity when the flow is described in terms of hydraulic head ($h = p/\rho g + z$), leading to the form $v = -K\nabla h$ [4].

For a steady-state flow of an incompressible fluid through a porous medium, the principle of mass conservation, expressed by the continuity equation, dictates that the divergence of the velocity field must be zero ($\nabla \cdot v = 0$). By combining this with the differential form of Darcy's law ($v = -(k/\mu)\nabla p$), we arrive at the governing equation for the pressure distribution within the medium: $\nabla \cdot (-(k/\mu)\nabla p) = 0$ when there is no source/sink to the system. In scenarios where the permeability (k) and the fluid viscosity (μ) are spatially uniform, this equation simplifies to Laplace's equation: $\nabla^2 p = 0$. However, in heterogeneous media where these properties vary from one point to another, the equation retains its more general elliptic partial differential form [7]. This mathematical similarity in form between the governing equation for pressure in Darcy flow and the steady-state diffusion equation ($\nabla \cdot (D\nabla u) = 0$), where D corresponds to k/μ and u to p , is a significant factor in considering the potential for adapting the existing diffusion solver.

Simulating Darcy flow typically involves the application of specific boundary conditions that define the state of the

flow at the edges of the computational domain. These commonly include: (i) Prescribed Pressure (Dirichlet boundary condition), where the pressure is set to a known value on certain boundaries, such as at the interface with a large fluid reservoir; (ii) Prescribed Flow Rate (Neumann boundary condition), where the rate at which fluid enters or leaves the porous medium across a boundary is specified such as no-flux boundary representing impermeable conditions; (iii) and Mixed Boundary Conditions, which involve applying different types of conditions on different segments of the domain's boundary. These are often used to model injection or production from wells. Well injection/extraction are typically dealt as source/sink conditions (non-zero RHS in the mass conservation equation). The ALE-AMR framework would need to be capable of implementing these types of boundary conditions, which might differ from those typically used in simulations of heat conduction and radiation transport.

One important point to note about the ALE-AMR framework however is the fact that because of the AMR refinement levels, one can often make the computational domain so large that the boundary conditions play little role in determining the early time behavior of the system. This is particularly beneficial for field-site applications since the modelers often make the domain large enough to minimize the effect of uncertain boundary conditions estimated in the field using geophysics or sparse field data sets. One can wrap the problem in a largely non-participatory airmesh for certain dynamical situations and effectively remove the boundary effects for the problem at hand.

Darcy's law and the associated governing equations are used extensively in modeling a wide variety of phenomena. These include the flow of groundwater in aquifers, encompassing scenarios such as flow towards extraction wells for freshwater supply, the regional movement of groundwater for contaminant remediation, and the interaction between groundwater and surface water bodies for flooding and drought risk mitigation. Notably, ALE methods have been successfully applied to simulate groundwater flow in situations involving free surfaces, which are characterized by moving boundaries [8]. In the field of petroleum engineering, Darcy's law is fundamental for simulating the flow of hydrocarbons (oil and gas) and water within subsurface reservoirs, enabling the prediction of production rates and the design of effective recovery strategies. The extraction of geothermal energy from the Earth's internal heat relies on the flow of fluids through porous rock formations, a process that can be modeled using Darcy's law, often in conjunction with equations governing heat transfer. Beyond these primary applications, Darcy's law is also utilized in modeling flow through various types of filters, membranes, and porous electrodes in devices like fuel cells.

While Darcy's law is a powerful tool, it is predicated on certain assumptions that limit its applicability to specific flow regimes and porous media characteristics. A fundamental assumption is that the flow is laminar, a condition typically met at low flow velocities and within media having small pore sizes, resulting in low Reynolds numbers (generally below

1 to 10). At higher velocities, inertial forces become non-negligible, and the flow transitions to turbulence, a regime where Darcy's law in its basic form is no longer accurate. In such cases, modifications like the Forchheimer equation, which incorporates a term proportional to the square of the velocity, are employed to account for these inertial effects. Furthermore, Darcy's law typically assumes a homogeneous and isotropic porous medium, meaning that the properties of the medium (like permeability) are uniform throughout and are the same in all directions. In reality, many geological formations and engineered materials exhibit heterogeneity, where properties vary spatially, and anisotropy, where permeability differs depending on the direction of flow. While the tensorial form of Darcy's law can accommodate anisotropy, significant heterogeneity might necessitate finer spatial discretization with more advanced modeling techniques. In very low permeability media under extremely small pressure gradients, deviations from the linear relationship described by Darcy's law have been observed, a phenomenon known as pre-Darcy flow. This behavior is thought to be due to factors such as the presence of immobile fluid layers at the pore walls. Lastly, the basic formulation of Darcy's law is for single-phase flow, where only one fluid is present in the porous medium. For scenarios involving multiple immiscible fluids (like oil and water in a reservoir), generalized forms of Darcy's law are used, which introduce the concept of relative permeabilities for each fluid phase. The dependence of fluid/rock density and permeability on the temperature, pressure, and displacement further complicate the modeling effort. We first consider what is involved in adapting the ALE-AMR code for Darcy flow based on its current diffusion solver for simulating laminar, single-phase flow in porous media. Modeling more complex flow regimes or multiphase scenarios would require additional developments.

IV. MATHEMATICAL PARALLELS AND DIVERGENCES: DIFFUSION VS. DARCY FLOW

The general form of the diffusion equation is given by $\partial u / \partial t = \nabla \cdot (D \nabla u) + S$, where u represents the quantity undergoing diffusion (such as temperature or concentration), D is the diffusion coefficient (e.g., thermal diffusivity or mass diffusivity), and S denotes any sources or sinks of the quantity u . In a steady-state scenario, where the conditions do not change with time, the time derivative becomes zero, resulting in the equation $\nabla \cdot (D \nabla u) + S = 0$. If, in addition, there are no sources or sinks within the domain, the equation further simplifies to $\nabla \cdot (D \nabla u) = 0$. In the special case where the diffusion coefficient D is also constant throughout the domain, the equation reduces to Laplace's equation, $\nabla^2 u = 0$. Within the ALE-AMR hydrocode, for the simulation of heat conduction, u corresponds to the temperature (T), and D is related to the thermal conductivity of the material. The source term S in this context can represent the generation or absorption of heat. For the modeling of radiation transport using the diffusion approximation, the equations involve the radiation energy density (E_R) and the material temperature (T), with the

"diffusion coefficient" being a function of radiation-specific properties such as opacities and the speed of light. Thus, the diffusion equation fundamentally describes the transport of a scalar quantity driven by its own spatial gradient.

As previously discussed, the steady-state flow of an incompressible fluid in a heterogeneous porous medium under Darcy's law is governed by the equation $\nabla \cdot (-(k/\mu) \nabla p) = 0$, where p is the pressure, k is the permeability tensor of the medium, and μ is the viscosity of the fluid.

Despite their different physical contexts, both the steady-state diffusion equation and the governing equation for Darcy flow with homogeneous medium share significant mathematical similarities. First, they are both second-order partial differential equations of the elliptic type. This classification implies that the solution at any given point within the domain is influenced by the conditions imposed at all the boundaries of the domain. Second, both equations describe a flux—be it heat flux or radiation flux in the case of diffusion, or Darcy velocity in the case of Darcy flow—that is directly proportional to the gradient of a scalar potential. For diffusion, this potential is temperature or radiation energy density, while for Darcy flow, it is pressure or hydraulic head. The constant of proportionality is a transport property, which is the diffusion coefficient in the diffusion equation and the permeability (divided by viscosity) in Darcy's law. Third, both types of equations are amenable to solution using similar numerical techniques, most notably the finite element method. This method involves discretizing the continuous domain into a mesh of smaller elements and then approximating the solution within each element using a set of basis functions. This suggests that the numerical methodologies already in place within the ALE-AMR diffusion solver could be adapted to address problems involving Darcy flow.

However, there are also key differences between these two types of physical processes and their mathematical representations. A primary divergence lies in the nature of the primary variable and the desired output. The diffusion solver in ALE-AMR is designed to solve for a scalar quantity, such as temperature or energy, which is also the main result of the simulation. In contrast, while the governing equation for Darcy flow is often solved for pressure, which is a scalar, the quantity of principal interest is frequently the Darcy velocity, which is a vector quantity representing the rate and direction of fluid flow [9]. To obtain this velocity, Darcy's law itself must be applied to the computed pressure gradient, either as a post-processing step or through a different formulation. Another significant difference pertains to the physical properties involved. The diffusion equation utilizes properties like thermal conductivity or radiation opacities, which are typically scalar quantities, although they can exhibit anisotropic behavior in some materials. Darcy flow, however, is characterized by permeability, which in anisotropic porous media is inherently a tensor, reflecting the fact that fluid flow can be more or less restricted depending on the direction. Another more complicated issue is spatial variability. One may need to discretize the domain with finer resolution, for example, the number of the cells with different permeability in the domain can be $100 \times 100 \times 100$

or $1000 \times 1000 \times 1000$ for 3D. In many cases, modelers assign only horizontal and vertical permeability in a cell instead of full tensor values. The full tensor is typically needed for coarse-scale discretization to simulate directional flow. With finer discretization, it can handle directional flow with small-scale horizontal/vertical permeability. The ALE-AMR solver would need to be capable of handling such tensorial properties for permeability. Additionally, the viscosity (μ) of the fluid is a critical parameter in Darcy's law, whereas it does not explicitly appear in the standard heat or radiation diffusion equations. While both types of equations can include source terms, their physical interpretations differ. Sources in diffusion problems represent the generation or absorption of the diffusing quantity (heat or energy), whereas in Darcy flow, sources typically correspond to the injection or extraction of fluid from the porous medium. Finally, the boundary conditions commonly encountered in Darcy flow simulations, such as prescribed flow rates or impermeable boundaries, might not have direct analogs in heat conduction or radiation transport problems, necessitating the implementation of new types of boundary conditions within the ALE-AMR framework.

V. ADAPTABILITY OF THE FINITE ELEMENT DIFFUSION SOLVER FOR DARCY FLOW

The current finite element diffusion solver within ALE-AMR employs a standard Galerkin method using first-order H1 elements. This choice of numerical method and element type is also prevalent in the solution of Darcy's equation, particularly when the primary variable being solved for is pressure. The solver's inherent capability to manage complex meshes arising from AMR, including the use of transition elements and specialized quadrature rules, presents a considerable advantage for potentially simulating fluid flow in porous media that exhibit geometric complexity or heterogeneity requiring localized mesh refinement. Furthermore, the utilization of an implicit solver (GMRES with a preconditioner) within the ALE-AMR framework suggests its suitability for handling elliptic partial differential equations, such as the steady-state form of the Darcy flow equation.

However, several potential challenges need to be addressed to adapt the existing diffusion solver for Darcy flow simulations. The solver would need to be modified to correctly interpret the "diffusion coefficient" in the governing equation as the permeability tensor of the porous medium divided by the viscosity of the fluid (k/μ). This adaptation would likely involve changes to the process of assembling the element stiffness matrix, particularly if the permeability is anisotropic, requiring the solver to handle tensor properties. The current diffusion solver is designed to output a scalar field (temperature or energy). For Darcy flow, the Darcy velocity vector is a key quantity that needs to be determined. This could be achieved through a post-processing step, where the gradient of the computed pressure field is calculated at the nodes or element centers, and then Darcy's law is applied to derive the velocity [9]. Alternatively, more significant modifications could involve exploring the implementation of mixed finite

element methods, which are formulated to solve for both pressure and velocity simultaneously, potentially offering a more direct and accurate way to obtain the velocity field [7]. The ALE-AMR framework would also require the implementation of boundary conditions that are specific to Darcy flow, such as the ability to prescribe flow rates at boundaries, which would necessitate adding new functionalities to both the solver and the overall framework. Finally, the way in which the pressure and velocity fields interact with the ALE mesh movement and remapping processes would need to be carefully designed and implemented, drawing upon the existing strategies used for coupling temperature and energy with the hydrodynamics.

VI. POTENTIAL MODIFICATIONS AND ADDITIONS TO THE ALE-AMR FRAMEWORK

To enable the simulation of Darcy flow within the ALE-AMR hydrocode, several modifications and additions to the existing framework would be necessary. A dedicated physics module for Darcy flow should be developed to encapsulate the governing equations and the specific parameters associated with fluid flow in porous media [7]. This module would be responsible for handling the input of spatially varying permeability (which could be a scalar or a tensor field), the viscosity of the fluid, and potentially porosity if more complex scenarios such as transient or compressible flow are to be considered. It would also manage the definition of source and sink terms that represent the injection or extraction of fluid from the porous medium.

The existing finite element solver would require several adaptations. Firstly, it needs to be capable of accepting the permeability tensor (divided by the fluid viscosity) as the transport property in the governing equation, instead of the thermal conductivity or radiation-related parameters it currently uses. This would likely necessitate modifications to the process by which the element stiffness matrix is assembled. Secondly, the solver should be configured to solve for pressure (or hydraulic head) as the primary unknown variable at the mesh nodes. Thirdly, a post-processing function should be incorporated to calculate the Darcy velocity vector at each node or within each element, based on the computed pressure gradient and Darcy's law. For potential future extensions to model transient Darcy flow, which would be relevant for applications such as groundwater flow with time-varying boundary conditions or sources, e.g., extraction in wells, the solver would need to include a time-stepping scheme. Furthermore, consideration should be given to exploring the implementation of mixed finite element methods, which employ different basis functions for pressure and velocity and solve for both simultaneously. This approach can often yield more accurate velocity fields directly, which is particularly important for problems where flow is coupled with transport processes. Various finite element methods, including continuous Galerkin (CG), discontinuous Galerkin (DG), weak Galerkin (WG), and mixed finite element methods (MFEM), are used for Darcy flow simulations, and the choice would depend on the desired accuracy and computational cost.

The ALE-AMR framework must also be extended to support boundary conditions that are specific to fluid flow in porous media. This involves implementing: the ability to prescribe flow rates at selected boundaries; the option to define impermeable boundaries where no flow occurs; and ensuring that these new boundary condition types can be applied and handled correctly by the finite element solver, while also being compatible with the ALE and AMR features of the code.

The mechanisms for mapping data between the nodal finite element representation and the cell-centered representation used in the ALE hydrodynamics need to be adapted to handle pressure and velocity fields. This is crucial for scenarios where Darcy flow might be coupled with other physical processes already modeled in ALE-AMR, such as thermal effects in geothermal reservoirs, or where the flow interacts with the moving mesh in ALE simulations. For example, changes in fluid pressure might induce deformation of the porous medium that in turn affect the rock permeability, or the movement of the computational mesh could affect the flow domain.

Finally, a rigorous program of verification and validation is essential. This includes developing a comprehensive suite of unit tests to ensure the correct implementation of the Darcy flow physics module and the modifications made to the finite element solver. Additionally, the adapted code should be thoroughly validated against analytical solutions for standard Darcy flow problems, such as flow in simple geometries or radial flow towards a well, and against established benchmark problems reported in the literature to confirm the accuracy and reliability of the new simulation capabilities.

VII. SUITABILITY ASSESSMENT AND RECOMMENDATIONS

To simulate Darcy flow within the framework of the ALE-AMR hydrocode, a series of key steps would be necessary. The most critical of these would be the development of a dedicated Darcy flow physics module. This module would be responsible for managing the input of parameters specific to porous media flow, such as permeability and fluid viscosity, as well as defining the source and sink terms relevant to fluid flow. The existing finite element solver would need to be modified to correctly interpret these parameters, to solve for pressure (or hydraulic head) as the primary variable, and to provide the Darcy velocity as a key output, either directly or through a post-processing calculation. Furthermore, the solver would need to be enhanced to handle permeability as a tensor to accurately model anisotropic porous media. The implementation of boundary conditions specific to fluid flow in porous media, such as prescribed flow rates and impermeable boundaries, would also be a necessary addition to the framework. Careful design and implementation of the coupling mechanisms between the new Darcy flow module and the existing ALE hydrodynamics would be crucial, drawing upon the experience gained from coupling heat conduction and radiation transport. Finally, a thorough and rigorous program of verification and validation, using both (semi-)analytical solutions and established benchmark problems from the liter-

ature, would be essential to ensure the accuracy and reliability of the newly implemented Darcy flow simulation capabilities.

In summary, we believe that extending the PISALE ALE-AMR framework to simulate Darcy flow is not only feasible but also holds significant promise for advancing subsurface models for geothermal applications. As high-performance computing increasingly relies on GPU architectures for acceleration, adapting a proven, fully parallel AMR-capable hydrocode like PISALE is a critical step toward next-generation modeling. We give details here of a direct path to building simulation tools that can exploit modern HPC architectures and software, and thus enable discovery of critical new important details in geothermal reservoir modeling.

VIII. ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy (DOE) by the University of Hawai'i at Mānoa and by Lawrence Livermore National Laboratory supported by DOE Science Foundations for the Energy Earthshots Award Number DE-SC0024728. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. Work at LLNL also supported by Contract DE-AC52-07NA27344. LLNL-CONF-2008557

REFERENCES

- [1] A. Koniges *et al.*, "Multi-material ale with amr for modeling hot plasmas and cold fragmenting materials," *Plasma Sci. and Technol.*, vol. 17, pp. 117–128, 2015.
- [2] A. Koniges *et al.*, "A survey of recent applications of the pisale code and pde framework," *ADVCOMP 2023 : The Seventeenth International Conference on Advanced Engineering Computing and Applications in Sciences*, pp. 20–25, 2023.
- [3] A. Fisher *et al.*, "An amr capable finite element diffusion solver for ale hydrocodes," *Plasma Science and Technology*, vol. 17, no. 2, p. 109, 2015.
- [4] G. Brown, "Henry darcy and the making of a law," *Water Resources Research*, vol. 38, no. 7, pp. 11–1, 2002.
- [5] Z. Chen and R. E. Ewing, *Fluid Flow and Transport in Porous Media, Mathematical and Numerical Treatment: Proceedings of an AMS-IMS-SIAM Joint Summer Research Conference on Fluid Flow and Transport in Porous Media, Mathematical and Numerical Treatment, June 17-21, 2001, Mount Holyoke College, South Hadley, Massachusetts. American Mathematical Soc., 2002*, vol. 295.
- [6] Y. Seo, J. Lee, A. Koniges, and A. Fisher, "Development of the pisale codebase for simulating flow and transport in large-scale coastal aquifer," *Eleventh International Conference on Computational Fluid Dynamics (ICCFD11)*, 2022, Paper 1502.
- [7] J. Liu, L. Mu, and X. Ye, "A comparative study of locally conservative numerical methods for darcy's flows," *Procedia Computer Science*, vol. 4, pp. 974–983, 2011.
- [8] Y. Jin, E. Holzbecher, and M. Sauter, "A novel modeling approach using arbitrary lagrangian-eulerian (ale) method for the flow simulation in unconfined aquifers," *Computers & Geosciences*, vol. 62, pp. 88–94, 2014.
- [9] R. W. Zimmerman, *The Imperial College Lectures in Petroleum Engineering: Volume 5: Fluid Flow in Porous Media*. World Scientific, 2018.

TABLE I. COMPARISON OF DIFFUSION AND DARCY FLOW EQUATIONS

Feature	Diffusion Equation	Darcy Flow Equation
Governing Equation	$\partial u/\partial t = \nabla \cdot (D \nabla u) + S$ (General)	$\nabla \cdot (-(k/\mu) \nabla p) = S$ (Steady-State)
Primary Unknown	Temperature (T), Energy (E_R)	Pressure (p), Hydraulic Head (h)
Key Properties	Conductivity, Diffusivity, Opacities	Permeability (k), Viscosity (μ)
Applications in ALE-AMR	Heat Conduction, Radiation Transport	Groundwater, Petroleum, Geothermal

TABLE II. SUMMARY OF REQUIRED MODIFICATIONS FOR DARCY FLOW IMPLEMENTATION

Category	Specific Required	Action	Purpose	Potential Challenges
New Physics Module	Develop a dedicated Darcy flow module.	Organize parameters (permeability, viscosity, and equations).		Seamless integration with ALE-AMR architecture.
Solver Adaptations	Modify solver for permeability (tensor) and viscosity; solve for pressure; post-process for velocity. Consider mixed FEM.	Accurately represent Darcy's law; provide velocity output; handle anisotropy and transient flow.		Handling tensor properties; ensuring velocity accuracy; significant code changes for mixed FEM.
Boundary Conditions	Implement prescribed flow rate and impermeable boundaries.	Model physical conditions at domain edges; ensure compatibility with ALE/AMR.		Correctly imposing conditions on AMR meshes with hanging nodes.
Data Mapping	Define pressure/velocity interaction with ALE hydrodynamics and evolving media properties.	Enable coupling with other physics (thermal, structural); ensure consistent data transfer.		Designing robust and conservative mapping strategies.
Verification	Develop unit tests and benchmark problems. Validate against analytical solutions.	Ensure correctness and accuracy of the new simulation capabilities.		Identifying appropriate validation cases and benchmark problems.

Modeling the Interaction of Laser-Produced Proton Beams with Matter

Jack McKee

Mathematics
University of Hawai'i at Mānoa
Honolulu, HI, USA
email: jmckee@math.hawaii.edu

David Eder

Physics and Astronomy
University of Hawai'i at Mānoa
Honolulu, HI, USA
email: dceder@hawaii.edu

Aaron Fisher

Information Technology Services
University of Hawai'i at Mānoa
Honolulu, HI, USA
email: fallen@andcheese.org

Alice Koniges

Information and Computer Sciences
University of Hawai'i at Mānoa
Honolulu, HI, USA
email: koniges@hawaii.edu

Claudia Parisuaña

Department of Mechanical Engineering
Stanford University
Stanford, CA USA
email: cparisua@stanford.edu

Maxence Gauthier

SLAC National Accelerator Laboratory
Stanford University
Menlo Park, CA USA
email: gauthier@slac.stanford.edu

Emma Elizabeth McBride

SLAC and School of Math & Physics
Queen's University Belfast
Belfast BT7 1NN, UK
email: e.mcbride@qub.ac.uk

Frank Seiboth

Deutsches Elektronen-Synchrotron DESY
Research Centre of the Helmholtz Association
22607 Hamburg, Germany
email: frank.seiboth@desy.de

Chandra Breanne Curry

SLAC National Accelerator Laboratory
Stanford University
Menlo Park, CA USA
email: ccurry@slac.stanford.edu

Mungo Frost

SLAC National Accelerator Laboratory
Stanford University
Menlo Park, CA USA
email: md frost@slac.stanford.edu

Eric Galtier

SLAC National Accelerator Laboratory
Stanford University
Menlo Park, CA USA
email: egaltier@slac.stanford.edu

Siegfried Glenzer

SLAC National Accelerator Laboratory
Stanford University
Menlo Park, CA USA
email: glenzer@slac.stanford.edu

Abstract—The paper reports on efforts to significantly increase our understanding of isochoric heating of matter using laser-produced proton beams, and the associated High Energy Density (HED) and Warm Dense Matter (WDM) regimes generated. This will benefit research fields such as planetary science, fusion energy, plasma physics, long-term battery storage, qubit synthesis, and material science. We discuss our experiments that irradiated Si targets with proton beams generated by the 20 TW-laser at the SLAC MEC end-station, Linac Coherent Light Source (LCLS), SLAC National Accelerator Laboratory. The HED/WDM states are probed using the 50 fs hard X-rays using the fundamental undulator harmonic of LCLS. We compare our results from the phase contrast X-ray imaging, which shows the generation of compression waves that produces rear surface spallation, to modeling results from the 3D multi-physics multi-material code, PISALE, that combines Arbitrary Lagrangian-Eulerian (ALE) hydrodynamics with Adaptive Mesh Refinement (AMR). This comparison required modifications to several physics models in the PISALE (Pacific Island Structured-AMR with ALE) code.

Keywords—Adaptive Mesh Refinement; Computational Fluid Dynamics; Arbitrary Lagrangian Eulerian Methods; High Performance Computing; isochoric heating; Phase Contrast Microscopy; X-ray Diffraction; X-ray Free Electron Lasers.

I. INTRODUCTION

Increasing our understanding of isochoric (constant volume) heating of matter and the associated High Energy Density (HED) and Warm Dense Matter (WDM) [1][2][3] generated regimes using laser-produced proton beams [4][5][6] will benefit research in many fields, such as planetary science

[7], fusion energy [8][9][10][11][12][13][14], plasma physics [15][16][17][18], long-term battery storage, qubit synthesis [19], and material science [20]. For example, it will enhance our understanding of WDM properties of iron and silica under conditions encountered in planetary interiors and diagnostic components in fusion devices exposed to high fluxes of energetic plasma ions. The work is also relevant to long-term battery storage using Si. Short-pulse laser-produced ion beams can create chains of closely coupled qubits based on "color centers" in diamond and other material [19].

We present results of experiments that irradiated Si targets, at near constant density to eV temperatures, with proton beams generated by the 20 TW-laser at the Matter in Extreme Conditions (MEC) end-station at the Linac Coherent Light Source (LCLS) located at the SLAC National Accelerator Laboratory. This short-lived high-temperature, high-density state is probed using the 50 fs hard X-rays using the fundamental undulator harmonic of LCLS. The results of the phase-contrast X-ray imaging indicate ion energy deposition inside the 20 μm thick Si target, and the subsequent evolution reveals the generation of compression waves that reflect off the rear surface resulting in intense tensile stress and spallation.

In this paper, we also discuss the development and use of the PISALE code to model our experiments. The 3D multi-physics multi-material code, PISALE, combines Arbitrary Lagrangian-Eulerian (ALE) hydrodynamics with Adap-

tive Mesh Refinement (AMR). The PISALE (Pacific Island Structured-AMR with ALE) code has physics models that include laser/ion deposition, radiation hydrodynamics, thermal diffusion, anisotropic material strength with material time history, advanced models for fragmentation, and surface tension models. The PISALE code has an ion beam deposition package that has been used to model Li ions. In some earlier papers (e.g., [21]), the PISALE code is called ALE-AMR as it was one of the first codes to combine those two methods.

Modeling new experimental configurations generally requires modifications to the code in addition to determining new parameters for the various models that are used in the simulations. One advantage of PISALE over some codes used to model US Department of Energy (DOE) funded experiments, is that graduate students and postdocs can modify the source for a particular experiment [22].

In section II, we provide a brief background of the PISALE code and the equations that are solved. In section III, we discuss improvements to the code that are needed to model ion beam experiments with a range of ion energies. In section IV, we provide details of our experiment using laser-produced protons to heat a Si target. In section V, we discuss PISALE simulations of our experiment. In section VI, we provide conclusions and discuss future work.

II. PISALE CODE BACKGROUND

PISALE operates on top of the scalable Structured-AMR Application Infrastructure (SAMRAI) library [23]. The PISALE code contains a general purpose PDE solver that uses a staggered-grid, Lagrangian formulation, written for coupled plasma/fluids with position and velocity being nodal variables and density, internal energy, temperature, pressure, strain, and stress being zonal (cell centered) variables. These plasma/fluid equations in a Lagrangian formulation (in vector and indicial notation $i, j, k = 1, 2, 3$) are:

$$\frac{D\rho}{Dt} = -\rho \nabla \cdot \vec{U} = \rho U_{i,i} \quad (1)$$

$$\frac{D\vec{U}}{Dt} = \frac{1}{\rho} \nabla \cdot \boldsymbol{\sigma} = \frac{1}{\rho} \sigma_{ij,j} \quad (2)$$

$$\frac{De}{Dt} = \frac{1}{\rho} V s : \dot{\boldsymbol{\epsilon}} - P \dot{V} = \frac{1}{\rho} V (s_{ij} \dot{\epsilon}_{ij}) - P \dot{V} \quad (3)$$

where $\frac{D}{Dt} = \frac{\partial}{\partial t} + \vec{U} \cdot \nabla$ is the substantial derivative, ρ is the density, $\vec{U} = (u, v, w)$ is the material velocity, t is time, $\boldsymbol{\sigma}$ is the total stress tensor, P is the pressure, e is the internal energy, V is the relative volume ($\rho V = \rho_0$ where ρ_0 is the reference density), s is the deviatoric stress defined as $s_{ij} = \sigma_{ij} + P \delta_{ij}$ where δ is the Kronecker delta and $\dot{\boldsymbol{\epsilon}}$ is the strain rate tensor. PISALE has a range of different strength and failure models.

Thermal conduction and radiation transport coupled to the basic conservation law equations are solved by implementing the diffusion approximation, which uses a nodal radiation energy and a zone-averaged nodal temperature [24]. The original PISALE Finite Element Method (FEM) package includes only first order H^1 quadrilateral and hexahedron elements in 2D and 3D. This works well for the diffusion equation solvers

utilized in the heat conduction and radiation diffusion modules. The original FEM package in PISALE only supports AMR in 2D. We are currently coupling PISALE to the Modular Finite Element Methods (MFEM) library [25]. This will provide a wide variety of powerful FEM-based PDE solvers, with the added benefit of support for GPU acceleration, automatic differentiation, high-order methods and more including the support of AMR in 3D simulations.

III. RECENT IMPROVEMENTS TO PISALE TO MODEL ION-BEAM EXPERIMENTS

In order to model our experiments using laser-produced ion beams, we evaluated and extended the ion-beam package in PISALE. Previous work using the package was restricted to a monoenergetic beam of Li ions. We have modified the ion-beam package in PISALE to have multiple ion beams with different ion properties, including ion energy, spot size, and fluences as a function of time. The PISALE code exploits complex numerical techniques for fully anisotropic stress tensors, interface reconstruction for multiple materials, and a flexible strength/failure infrastructure for analytic or tabulated material models and equations of state. We have determined the appropriate equation of state and material strength/failure model for the Si material used in our experiments. In the PISALE code, a void material with an associated volume fraction is introduced when a material failure occurs. This can result in spall planes that will be compared to measured spall generation.

IV. EXPERIMENTS DETAILS USING LASER-PRODUCED PROTON BEAMS




Figure 1. Experimental configuration using proton beams generated by the 20 TW-laser at the SLAC MEC end-station to heat a Si target.

Laser-produced high-energy (MeV) proton beams can be used for isochoric heating of solid targets [26][27][28]. In our experiment, the 0.66 J, ~50 fs, 800 nm Ti:Sapphire MEC pump laser (P-polarized) was focused onto a 4 μm-thick propylene foil using an f/6 off-axis parabola, reaching peak intensities of $\sim 2 \times 10^{19}$ W/cm². Interaction with the solid-density target generated a population of hot electrons, a fraction of which escaped from the rear surface, forming a strong electrostatic sheath field. This field accelerates surface ions via the Target Normal Sheath Acceleration (TNSA) mechanism [29], producing a well-collimated proton beam normal to




Figure 2. Measured proton spectrum generated by the 20 TW-laser striking polypropylene target.




Figure 3. Measured response in Si target from proton beam arriving from the left.

the target. A secondary sample—20 μm -thick single-crystal silicon—was placed 300–500 μm from the proton source, where the ions deposited most of their energy at the end of their range, resulting in isochoric heating. The experimental set-up is shown in Figure 1. This short-lived warm dense matter state sample is imaged using the 50 fs, 11 keV X-ray LCLS beam. The proton spectrum and angular distribution produced during the interaction were measured using an absolutely calibrated Thomson parabola spectrometer and a stack of RadioChromic Films (RCFs). The measured proton spectrum is shown in Figure 2. The combination of the high intensity short pulse laser system at MEC, able to create MeV proton beams to isochorically heat matter, with the LCLS, allows unprecedented access to the creation and interrogation of this exotic state of matter.

The heated target was spatially and temporally diagnosed using the MEC X-ray Imager (MXI), operated in a Phase-Contrast Imaging (PCI) configuration [30]. The collimated 11 keV LCLS beam was prefocused upstream of the sample using

a stack of 40 beryllium Compound Refractive Lenses (CRLs), generating a secondary x-ray source located 214 mm before the sample. The transmitted x-rays were recorded on a YAG phosphor screen positioned 4901 mm downstream. This screen was re-imaged by a high-resolution optical system coupled to an sCMOS camera using a 1.8 x magnification objective, yielding a total geometric magnification of 42 and an effective pixel size of approximately 150 nm.

Typical experimental results are shown in Figure 3. The MeV proton beam arrives at the target from the left. These measurements visualize the density evolution as a function of time starting with the firing of the short-pulse laser. The generated protons travel to the Si target with the highest energy ones arriving first. One nanosecond after the pump laser irradiates the polypropylene target and starts producing the proton beam, we observe the effects of proton-induced energy deposition (left image). At 5 ns (center image) we see the generation of a compression wave. At 17.5 ns (right image), we see material failure and spallation. In another experiment, we observed similar spallation in Ge foils using a 20 ns laser pulse to generate the shock [30].




Figure 4. Calculated internal energy per unit mass of 20 μm thick Si target heated by proton beam arriving from the left at four times.

V. PISALE SIMULATION RESULTS

Simulations of our experiments can provide insight into the properties of the Si target and the cause of the observed spallation. We use the measurements of the proton spectrum, shown in Figure 2, as input for our PISALE simulations. We split energies into fixed-width ranges based on the individual proton energies, and for each bin we calculated a beam intensity that matches the total deposited energy of that bin determined by integrating an exponential fit of the curve in Figure 2. The arrival time at the Si sample is a function of their energy, with the highest energy ones arriving first. Given

the relatively small distance between the foil where the protons are produced and the Si targets, all the protons arrive before any significant hydrodynamics motion occurs. This results in isochoric heating by the protons. We divide the protons into different energy groups with corresponding different arrival times. In addition to arriving at different times, the ions with different energies deposit their energy at different locations in the Si sample. We explored using different numbers of energy groups, and the results shown in this report are for 17 energy groups going from 0.05 to 1.8 MeV with almost 60% of the energy in the proton beam being in the 0.2 to 0.6 MeV energy range. Less than 10% of the proton beam energy is below 0.2 MeV.

In general, the most important aspect of proton beams having a wide range of energies is the location of the deposited energy. In Figure 4, we show results from the PISALE simulation for the internal energy per unit mass inside the 20 μm thick Si target heated by the 50 μm radius proton beam coming from the left. The proton beam is centered on 0 in the Y axis. We model just the top half the Si target given that the target and proton beam are symmetric in Y. [These results are for a distance of 4 mm between the polypropylene foil and the Si sample, which is larger than the distance of 300 μm discussed in the previous experimental section.] The first image is at 250 ps when only protons with energy greater than 1.3 MeV have started to arrive at the target. One can see a small amount of heating in the center of the beam. Note the log scale. The next image is at 310 ps, with heating seen in the entire beam with the deposition relatively uniform through the thickness of the Si target. This is expected for these high-energy protons. At 431 ps, some of the lower-energy protons with shorter stopping distances start to arrive, and we see a higher internal energy on the front side of Si target. In the far right image, we show results at 713 ps when protons in the lowest-energy group have started to arrive, and most higher-energy protons have deposited their energy or have passed through the Si target.

The high internal energy near the front side of the target causes a shock to form and move to the right before significant hydrodynamic motion. The calculated response of the Si target to proton-beam heating is shown in Figure 5, where we show the density, with a linear scale, at four times. All the energy deposition has essentially completed by 1 ns. In the far left image at 1.4 ns, we see the shock moving to the right with a narrow density enhancement location about 3 μm into the Si target in the center of the beam path. In the next image at 2 ns, the density enhancement has broadened and is about 7 μm into the Si target. This corresponds to a velocity of 6.7 $\mu\text{m}/\text{ns}$. This is consistent with the observed velocity obtained from images at 1 and 5 ns in Figure 3. The start of rear surface motion in the simulation is seen in the 6 ns image in Figure 4. Rear-surface spall is seen in the image on the far right at 9 ns with some melting of the Si target. The rear-surface spall is calculated to start earlier than is observed. These PISALE simulations are in good agreement with the experimental data for the location of deposition and the measured shock speed.

The simulation and the experiment both have rear-surface spall with the simulation having onset of spall approximately 2X sooner than the experiment.

As already stated, these simulation results are for a flight distance of 4 mm for protons, as compared to 300 μm in the experiment. For a closer distance of 300 μm , all of the proton energy will be deposited in approximately 100 ps. Given that there is little volume change in the first nanosecond, we expect the Si response, including the spall, to be similar. Another reason for expecting a similar Si response is that the location of proton deposition within the Si target is primarily a function of the proton energy, which is independent of the flight distance of the protons.

VI. CONCLUSIONS AND FUTURE WORK

The project has expanded the capability of the PISALE code to model complex ion beams with varying energies, arrival times, and other beam properties. This is beneficial for ongoing modeling efforts in modeling laser-produced ion for the formation of color centers relevant to photon qubits and studying new materials for long-term battery storage. We plan to study the effect of smaller flight distance and subsequent shorter energy deposition times. We also plan to explore the effects of different equation-of-state tables/models and strength models for Si on the simulation of spallation.




Figure 5. Calculated response of 20 μm thick Si target following heating by proton beam at four times. Measured response in Si target from MeV proton beam arriving from the left.

ACKNOWLEDGMENT

We would like to acknowledge the support and contributions of Z. Chen, A. Descamps, G. Dyer, L.B. Fletcher, G.D. Glenn, J. Hastings, P. Heimann, J.B. Kim, M. Mo, B. Ofori-Okai, and F. Treffert from SLAC National Accelerator Laboratory, and C. Spindloe from the Central Laser Facility, for their

involvement in the experimental campaign. This work is supported by the U.S. Department of Energy, Office of Science, under Fusion Energy Sciences Research Division Award Numbers DE-SC0023475 and DE-SC0024403, and through through FWP100182, as well as under Science Foundations for the Energy Earthshots Award Number DE-SC0024728. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. This work was also supported by the UK Research & Innovation Future Leaders Fellowship (MR/W008211/1).

REFERENCES

- [1] P. Patel *et al.*, “Isochoric heating of solid-density matter with an ultrafast proton beam,” *Physical review letters*, vol. 91, no. 12, p. 125 004, 2003.
- [2] R. Snavely *et al.*, “Laser generated proton beam focusing and high temperature isochoric heating of solid matter,” *Physics of Plasmas*, vol. 14, no. 9, p. 092 703, 2007.
- [3] S. Malko *et al.*, “Proton stopping measurements at low velocity in warm dense carbon,” *Nature Communications*, vol. 13, no. 1, p. 2893, 2022.
- [4] S. P. Hatchett *et al.*, “Electron, photon, and ion beams from the relativistic interaction of petawatt laser pulses with solid targets,” *Physics of Plasmas*, vol. 7, no. 5, pp. 2076–2082, 2000.
- [5] E. Clark *et al.*, “Energetic heavy-ion and proton generation from ultraintense laser-plasma interactions with solids,” *Physical Review Letters*, vol. 85, no. 8, p. 1654, 2000.
- [6] S. Wilks *et al.*, “Energetic proton generation in ultra-intense laser–solid interactions,” *Physics of plasmas*, vol. 8, no. 2, pp. 542–549, 2001.
- [7] R. Smith *et al.*, “Ramp compression of diamond to five terapascals,” *Nature*, vol. 511, no. 7509, pp. 330–333, 2014.
- [8] D. A. Callahan-Miller and M. Tabak, “Progress in target physics and design for heavy ion fusion,” *Physics of Plasmas*, vol. 7, no. 5, pp. 2083–2091, 2000.
- [9] I. Hofmann, “Review of accelerator driven heavy ion nuclear fusion,” *Matter and Radiation at Extremes*, vol. 3, no. 1, pp. 1–11, 2018.
- [10] M. Roth *et al.*, “Fast ignition by intense laser-accelerated proton beams,” *Physical review letters*, vol. 86, no. 3, p. 436, 2001.
- [11] J. C. Fernandez *et al.*, “Fast ignition with laser-driven proton and ion beams,” *Nuclear fusion*, vol. 54, no. 5, p. 054 006, 2014.
- [15] M. Borghesi *et al.*, “Electric field detection in laser-plasma interaction experiments via the proton imaging technique,” *Physics of Plasmas*, vol. 9, no. 5, pp. 2214–2220, 2002.
- [16] L. Romagnani *et al.*, “Dynamics of electric fields driving the laser acceleration of multi-mev protons,” *Physical review letters*, vol. 95, no. 19, p. 195 001, 2005.
- [12] M. Key *et al.*, “Proton fast ignition,” *Fusion science and technology*, vol. 49, no. 3, pp. 440–452, 2006.
- [13] M. Roth *et al.*, “Proton acceleration experiments and warm dense matter research using high power lasers,” *Plasma Physics and Controlled Fusion*, vol. 51, no. 12, p. 124 039, 2009.
- [14] O. Hurricane *et al.*, “Approaching a burning plasma on the nif,” *Physics of Plasmas*, vol. 26, no. 5, p. 052 704, 2019.
- [17] J. Fuchs *et al.*, “Comparative spectra and efficiencies of ions laser-accelerated forward from the front and rear surfaces of thin solid foils,” *Physics of plasmas*, vol. 14, no. 5, p. 053 105, 2007.
- [18] K. Quinn *et al.*, “Laser-driven ultrafast field propagation on solid surfaces,” *Physical review letters*, vol. 102, no. 19, p. 194 801, 2009.
- [19] W. Redjem *et al.*, “Defect engineering of silicon with ion pulses from laser acceleration,” *Communications Materials*, vol. 4, no. 1, p. 22, 2023.
- [20] M. Roth *et al.*, “Energetic ions generated by laser pulses: A detailed study on target properties,” *Physical Review Special Topics-Accelerators and Beams*, vol. 5, no. 6, p. 061 301, 2002.
- [21] A. Koniges *et al.*, “Multi-material ale with amr for modeling hot plasmas and cold fragmenting materials,” *Plasma Sci. and Technol.*, vol. 17, pp. 117–128, 2015.
- [22] A. Koniges *et al.*, “A survey of recent applications of the pisale code and pde framework,” *ADVCOMP 2023, The Seventeenth International Conference on Advanced Engineering Computing and Applications in Sciences*, p. 20, 2023.
- [23] B. T. N. Gunney and R. W. Anderson, “Advances in patch-based adaptive mesh refinement scalability,” *Journal of Parallel and Distributed Computing*, vol. 89, pp. 65–84, 2016.
- [24] A. C. Fisher *et al.*, “Modeling heat conduction and radiation transport with the diffusion equation in NIF ALE-AMR,” *Journal of Physics: Conference Series*, vol. 244, p. 022 075, 2010.
- [25] Mfem, <https://mfem.org/>, Accessed: 2025-07-11.
- [26] C. McGuffey *et al.*, “Focussing protons from a kilojoule laser for intense beam heating using proximal target structures,” *Scientific reports*, vol. 10, no. 1, p. 9415, 2020.
- [27] A. Mancic *et al.*, “Isochoric heating of solids by laser-accelerated protons: Experimental characterization and self-consistent hydrodynamic modeling,” *High Energy Density Physics*, vol. 6, no. 1, pp. 21–28, 2010.
- [28] M. Carrié, P. Combis, and E. Lefebvre, “Isochoric heating with laser-accelerated proton beams,” *Physics of Plasmas*, vol. 17, no. 12, p. 122 707, 2010.
- [29] R. Snavely *et al.*, “Intense high-energy proton beams from petawatt-laser irradiation of solids,” *Physical review letters*, vol. 85, no. 14, p. 2945, 2000.
- [30] F. Seiboth *et al.*, “Simultaneous 8.2 kev phase-contrast imaging and 24.6 kev x-ray diffraction from shock-compressed matter at the lcls,” *Applied Physics Letters*, vol. 112, no. 22, 2018.

Artificial Intelligence-Based Local Weather Forecasting for Agricultural Digital Twins

Miguel Zaragoza-Esquerdo¹, Alberto Ivars-Palomares¹, Irene Eiros-Fonseca², Sandra Sendra¹, Jaime Lloret¹

¹Instituto de Investigación para la Gestión Integrada de Zonas Costeras, Universitat Politècnica de València

C/Paranimf, 1, 46730 Grao de Gandia, Spain

²Universidad CEU San Pablo

C. de Julián Romea, 23, Chamberí, 28003 Madrid

mizaes2@epsg.upv.es, aivapal@epsg.upv.es, i.eiros@usp.ceu.es, sansenco@upv.es, jlloret@dcom.upv.es

Abstract—This work presents the development of a local weather forecasting system integrated into an agricultural digital twin, leveraging classical machine learning. Data were collected from ESP32-based weather stations equipped with temperature, relative humidity, and atmospheric pressure sensors. The acquired measurements were processed through a Node.js server and used to train predictive models, including Random Forest, Gradient Boosting, Ridge Regression, Lasso Regression and K-Nearest Neighbors. A sliding window approach was applied to structure the input data for short-term forecasting. Experimental results show that Gradient Boosting achieved the best performance among classical methods for atmospheric pressure but exhibited overfitting for temperature and humidity. These findings highlight the potential of Artificial Intelligence (AI)-powered digital twins to enhance precision agriculture by providing accurate, localized, and up-to-date weather forecasts.

Keywords—Digital twin; local weather forecasting; machine learning; deep learning; agriculture; Gradient Boosting.

I. INTRODUCTION

In recent decades, technological advancements have driven the development of digital twins, virtual replicas of real-world physical systems that enable real-time monitoring, simulation, and control. In the agricultural sector, these digital twins integrate data from sensors, weather stations, and Internet of Things (IoT) devices to remotely model and optimize complex processes, facilitating the management of resources such as water, fertilizers, and pesticides [1].

The reliability of these models largely depends on the underlying sensing infrastructure. Well-designed Wireless Sensor Network (WSN) architectures, such as those proposed by Lloret et al. [2], and Hussein et al. [3] have proven essential for ensuring coverage, scalability, and efficiency in data collection. Such infrastructures enable parallel and organized communication between multiple nodes, optimizing network topology and reducing latency in transmitting critical data.

By replicating plants and cultivation environments, digital twins offer farmers decision support systems that reduce resource consumption and improve productivity [4]. These virtual representations facilitate scenario evaluation and predictive analysis without extensive physical trials, accelerating the digital transformation of the agricultural

sector [5]. Incorporating artificial intelligence and machine learning expands the capabilities of these models, allowing, for example, adaptive irrigation scheduling based on real-time soil moisture data and weather forecasts, significantly reducing water waste [6].

Artificial intelligence techniques have emerged as a key component in high-precision weather forecasting, leveraging convolutional and recurrent neural networks to model complex atmospheric phenomena [7]. These approaches have been shown to improve the estimation of precipitation and extreme temperatures, enabling digital twins to anticipate adverse conditions and proactively adapt their crop management strategies [8].

Moreover, digital twins support sustainable agricultural practices through continuous environmental monitoring and adaptive management strategies [9]. The combination of cloud computing technologies and edge devices increases data processing capacity and allows faster, more precise responses to changes in the field [10].

This work presents the design and implementation of a local weather forecasting system integrated into an agricultural digital twin. Unlike previous studies, our proposal combines IoT-based sensing infrastructure with classical machine learning models to generate short-term, high-resolution forecasts directly tailored to the conditions of a specific agricultural plot. The system demonstrates the feasibility of deploying low-cost weather stations with real-time data processing and highlights the comparative advantages and limitations of different predictive approaches. This contribution provides a practical framework for enhancing decision-making in precision agriculture through accurate, localized, and continuously updated forecasts for digital twins data input.

The remainder of this article is structured as follows: Section II reviews related works on digital twins, artificial intelligence, and weather forecasting in agricultural contexts. Section III details the architecture and operation of the proposed system, including data acquisition, preprocessing, and model training. Section IV presents and discusses the experimental results, covering machine learning approaches. Finally, Section V summarizes the conclusions and outlines potential future research directions.

II. RELATED WORK

Several studies have reviewed the adoption of digital twins in agriculture, focusing on controlled environments such as greenhouses. A systematic review highlights how these models improve horticultural productivity and sustainability, emphasizing their role in microclimate control, crop growth monitoring, and resource use efficiency [11].

Integrating sensors to measure humidity, temperature, CO₂ levels, and light, along with IoT platforms, increases the accuracy of growth simulations and optimizes climate control in greenhouses [12]. As documented by Bri et al. [13], real deployments of wireless sensor networks demonstrate the feasibility and challenges associated with large-scale agricultural monitoring, including node resilience in adverse environments and energy management to ensure continuity of real-time measurements.

Virtual and augmented reality technologies have strengthened digital twin platforms by enabling immersive interactions with virtual crop models. Examples like the “Virtual Breeding Nursery” allow farmers to explore virtual plots, manipulate environmental variables, and simulate stress or pest infestation scenarios [14]. These intuitive data visualization interfaces combine sensor information with 3D models of plants and structures, supporting more informed decision-making [15].

In open-field contexts, pilot projects demonstrate that platforms equipped with IoT sensor networks—streaming real-time soil moisture, nutrient content, and weather data—can dynamically adapt management practices, optimizing fertilizer use and pest control while minimizing environmental impact [16]. Cooperative group-based solutions, such as those presented by García et al. [17], reduce energy consumption and improve communication efficiency in WSNs, increasing the viability of these platforms in rural areas with limited infrastructure.

Advanced Deep Learning techniques have been successfully applied to environmental monitoring, enabling the early detection of anomalous patterns in air quality, humidity, and temperature [18]. This automated analysis facilitates the integration of predictive systems within digital twins, enhancing responsiveness to emerging weather events or pest outbreaks. Nevertheless, challenges remain, such as data interoperability, real-time synchronization, and affordability for smallholder farmers, which currently limit widespread adoption [19].

III. SYSTEM PROPOSAL

This section presents a proposed system for implementing local weather forecasting based on artificial intelligence within a digital twin environment.

A. System Description

The proposed system is based on weather stations developed on the ESP32 platform, as illustrated in Figure 1. These stations are equipped with sensors capable of measuring key meteorological variables such as air temperature, relative humidity, and atmospheric pressure.

The data captured by the sensors is transmitted and managed through a server implemented with Node.js, which allows for their storage, processing, and subsequent use for modeling. The locally obtained meteorological variables constitute the inputs to the prediction model based on a machine learning models.

The management of weather forecast requests and the visualization of real-time data and model results is carried out through a user interface developed as part of an application for monitoring agricultural fields. This application serves a dual purpose: to facilitate interaction with the system and as a key component in creating a digital twin of the agricultural environment.



Figure 1. Photograph of the weather station.

B. Operating Algorithm

The weather station's operating algorithm operates in a loop with an execution frequency of once per hour. After this interval, the node automatically connects to the wireless network, establishes communication with the sensors using the Inter-Integrated Circuit (I2C) or Universal Asynchronous Receiver-Transmitter (UART) protocols, and measures the meteorological variables.

Once the data is obtained, the system attempts to connect with the server. If the connection is successful, the data is sent, and the node again enters a standby state until the next iteration of the cycle.

Regarding the operation of the digital twin, the interaction begins when the user clicks the forecast button on the interface associated with a specific weather station. This action activates the weather prediction model, which generates a forecast for the next 24 hours.

Additionally, the data collected daily by the weather stations is used to update the model. In this way, a new model version is trained daily, ensuring that predictions are based on the most recent local weather records, thereby increasing the system's accuracy and adaptability to environmental conditions. This data is used for input for digital twins.

C. Dataset Description

The dataset used in this work corresponds to a time series of local meteorological observations recorded from February 1, 2025 to August 1, 2025. During this period, atmospheric conditions were assessed intermittently over a specific agricultural plot. The dataset is composed of the following variables:

- Air temperature (°C)
- Relative humidity (%)
- Atmospheric pressure (hPa)
- Date
- Time

Each observation is associated with a timestamp (date and time) indicating the exact moment the measurement was taken. These variables constitute the basis for generating weather forecasts in the proposed system, allowing for modeling short-term local climate evolution within the context of the digital twin.

D. Model Training and Testing

The weather forecast model was generated following the steps illustrated in Figure 2.

First, the dataset is loaded into a DataFrame, where the temperature, relative humidity, and atmospheric pressure variables are converted to numeric format. The date and time column is transformed into datetime objects and sorted in natural chronological sequence.

Next, a data cleaning process is carried out. Invalid samples are eliminated, including those generated during testing (e.g., records with a frequency of one second) and those with temporal discontinuities greater than one hour.

Once the dataset is cleaned, an exploratory analysis of the variables of interest is performed. To do this, histograms are used to visualize the distribution of the variables extracted from the database.

Subsequently, the input structure for the model is built based on a sliding data window. The input to the forecasting models was defined using a sliding-window approach over the time series. Each input sequence consisted of a window of 24 consecutive hourly samples (corresponding to a 24-hour period), with a stride of 3 hours between windows. This resulted in partially overlapping input sequences, capturing both short-term dynamics and daily patterns. The sampling frequency within each window was fixed at one hour, and the target variable was defined as the meteorological condition in the subsequent 24 hours. Thus, each input had a shape of (24, number of features), where the features corresponded to air temperature, relative humidity, and atmospheric pressure. With this structure, the input tensor is generated, on which statistical analysis is performed to obtain:

- The tensor dimension.
- The total number of elements.
- The number and percentage of null values.
- The percentage of valid data.
- The ranges of values before normalization.
- The identification of potential outliers.

Before training the models, the variables are normalized, and the correlation matrix between them is analyzed to identify relevant relationships between the different meteorological variables.

Different machine learning and deep learning algorithms are trained and comparatively evaluated in the final stage. In

the machine learning field, the following models were considered:

- Random Forest
- Gradient Boosting
- Ridge Regression
- Lasso Regression
- K-Nearest Neighbors

The primary evaluation metric employed in this study is the Coefficient of Determination (R^2), which provides a measure of how well the predicted values approximate the actual data. In addition to R^2 , the Mean Squared Error (MSE) and Mean Absolute Error (MAE) are also analyzed to gain deeper insights into the model's predictive performance. These complementary metrics allow us to assess aspects such as robustness, generalization capability, and potential overfitting. By considering multiple evaluation criteria, we ensure a more comprehensive understanding of the model's behavior across different scenarios.

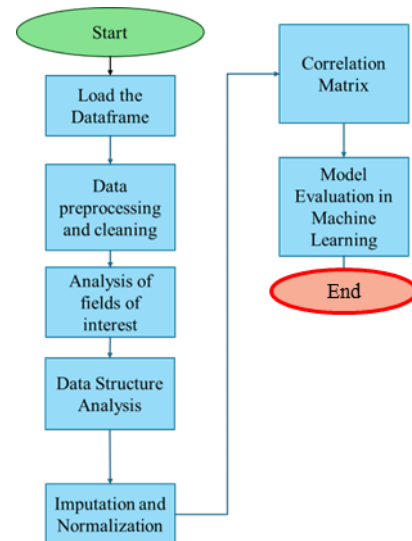


Figure 2. Flowchart of the artificial intelligence model selection, training, and testing process.

E. Computational Tools and Libraries

The data processing, analysis, and modeling pipeline was implemented in Python 3.10. For data handling and preprocessing, the libraries Pandas and NumPy were employed to manage time series, perform data cleaning, and generate sliding windows for model input. Exploratory data analysis and visualization were conducted using Matplotlib and Seaborn, which allowed the inspection of statistical distributions, correlations, and temporal trends.

Traditional machine learning models, including Random Forest, Gradient Boosting, Ridge, Lasso, and K-Nearest Neighbors, were implemented with Scikit-learn, which was also used to compute the performance metrics (R^2 , MSE, and MAE).

IV. RESULTS

This section will detail the results obtained from the training and testing of the different models proposed for the meteorological dataset.

A. Histogram of dataset variables

Figure 3 presents the frequency distributions of the temporal and meteorological variables in the dataset. The temporal variables (minutes, day, and month) display distinct sampling patterns, with the distribution of minutes being nearly uniform across the recorded range, the distribution of days showing alternating frequencies, and the distribution of months indicating data collection concentrated in two main periods.

Regarding the meteorological variables, the air temperature histogram shows a slightly left-skewed distribution, with most observations ranging between 27 °C and 30 °C, and fewer occurrences at extreme values. Relative humidity exhibits a bimodal distribution, with peaks near 70–75 % and around 100 %, indicating frequent saturation events. Atmospheric pressure values are clustered mainly between 1012 hPa and 1016 hPa, following a near-normal distribution with moderate variability.

These histograms provide insight into the statistical characteristics of the dataset, highlighting the predominant conditions recorded by the meteorological station and potential patterns relevant for model training.




Figure 3. Histograms of the weather and time variables from the dataset.

B. Correlation Matrix

Figure 4 illustrates the relationships between the temporal and meteorological variables in the dataset. When considering the absolute values of the correlation coefficients, the highest associations are observed between relative humidity and both the day and month variables ($|r| = 0.60$ and 0.51 , respectively) negative in sign and between atmospheric pressure and the same variables ($|r| = 0.53$ and 0.51 , respectively) positive in sign. Air temperature also shows a relatively strong correlation ($|r| = 0.63$) with the minute of the day, indicating a marked diurnal pattern. These high-magnitude correlations, regardless of their direction, suggest that temporal factors exert a significant influence on the measured meteorological variables, and these relationships can be exploited to improve the performance of the forecasting models.

These results indicate the presence of significant relationships among certain variables, which can be exploited to enhance the training process of the artificial intelligence model for local weather forecasting.




Figure 4. Correlation matrix of temporal and meteorological variables.

C. Machine learning analysis

Figure 5 compares the results of the different models for temperature prediction. Gradient Boosting achieved the best overall performance among the classical machine learning methods, with the lowest MAE (0.37) and MSE (0.30), and the highest R^2 (0.656). K-Nearest Neighbors also provided competitive results ($R^2 = 0.640$), followed by Lasso and Ridge regressions with similar performance levels ($R^2 \approx 0.60$). Random Forest yielded the lowest R^2 (0.590) and slightly higher error values.

However, Figure 6 reveals that, despite the relatively high R^2 value, the Gradient Boosting model exhibits substantial dispersion in the test set predictions, not only at extreme temperatures but also in central value ranges. This discrepancy between training and testing performance suggests overfitting, indicating that the model may be capturing noise or dataset-specific patterns rather than generalizable relationships. This limitation will later be addressed in the deep learning experiments, where regularization techniques are introduced to reduce overfitting and improve generalization.




Figure 5. Performance comparison of classical machine learning models for air temperature prediction, evaluated using MSE, MAE, R^2 .




Figure 6. Predicted versus actual air temperature values for the Gradient Boosting model, including both training and testing datasets, compared to the ideal prediction line.

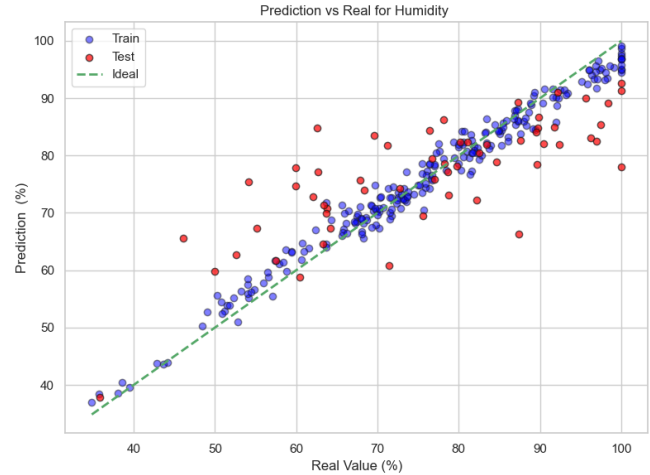


Figure 7. Predicted versus actual relative humidity values for the Gradient Boosting model, including both training and testing datasets, compared to the ideal prediction line.

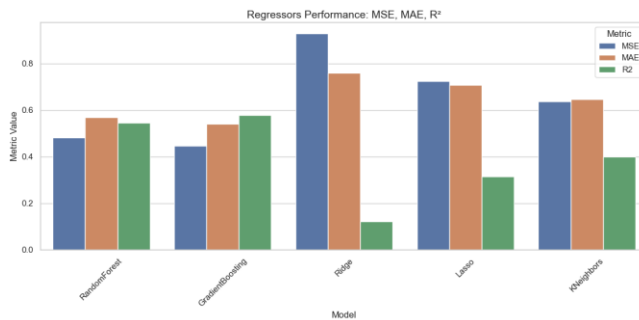


Figure 8. Performance comparison of classical machine learning models for relative humidity prediction, evaluated using MSE, MAE, R^2 .

Figure 7 presents the performance comparison for relative humidity prediction. As in the case of temperature, Gradient Boosting outperformed the other models, achieving the highest R^2 (0.578) and the lowest error metrics (MAE = 0.54, MSE = 0.45). Random Forest showed slightly lower performance ($R^2 = 0.546$), while Ridge Regression performed poorly ($R^2 = 0.123$), indicating difficulty in modeling the underlying relationships. Lasso Regression and K-Nearest Neighbors produced intermediate results.

Figure 8 depicts the predicted versus actual humidity values for the Gradient Boosting model, showing a strong alignment with the ideal prediction line for most observations, though with higher dispersion in the test set, particularly at mid-range humidity levels.

Figure 9 compares the performance of the evaluated machine learning models for atmospheric pressure prediction. Gradient Boosting achieved the highest coefficient of determination ($R^2 = 0.899$) and the lowest error metrics (MAE = 0.26, MSE = 0.10), demonstrating its superior accuracy and generalization capacity. Random Forest also performed well, with an R^2 of 0.845 and slightly higher errors (MAE = 0.34, MSE = 0.16). In contrast, Ridge Regression and Lasso Regression yielded moderate results ($R^2 = 0.618$ and 0.593 , respectively), while K-Nearest Neighbors showed the weakest performance ($R^2 = 0.457$).

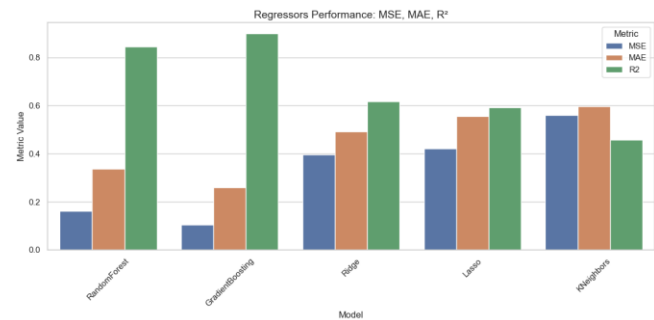


Figure 9. Performance comparison of classical machine learning models for atmospheric pressure prediction, evaluated using MSE, MAE, R^2 .

Figure 10 illustrates the predicted versus actual atmospheric pressure values using the Gradient Boosting model. The predictions closely follow the ideal line for both training and testing datasets, with minimal dispersion, particularly in the test set. These results confirm that Gradient Boosting provides the most reliable predictions for atmospheric pressure in the given dataset, outperforming all other tested models.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a system successfully integrates artificial intelligence into a digital twin for local weather forecasting in agricultural environments. The results demonstrate that Gradient Boosting offers the most accurate predictions among the classical machine learning models for all three meteorological variables, with particularly strong performance in atmospheric pressure forecasting.

These outcomes validate the feasibility of combining IoT-based sensing infrastructure with advanced predictive models to enhance decision-making in precision agriculture. By providing localized and timely forecasts, the system can support improved resource allocation, crop management, and environmental sustainability.

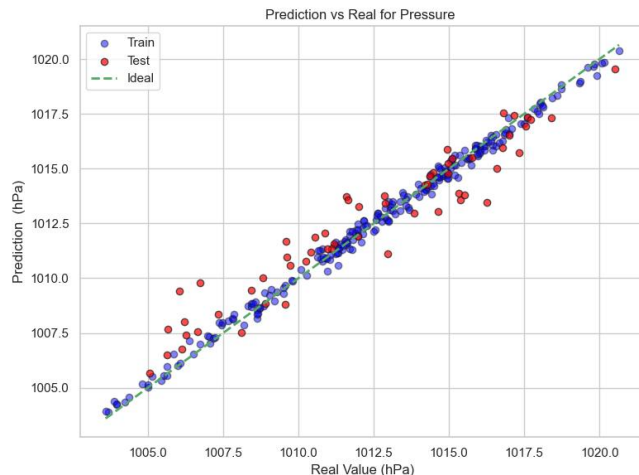


Figure 10. Predicted versus actual atmospheric pressure values for the Gradient Boosting model, including both training and testing datasets, compared to the ideal prediction line.

Future work will focus on expanding the dataset to cover multiple seasons and diverse climatic conditions, integrating additional variables such as wind speed, solar radiation, and soil moisture. Model optimization will include hyperparameter tuning, advanced regularization techniques, and the exploration of hybrid architectures that combine statistical and neural approaches. Furthermore, deploying the forecasting system in real-time operational scenarios and integrating it with automated control mechanisms in the digital twin will be key steps towards its practical adoption in smart farming applications.

ACKNOWLEDGMENT

This work is funded by the Ministry of Science, Innovation and Universities through the State Research Agency (AEI) with the projects PID2020-114467RR-C33/AEI/10.13039/501100011033, TED2021-131040B-C31.

REFERENCES

- [1] L. Wang, "Digital twins in agriculture: a review of recent progress and open issues," *Electronics*, vol. 13, no. 11, p. 2209, 2024.
- [2] J. Lloret, M. Garcia, D. Bri, and J. R. Diaz, "A cluster-based architecture to structure the topology of parallel wireless sensor networks," *Sensors*, vol. 9, no. 12, pp. 10513–10544, 2009.
- [3] A. Hussein, A. Challoor, and M. Yousif, "Optimizing network topologies for enhanced communication and scalability in distributed networks," *J. Sci. Res. Teach.*, vol. 3, no. 11, 2024.
- [4] W. Purcell, T. Neubauer, and K. Mallinger, "Digital twins in agriculture: challenges and opportunities for environmental sustainability," *Curr. Opin. Environ. Sustain.*, vol. 61, p. 101252, 2023.

- [5] P. Angin, M. H. Anisi, F. Göksel, C. Gürsoy, and A. Büyükgülcü, "Agrilora: a digital twin framework for smart agriculture," *J. Wirel. Mob. Netw. Ubiquitous Comput. Dependable Appl.*, vol. 11, no. 4, pp. 77–96, 2020.
- [6] N. Peladarinos, D. Piromalis, V. Cheimaras, E. Tserepas, R. A. Munteanu, and P. Papageorgas, "Enhancing smart agriculture by implementing digital twins: A comprehensive review," *Sensors*, vol. 23, no. 16, p. 7128, 2023.
- [7] H. Zhang, Y. Liu, C. Zhang, and N. Li, "Machine learning methods for weather forecasting: A survey," *Atmosphere*, vol. 16, no. 1, p. 82, 2025.
- [8] R. Ahsen et al., "Harnessing digital twins for sustainable agricultural water management: A systematic review," *Appl. Sci.*, vol. 15, no. 8, p. 4228, 2025.
- [9] M. Awais, X. Wang, S. Hussain, F. Aziz, and M. Q. Mahmood, "Advancing precision agriculture through digital twins and smart farming technologies: A review," *AgriEngineering*, vol. 7, no. 5, p. 137, 2025.
- [10] A. Khanna and S. Kaur, "Evolution of Internet of Things (IoT) and its significant impact in the field of precision agriculture," *Comput. Electron. Agric.*, vol. 157, pp. 218–231, 2019.
- [11] N. Ariesen-Verschuur, C. Verdouw, and B. Tekinerdogan, "Digital twins in greenhouse horticulture: A review," *Comput. Electron. Agric.*, vol. 199, p. 107183, 2022.
- [12] J. A. López, A. J. Garcia-Sanchez, F. Soto, A. Iborra, F. Garcia-Sanchez, and J. Garcia-Haro, "Design and validation of a wireless sensor network architecture for precision horticulture applications," *Precis. Agric.*, vol. 12, no. 2, pp. 280–295, 2011.
- [13] D. Bri, M. Garcia, J. Lloret, and P. Dini, "Real deployments of wireless sensor networks," in *Proc. 2009 3rd Int. Conf. Sensor Technol. Appl. (Sensorcomm 2009)*, Athens, Greece, Jun. 2009, pp. 415–423.
- [14] M. M. Sakha, F. Daiber, C. Tieben, and M. Enders, "Virtual breeding nursery: Towards a VR digital twin for plant breeding," in *Proc. 2024 ACM Symp. Spatial User Interact.*, 2024, pp. 1–2.
- [15] O. Spyrou, M. Ariza-Sentís, and S. Vélez, "Enhancing education in agriculture via XR-based digital twins: A novel approach for the next generation," *Appl. Syst. Innov.*, vol. 8, no. 2, p. 38, 2025.
- [16] M. Awais, X. Wang, S. Hussain, F. Aziz, and M. Q. Mahmood, "Advancing precision agriculture through digital twins and smart farming technologies: A review," *AgriEngineering*, vol. 7, no. 5, p. 137, 2025.
- [17] M. Garcia, S. Sendra, J. Lloret, and A. Canovas, "Saving energy and improving communications using cooperative group-based wireless sensor networks," *Telecommun. Syst.*, vol. 52, no. 4, pp. 2489–2502, 2013.
- [18] Y. Zhang, J. Li, X. Wang, and P. Zhao, "Deep learning-based environmental monitoring: A survey," *Environ. Model. Softw.*, vol. 156, p. 105495, 2022.
- [19] S. Cesco, P. Sambo, M. Borin, B. Basso, G. Orzes, and F. Mazzetto, "Smart agriculture and digital twins: Applications and challenges in a vision of sustainability," *Eur. J. Agron.*, vol. 146, p. 126809, 2023.
- [20] A. A. Hapsari, D. J. Vresdian, M. Aldiansyah, B. W. Dionova, and A. C. Windari, "Indoor air quality monitoring system with Node.js and MQTT application," in *Proc. 2020 1st Int. Conf. Inf. Technol., Adv. Mech. Electr. Eng. (ICITAMEE)*, Yogyakarta, Indonesia, 2020, pp. 144–149.