Bridging Domains: AI-Driven Agnostic Reconstruction Frameworks for Irregular High-Dimensional Data (BRAID)

Jan Kieseler (speaker) KIT, **Nicolas R. Gauger** RPTU,
Wolfgang Rhode TU Dortmund, **Christian Glaser** Uppsala (associated),
Anastasios Belias GSI,

Abstract

Recent breakthroughs in artificial intelligence (AI) have demonstrated transformative power across a wide spectrum of applications, from object classification in computer vision to generative modeling of images and natural language. This project seeks to leverage cutting-edge AI techniques to revolutionize object detection (particle reconstruction) within particle physics and related domains. By developing a flexible, detector-agnostic framework capable of handling irregular, high-dimensional data, we aim to enhance scientific discovery and data-analysis efficiencies in experiments such as at the HL-LHC, FAIR, and IceCube-Gen2. We will bridge expertise from high-energy, hadron, nuclear, and astroparticle physics with computer science to create open-source datasets, domain-specific evaluation metrics, and robust AI-based reconstruction architectures. Emphasizing sustainability, interdisciplinary collaboration, and open science, the proposed methods adhere to physics constraints and incorporate advanced dimensionality reduction strategies for large-scale detector data. Our approach promises broader applicability beyond particle physics, with potential benefits for medical imaging and industrial monitoring, ultimately fostering paradigm changes in future detector design while aligning with BMBF ErUM-Data objectives.

Introduction and Objectives

Advancements in artificial intelligence (AI) have repeatedly transformed entire fields, example being the revolution in object classification and detection in computer vision as well as the rather novel breakthrough of generative modelling in image, video and natural language. Inspired by the massive transformative power of AI, we aim to facilitate a paradigm shift for object detection (particle reconstruction) in particle physics and related domains. By leveraging cutting-edge AI techniques, this project has the potential to significantly enhance data analysis capabilities, leading to unprecedented scientific discoveries and efficiencies, and pave the way to paradigm changes in future detector design. The proposed project aims to develop a flexible and diverse detector-agnostic reconstruction framework for irregular high-dimensional data by leveraging advanced artificial intelligence (AI) and machine learning (ML) techniques. By bridging multiple domains—high-energy physics (HEP), hadron and nuclear physics (HaN), astroparticle physics (Astro), and computer science (CS)—we seek to create innovative concepts, public datasets, and frameworks that are universally applicable across different experimental setups. Our primary objective is to address the challenges posed by the increasing complexity and volume of data generated by modern and future detectors, such as those at the High-Luminosity Large Hadron

References marked with † indicate contributions from a consortium member; references marked with ‡ indicate joint publications, with contributions from at least two consortium members.
Generative AI has been used for editorial purposes when creating this document.

Collider (HL-LHC), the international Facility for Antiproton and Ion Research (FAIR), and the IceCube Neutrino Observatory. Just as the shift from handcrafted feature extraction to convolutional neural networks (CNNs) revolutionized computer vision by unlocking unprecedented potential through automated feature learning [**1**, **2**, **3**], we aim to bring a similar paradigm shift to data reconstruction in physics experiments. CNNs have enabled significant advancements in image recognition by learning complex patterns directly from raw data [**4**], leading to breakthroughs in various applications discussed e.g. in Refs. [**4**, **5**, **6**, **7**]. At the core of CNNs are translation equivariant kernels that learn local features, combined with dimensionality reduction and abstraction. We believe that adopting analogous AI techniques in our domains will similarly transform the way we process and interpret experimental data.

To achieve this, we propose the following objectives:

- **Develop Detector-Agnostic Reconstruction Methods:** Create new graph or transformer neural network based architectures that adhere to physics constraints, particularly spatial and temporal locality. We aim to efficiently handle irregular, high-dimensional data common in modern detectors.
- **Implement Adaptive Dimensionality Reduction Techniques:** Design algorithms that reduce the dimensionality of data dynamically without violating locality requirements. This will facilitate the reduction of a large number of detector signals into reconstructed particles used for further processing.
- **Establish Common Frameworks and Public Datasets:** Generate open-source datasets for community-wide use and beyond alongside the domain-specific evaluation metrics to evaluate the performance of the algorithms. This will promote standardization, facilitate benchmarking, and encourage collaboration across different domains.
- **Foster Interdisciplinary Collaboration:** Strengthen connections between physics domains and CS by incorporating expertise from all consortium partners. This collaboration will promote knowledge transfer, ensure robustness of the developed methods, and optimize them for practical deployment.
- **Contribute to Open Science and Sustainability:** Adhere to FAIR principles and open-access policies to ensure wide accessibility of our results. By focusing on efficient hardware utilization and potential reductions in energy consumption (Watts per event), we aim to contribute to sustainable computing practices in scientific research.

By achieving these objectives, the project will not only enhance the physics reach of current and future experiments but also address the significant computing challenges associated with processing and analysing large-scale, complex datasets. The development of detector-agnostic methods will have broad applicability, potentially benefiting other domains that deal with irregular high-dimensional data, such as medical signal processing and imaging and industrial monitoring. Although the resulting algorithms have the potential to become part of the real-time selection systems (triggers) in the experiments, in this first step we will focus on less stringent runtime constraints imposed by processing data “offline” that has already been recorded. This project aligns with the goals of the BMBF ErUM-Data call, focusing on software and algorithms with an emphasis on AI and ML, research data management, and federated digital infrastructures. By bridging multiple communities and fostering interdisciplinary collaboration, we aim to develop solutions that go beyond the capabilities of individual domains, thereby maximizing scientific impact and promoting sustainable, open science practices. This project is not within the scope of NFDI, EOSC, or WLCG.

State of the Art

In the high-luminosity phase of the HL-LHC, experiments like ATLAS and CMS aim to push the boundaries of our understanding of fundamental particles and their interactions. Among the key objectives are precision measurements of the Higgs boson's properties and its couplings to other particles, extensive searches for new physics phenomena beyond the Standard Model (e.g., supersymmetric particles or additional spatial dimensions), and improved constraints on models of dark matter production. By collecting more than an order of magnitude more data than previous runs, the HL-LHC will enable physicists to explore subtle effects in rare processes, probe higher energy scales, and potentially uncover entirely new sectors of particle physics.

Experiments like the IceCube Neutrino Observatory focus on understanding the high-energy universe by detecting cosmic neutrinos. Neutrinos, as messengers of astrophysical processes, offer insights into the origins of ultra-high-energy cosmic rays and the nature of extreme astrophysical accelerators such as blazars, active galactic nuclei, and supernova remnants. The physics goals include, charting the energy spectrum of astrophysical neutrinos, studying neutrino flavor composition and oscillations over astronomical scales and pinpointing individual cosmic neutrino sources. By elucidating where and how nature's most powerful accelerators operate, IceCube and forthcoming detectors such as KM3NeT, Baikal-GVD, P-ONE and IceCube-Gen2 aim to answer long-standing questions about the birthplaces of the universe's highest-energy particles. Furthermore, cosmic sources provide a beam of neutrinos at energies thousandfold higher and over vaster distances than accelerator experiments, allowing us to probe fundamental particle physics in ways not otherwise possible

At facilities like FAIR, experiments such as PANDA and CBM are designed to probe the behaviour of strongly interacting matter under extreme conditions. By colliding heavy ions at high energies, researchers create hot and dense environments that mimic the early universe moments after the Big Bang or the interiors of neutron stars. The overarching goals include determining how matter transitions between different phases—such as hadronic matter and the quark-gluon plasma—exploring exotic states of matter like glueballs, and clarifying the properties and dynamics of the strong force. Through these studies, FAIR experiments aim to deepen our understanding of quantum chromodynamics, the origin of mass, and the fundamental structure of matter itself.

The detectors in all three domains share the common goal of extracting physics information by analysing final-state particles. These particles, produced by the fundamental processes under study, remain stable over the timescales of their interaction with the detector. Each detector comprises thousands to millions of sensors distributed across various subsystems, which either directly detect the particles or record the signals left as the particles traverse materials within the detector. Depending on the technology, these sensors can provide outputs ranging from simple binary data to intricate pulse shapes. Arranged in three-dimensional patterns that rarely conform to grid structures, these sensors necessitate an initial processing step: identifying the final-state particles responsible for the observed pattern of sensor responses.

Depending on the domain, there are specific challenges attached to the pattern recognition problem. For the HL-LHC data, the unprecedented collision rates and resulting event complexity lead to an environment where hundreds of thousands of detector hits must be associated to final-state particles in a single collision (event). These hits stem from multiple detector subsystems using very different technologies. For example the tracking system aims to measure trajectories without changing the particle momenta, and the calorimeters aim to measure the particle energy by stopping them entirely. In addition, multiple simultaneous particle interactions (up to 200) often overlap, making it challenging to disentangle signals belonging to different particles or interactions. Moreover, the increase in data complexity and size puts significant pressure on the world-wide LHC computing infrastructure, demanding new solutions [8]. For IceCube and other neutrino observatories, the irregular geometry and directional acceptance of the sensors embedded in natural and inhomogeneous media like glacier ice, the high data dimensionality of time series per sensor, and the large stochasticity of energy deposits within the cubic-kilometre-big detector volume complicate the reconstruction of the neutrino's direction and energy. In HaN experiments at FAIR, the wide variety of beam energies, target materials, and dynamically changing detector configurations demand flexible reconstruction algorithms that can adapt to evolving conditions while maintaining high efficiency and accuracy.

The algorithms currently used in HEP and HaN have their roots in techniques developed decades ago when data complexity and volume were orders of magnitude smaller [9, 10, 11]. Classical methods rely on predefined sequences of threshold-based filters and seed-finding procedures, leading to the premature discarding of valuable information and requiring extensive manual tuning. These algorithms typically exploit symmetries and are strictly local. Maintaining strict locality ensures that the calibration of the simulation against real data remains tractable, as it relies on the independence of widely separated parts of the detector. While incremental improvements have been introduced over time, these algorithms struggle to keep pace with complexity and computing demands, especially as the HL-LHC and other new facilities drastically increase data rates. In particular with the rise of timing capabilities of the sensors, but also in general, the inflexibility of these algorithms in adapting to changing detector conditions further strain resources, both in terms of computing time and power consumption, but also due to the fact that these algorithms need to be tuned by hand, resulting in person-power intense, yet slow and suboptimal tuning; a tuning that is performed for stable beam conditions and static detector configurations. For Cherenkov-based neutrino observatories such as IceCube, recent years have shown that classical likelihood approaches have reached their limits: systematic effects and high-dimensional data have led to a saturation in angular and energy resolution for all event classes. Multi-PMT sensors, as e.g. used by KM3Net and planned for IceCube-Gen2, will add additional data dimensionality by providing 31 (in case of KM3Net) directional PMTs per sensor compared to just one in the case of IceCube, and will require more complex algorithms to exploit the additional information. A futher improvement of the angular resolution is particularly interesting as IceCube see several 3-sigma excesses. Hence, an improved reconstruction can lead to the discovery of new cosmic neutrino sources.

Recent developments in ML-based reconstruction are beginning to address these problems. In HEP, GravNet [12†], a graph neural network (GNN), with Object Condensation [13†] as a training formalism have demonstrated their versatility, achieving success in environments ranging from CMS-like calorimeter data [14†, 15†] to Belle II calorimetry and tracking [16] and future collider environments such as FCC [17†]. These architectures represent detector hits as points in a learned latent space, allowing dynamic neighbourhood building and flexible clustering without rigid assumptions about detector geometry. A similar, yet less rigorous approach is followed by MLPF [18], where classic clustering is performed in subdetectors and then the information is combined using ML. Methods like HGPflow [19] can estimate particle multiplicities either within jets or localized regions of the detector, and then assign properties to them using hypergraphs. In astroparticle physics, CNNs have been applied to IceCube data by approximating detector layouts on coarse grids, surpassing traditional methods for certain event topologies [20†]. Furthermore, integration of neural network surrogate models into a Likelihood reconstruction led to further improvements [21†], which has already led to the recent discoveries of neutrino emission from the galactic plane [22†].

All these techniques leverage industry-standard software stacks and are highly parallelizable on GPUs, offering the potential for better physics performance per Watt and per second of processing time, thereby enhancing both efficiency and sustainability.

However, despite these advances, some critical challenges remain unresolved. Many existing algorithms in HEP still rely on global operations (e.g., counting total particles upfront or segmenting the detector into patches), which conflicts with the stringent locality requirements demanded by precise physics calibration and factorization assumptions or the necessity to avoid possible non-physical artifacts at the patch edges. Other ML approaches, such as MLPF or GravNet with object condensation, either still rely on classical methods as a first step or defer dimensionality reduction until the late stages of inference, missing opportunities for early, efficient filtering that could dramatically improve speed and resource utilization. This also applies to the advanced architectures under study at IceCube, which face severe scalability issues, and innovative strategies are needed that maintain interpretability and manage computational costs while providing maximally robust results. For IceCube, ML methods have so far only been successful for cascade topologies. However, an application of the established ML methods to track topologies, which are the driving events for identifying neutrino sources and IceCube's real-time alert stream, has so far been unsuccessful, and it seems clear that cascades have been the "low-hanging fruit". Furthermore, the sensor positions of in-water detectors are constantly changing, requiring algorithms that take the current positions dynamically into account, which requires going beyond CNNs. An application of network architectures such as Transformers and Graph Neural Networks tailored to the needs in Cherenkov-based neutrino observatories, and not just CNNs, is expected to yield further improvements on all event topologies and have shown promising initial results (e.g. [**23**, **24**, **25**†] plus initial unpublished work by the applicants).

To overcome these limitations, we need local dynamic dimensionality reduction methods that can smoothly compress information at an early stage, guided solely by local densities and detector signals. Implementing such capabilities will not only improve "offline" data processing - processing data already recorded - but also opens the door to deploying these techniques in real-time online trigger systems, thereby enabling more selective and efficient event filtering at the very first level of data selection. In the case of neutrino astronomy, this will enable improved real-time event classification and reconstruction to alert the astronomical community for follow-up observations. Moreover, moving beyond static, hand-tuned algorithms toward ML-based methods that are designed to be more generic and can be developed to respond automatically to varying conditions, offers a huge potential for complex and dynamic detector environments such as the FAIR experiments or future neutrino observatories in water.

Off-the-shelf techniques from CS cannot close this gap. The current state-of-the art ML techniques used for object detection rely on CNN architectures, recently often combined with transformer architectures, e.g. [**26**, **27**, **28**]. The data from detectors is typically sparse and irregular, and contains up to hundred-thousands of inputs, ruling out this otherwise successful combination. As an alternative, GNNs have shown promise for processing complex relational data in areas like chemistry, combinatorial optimization, and autonomous driving. For the point-cloud like data from detectors, the GNN architectures are required to also perform a graph topology building, creating connections between the points. Methods like SAGPooling [**29**], EdgePool [**30**], DiffPool [**31**], and architectures inspired by U-Net or Graph U-Net [**32**] can produce hierarchical representations and handle multi-scale structures. Similarly, PointNet++ [**33**], EdgeConv [**34**], and GravNet [**12**†] provide ways to dynamically build and rebuild topologies.

Yet, these techniques typically assume simpler forms of locality and do not address the non-differentiability of discrete thresholding operations at the heart of dimensionality reduction, making them less suitable for complex, structured transformations especially in cases with obsolete or noisy information. As these networks are trained with gradient-descent techniques, any operation that should be learnable must be implemented in a differentiable way. Current state-of-the-art techniques also do not incorporate domain-specific constraints like those seen in particle detectors, where local invariances and symmetries must be preserved. To meet the demands of current and next-generation detectors, we need new ML solutions that integrate dynamic dimensionality reduction, locality, and physics constraints within a single differentiable framework. Instead of relying on global operations or fixed thresholds, architectures must learn to perform pooling and coarsening based solely on local information- such as density-based criteria that can be differentiated smoothly.

Achieving this will require developing new layers or modules that ensure the number of outputs emerges naturally from the data, and that the network respects data symmetries (equivariances) and locality. These developments must be made in close contact with and adhering to the domain-specific requirements of HEP, Astro, and HaN detectors, while containing generic enough concepts to serve the whole community. To stimulate cross-domain advancements, public datasets in widely accepted formats are essential to ensure broader recognition, as demonstrated by the limited impact of datasets like Ref. [**19**]. These datasets must include relevant annotations, such as physics constraints and event characteristics, and provide clear success metrics for algorithm evaluation—key elements that enabled transformative efforts like AlphaFold(2).

The work proposed here explicitly tackles these challenges, focusing on offline applications first to develop concepts with less stringent resource constraints. In addition, we intend to utilize the developed algorithms to intraoperative neurophysiological monitoring (IONM) through our associated industry partner inomed, demonstrating the applicability of our methods to other domains. IONM is vital for monitoring the functional integrity of nerves during surgeries, avoiding temporary or persistent damage resulting in neurological deficits. Classical manual interpretation of the signals is prone to inter-expert variability as well as human errors, leading to a flourishing research landscape proposing AI solutions to support surgical staff [**35**, **36**†]. Many solutions for analysing continuous IONM signals rely on recurrent architectures, which maintain "memory" of previous signals. However, these approaches often struggle to capture long-term dependencies effectively and are not well-equipped to handle signal artifacts or outliers that may corrupt the hidden state of the network.

We aim to incorporate the advanced message-passing architectures and hierarchical dimensionality reduction techniques developed in our consortium into the signal processing pipeline. By a re-interpretation of spatial locality and dynamic reduction to temporal locality and reduction, we aim to develop new methods that are more resilient to noise and signal variability, while exploiting existing long term and hierarchical structures in IONM data, that are unaccounted by previous approaches. Furthermore, by interpreting the reduction and learned locality patterns, this approach opens up new possibilities for explainable AI solutions, offering insights at microscopic and macroscopic levels.

Expertise of the Consortium Partners

KIT: Jan Kieseler has established himself as a leader in machine learning (ML)-based particle reconstruction in high-energy physics (HEP), with contributions that have significantly advanced the field. His development of the DeepJet algorithm, a deep learning-based approach to jet identification [**37**†], has enabled the CMS experiment to enhance the sensitivity of key physics searches and measurements - in some cases equivalent to decades of additional data taking. Its wide adoption has already improved over 300 physics analyses and set a new standard for ML applications in HEP. As the co-founder of the CMS ML4RECO initiative, Dr. Kieseler has also pioneered efforts to integrate ML into particle reconstruction workflows. His state-of-the-art algorithms, such as GravNet and Object Condensation, represent the cutting edge of ML-based particle reconstruction in HEP. These methods have demonstrated remarkable versatility, with successful applications ranging from CMS-like data to Belle II and future collider studies such as the FCC. In addition, Dr. Kieseler has held pivotal leadership roles within CMS, including as convener of the TOP physics group and the Upgrade Physics Studies group, overseeing data analyses of up to 100 scientists. With extensive experience in detector upgrades and a deep understanding of particle physics, Dr. Kieseler is uniquely positioned to address the challenges outlined in this proposal, combining domain expertise with cutting-edge computational techniques.

RPTU: Nicolas R. Gauger is a leading expert in Scientific Computing, with a strong focus on the mathematical foundations of machine learning and artificial intelligence. He is a full professor for Computer Science as well as Applied Mathematics at RPTU Kaiserslautern-Landau, where he currently holds the Chair for Scientific Computing and serves as Director of the Computing Center (RHRZ). Prof. Gauger's research is deeply rooted in the application of advanced mathematical techniques to real-world problems, with significant contributions to turbulence modelling in aerospace, as well as the application of machine learning (ML) in various domains, including high-energy physics [38†, 39†, 40‡], climate science, and medicine. His work has provided innovative solutions for data-driven modelling of complex physical systems, particularly in the context of turbulent flows, which requires sophisticated AI-driven methods for accurate predictions. His recent research on the mathematical foundations of machine learning, particularly in the training of neural networks, is highly relevant to the proposed project. Prof. Gauger's work in this area aims to refine and optimize ML algorithms, ensuring their robustness and efficiency when applied to complex, high-dimensional data—precisely the kind of challenge the project seeks to address in the context of high-energy physics. His expertise in developing new algorithms and optimizing their performance in computationally demanding applications positions him well to contribute to the development of detector-agnostic reconstruction methods for irregular high-dimensional data, as outlined in the project proposal. Prof. Gauger's deep expertise in both the theoretical and applied aspects of machine learning, combined with his interdisciplinary approach, will be valuable in helping achieve the project's objectives. His involvement will not only advance the development of cutting-edge reconstruction algorithms but also foster collaboration across diverse domains, contributing to the broader goals of the proposal, such as establishing open-source frameworks and promoting sustainable computing practices in scientific research.

TU-Dortmund: Wolfgang Rhode is a distinguished physicist specializing in astroparticle physics at TU Dortmund University. His research focuses on understanding the origins of high-energy cosmic rays through the detection and correlation of cosmic messenger particles using Cherenkov telescopes (MAGIC, FACT, CTA/LST), neutrino telescopes (AMANDA, IceCube/GEN2), and radio telescopes (MeerKAT, GLOW). By developing theoretical models, his team calculates expected signatures in these telescopes and creates simulation programs like PROPOSAL and contributes to CORSIKA to identify these processes. A significant aspect of Prof. Rhode's work involves advancing data analysis methods, particularly machine learning (ML) within artificial intelligence. Collaborating with computer scientists in the Collaborative Research Center 876 and as Area Chair Physics at the Lamarr Institute for Machine Learning and Artificial Intelligence, he has applied ML techniques to enhance neutrino detection and analysis. Notably, his contributions have led to breakthroughs such as the first detection of neutrinos from the Milky Way using the IceCube detector, achieved through ML methods that improved event identification and reconstruction [22†]. Additionally, Prof. Rhode's interdisciplinary approach extends to exploring the epistemological aspects of ML in scientific knowledge acquisition, in cooperation with the field of philosophy.

Uppsala (associated): Christian Glaser is an astroparticle physicist specializing in high-energy neutrinos, cosmic rays, and air showers, with expertise in experimental data analysis, detector technologies, and deep learning. He plays a key role in major international collaborations, including IceCube, ARIANNA, RNO-G, and MODE, and co-leads the IceCube-Gen2 radio working group. His research has advanced in-ice radio detection techniques and machine learning applications through conditional normalizing flows and transformers to improve neutrino detection and event reconstruction. As the leader of open-source projects like NuRadioMC, widely used in the field, and the ERC-funded NuRadioOpt project, he is pioneering innovations in detector optimization through deep learning and differential programming. Additionally, he has developed popular deep-learning courses that integrate advanced methods with innovative teaching strategies, fostering interdisciplinary collaboration and advancing neutrino astronomy.

GSI: Anastasios Belias is a physicist at GSI Helmholtzzentrum für Schwerionenforschung, specializing in the development and coordination of detector systems for the PANDA experiment at the Facility for Antiproton and Ion Research (FAIR). His work focuses on the technical realization of detectors for the PANDA experiment and more recently also of detectors for the CBM experiment, contributing significantly to their design and implementation. A notable achievement in Dr. Belias's career is his role in repurposing the former Outer Tracker of the LHCb experiment at CERN for use in PANDA and CBM. This endeavor involved the complex task of transporting the gas-filled straw-tube detector from CERN to GSI and upgrading its front-end electronics to interface with modern readout systems. His efforts have been instrumental in adapting this advanced tracking detector beyond its original scope for future experiments at FAIR, demonstrating his commitment to enhancing experimental capabilities in hadron physics. Dr. Belias has also engaged in collaborations exploring the integration of machine learning techniques into detector optimization and data analysis workflows. Through his participation in initiatives such as the MODE collaboration, the EUCAIF and the AI Workgroups at GSI and NuPECC, he contributes to the application of advanced algorithms to challenges in tracking and event reconstruction. His involvement in leveraging ML methods aligns with efforts to improve the efficiency and adaptability of reconstruction processes in dynamic experimental environments like FAIR.

All scientific partners have a track record of collaborating with experts of other fields, as well as with each other, in particular in the MODE collaboration (Belias, Gauger, Glaser, Kieseler) [**41**†, **42**†, **43**‡, **44**‡, **45**‡, **46**‡] and within the EUCAIF WG2-Codesign (Belias, Gauger, Glaser, Kieseler). The long term goal of both is to develop methods and concepts for holistic optimisation of future detectors using tools from machine learning and differentiable programming. A detector agnostic reconstruction is an important stepping stone towards this goal, with direct applications to running experiments.

Industry partner: inomed Medizintechnik GmbH (Emmendingen, founded in 1991) is a privately owned medical technology company specialising in research, development, manufacture and sale of innovative devices for intraoperative monitoring of the nervous system. With around 250 employees, including 65 in research and development, the company actively participates in national and international research projects. Close collaboration between research and development teams ensures project outcomes influence future product innovation. inomed's innovative strength has been recognized with awards such as the 2016 Industry Award for the Raabe mapping suction device, the VDE Patient Safety Award in 2019, and the Top100 Innovation Award in 2022. Continuous improvement of products and processes is part of the company philosophy.

Detailed Description of the Work Plan

The project is structured into four main Work Packages (WPs) designed to achieve the overarching goal of developing detector-agnostic reconstruction frameworks for irregular high-dimensional data. These WPs involve close collaboration among participating institutes, leveraging expertise in CS, HEP, Astro, and HaN.

WP1 covers the more CS focused aspects, such as the new ML concepts and methods to be developed. In WPs 2-4, these concepts are applied to domain specific problems, as well as validated and further refined focusing on particular key aspects with help of the CS experts of WP1. These aspects are aligned with the main challenges and needs of the respective community. This setup makes WP1 a central hub for the consortium. Each of the WPs 2-4 will provide a public (simulated) dataset alongside evaluation metrics that are representative for the physics goals of the specific domain. This is a mid-term goal to facilitate possible input and collaboration beyond the consortium, e.g. with a larger group of computer scientists. For this purpose, we are planning to release the datasets in formats that are widely accepted in the CS domain, and will not rely on physics-domain specific data formats.

The milestones of each work package are aligned with annual meetings and workshops also detailed in this Section. The first ones facilitate the information exchange across the consortium, while the ones towards the end of the project target the dissemination and publication of results. A general overview is provided in the next Section and Fig. 1.

Since training ML models in all WPs relies on dedicated hardware, we have also carefully evaluated the requirements and determined that the consortium's current hardware resources are sufficient to support the planned work packages.

Work Package 1: ML Concepts

Participating Institutes: RPTU, KIT
Associated partner: inomed Medizintechnik GmbH
FTE Provided:

- RPTU: PI (10%): Nicolas R. Gauger, PostDoc (30%): Long Chen
- KIT: PI (30%): Jan Kieseler, PhD student (20%): Alessandro Brusamolino

FTE Requested: 1 PostDoc (WP1a, KIT), 1 PhD student (WP1b, RPTU)

WP1 is dedicated to developing advanced ML concepts that are universally applicable to irregular high-dimensional data across diverse domains. The work is organized into two closely interconnected sub-packages (WP1a and WP1b), each addressing key aspects of object detection in high-dimensional data from complementary perspectives. WP1a focuses on designing physics-informed architectures and embedding domain-specific constraints, while WP1b explores novel optimization techniques and solvers tailored to the inherently discrete nature of the problem. This strong interplay between the sub-packages fosters close collaboration between HEP and CS institutes, enhancing interdisciplinary synergies within the consortium. A key focus of WP1 is also the transfer of its findings to other work packages, positioning it as a central hub for ML-related support and feedback across the project. Additionally, this WP includes the application of developed techniques to novel architectures for analysing high-dimensional intraoperative time-series signals, recorded as part of routine nerve function monitoring during surgery. The data, provided by our associated industry partner inomed, will be accessed via their machines to ensure compliance with their privacy requirements if direct copying is not feasible. Applying the developed algorithms in this context will not only demonstrate their applicability beyond academia, benefiting real-world applications such as those in industry, but will also serve as a vital test bed for their generalisability to the analysis of time-series sensor readouts, event building, and online monitoring in scientific experiments.

WP1a

Participating Institute: KIT, PI: Jan Kieseler
Scope:

- **Coordination** between WP1a and WP1b, and information exchange with WP2-4
- **Integration of Physics Constraints:** Incorporate physics knowledge from WP2-4 into ML models to improve interpretability and performance.
- **Graph Topology Building Message Passing Layers:** Develop new message-passing architectures that construct graph topologies adhering to physical locality constraints, e.g. integrating dot-product attention mechanisms that are the key to the success of large language models.
- **Adaptive Point-Cloud Reduction Architectures:** Design point-cloud neural network architectures for adaptive dimensionality reduction without violating locality requirements (relies on WP1b).
- **Scalability and Efficiency:** Optimize models for computational efficiency, targeting reductions in energy consumption (Watts per event) through efficient hardware utilization.
- **Application** of developed concepts to inomed data (shared with WP1b)

WP1b

Participating Institute: RPTU, PI: Nicolas R. Gauger
Scope:

- **Adaptive Dimensionality Reduction Techniques:** Implement techniques that dynamically reduce data dimensionality, facilitating efficient processing without significant information loss. This includes studying the integration of classical algorithms with modern ML techniques to create hybrid models. Examples include combining density-based clustering methods (e.g., DBSCAN) with neural networks to effectively handle outliers and non-informative signals.
- **Machine Learning for Non-Differentiable Thresholds:** Develop machine learning algorithms and learned heuristics to handle non-differentiable operations in data reconstruction, such as thresholding, combinatorics and clustering, common in particle detection. This is an integral ingredient into dynamic dimensionality reduction.
- **Application** of developed concepts to inomed data (shared with WP1a)

Collaboration and Coordination

Both sub-packages will maintain close collaboration, ensuring compatibility and synergy between architectures and ML driven discrete optimization techniques and heuristics. In particular the U-Net-inspired reduction architectures will rely on the findings of the dimensionality reduction techniques developed in WP1b as well as the approaches to learn the reduction thresholds. The architectures developed in WP1a will ensure information propagation throughout the input data adhering to and exploiting physics constraints (inductive bias) providing the basis for the dimensionality reduction developed in WP1b.

Joint Milestones for WP1

- **M1.1** (Month 6): Initial method concepts; definition of specific neural network architecture requirements.
- **M1.2** (Month 12): Prototype implementation of hybrid ML algorithms integrating classical methods and ML as input to other WPs.
- **M1.3** (Month 18): Joint workshop on ML concepts developed in WP 1 and definition of a path forward for combination of findings.
- **M1.4** (Month 18): data agreement with inomed
- **M1.5** (Month 30): Joint publication summarizing ML concepts developed in WP 1, accompanied by public release of code and documentation.
- **M1.6** (Month 36): Organization of an international workshop on detector-agnostic ML methods (including results from other WPs) and application of ML concepts to inomed data.
- **M1.7** (Month 36): Final report on WP1 outcome (ML developments and inomed application)

Work Package 2: Low-Level ML Reconstruction for High-Energy Physics

Participating Institute: KIT, PI: Jan Kieseler
FTE Provided: PI (10%): Jan Kieseler, PhD student (30%): Alessandro Brusamolino
FTE Requested: 1 PhD student

WP2 focuses on applying the ML concepts developed in WP1 to low-level data reconstruction in HEP experiments, particularly for the LHC and the forthcoming HL-LHC phase. These experiments face two critical, potentially competing demands: algorithms must efficiently process Exabyte-scale datasets originating from unprecedented data rates while accurately identifying thousands of particles in highly complex, overlapping topologies.

The primary goal of WP2 is to extend the robustness of the dimensionality reduction techniques and networks from WP1 to handle high multiplicities and very high-dimensional input data within stringent computing constraints. Addressing these challenges will involve close collaboration with WP3. Additionally, the diverse nature of signals from tracking and calorimeter subsystems provides a unique opportunity to validate the agnostic capabilities of the algorithms from WP1. WP2 will integrate information from these subsystems directly, serving as a comprehensive evaluation of the methods developed.

Given the project's timeline, full integration of the algorithms into the CMS or ATLAS frameworks is deferred to future work. Instead, a generic HEP-detector simulation will be used as a benchmark, providing a robust foundation for future extensions to operational experiments. Alongside published code and datasets, this approach ensures a solid framework for broader adoption.

Scope:

- **Definition of Physics Goals and Metrics:** Establish clear physics objectives and performance metrics for particle reconstruction, considering increased data complexity and volume.
- **Public Dataset Creation:** Develop and release a generic HEP-detector simulation dataset using tools like COCOA for community-wide use and benchmarking.
- **Implementation of ML Methods:** Apply message-passing neural networks and adaptive dimensionality reduction techniques from WP1 to the HEP detector dataset, including

different subsystems.

Milestones:

- **M2.1** (Month 6): Establishment of physics objectives and performance metrics.
- **M2.2** (Month 18): Public release of the generic HEP simulation dataset.
- **M2.3** (Month 33): Publication of results.
- **M2.4** (Month 36): Final report and dissemination of methods, documentation and public release of code and models.

Work Package 3: Low-Level ML Reconstruction for Astroparticle Physics

Participating Institute: TU Dortmund, PI: Wolfgang Rhode
Associated Partner: Uppsala University, PI: Christian Glaser
FTE Provided: PI (10%): Wolfgang Rhode, co-PI (10%): Christian Glaser, PostDoc (20%): Thorsten Glüsenkamp, PhD student (20%): Marin Ravn
FTE Requested: 1 PhD student

WP3 focuses on developing and deploying advanced ML-based event reconstruction methods for optical Cherenkov neutrino detectors, with an emphasis on the IceCube Neutrino Observatory at the South Pole. These methodologies will also extend to next-generation detectors such as KM3NeT in the Mediterranean and P-ONE in Canada, both crucial for advancing high-energy neutrino astronomy.

A central challenge in IceCube is reconstructing neutrino directions with sub-degree precision, often based on signals from a single particle traversing the detector. This is further complicated by the detector medium—natural Antarctic ice—with its inhomogeneous, location-dependent asymmetries that disrupt conventional algorithms. WP3 will adapt and validate the techniques from WP1 to address these challenges, enhancing reconstruction accuracy and scalability. This work is particularly timely, as the upcoming IceCube Upgrade, planned for the 2025/26 polar season, will deploy advanced instrumentation to refine ice characterization – currently the dominant source of systematic uncertainty – necessitating enhanced reconstruction methods to fully exploit the new and existing IceCube data.

IceCube neutrino events fall into two topologies: tracks and cascades. Tracks, with their elongated light patterns, offer excellent directional resolution and are key to identifying astrophysical neutrino sources and generating real-time alerts for multi-messenger follow-ups. Cascades, characterized by spherical light patterns, excel in energy resolution but have poorer directional accuracy. WP3 will optimize ML-based reconstruction for both topologies, leveraging their strengths to maximize scientific output. Close collaboration with WP2 will ensure that computational efficiency and scalability are prioritized during deployment, opening up the option to also profit from the new algorithms in IceCube's real-time alert system for cosmic neutrinos. By leveraging the advanced ML concepts from WP1, providing feedback to them, and integrating them into the unique challenges of IceCube, WP3 will deliver essential tools for identifying astrophysical neutrino sources and advancing the capabilities of large-scale neutrino observatories.

Scope:

- **Definition of Physics Goals and Metrics:** Set objectives for improving cosmic neutrino identification from a large background and improving the precision in pinpointing neutrino origins, aiding in identifying ultra-high-energy cosmic ray sources.
- **Public Dataset Creation:** Generate and release a public dataset for an IceCube-like detector using new open-source Monte Carlo simulation software to facilitate broader collaboration and benchmarking.

- **Implementation of ML Methods:** Apply message-passing neural networks and adaptive dimensionality reduction techniques from WP1 to reconstruction in IceCube-like detectors. The emphasis will be on muon tracks (generated by muon neutrino charged-current interactions) which is the most challenging event topology with the largest potential to improve the localization of neutrino sources.

Milestones:

- **M3.1** (Month 6): Definition of physics objectives and performance metrics.
- **M3.2** (Month 18): Public release of the astroparticle physics simulation dataset.
- **M3.3** (Month 33): Publication of improved reconstruction algorithms.
- **M3.4** (Month 36): Final report and dissemination of methods, documentation and public release of code and models.

Work Package 4: Low-Level ML Reconstruction for Hadrons and Nuclei

Participating Institute: GSI (FAIR), PI: Anastasios Belias
FTE Provided: PI (20%): Anastasios Belias
FTE Requested: 1 PhD student

The new FAIR accelerator complex is designed to support a diverse range of hadron and nuclear physics experiments across multiple beamlines and storage rings. These experiments rely on various detector systems, including both compact movable setups and large permanent installations, to explore the behavior of strongly interacting matter under extreme conditions. Central to the physics goals are large-area detectors for charged-particle tracking and hadronic calorimetry, critical for fixed-target experiments utilizing particle beams, ions and protons, at the new synchrotron SIS100 and High Energy Storage Ring (HESR), as exemplified by the CBM and PANDA experiments.

WP4 focuses on leveraging ML techniques to enhance event reconstruction in hadron and nuclear physics experiments, aiming to optimize detector performance and adapt to dynamic changes in experimental conditions. This includes addressing both small operational variations (e.g., minor fluctuations in beam optics) and more significant disruptions (e.g., magnet quenches or stochastic cooling instabilities). The work emphasizes fast, adaptive, and flexible reconstruction methods that can maintain high efficiency and accuracy under varying conditions.

A key use case is the recent donation of the Outer Tracker (OT) straw modules from the LHCb experiment to GSI, providing a unique opportunity to apply ML-based reconstruction techniques. These modules are being integrated into PANDA, with potential applicability to CBM, extending their capabilities beyond their original design. This effort also facilitates testing and optimization of different detector configurations, directly aligning with WP4's focus on enhancing the adaptability and robustness of reconstruction methods.

By applying ML-driven approaches to detector data, WP4 seeks to improve physics output while addressing the complex and dynamic environments of FAIR experiments. The methods developed here will directly support the objectives outlined in the "Scope of Work," including the efficient reconstruction of events, creation of public datasets, and exploration of detector geometry optimizations.

Scope:

- **Definition of Physics Goals and Metrics:** Establish objectives for improving reconstruction in charged particle tracking coupled to hadronic calorimeters and other relevant detectors.

- **Public Dataset Creation:** Produce and share datasets relevant to hadron and nuclear physics for community-wide use. Including the various beam parameters and collisions of protons, heavy-ions from SIS100 and anti-protons with protons and nuclear-targets at HESR.
- **Implementation of ML Methods:** Apply algorithms for fast, adaptive, and dynamic reconstruction, with focus on robustness and adaptiveness against changes in detector conditions. This includes small changes during operations and larger changes due to instabilities in beam optics (stochastic cooling), and magnet quenches.
- **Detector Geometry Optimization:** Utilise the ML methods developed in WP1 to efficiently study proposals for detector geometries for improved physics output for experiments in hadrons and nuclei such as in the proton-proton collisions, heavy-ion collisions and anti-proton annihilation at CBM and PANDA respectively.

Milestones:

- **M4.1** (Month 6): Definition of physics objectives and performance metrics.
- **M4.2** (Month 18): Public release of hadron and nuclear physics simulated datasets.
- **M4.3** (Month 33): Application of methods to current experiments and geometry comparisons.
- **M4.4** (Month 36): Final report and dissemination of results, documentation and public release of code and models.

Overview, Cross-Cutting Activities, and Collaboration

An overview of the work plan and milestones can be found in Fig.1.

The work packages are highly interconnected, offering students and postdocs opportunities to build expertise in their respective tasks while fostering a sense of purpose through contributions to the broader project goals. In WP1, initial work on non-differentiable thresholds serves as a foundation for developing adaptive U-Net-like architectures. Similarly, the physics metrics identified during the early milestones will be integrated into the neural network designs, creating a feedback loop between WP1 and the application domains. As the project progresses, the methods developed in WP1 will be refined based on input and challenges identified by the application domains. This collaborative process, coordinated through WP1a, ensures that the algorithms remain domain-agnostic while addressing specific needs across the different fields, fostering strong cross-domain collaboration.

The publication timeline is carefully structured to maintain logical causality: the ML concepts and code will be released in conjunction with the domain application results, ensuring a clear progression in the research outputs. Additionally, for work packages involving PhD students, sufficient time has been allocated to support thesis preparation, aligning with standard timelines and practices in the respective research communities.

The milestones for each work package are strategically aligned to coincide with key points where interdependencies between WPs require coordinated input, creating clear deadlines for deliverables. Most of these milestones also coincide with the annual in-person consortium meetings, held in spring at one of the partner institutes. These meetings are designed as comprehensive workshops, directly tied to the milestones, ensuring a focused and collaborative environment. The first workshop, in month 6, will focus on introducing initial ML concepts, plans, and a discussion of physics constraints and metrics. By month 18, a workshop will review the ML methods developed in WP1, incorporating feedback from WPs 2-4. The final workshop, in month 30, will emphasize application results and lessons learned from cross-domain collaboration, ensuring alignment in preparing and synchronizing public results from WPs 2-4.

To ensure continuous and effective collaboration within the consortium beyond the defined milestones, we plan to hold bi-weekly virtual meetings, such as via Zoom, with all partners. These brief sessions will focus on addressing low-level issues and providing quick status updates, facilitating a fast turnaround of information and progress tracking. They will be organised through WP1.

As a culmination of the project, we will host an international workshop on detector-agnostic ML reconstruction methods at the end of the funding period. This event will serve as a platform to share project outcomes and gather valuable feedback on the public datasets released in month 18, further strengthening community engagement and collaboration. The organisation will be responsibility of the PIs and WP1.

Exploitation plan

The exploitation plan outlines how the outcomes of this project will be utilized and disseminated to maximize scientific impact, ensure sustainability, and contribute to both the scientific community and society at large. Our approach aligns with the FAIR principles (Findable, Accessible, Interoperable, Reusable) as described in Ref. [**47**], ensuring that data and tools developed are openly accessible and usable by others.

Commercialization Prospects

While the primary focus of this project is fundamental research, there are potential avenues for technology transfer to industry. Collaboration with our associated industry partner, inomed, provides an opportunity to apply our algorithms in medical settings, specifically in IONM as described in WP1. The adaptable ML techniques developed could benefit industries dealing with irregular, high-dimensional data, such as medical signal processing and imaging, telecommunications, and autonomous systems. Although commercialization is not the main objective, the potential for industrial application enhances the project's value and may lead to future collaborations or technology transfer opportunities.

In addition to benefiting from our direct associated industry partner, we will leverage the network and resources of the Lamarr institute, where Wolfgang Rhode leads the physics domain. This connection provides access to a dedicated transfer area within the institute, allowing us to continuously expand our industry-facing portfolio and actively engage with additional potential industrial collaborators throughout the project.

Scientific and Technological Prospects

The project is expected to make significant contributions to the fields of high-energy physics, astroparticle physics, hadron and nuclear physics, and computational science.

- **Publications:** We anticipate publishing our findings in high-impact, peer-reviewed journals across physics and computer science disciplines.
- **Conference Contributions:** Results will be showcased at prominent national and international conferences, collaboration meetings, and workshops dedicated to machine learning applications in physics. To foster early engagement and potential collaborations beyond the consortium, we aim to present our plans at the outset of the funding period. Furthermore, we plan for all funded personnel and PIs to participate in an international conference in both years 2 and 3, ensuring broad dissemination of our progress and findings while strengthening connections within the global scientific community.
- **Open-Source Release:** All developed software, algorithms, and models will be released under open-source licenses (e.g., MIT or GPL) on platforms like GitHub. This ensures broad accessibility and encourages community engagement.
- **Public Datasets:** As one of the primary goals of this project, we will generate and share high-quality, labeled datasets for HEP, astroparticle physics, and hadron and nuclear physics. These datasets will be accompanied by domain-specific metrics and will serve as benchmarks for the community. Furthermore, they will be provided in commonly accepted formats to facilitate close future collaboration with computer scientists.

Scientific and Commercial Impact

Scientific Impact:

- **Advancement of Reconstruction Techniques:** Developing detector-agnostic ML methods will significantly enhance data analysis capabilities, improving physics reach in experiments like ATLAS, CMS, IceCube, PANDA, CBM, and beyond. In addition to strengthening existing experiments, their adaptability will significantly ease the development of new detectors by connecting the detector design more closely to the physics objective and allowing significantly faster turn-around (days instead of years) in case of new proposals, compared to classic algorithms.
- **Interdisciplinary Collaboration:** Strengthening ties between physics domains and computer science will promote knowledge transfer and foster innovative solutions to complex problems.
- **Educational Advancement:** Training young researchers in cutting-edge ML techniques contributes to building a skilled workforce equipped to tackle future scientific challenges.

Commercial Impact:

- **Industry Applications:** The algorithms developed could be adapted for use in industries requiring analysis of complex, high-dimensional data, potentially leading to new products or services.
- **Innovation Stimulation:** Demonstrating the applicability of advanced ML methods may inspire further research and development within the industrial sector.
- **Educational Advancement:** Skills acquired in cutting-edge ML techniques prepare young researchers also for high-ranked jobs outside academia.

Subsequent Use

- **Integration into Experimental Frameworks:** The methods and tools can be integrated into the data analysis pipelines of ATLAS, CMS, IceCube (KM3NeT or similar), CBM, PANDA, and other experiments, enhancing their long-term utility.
- **Foundation for establishing new detector designs:** A detector agnostic, adaptable reconstruction may offer completely new avenues for detector design, and will additionally decrease the turn-around time when estimating proposals for new detectors by orders of magnitude.
- **Foundation for Future Research:** By adhering to open science principles, the project lays a foundation for future studies, allowing other researchers to build upon our work.
- **Educational Resources:** The developed materials and datasets can be used in academic courses, workshops, and training programs, extending their impact beyond the project's duration.

Sustainability

The project actively contributes to several United Nations Sustainable Development Goals (SDGs) and is in line with the ErUM-Data sustainability recommendations (Mx, Sx) [**48**], as well as the FAIR Guiding Principles for scientific data management and stewardship:

Industry, Innovation, and Infrastructure (SDG 9):
Advancing computational methods contributes to technological innovation and strengthens research infrastructure. Building on FAIR principles, we will make our data and workflows as

reusable and accessible as possible, directly addressing the recommendation to make data FAIR (M1) and to use workflow management to make processing FAIR (M5). By **explicitly publishing datasets alongside metrics and ML models as part of our WPs**, we enhance transparency, reproducibility, and long-term accessibility of our research outputs. This approach improves interoperability, facilitates community-wide benchmarking, and ultimately maximizes the scientific return on invested resources.

Responsible Consumption and Production (SDG 12):
Optimizing algorithms for efficiency reduces computational resource consumption, thus promoting sustainable practices. **We explicitly research data reduction strategies within the algorithms** to retain maximum scientific value with minimal resource usage, such as reducing and compressing data (M2) and thereby opening opportunities for optimizing storage vs. re-calculation choices (M3). Additionally, our emphasis on maintainable and reusable code aligns with good software development practices (M6). To ensure software longevity and adaptability, **we will provide a well-documented public code release with fine-grained test cases integrated into a continuous integration pipeline**, reducing the need for redevelopment. Furthermore, by using parallelizable algorithms on energy-efficient hardware (e.g., GPUs), we aim to achieve performance per watt improvements (M7, M8) and address the trade-off between research benefit and climate impact (S5). Beyond computational efficiency, by facilitating rapid exploration of new detector concepts, the outcome of our project leads to more sustainable experimental designs in the future, potentially reducing material use and construction efforts (S4). This is also explicitly reflected in enabling re-use of LHCb tracker modules in WP4. In general, this forward-looking stance encourages continuous assessment and improvement, ensuring that our methods and infrastructure remain both scientifically valuable and environmentally responsible.

Gender Equality (SDG 5):
We are committed to fostering gender equality and inclusivity, actively promoting equal opportunities for individuals of all genders and fostering a welcoming and supportive environment for LGBTQ+ researchers. Recognizing the persistent underrepresentation of women and other marginalized groups in STEM fields, we will implement proactive measures to ensure diversity and equity within the consortium. In line with these goals, we will prioritize inclusivity in the hiring process by encouraging applications from candidates of underrepresented groups and ensuring transparent, fair evaluation criteria. To further strengthen this commitment, the Principal Investigators (PIs) will take responsibility for cultivating a collaborative and equitable working environment, actively supporting diversity in participation and leadership. These principles will extend to our outreach and training efforts, ensuring that inclusivity and respect are embedded in every aspect of the project.

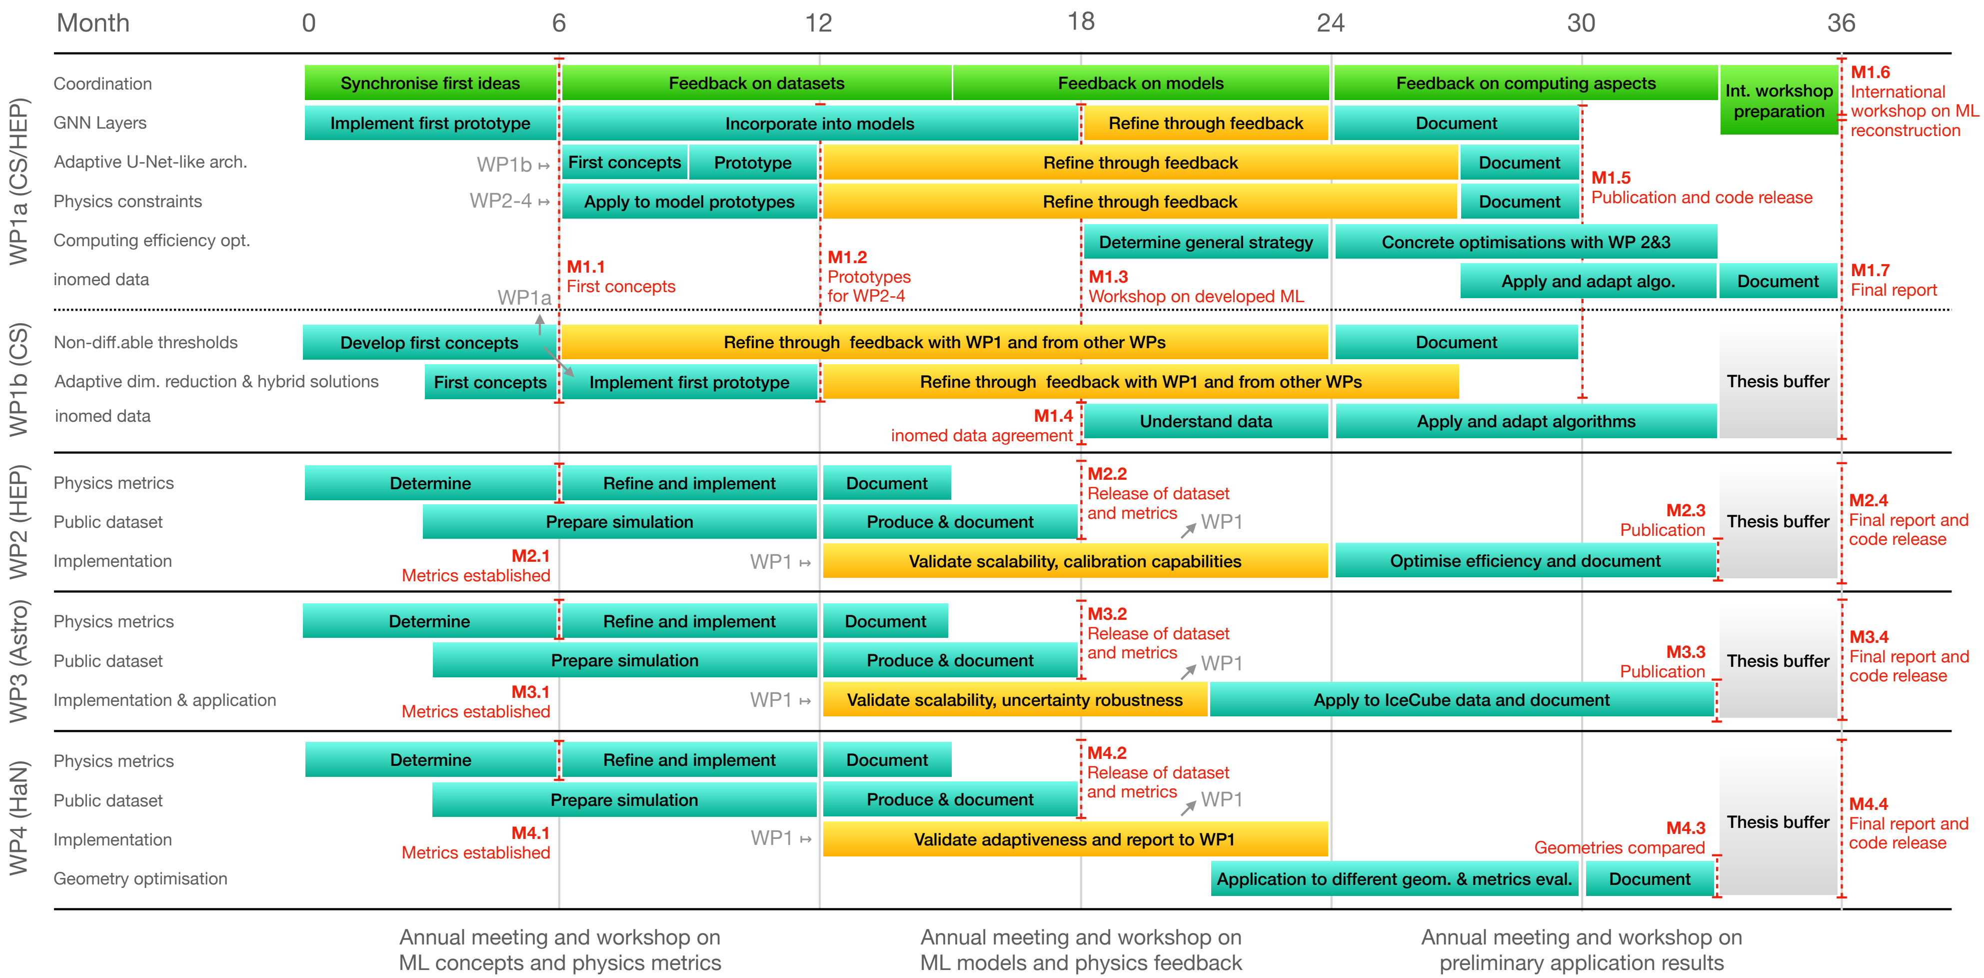

Figure 1: Overview of the project and its milestones. The colours indicate coordination activities (green), development and documentation (teal), and phases of particularly close collaboration between the work packages (yellow).

References

[1] Y. LeCun et al. “Backpropagation Applied to Handwritten Zip Code Recognition”. In: *Neural Computation* 1.4 (1989), pp. 541–551. DOI: 10.1162/neco.1989.1.4.541.

[2] Yann LeCun, Y. Bengio, and Geoffrey Hinton. “Deep Learning”. In: *Nature* 521 (May 2015), pp. 436–44. DOI: 10.1038/nature14539.

[3] L. Alzubaidi, J. Zhang, A.J. Humaidi, et al. “Review of deep learning: concepts, CNN architectures, challenges, applications, future directions”. In: *J Big Data* 8.53 (2021). DOI: 10.1186/s40537-021-00444-8.

[4] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. “ImageNet Classification with Deep Convolutional Neural Networks”. In: *Advances in Neural Information Processing Systems*. Ed. by F. Pereira et al. Vol. 25. Curran Associates, Inc., 2012.

[5] Ross Girshick et al. “Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation”. In: *2014 IEEE Conference on Computer Vision and Pattern Recognition*. 2014, pp. 580–587. DOI: 10.1109/CVPR.2014.81.

[6] Joseph Redmon et al. “ You Only Look Once: Unified, Real-Time Object Detection ”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, June 2016, pp. 779–788. DOI: 10.1109/CVPR.2016.91. URL: https://doi.ieeecomputersociety.org/10.1109/CVPR.2016.91.

[7] Hanxi Li, Yi Li, and Fatih Porikli. “DeepTrack: Learning Discriminative Feature Representations Online for Robust Visual Tracking”. In: *IEEE Transactions on Image Processing* 25.4 (Apr. 2016), pp. 1834–1848. ISSN: 1941-0042. DOI: 10.1109/tip.2015.2510583. URL: http://dx.doi.org/10.1109/TIP.2015.2510583.

[8] CMS Offline Software and Computing. *CMS Phase-2 Computing Model: Update Document*. Tech. rep. Geneva: CERN, 2022. URL: https://cds.cern.ch/record/2815292.

[9] A. M. Sirunyan et al. “Particle-flow reconstruction and global event description with the CMS detector”. In: *JINST* 12.10 (2017), P10003. DOI: 10.1088/1748-0221/12/10/P10003. arXiv: 1706.04965 [physics.ins-det].

[10] Morad Aaboud et al. “Jet reconstruction and performance using particle flow with the ATLAS Detector. Jet reconstruction and performance using particle flow with the ATLAS Detector”. In: *Eur. Phys. J. C* 77.7 (2017). 67 pages in total, author list starting page 51, 37 figures, 1 table, final version published in Eur. Phys. J C, all figures including auxiliary figures are available at http://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/PERF-2015-09/, p. 466. DOI: 10.1140/epjc/s10052-017-5031-2. arXiv: 1703.10485. URL: https://cds.cern.ch/record/2257597.

[11] Stefano Spataro. “The PandaRoot framework for simulation, reconstruction and analysis”. In: *J. Phys. Conf. Ser.* 331 (2011). Ed. by Simon C. Lin, p. 032031. DOI: 10.1088/1742-6596/331/3/032031.

[12†] Shah Rukh Qasim et al. “Learning representations of irregular particle-detector geometry with distance-weighted graph networks”. In: *Eur. Phys. J. C* 79.7 (2019), p. 608. DOI: 10.1140/epjc/s10052-019-7113-9. arXiv: 1902.07987 [physics.data-an].

[13†] Jan Kieseler. “Object condensation: one-stage grid-free multi-object reconstruction in physics detectors, graph and image data”. In: *Eur. Phys. J. C* 80.9 (2020), p. 886. DOI: 10.1140/epjc/s10052-020-08461-2. arXiv: 2002.03605 [physics.data-an].

[14†] Shah Rukh Qasim et al. “End-to-end multi-particle reconstruction in high occupancy imaging calorimeters with graph neural networks”. In: *Eur. Phys. J. C* 82.8 (2022), p. 753. DOI: 10.1140/epjc/s10052-022-10665-7. arXiv: 2204.01681 [physics.ins-det].

[15†] Shah Rukh Qasim et al. "Multi-particle reconstruction in the High Granularity Calorimeter using object condensation and graph neural networks". In: *EPJ Web Conf.* 251 (2021), p. 03072. DOI: 10.1051/epjconf/202125103072. arXiv: 2106.01832 [physics.ins-det].

[16] Lea Reuter et al. "End-to-End Multi-Track Reconstruction using Graph Neural Networks at Belle II". In: *arXiv* (Nov. 2024). arXiv: 2411.13596 [physics.ins-det].

[17†] Dolores Garcia et al. "Towards Particle Flow Event Reconstruction at the Future Circular Collider with GNNs". In: *NeurIPS GLFrontiers Workshop proceedings* (2023). URL: https://openreview.net/pdf?id=PYcp183GBL.

[18] Joosep Pata et al. "Improved particle-flow event reconstruction with scalable neural networks for current and future particle detectors". In: *Commun. Phys.* 7.1 (2024), p. 124. DOI: 10.1038/s42005-024-01599-5. arXiv: 2309.06782 [physics.data-an].

[19] Nilotpal Kakati et al. "HGPflow: Extending Hypergraph Particle Flow to Collider Event Reconstruction". In: *arXiv* (Oct. 2024). arXiv: 2410.23236 [hep-ex].

[20†] R. Abbasi et al. "A Convolutional Neural Network based Cascade Reconstruction for the IceCube Neutrino Observatory". In: *JINST* 16 (2021), P07041. DOI: 10.1088/1748-0221/16/07/P07041. arXiv: 2101.11589 [hep-ex].

[21†] Mirco Hünnefeld. "Observation of High-Energy Neutrinos from the Milky Way". PhD thesis. TU Dortmund, 2023.

[22†] R. Abbasi et al. "Observation of high-energy neutrinos from the Galactic plane". In: *Science* 380.6652 (2023), adc9818. DOI: 10.1126/science.adc9818. arXiv: 2307.04427 [astro-ph.HE].

[23] R. Abbasi et al. "Graph Neural Networks for low-energy event classification & reconstruction in IceCube". In: *JINST* 17.11 (2022), P11003. DOI: 10.1088/1748-0221/17/11/P11003. arXiv: 2209.03042 [hep-ex].

[24] Habib Bukhari et al. "IceCube – Neutrinos in Deep Ice: The top 3 solutions from the public Kaggle competition". In: *Eur. Phys. J. C* 84.6 (2024), p. 646. DOI: 10.1140/epjc/s10052-024-12977-2. arXiv: 2310.15674 [astro-ph.HE].

[25†] T. Glüsenkamp for the IceCube Collaboration. "Conditional normalizing flows for IceCube event reconstruction". In: *PoS(ICRC2023)1003* (). DOI: 10.22323/1.444.1003. arXiv: 2309.16380 [astro-ph.HE].

[26] Alexey Dosovitskiy et al. "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale". In: *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=YicbFdNTTy.

[27] Haiping Wu et al. "CvT: Introducing Convolutions to Vision Transformers". In: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. 2021, pp. 22–31. DOI: 10.1109/ICCV48922.2021.00009.

[28] Stéphane D'Ascoli et al. "ConViT: Improving Vision Transformers with Soft Convolutional Inductive Biases". In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, 18–24 Jul 2021, pp. 2286–2296. URL: https://proceedings.mlr.press/v139/d-ascoli21a.html.

[29] Junhyun Lee, Inyeop Lee, and Jaewoo Kang. "Self-Attention Graph Pooling". In: *Proceedings of the 36th International Conference on Machine Learning*. Sept. 2019.

[30] Frederik Diehl et al. "Towards Graph Pooling by Edge Contraction". In: 2019. URL: https://api.semanticscholar.org/CorpusID:208144770.

[31] Rex Ying et al. “Hierarchical graph representation learning with differentiable pooling”. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. NIPS’18. Montréal, Canada: Curran Associates Inc., 2018, pp. 4805–4815.

[32] Hongyang Gao and Shuiwang Ji. “Graph U-Nets”. In: *International Conference on Machine Learning*. 2019, pp. 2083–2092.

[33] Charles R Qi et al. “PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space”. In: *arXiv preprint arXiv:1706.02413* (2017).

[34] Yue Wang et al. “Dynamic Graph CNN for Learning on Point Clouds”. In: *ACM Transactions on Graphics (TOG)* (2019).

[35] Dougho Park and Injung Kim. “Application of Machine Learning in the Field of Intraoperative Neurophysiological Monitoring: A Narrative Review”. In: *Applied Sciences* 12.15 (2022). ISSN: 2076-3417. DOI: `10.3390/app12157943`. URL: `https://www.mdpi.com/2076-3417/12/15/7943`.

[36†] Tobias Kortus et al. “Automated Robust Interpretation of Intraoperative Electrophysiological Signals – A Bayesian Deep Learning Approach”. In: *Current Directions in Biomedical Engineering* 7.2 (2021), pp. 69–72. DOI: `doi:10.1515/cdbme-2021-2018`. URL: `https://doi.org/10.1515/cdbme-2021-2018`.

[37†] Emil Bols et al. “Jet Flavour Classification Using DeepJet”. In: *JINST* 15.12 (2020), P12012. DOI: `10.1088/1748-0221/15/12/P12012`. arXiv: `2008.10519 [hep-ex]`.

[38†] Tobias Kortus et al. “Towards Neural Charged Particle Tracking in Digital Tracking Calorimeters With Reinforcement Learning”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45.12 (2023), pp. 15820–15833. DOI: `10.1109/TPAMI.2023.3305027`.

[39†] Tobias Kortus, Ralf Keidel, and Nicolas R. Gauger. *Exploring End-to-end Differentiable Neural Charged Particle Tracking – A Loss Landscape Perspective*. 2024. arXiv: `2407.13420 [physics.comp-ph]`. URL: `https://arxiv.org/abs/2407.13420`.

[40‡] Tobias Kortus et al. *Constrained Optimization of Charged Particle Tracking with Multi-Agent Reinforcement Learning*. 2025. arXiv: `2501.05113 [physics.comp-ph]`. URL: `https://arxiv.org/abs/2501.05113`.

[41†] Giles C. Strong et al. “TomOpt: differential optimisation for task- and constraint-aware design of particle detectors in the context of muon tomography”. In: *Mach. Learn. Sci. Tech.* 5.3 (2024), p. 035002. DOI: `10.1088/2632-2153/ad52e7`. arXiv: `2309.14027 [physics.ins-det]`.

[42†] M. Aehle et al. “Optimization Using Pathwise Algorithmic Derivatives of Electromagnetic Shower Simulations”. In: (2024). URL: `http://arxiv.org/abs/2405.07944`. published.

[43‡] Tommaso Dorigo et al. “Toward the End-To-End Optimization of the SWGO Array Layout”. In: *submitted to Nuclear Physics B* (Oct. 2023). arXiv: `2310.01857 [astro-ph.IM]`.

[44‡] Max Aehle et al. *Efficient Forward-Mode Algorithmic Derivatives of Geant4*. 2024. arXiv: `2407.02966 [physics.comp-ph]`. URL: `https://arxiv.org/abs/2407.02966`.

[45‡] Max Aehle et al. “Progress in End-to-End Optimization of Detectors for Fundamental Physics with Differentiable Programming”. In: *arXiv* (Sept. 2023). arXiv: `2310.05673 [physics.ins-det]`.

[46‡] Tommaso Dorigo et al. “Toward the end-to-end optimization of particle physics instruments with differentiable programming”. In: *Rev. Phys.* 10 (2023), p. 100085. DOI: `10.1016/j.revip.2023.100085`. arXiv: `2203.13818 [physics.ins-det]`.

[47] M. Wilkinson, M. Dumontier, I. Aalbersberg, et al. “The FAIR Guiding Principles for scientific data management and stewardship.” In: *Sci Data* 3 (2016). DOI: 10.1038/sdata.2016.18.

[48] Ben Bruers et al. “Resource-aware Research on Universe and Matter: Call-to-Action in Digital Transformation”. In: Nov. 2023. arXiv: 2311.01169 [physics.comp-ph].

Entries marked with † indicate contributions from consortium members; entries marked with ‡ indicate joint publications, with contributions from at least two members of the consortium.