Unleashing the Power of Machine Learning and Artificial Intelligence: Advancing the Production of Official Statistics through MLOps implementation

R. Piela[1] and H. Luomaranta-Helmivuo[2]

[1] Statistical Office, Finland
e-mail: riitta.piela@stat.fi

[2] Statistical Office, Finland
e-mail: henri.luomaranta@stat.fi

Abstract

1. Introduction

1.1. Background and context

Official statistics provide objective and reliable information about various aspects of society. Accurate statistics help identify societal challenges, assess policy effectiveness, and shape evidence-based strategies to address issues such as unemployment, and inequality. Official statistics serve as a critical tool for monitoring economic performance, tracking demographic changes, and assessing the impact of government policies. Reliable data allows governments to identify areas that require intervention, assess the success of existing programs, and allocate resources efficiently to achieve societal goals.

The availability of official statistics should ensure transparency and accountability in decision-making processes. Transparent statistics enhance public trust in government actions and foster open discussions on public policy matters. Implementing ML/AI in statistical processes need the implementation of MLOps (a paradigm for the operationalization of ML efficiently and transparently) principles and the infrastructure and platform to implement these principles. The upcoming AI Act from EU may also set new requirements for statistical production or at least raise people's awareness of the requirements for information produced with the help of artificial intelligence.

1.2. Problem statement

Statistical institutes around the world are increasingly adopting machine learning (ML) techniques to enhance official statistics production. However, globally many statistical institutes encounter difficulties in implementing MLOps principles, such as transparency and reproducibility or they are not yet aware what kind of elements should exist when using AI in statistical production. In the future, decision-makers and the public may lack confidence in the validity and accountability of the ML-driven statistics and its impact on policymaking and governance if the AI driven processes are not transparent enough.

The importance of official statistics in guiding policy formulation, resource allocation, and governance decisions cannot be overstated. In the era of AI and ML, leveraging these advanced technologies has the potential to revolutionize statistical production, improve accuracy, and speed up the ability of statistics to quickly describe the changes taking place in society. Nevertheless, the lack of transparency and reproducibility in ML models raises concerns about their integrity and potential biases.

To gain the trust of decision-makers and the public, statistical institutes must prioritize Responsible AI and MLOps principles, transparency, and reproducibility. Transparency ensures that the decision-making process and ML algorithms are open and explainable, providing stakeholders with insight into how conclusions are reached. Reproducibility guarantees that results can be independently verified, increasing the confidence in the accuracy and reliability of statistical outputs. Concepts such as

Responsible AI and MLOps are not fully established. They contain many common principles, but usually Responsible AI includes also principles related to data protection and data security.

MLOps, an amalgamation of various definitions, represents a paradigm that encompasses best practices, essential concepts, and a development culture aimed at facilitating the end-to-end lifecycle of machine learning products. At its core, MLOps merges principles, tools, and techniques from both machine learning and traditional software engineering to design and construct complex computing systems. However, MLOps is not merely a collection of principles or a stack of technological elements; it extends its reach to encompass data, processes, roles, and methods. Embracing MLOps presents challenges due to its multidimensional nature. One of the prominent challenges lies in navigating the diverse dimensions it encompasses. The implementation of MLOps can be intricate, necessitating the adoption of new roles, such as ML engineers, data scientists, and data engineers, along with the establishment of robust collaboration mechanisms among these roles. Despite the complexities, embracing MLOps is essential for streamlined and accountable AI integration, empowering organizations to leverage the full potential of machine learning and artificial intelligence for advancing official statistics.

By addressing the challenges in implementing MLOps principles, statistical institutes can build a foundation of trust and accountability in their ML driven statistical processes. This trust is crucial for fostering evidence-based policymaking, efficient governance, and the public's acceptance of AI and ML advancements in official statistics. The motivation is to establish a robust and ethical framework that empowers decision-makers to make informed choices and the public to have confidence in the data driving policy and governance decisions.

2. AI in Statistical production

2.1. Current Landscape of ML adoption in Statistical Institutes

As mentioned, machine learning (ML) and artificial intelligence (AI) can bring improvements to the accuracy, efficiency, and timeliness of official statistics production compared to traditional manual processes, or even compared to rules-based systems of validation. Once properly established, ML models can process and analyze data in real-time, allowing statistical offices to produce more up-to-date and accurate data. This is especially valuable in rapidly changing situations of society. ML and AI algorithms can handle large-scale datasets efficiently, enabling statistical offices to process and analyze massive amounts of data quickly and reliably. Manual data validation and editing processes are also prone to human errors. ML-driven automation reduces these errors and is likely to lead to more accurate and reliable official statistics. ML and AI can automate labor-intensive tasks in statistical production processes. Many statistical institutes understand these potential benefits and have explored ML in automating classification and editing tasks, especially when the data are large and difficult. Some effort has been put into exploring new areas such as nowcasting or exploring new data sources such as satellite images. But systematic adoption of ML is not yet achieved. This is likely because these explorations have been done on a case-by-case basis, by few experts dealing simultaneously with many other tasks, and the benefits are not realized at the level of the organization, even less in the statistical system as a whole.

2.2. The Importance of MLOps Principles for Trustworthy Statistical Production

The systematic adoption of transparency, reproducibility and explainability in ML-driven processes play a crucial role in ensuring that decision-makers and the public to understand and trust the outcomes of AI and ML applications in statistical production and ensures that systemwide adoption can be possible. It fosters accountability, promotes ethical decision-making, and facilitates effective communication between stakeholders, leading to better-informed policies and governance decisions. Automated reports and visualization can be great benefit to users of statistics, but only if they are trusted.

In this paper, we underscore the significance of MLOps principles like transparency and reproducibility in the context of machine learning adoption in statistical institutes and aims to inspire statistical institutes to embrace MLOps principles, thereby fostering greater trust among decision-makers and the public in the credibility and impact of ML-driven statistical outputs. Additionally, the paper seeks to contribute to the responsible and ethical implementation of AI in official statistics, promoting evidence-based policymaking and efficient governance by critically discussing the experiences in Finland, noting that there is much work to be done.

MLOps principles represent a set of best practices that aim to streamline and standardize the entire ML model lifecycle, from development to deployment and maintenance. These principles are fundamental for ensuring that ML models are effectively managed, monitored, and governed, facilitating efficient collaboration among data scientists, decision-makers, and stakeholders.

2.3. Technology and processes implementing MLOps principles

Technology plays a central role in enabling MLOps principles and practices, as it provides the infrastructure, platforms, and frameworks necessary to implement these principles seamlessly. Tools and platforms (or frameworks) like TensorFlow Extended (TFX), MLflow, and Kubeflow offer a suite of tools for automating ML workflows, versioning models, and tracking experiments, promoting transparency and reproducibility. By integrating MLOps principles with appropriate technological solutions, statistical institutes can optimize data collection, validation, analysis, and dissemination processes, leading to more accurate, timely, and trustworthy official statistics.

ML platforms typically provide tools for managing data, model training, serving models, and monitoring, and support features that promote transparency and reproducibility, such as model tracking, artifact management, and version control. In fact, one may consider replacing the entire statistical production process pipelines with the available platforms. Countries such as France, Norway and Lithuania seem to be headed this way.

While choosing an ML platform for implementing MLOps principles, it's essential to consider the specific needs and requirements of your organization. Each platform may have different strengths and may be better suited for certain use cases or deployment environments. Additionally, integration with existing infrastructure and compatibility with the chosen ML frameworks should be considered.

3. Challenges and Ethical Considerations

3.1. Ethical implications of using AI in producing official statistics

Ethics is a critically important issue when using ML in producing data for official statistics due to several key reasons like fairness and equity. ML-driven statistical models have the potential to impact individuals and communities directly through policy decisions. If ML models are not designed with fairness in mind, they may inadvertently perpetuate biases, leading to inequitable outcomes and exacerbating existing societal disparities.

Official statistics influence critical decisions and policies that affect people's lives. Lack of transparency in ML models can lead to "black-box" decision-making, where it's challenging to understand how conclusions are reached. Transparent ML ensures accountability, enabling stakeholders to verify and challenge the basis of statistical outputs.

3.2. Notes regarding to training data

ML-driven official statistics depend on the quality and representativeness of the data used for training. Ethical practices ensure that data used in ML models accurately reflect the population they represent, avoiding skewed or inaccurate results. Establishing governance frameworks for ML use in official statistics ensures that decisions about model deployment and data utilization are made responsibly and with careful consideration of potential impacts on society.

Official statistics often involve sensitive and personal data. ML models must be designed with robust data privacy measures to protect individuals' confidentiality and adhere to data protection regulations. ML models can be susceptible to adversarial attacks, wherein malicious actors manipulate data inputs to produce inaccurate results. Ensuring the accuracy of ML-driven official statistics is essential for making informed and reliable decisions.

When using data from individuals, obtaining informed consent becomes vital to respecting their autonomy and privacy rights. ML practitioners must ensure that individuals are aware of how their data will be used and have the option to opt-out if desired. ML can make predictions and decisions based on patterns that are not easily understood by humans. This introduces the risk of unintended consequences if models are deployed without careful consideration of potential impacts. ML, as any modeling, can amplify biases present in the data it learns from. Ethical considerations mandate the detection and mitigation of biases to ensure that AI models produce unbiased and impartial results.

3.3. Maintaining public trust

Public trust is one the most important issues in implementing AI systems in official statistics. The adoption of ethical AI practices in official statistics fosters public trust and confidence in the results and decision-making processes. Transparency and fairness lead to greater acceptance of AI-driven insights by policymakers and the public.

The importance of ethics in using AI in producing data for official statistics cannot be overstated. Ethical considerations are crucial for ensuring fairness, transparency, accountability, data privacy, and the accuracy of AI-driven statistical outputs. By upholding ethical principles, statistical institutes can leverage AI to produce data that supports informed policymaking, governance, and decision-making while protecting individuals' rights and promoting societal well-being. Canadian Trust Centre is an excellent effort building trust between public and statistical institute.

3.4. Reproducibility: a cornerstone of scientific research and statistics

Reproducibility is paramount in scientific research and in official statistics. By promoting transparency and fostering collaboration and trust, reproducibility ensures the credibility of results. In official statistics, it serves as a reliable foundation for evidence-based decision-making and contributes to scientific progress.

Both ethical considerations and reproducibility are essential components when deploying AI in official statistics. Maintaining reproducibility ensures that the statistical outputs are verifiable, transparent, and trustworthy, providing a robust basis for effective decision-making. By embracing these principles, statistical institutes can navigate the challenges of AI in official statistics while maintaining public trust and promoting responsible and ethical use of AI technologies.

4.Implementing MLOps, steps towards success at Statistics Finland

4.1. Versioning

In the domain of MLOps, achieving reproducibility is of paramount importance to ensure the credibility and reliability of machine learning models. The main components that play a pivotal role in attaining reproducibility are versioning of the model, data, and code. Model versioning entails the systematic tracking and recording of changes made to the ML model throughout its development lifecycle. By preserving the model's version history in a model registry, it becomes possible to precisely recreate and compare results at different stages of the development process, enabling practitioners to understand the model's evolution and make informed decisions based on its performance. Similarly, data versioning involves meticulously documenting the datasets used for training and evaluation, including data preprocessing steps and any modifications made during the data preparation phase. Proper data versioning ensures that the same dataset can be reliably used in future iterations, allowing for the validation of results, and reducing the risk of introducing data-related discrepancies.

Lastly, code versioning involves the systematic management of the software code used to develop, train, and deploy the ML model. By maintaining a clear record of code changes and dependencies, data scientists and statisticians can reproduce experiments and workflows accurately, leading to consistent results and facilitating collaboration among team members. In combination, model, data, and code versioning in MLOps contribute to the establishment of a robust and transparent framework, fostering confidence in the reproducibility of machine learning outcomes.

Statistics Finland has implemented a comprehensive model register for ML models, ensuring that each model's version history is well-documented and traceable. Additionally, the use of Model Cards is adopted, which serve as repositories for storing essential metadata related to every model. This includes important information about the model's architecture, training data, hyperparameters, and performance metrics. By maintaining detailed version records and utilizing Model Cards, we enhance transparency, reproducibility, and accountability in our machine learning practices.

4.2. Automation of the workflow – the levels

As the demand for AI and ML solutions continues to soar, organizations face increasing pressure to streamline their ML workflows and enhance model deployment efficiency. Google Cloud's MLOps automation, which has become somewhat of an MLOps standard, has emerged as a transformative approach to address these challenges, presenting three distinct levels of maturity. Each level represents a significant milestone in the automation journey, enabling organizations to achieve higher levels of operational efficiency, model reproducibility, and collaboration among different ML roles.

At the foundational level of MLOps automation, organizations rely on manual processes to manage their ML workflows. Data scientists and ML engineers typically conduct each step of the ML lifecycle manually, from data preprocessing and feature engineering to model training and deployment. Code and data versioning are often managed through traditional version control systems, leading to potential inconsistencies and difficulties in tracking changes. **Level 0** automation serves as a starting point for organizations, providing valuable insights into the intricacies of their ML workflows and paving the way for further advancements in MLOps automation.

As organizations mature in their MLOps journey, they seek to address the challenges encountered at Level 0 by adopting MLOps tools and platforms. **At Level 1**, automation becomes more prevalent in data versioning, model training, and deployment processes. MLOps tools facilitate the automation of ML workflows, offering capabilities for versioning models, tracking experiments, and managing data pipelines. These tools enhance transparency, traceability, and collaboration among ML roles, enabling better reproducibility and efficiency in model development. However, certain elements of the ML process may still require manual intervention, limiting the level of end-to-end automation achieved at this stage.

The pinnacle of MLOps automation is reached at **Level 2,** where organizations achieve full end-to-end automation of their ML workflows. At this stage, advanced MLOps platforms seamlessly orchestrate the entire ML lifecycle, from data ingestion to model deployment and monitoring. AutoML solutions further enhance the efficiency of model development by automating hyperparameter tuning and architecture selection. Fully automated MLOps streamlines collaboration, fosters transparency, and ensures model reproducibility at scale. The robustness of Level 2 automation empowers data scientists and ML engineers to focus on high-value tasks, driving innovation and accelerating the delivery of AI solutions.

At Statistics Finland, we have made significant progress in automation, but the current level falls somewhere between 0 and 1 on the automation scale. While we have successfully implemented many components and have partially addressed versioning requirements, there is still ongoing development work required to achieve seamless automation across process states. Our team is actively working to bridge this gap and enhance our automation capabilities, with the goal of streamlining our processes, improving efficiency, and ensuring the reproducibility of our models and data. As we continue to invest in automation and embrace best practices, we aim to elevate our capabilities to a higher level. It is a difficult management exercise to determine the sweet spot lies where efficiency gains start to show throughout the organization and input costs start to sink in relation to output as compared to the situation without implementing MLOps. Nevertheless, these are necessary sometimes to convince higher leadership. Often, we must rely on demonstrating gains by showing an example and trying to analyze what cost savings are achieved in manual data editing phase which is sometimes offset with increased investments in technology when a machine learning is adopted, but this does not show the benefits of MLOps, which is a strategy that ensures that the one-off-ML model is maintained, transparent, and available for re-use. What makes it more difficult is that sometimes, short term gains are not favorable for the adoption of a single machine learning model, especially when evaluated case-by-case. However, with system-wide adoption of MLOps and with increased potential output, these gains should be easy to see. MLOps is a strategic investment, more than allowing some statistical team to develop a ML approach to a specific problem.

4.3. Monitoring the performance

Monitoring the performance of machine learning (ML) models is of utmost importance to ensure their continued effectiveness and avoid degradation over time. ML models are typically trained on historical data, and as the underlying patterns in the data change, the model's performance can deteriorate. Monitoring enables us to detect and address these changes promptly, maintaining the model's accuracy and reliability.

The monitoring process involves regularly evaluating the model's performance metrics, such as accuracy, precision, recall, and F1 score, on a representative dataset. By comparing these metrics against predefined thresholds, we can identify deviations and potential performance degradation. Additionally, tracking other relevant statistics, such as data distribution shifts and input-output correlations, can help spot anomalies that might affect the model's performance.

To monitor model performance effectively, a robust and scalable monitoring system is required. This system should be integrated into the MLOps process, enabling automated and continuous monitoring. Regularly retraining the model with fresh data can further help mitigate performance degradation, ensuring that the model adapts to changing patterns in the data.

Moreover, ongoing monitoring facilitates the detection of concept drift, where the relationships between features and target variables change over time. By identifying concept drift early, data scientists can take appropriate actions, such as retraining the model on more recent data or updating the feature set to account for new trends.

At Statistics Finland, our efforts to identify data drift and ensure model monitoring have been ongoing. We have conducted extensive studies using various libraries and off-the-shelf solutions to explore the best approaches for our needs. Notably, we have extensively tested the capabilities of the Alibi Detect library, which offers a collection of state-of-the-art algorithms and methods that can seamlessly integrate into our existing machine learning pipelines. While we have made significant progress, we are still in the process of determining the most suitable implementation strategy for monitoring our ML models effectively. Our commitment to continuous improvement ensures that we will make informed decisions to achieve optimal model performance and maintain the highest standards of data quality and accuracy.

4.4. Explainability

Although explainability and especially explainability of deep learning is not included as an independent principle in MLOps it is still part of transparency. When using deep learning models, achieving explainability is particularly challenging due to the complex and non-linear nature of these models. However, researchers have developed various techniques to provide some level of insight into how deep learning models make decisions.

It's important to note that while methods provide some insights into model behavior, they might not offer a complete understanding of why a deep learning model makes specific decisions. Achieving full explainability in deep learning models remains an ongoing research challenge. The choice of explainability method may depend on the specific use case and the level of interpretability required for the application.

Explainability of deep learning at Statistics Finland is still in its nascent stage. As we delve into the domain of deep learning, we recognize the significance of achieving interpretability for our models. To tackle this complex challenge, we have actively sought collaborations with academic partners, including those from the FCAI (Finnish Center for Artificial Intelligence) funded by the Research Council of Finland. Explainability in deep learning is a highly theoretical and intricate area, demanding a cautious and methodical approach. We are committed to taking small yet deliberate steps to embark on this journey, working collaboratively to unravel the inner workings of the deep learning models adopted and make them as transparent and interpretable as possible. Through these efforts, we aim to enhance the trustworthiness of our AI-driven statistical insights and pave the way for responsible and ethical implementation of deep learning in official statistics.

4.5. ML Platforms

At this crucial juncture, Statistics Finland is exploring various possibilities to optimize the whole data processing infrastructure (data stack) and leverage the full potential of AI. One option is to replace our in-house ML platform with an off-the-shelf machine learning platform, allowing us to benefit from the advancements and features offered by established solutions. By shifting away from in-house development, we could streamline our processes and reduce maintenance overheads. Additionally, in parallel, we are considering the integration of pre-trained models within the cloud infrastructure of a commercial cloud vendor. While preserving the advantages of the cloud environment, this approach enables us to tap into the power of AI and benefit from the pre-trained models provided by Microsoft/Open AI. However, as we progress towards these potential solutions, we remain mindful of our data's privacy and sensitivity and the general problems with pre-trained foundation models. Also, certain data may require on-premises solutions to ensure the utmost security and compliance with regulatory requirements.

While it may not be classified as a conventional success story, Statistics Finland has shown the foresight and strategic positioning to harness the opportunities presented by AI in official statistics. By being in the right place at the right time, the organization has successfully identified and acted upon the essential elements needed for the effective implementation of AI. This proactive approach has allowed Statistics Finland to make significant strides in adopting AI technology, setting the stage for potential future successes in enhancing the accuracy, efficiency, and impact of official statistics through responsible AI practices.

5. Potential future developments in AI and ML for official statistics

The advent of artificial intelligence (AI) has ushered in transformative opportunities for statistical institutes, promising to revolutionize various stages of the statistical process, from data collection and analysis to validation and dissemination. In this chapter, we delve into the exciting possibilities that AI holds for statistical data production, exploring potential applications, benefits, and challenges on the horizon.

One promising area where AI can make a significant impact is in automating and enhancing data collection processes. By harnessing advanced techniques like natural language processing (NLP) and computer vision, AI facilitates web scraping, social media analysis, and satellite imagery processing, providing real-time and large-scale data streams. This empowers statistical institutes to monitor and respond more effectively to dynamic changes in society. These can improve existing statistics by providing new data sources against which they can be analysed.

Machine learning algorithms, such as deep learning and neural networks, offer powerful tools for data analysis and modeling. AI-driven analytics can identify intricate patterns, correlations, and anomalies in large datasets, enabling statistical institutes to gain deeper insights into complex socio-economic phenomena and forecast trends with improved accuracy. These can lead to new statistical products to complement more established statistical outputs that are based on well-established frameworks.

AI's influence extends to data validation and quality control processes, streamlining the often-time-consuming manual validation and editing tasks. With automated anomaly detection algorithms, potential errors, inconsistencies, and outliers can be swiftly identified, allowing statisticians and subject matter experts to focus on data verification and ensuring the accuracy of statistical outputs.

Furthermore, AI-powered data visualization tools pave the way for real-time and interactive data dissemination. Dynamic visualizations, dashboards, and geospatial mapping enhance data accessibility and understanding for decision-makers and the public, fostering greater engagement with official statistics.

In survey design and sampling methodologies, AI optimization can contribute once scientifically validated. AI-enabled adaptive surveys can dynamically adjust questions based on respondents' previous answers, ensuring personalized and relevant data collection experiences. This leads to more efficient and representative data collection, enhancing the quality of statistical insights.

It is crucial to recognize that AI should not completely replace human expertise but rather complement it in highly specialized field such as official statistics, once the low-hanging fruit of eradicating repetitive tasks is achieved. Statistical institutes can adopt a human-in-the-loop approach, where AI collaborates harmoniously with statisticians, subject matter experts, and decision-makers. This collaborative synergy facilitates more robust analysis, validation, and interpretation of AI-driven statistical outputs.

6. Challenges and Ethical Considerations of the use of foundation models

Foundation models represent a significant shift in the paradigm of AI because they introduce a new approach to language understanding and knowledge representation. Traditionally, AI models were designed with specific tasks in mind and required extensive fine-tuning to perform well on those tasks. However, foundation models, such as large language models like GPT models, are pre-trained on vast amounts of data from the internet, learning the structure of language and general knowledge in an unsupervised manner.

The shift towards foundation models has democratized access to advanced AI capabilities, allowing developers to access state-of-the-art language understanding with minimal effort. It reduces the barriers to entry for AI development and accelerates the pace of innovation in natural language processing tasks.

While leveraging pre-trained models can significantly enhance the efficiency of statistical analysis and insights generation, it's essential to note that pre-trained models may not fully align with the specific needs of official statistics. Customization options might be limited and adapting the models to the specific nuances of official statistical domains could become challenging.

The lack of knowledge about the training data and process raises transparency and explainability concerns. Especially for statistical offices, it is crucial to ensure that the models' underlying data and algorithms align with ethical guidelines and produce interpretable outputs. Hosting and training models on external platforms raise data privacy and security considerations, particularly since official statistics often deal with sensitive data. Sharing such data with external providers requires robust data protection measures.

Dependence on external platforms and models necessitates consideration of the long-term sustainability of such arrangements. It is important to assess potential risks and formulate contingency plans in case the platform or provider becomes unavailable or discontinues services. Delegating model training to external entities requires clear accountability and data governance mechanisms to ensure that the models adhere to official statistical standards and legal frameworks.

Mitigating biases in the usage of foundation models is crucial for ensuring equitable and unbiased official statistics. Efforts to improve the interpretability and fairness of foundation models are essential. Investing in research and development for domain-specific pre-trained models tailored to official statistics can also help address some of the limitations.

To address these challenges effectively, statistical offices should strike a balance between using pre-trained models and maintaining control over the process. Collaborations with external providers should involve transparent agreements on data sharing, data privacy, and explainability. Prioritizing Responsible AI and MLOps principles like transparency, reproducibility, and accountability is essential to ensure the ethical implementation of AI in official statistics. Adopting open-source frameworks and platforms allows for better scrutiny and adaptation to specific needs while ensuring adherence to best practices.

Conclusions

The implementation of AI in official statistics presents both exciting opportunities and critical challenges. The paradigm shift brought by foundation models has democratized access to advanced AI capabilities, accelerating innovation in natural language processing tasks and other AI applications. However, as we venture into this new frontier of AI, it is crucial to address the ethical implications and challenges that arise. Fairness, transparency, explainability, and data privacy must be prioritized to ensure AI-driven statistical outputs are reliable, equitable, and unbiased. The needs of official statistics can be more specific or differ from what the developers of AI models have in mind. It remains a challenge to adopt what works for official statistics and discard the rest. It is certain however, that high speed of development of AI will continue to shape how we work and what is expected from us.

In the future, the responsible and ethical implementation of AI in official statistics will continue to be a priority. As we navigate the complexities and challenges, we must remain committed to upholding the highest ethical standards, fostering public trust, and contributing to evidence-based decision-making. By leveraging the power of AI responsibly with MLOps, we can shape a brighter and more AI-driven future for official statistics.

Keywords: MLOps, transparency, Artificial Intelligence, Machine Learning

References

Kreuzberger, D., Kühl, N., Hirschl, S. (2023) *Machine Learning Operations (MLOps): Overview, Definition, and architecture.* IEEE Xplore Full-Text PDF:

Google Cloud MLOps. MLOps: Continuous delivery and automation pipelines in machine learning | Cloud Architecture Center | Google Cloud

ml-ops.org ML Ops: Machine Learning Operations (ml-ops.org)