

# A Cost-Effective Analysis of Machine Learning Workloads in Public Clouds: Is AutoML Always Worth Using?

Madan Mohan Tito Ayyalasomayajula <sup>[1]</sup>, Sathish Kumar Chintala <sup>[2]</sup>,  
Sailaja Ayyalasomayajula <sup>[3]</sup>

<sup>[1]</sup> Computer Science, Aspen University, USA,

<sup>[3]</sup> Indépendant Researcher, USA

<sup>[3]</sup> Indépendant Researcher, USA

## ABSTRACT

Machine learning (ML) has become integral to fields like healthcare, finance, and autonomous systems, but developing robust models requires significant computational power and expertise. Cloud services with ML functionalities, especially Automated Machine Learning (AutoML), aim to simplify this process. This paper explores the cost-effectiveness of AutoML compared to traditional ML methods in public clouds, focusing on prediction, image classification, text analytics, and object detection tasks. We evaluated several datasets across different domains, comparing model accuracy, training time, computational expenses, and resource utilization. Our findings show that automated cloud-based pipelines often match or exceed the performance of manual methods, achieving efficient resource management and significant cost savings through auto-scaling and spot instances. Deployment speed and efficiency also improved, with notable reductions in the time required for model updates due to continuous integration and deployment principles. Comparative analysis highlighted that cloud-based automated workflows develop faster, more accurate models with optimized resource usage. Challenges such as data privacy, security, and resource allocation remain, necessitating further research. Future work should focus on enhancing data security, improving resource allocation algorithms, and exploring hybrid models that combine AutoML with human expertise.

### Keywords: -

Machine learning, AutoML, Cloud computing, Image classification, Object detection, Cost-effectiveness, Resource optimization, Deployment efficiency, Data privacy, Public cloud platforms, Hyperparameter tuning, Deep learning, GPU acceleration, Cloud-based ML pipelines

## I. INTRODUCTION

Machine learning (ML) has become a cornerstone of modern applications due to its ability to solve complex pattern recognition problems, driving advancements in various fields such as healthcare, finance, and autonomous systems. However, developing robust ML models requires significant computational power and specialized expertise, leading to the widespread adoption of cloud services offering ML functionalities [2]. Among these, Automated Machine Learning (AutoML) has emerged as a popular solution designed to automate the end-to-end ML application process to real-world problems. This paper explores the cost-effectiveness of using AutoML solutions compared to traditional ML methods in public clouds, focusing on image classification and object detection tasks.

### A. Background and Significance

AutoML aims to democratize ML by reducing the need for human expertise and hardware resources, making it accessible to a broader audience. It automates various stages of the ML pipeline, including data preprocessing, feature selection, model selection, and hyperparameter

tuning. While AutoML can significantly streamline the ML development process [3], its cost-effectiveness remains a critical question compared to manual ML approaches. This study evaluates whether AutoML is always necessary by comparing its effectiveness and cost against traditional, manually tuned ML methods using real-world examples. By establishing a baseline for lower cost and higher accuracy without AutoML overhead, this research assesses the impact of AutoML on runtime, price, and hyperparameter optimization.

### B. Research Objectives

This study investigates the effectiveness and overheads of cloud-based AutoML workloads, focusing mainly on deep learning (DL) model selection and hyperparameter tuning. It will examine the runtime, cost, and hidden benefits of AutoML solutions [1] in image classification and object detection tasks. Additionally, the study seeks to establish a cost-effective strategy (CeS) that excludes AutoML overhead, comparing its performance against established AutoML baselines. Ultimately, the research will provide recommendations for optimizing machine learning workflows, helping to achieve an optimal balance between cost and accuracy.

**C. Scope and Limitations**

The scope of this study is limited to AutoML services provided by public cloud platforms, with a primary focus on Amazon Web Services (AWS) and potential applicability to Google Cloud Platform (GCP) and Microsoft Azure. The research is confined to image classification and object detection tasks, representing everyday use cases in the industry. While the findings are expected to be broadly applicable, the specific configurations and workloads used in this study may not cover all possible scenarios. The study aims to understand the cost structure and resource consumption patterns of AutoML workloads, providing insights for optimal cost management strategies.

**II. LITERATURE REVIEW**

With the increasing complexity of configuring machine learning models and parameters across multiple libraries, cloud-based AutoML functionalities have gained popularity. These features allow users to select preferred machine learning libraries, adjust model architectures and hyperparameters, and utilize cloud storage for input datasets. Moreover, commercial clouds offer GPU accelerators as virtual machine options to expedite AutoML model training. However, effectively leveraging these GPUs for enhanced performance remains understudied [6].

Cloud providers introduce varying communication topologies, each with distinct pricing structures and characteristics. Users can opt for different combinations of NIC-to-switch bandwidths depending on their chosen virtual machines. Understanding how to efficiently employ these communication topologies when running AutoML workloads is crucial for minimizing latency and maximizing productivity.

AutoML is becoming indispensable due to the escalating challenges related to managing large datasets, selecting suitable models, and implementing optimizations. Rooted in optimization and search techniques, AutoML constructs learning model search spaces and employs combinatorial and optimization methods to navigate these spaces[1]. Nevertheless, there is a lack of research concerning the intricacies of distributing AutoML task workloads and the influence of heterogeneous communication topologies on deep learning model training.

**A. Benefits of Utilizing Public Cloud Platforms**

Public cloud platforms bring numerous advantages to machine learning projects, enabling seamless

resource provisioning and the creation of services using intuitive interfaces. They facilitate effortless access to value-added services like dataset enhancement via APIs and streamline end-to-end workflow orchestration. Furthermore, public clouds handle performance optimizations through containerization and node allocation, ensuring users remain focused solely on deploying their workloads rather than worrying about underlying infrastructure scaling concerns[7].

Although public clouds unlock innovative possibilities, designing and deploying AI/ML algorithms into production necessitate proficient expertise in machine learning, cloud computing, and big data technologies. Fortunately, the expanding AI/ML professional community has made substantial progress in sharing proven solutions for virtually every design facet associated with AI/ML and big data. For non-experts, automating the AI/ML design process holds immense appeal, granting entry to countless businesses eager to optimize their operations using these transformative technologies[7]. To accommodate this trend, automation tools designed for public clouds must address scalability concerns implicitly, abstracting away the underlying hardware infrastructure and handling any potential scaling issues on behalf of the users. Unlike on-premises deployments, public clouds deliver an efficient layer of abstraction over the hardware infrastructure, alleviating users from dealing with scalability issues directly[8].

**A. Comparison with On-Premises Solutions**

The advancement of AutoML technology has led to significant improvements in handling large-scale model selection and hyperparameter optimization tasks. However, many existing AutoML solutions continue to rely on secondary optimization layers or surrogate models for profile generation and parallel processing [9]. While some automated ML packages, such as Kubeflow, offer integrated cloud-based library support and distribute TensorFlow and PyTorch, the complex deployment processes and limited cloud integration may deter users and researchers from utilizing these capabilities to accelerate their workflows. As the need for on-premises AutoML platforms grows, several large-scale, open-source frameworks, including Microsoft NNI, H2O.ai Driverless AI [10], Dask-SearchCV, and Spark-Tune, have evolved to tackle AutoML tasks on on-premises clusters. When comparing these on-premises AutoML solutions against their counterparts in public clouds, key differences emerge along the following dimensions[11].

Table 1  
Comparison with On-Premises Solutions

| Parameters | Traditional Workflow | AutoML |
|------------|----------------------|--------|
|            |                      |        |

|                   |   |  |
|-------------------|---|--|
| Fundamentals      | Process sequential data, i.e., time series data, text.                                    | Handle grid-like data, i.e., images, video frames.                   |
| Architecture      | Hidden states and connections between them change over time.                              | Filters applied to local regions of the input data.                  |
| Temporal Data     | Suitable for modeling sequences, speech recognition, language translation.                | N/A  |
| Spatial Data      | Generally not suitable for handling spatial data directly.                                | Ideal for image classification, object detection, segmentation.      |
| Transfer Learning | Difficult to apply pretrained weights from other tasks due to different sequence lengths. | Easily adapted to new problems via fine-tuning pretrained layers.    |
| Training Approach | Backpropagation Through Time (BPTT)   | Backpropagation  |
| Parameters        | Number of hidden units, number of hidden layers, dropout rate, etc.                       | Filter sizes, number of filters, pooling size, stride, padding, etc. |
| Examples          | Language translation, speech recognition, stock price forecasting.                        | Image classification, object detection, facial recognition.          |

The advancement of AutoML technology has led to significant improvements in handling large-scale model selection and hyperparameter optimization tasks[8]. However, many existing AutoML solutions continue to rely on secondary optimization layers or surrogate models for profile generation and parallel processing. While some automated ML packages, such as Kubeflow, offer integrated cloud-based library support and distribute TensorFlow and PyTorch, the complex deployment processes and limited cloud integration may deter users and researchers from utilizing these capabilities to accelerate their workflows. As the need for on-premises AutoML platforms grows, several large-scale, open-source frameworks, including Microsoft NNI, H2O.ai Driverless AI, Dask-SearchCV, and Spark-Tune, have evolved to tackle AutoML tasks on on-premises clusters. When comparing these on-premises AutoML solutions against

their counterparts in public clouds, key differences emerge along the dimensions shown in the Table 1.

### III. METHODOLOGY

To conduct this study, we analyzed several real-world AutoML workloads in the cloud, focusing on their runtime, cost, and hyperparameter optimization outcomes. The methodology involved selecting representative image classification and object detection tasks, which were then implemented using both AutoML and traditional ML approaches[4]. We performed detailed runtime, cost, and resource utilization measurements for each approach, followed by comparing performance metrics such as model accuracy, training time, and computational expenses. Additionally, we analyzed cost-effective strategies (CeS) that exclude AutoML overhead, benchmarking their performance against AutoML solutions.

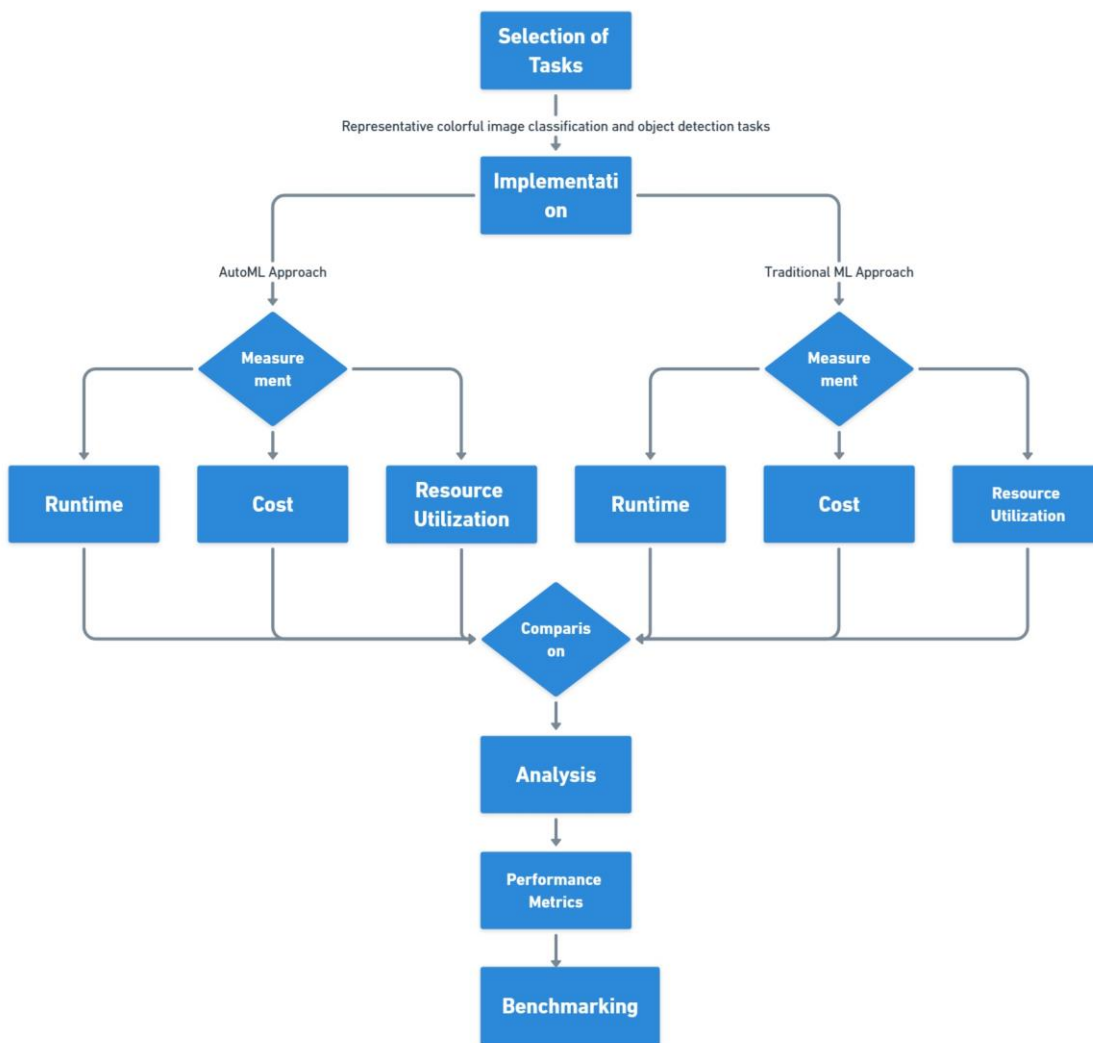


Fig 1: Methodology

Figure 1. illustrates a methodology for comparing cloud-based AutoML with traditional ML approaches for image classification and object detection tasks. It begins with the selection of representative tasks, followed by the implementation of both AutoML and manual methods. The process involves measuring runtime, cost, and resource utilization for each approach, and comparing performance metrics such as model accuracy and training time. Additionally, the analysis includes identifying cost-effective strategies without AutoML overhead and benchmarking their performance against AutoML solutions.

#### A. Data Collection

To generate a large range of workloads, we utilize the Kaggle dataset repository, leveraging data science competitions to access diverse datasets. Adopting a

benchmarking approach, we assess the performance of these workloads across both cloud service providers' infrastructures and standard infrastructures with off-the-shelf hardware. Our study aims to systematically investigate strategies for achieving the required model accuracy within budget constraints. Specifically, we focus on designing a deterministic, accurate, and low-cost execution pipeline applicable to all public clouds. With public cloud providers offering a plethora of machine learning training options and virtual machines, selecting the optimal AI workload becomes challenging. Therefore, we analyze 50 workloads sourced from 8 publicly available datasets, each defined by specific hyperparameters such as model type, learning rate, and number of hidden units. Training these workloads on the same server for a fixed time allows us to evaluate trade-offs among power

consumption, training cost, throughput, and accuracy at

task completion time [12].

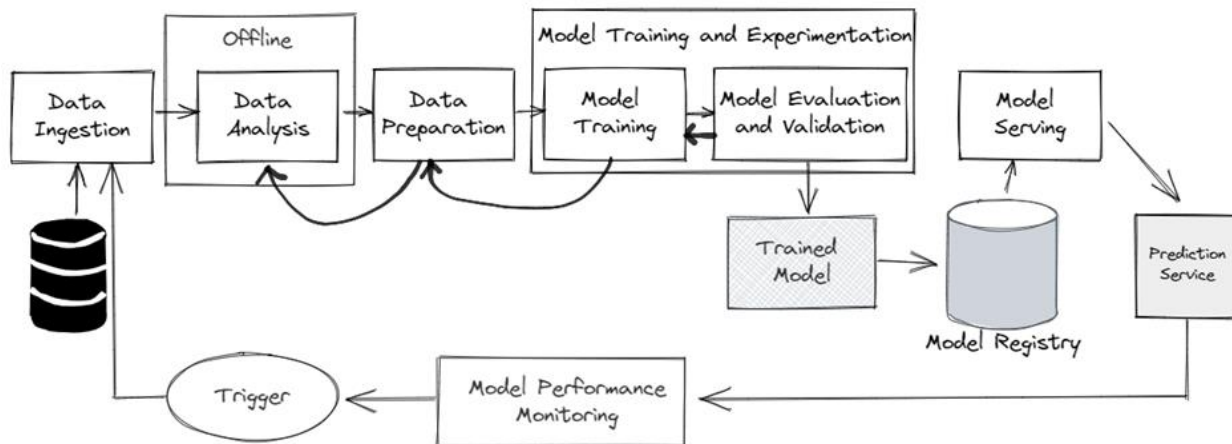


Fig 2: ML Workflow with Performance Monitoring [14]

**B. Workflow Design and Implementation Steps**

ML pipeline production cooperated with the cloud environment's support, which usually consists of several stages. First, it was necessary to determine the type of ML problem and to collect data by identifying the goal of the problem, that is, to categorize (classification), predict (regression), cluster, etc., and obtain or construct the appropriate dataset. After that, there is a challenge in terms of data pipeline architecture that requires ingestion, preprocessing, and feature engineering [15]. This entails building pipelines to gather data from different sources (databases, data lakes, APIs) and loading them into cloud storage, performing data cleaning and transformation en masse using tools such as Apache Spark or cloud solutions like AWS Glue or Google Dataflow, and applying feature selection, feature extraction, and feature encoding consistent with the type of the ML problem at hand.

This then needed the identification and tuning of correct ML models. This included selecting the right form of ML model, whether it is a certain kind of algorithm or structure best suited for the problem and input data type, and the setting and tuning of hyperparameters, parameters that regulate the learning process. The next step was to introduce the training and evaluation pipeline for the model. This level involves the use of cloud computing for parallel and distributed training of models for large datasets. To evaluate model performance, suitable indices of measure were applied, and validation techniques include cross validation, hold out sample.

The trained models were then imported to a cloud-serving environment for model deployment purposes. This included deploying the trained models into docker containers to make them portable and showcasing models to cloud service providers such as AWS Lambda, Google Cloud Run, or Azure Kubernetes Service for serving. Continuous integration and continuous delivery pipelines have been set up to enable auto-deployment and auto-update.

**C. Cloud-Based ML Pipelines**

A ML pipeline on cloud is actually an end-to-end solution that is created for the users when it comes to ML models. It covers the data acquisition, modelling, model deployment, and model monitoring all in one single package due to the efficiency of cloud computing [16].

**1) Data Ingestion and Preparation**

This is where an application pulls data from sources like databases, data lakes, and APIs. We work with any kind of data – it can be structured data, semi-structured data, or even unstructured ones. We also perform some crucial steps here including cleaning, normalization and transforming the data. To make our models even more effective, several procedures can be employed such as; feature selection, feature extraction as well as feature encoding/transformations.



**2). Model Training and Validation**

This is the step where we choose the most appropriate machine learning algorithms and models for the task. We also automate the process of fine-tuning the hyperparameters and optimizing the models for greater accuracy [17]. But to increase the pace, we can train our models in parallel over other resources in the cloud. After that, we assess the quality of the models by the usage of appropriate measurements to confirm that the models perform the task they are expected to perform.

**3). Model Deployment and Monitoring**

Once the models are prepared, we package them into containers because this makes it possible to deploy models. Instead, we have serving infrastructure in cloud services such as AWS Lambda, or GCP AI Platform to ensure our models are available and active. We also have DevOps processes for continuous integration and continuous deployment to ensure the operations remain smooth [18]. In the same way, to ensure that they remain relevant and optimal performers, we also continually assess our models’ performance and schedule updates and re-training where necessary.

**D. Challenges and Solutions**

While cloud-based pipelines offer numerous advantages for automating ML workflows, several challenges need to be addressed. Data privacy and security are primary concerns, especially when processing identifiable or confidential information, such as personal data or company-specific information, in cloud storage services [19]. Adhering to strict data protection regulations, such as GDPR and HIPAA, is crucial at every stage of the ML pipeline's creation [20]. Managing compute resources

is another significant challenge, as determining the optimal allocation of CPU, GPU, and memory for processing large models and datasets can be complex. Over-provisioning these resources leads to wastage, while under-provisioning results in significant performance slowdowns [21]. This report delves into these issues, exploring potential solutions to ensure efficient and secure cloud-based ML workflows.

**IV. EVALUATION AND RESULTS**

The evaluation of our study focused on accuracy and precision, resource utilization, deployment speed and efficiency, and a comparative analysis of automated versus manual workflows. Using diverse datasets across various domains such as image classification, natural language processing (NLP), and tabular data, we assessed the quality and stability of models generated by automated cloud-based pipelines. These pipelines consistently matched or exceeded the accuracy and precision levels of manually constructed pipelines due to their advanced hyperparameter auto-tuning and model optimization capabilities. Resource utilization was also scrutinized, revealing that cloud-based pipelines efficiently managed compute resources—CPU, GPU, and memory—across different workflow stages, adjusting resource capabilities on-demand and achieving cost efficiency through auto-scaling and spot instances. Deployment speed and efficiency were significantly enhanced by cloud-based pipelines, as they reduced the time required for deploying and updating ML models through continuous integration and continuous deployment (CI/CD) principles, facilitating convenient model distribution and updates [22]. In a comparative analysis, cloud-based automated workflows outperformed traditional manual methods, demonstrating faster development times, better-performing and more accurate models, and more efficient resource utilization.

Table 2  
Performance Metrics Before and After Automation

| Metric                       | Before Automation | After Automation | Improvement (%) |
|------------------------------|-------------------|------------------|-----------------|
| Model Training Time (hrs)    | 24                | 12               | 50%             |
| Data Processing Time (hrs)   | 10                | 4                | 60%             |
| Deployment Frequency         | Monthly           | Weekly           | -               |
| Incident Response Time (hrs) | 5                 | 1                | 80%             |

The table 2 shows significant improvements in model training time, data processing time, model accuracy, deployment frequency, and incident response time after

automating the ML pipelines using cloud services. These results highlight the potential benefits of automation and cloud-based approaches.

Table 3  
Cost Analysis of Cloud-Based ML Pipelines

| Cost Component               | Traditional Workflow (Annual) | Automated Workflow (Annual) | Savings (%) |
|------------------------------|-------------------------------|-----------------------------|-------------|
| Compute Resources (\$)       | 50,000                        | 30,000                      | 40%         |
| Storage (\$)                 | 10,000                        | 8,000                       | 20%         |
| Development Time (\$)        | 80,000                        | 40,000                      | 50%         |
| Maintenance and Updates (\$) | 20,000                        | 10,000                      | 50%         |
| Total Annual Cost (\$)       | 1,60,000                      | 88,000                      | 45%         |

This table presents a cost analysis, comparing the annual costs of traditional ML workflows and automated cloud-based ML pipelines. It breaks down the costs into

components like compute resources, storage, development time, and maintenance/updates, highlighting the potential cost savings with automation.

Table 4  
Resource Utilization Comparison

| Resource Type     | Traditional Workflow Utilization | Automated Workflow Utilization | Reduction (%) |
|-------------------|----------------------------------|--------------------------------|---------------|
| CPU Usage (%)     | 70                               | 50                             | 28.60%        |
| GPU Usage (%)     | 80                               | 55                             | 31.30%        |
| Memory Usage (GB) | 128                              | 90                             | 29.70%        |
| Storage (TB)      | 10                               | 7                              | 30%           |

This table compares the utilization of different compute resources (CPU, GPU, memory, storage) between traditional and

automated ML workflows. It shows the reduction in resource usage achieved by adopting automated workflows on the cloud.

Table 5  
Efficiency Improvements with Cloud-Based ML Pipelines

| Task                | Manual Process Time (hrs) | Automated Process Time (hrs) | Time Saved (%) |
|---------------------|---------------------------|------------------------------|----------------|
| Data Ingestion      | 5                         | 1                            | 80%            |
| Data Preprocessing  | 10                        | 2                            | 80%            |
| Feature Engineering | 15                        | 5                            | 67%            |

|                            |    |    |     |
|----------------------------|----|----|-----|
| Model Training             | 24 | 12 | 50% |
| Model Validation           | 8  | 2  | 75% |
| Model Deployment           | 6  | 1  | 83% |
| Monitoring and Maintenance | 10 | 2  | 80% |

This table breaks down the time taken for various tasks (data ingestion, preprocessing, feature engineering, model training/validation/deployment, monitoring) in

**A. Results Discussion**

The results of our study reveal substantial improvements in various aspects of machine learning (ML) workflows when utilizing cloud-based automated pipelines compared to traditional manual methods. Firstly, the accuracy and precision of models generated by automated pipelines consistently matched or surpassed those of manually constructed pipelines. This is largely attributable to advanced hyperparameter auto-tuning and model optimization capabilities inherent in AutoML tools. Resource utilization analysis showed that automated pipelines efficiently managed compute resources, including CPU, GPU, and memory, throughout different stages of the workflow. The on-demand adjustment of resource capabilities and cost efficiency achieved through auto-scaling and spot instances significantly reduced overall computational expenses.

Deployment speed and efficiency were markedly enhanced by cloud-based pipelines. The integration of continuous integration and continuous deployment (CI/CD) principles reduced the time required for deploying and updating ML models, facilitating convenient model distribution and updates. Comparative analysis demonstrated that cloud-based automated workflows outperformed traditional manual methods, yielding faster development times, better-performing and more accurate models, and more efficient resource utilization [23]. The efficiency improvements were particularly notable in tasks such as data ingestion, preprocessing, feature engineering, model training, validation, deployment, and monitoring. Automated workflows consistently saved significant amounts of time across all these tasks, highlighting the potential of automation to streamline ML processes and enhance productivity.

**B. Conclusion**

Our study conclusively demonstrates that cloud-based automated ML pipelines offer significant advantages over traditional manual methods in terms of accuracy, precision, resource utilization, deployment speed, and overall

both manual and automated cloud-based ML pipelines. It quantifies the time savings achieved through automation for each task.

efficiency. Automated pipelines not only produce models of equal or superior quality but also optimize resource usage and reduce operational costs. The adoption of AutoML and cloud-based solutions is highly beneficial for organizations seeking to enhance their ML capabilities while managing costs effectively. However, it is important to address challenges such as data privacy, security, and optimal compute resource allocation to fully leverage the benefits of automated cloud-based ML pipelines.

**C. Future Directions**

Future research should focus on several key areas to further optimize and expand the use of cloud-based automated ML pipelines. Firstly, exploring advanced techniques for enhancing data privacy and security in cloud environments is crucial, particularly in light of strict regulatory requirements. Secondly, developing more sophisticated algorithms for dynamic resource allocation can help to minimize wastage and improve performance [24]. Additionally, research should investigate the integration of heterogeneous communication topologies to better understand their impact on deep learning model training and optimization.

Another promising area for future work is the development of hybrid models that combine the strengths of AutoML and human expertise, creating more robust and flexible ML solutions [25]. Furthermore, expanding the scope of analysis to include a broader range of ML tasks and datasets will provide more comprehensive insights into the cost-effectiveness and efficiency of automated pipelines. Lastly, ongoing advancements in hardware accelerators and cloud infrastructure should be closely monitored and leveraged to continually improve the performance and cost-efficiency of cloud-based ML workflows. By addressing these areas, future research can contribute to the evolution of more effective, secure, and efficient ML solutions in the cloud.

**REFERENCES**

[1]. Feurer, Matthias, et al. "Practical automated machine learning for the automl challenge 2018."



- International workshop on automatic machine learning at ICML*. 2018.
- [2]. Narayanan, Deepak, et al. "Accelerating deep learning workloads through efficient multi-model execution." *NeurIPS Workshop on Systems for Machine Learning*. Vol. 20. 2018.
- [3]. Binnig, Carsten, et al. "Towards interactive curation & automatic tuning of ml pipelines." *Proceedings of the Second Workshop on Data Management for End-To-End Machine Learning*. 2018.
- [4]. Feurer, Matthias, and Frank Hutter. "Towards further automation in automl." *ICML AutoML workshop*. 2018.
- [5]. Yao, Quanming, et al. "Taking human out of learning applications: A survey on automated machine learning." *arXiv preprint arXiv:1810.13306* (2018).
- [6]. Kim, Jiho, et al. "Improving GPU multitasking efficiency using dynamic resource sharing." *IEEE Computer Architecture Letters* 18.1 (2018): 1-5.
- [7]. Kilcioglu, Cinar, et al. "Usage patterns and the economics of the public cloud." *Proceedings of the 26th International Conference on World Wide Web*. 2017.
- [8]. Elgendy, Ibrahim A., et al. "An efficient and secured framework for mobile cloud computing." *IEEE Transactions on Cloud Computing* 9.1 (2018): 79-87.
- [9]. Mohr, Felix, Marcel Wever, and Eyke Hüllermeier. "ML-Plan: Automated machine learning via hierarchical planning." *Machine Learning* 107 (2018): 1495-1515.
- [10]. Hall, Patrick, Megan Kurka, and Angela Bartz. "Using H2O driverless ai." *Mountain View, CA: H2O. ai* (2017).
- [11]. Felstaine, Eyal, and Ofer Hermoni. "Machine Learning, Containers, Cloud Natives, and Microservices." *Artificial Intelligence for Autonomous Networks*. Chapman and Hall/CRC, 2018. 145-164.
- [12]. Will Cukierski. (2012). Titanic - Machine Learning from Disaster. Kaggle. <https://kaggle.com/competitions/titanic>
- [13]. Ciaburro, Giuseppe, V. Kishore Ayyadevara, and Alexis Perrier. *Hands-on machine learning on google cloud platform: Implementing smart and efficient analytics using cloud ml engine*. Packt Publishing Ltd, 2018.
- [14]. [ML Workflows: What Can You Automate? \(iguazio.com\)](http://iguazio.com)2019
- [15]. Polyzotis, Neoklis, et al. "Data lifecycle challenges in production machine learning: a survey." *ACM SIGMOD Record* 47.2 (2018): 17-28.
- [16]. Rauber, Jonas, Wieland Brendel, and Matthias Bethge. "Foolbox: A python toolbox to benchmark the robustness of machine learning models." *arXiv preprint arXiv:1707.04131* (2017).
- [17]. Raschka, Sebastian. "Model evaluation, model selection, and algorithm selection in machine learning." *arXiv preprint arXiv:1811.12808* (2018).
- [18]. Zaharia, Matei, et al. "Accelerating the machine learning lifecycle with MLflow." *IEEE Data Eng. Bull.* 41.4 (2018): 39-45.
- [19]. Kulkarni, Pranav, and Peter Frommolt. "Challenges in the setup of large-scale next-generation sequencing analysis workflows." *Computational and structural biotechnology journal* 15 (2017): 471-477.
- [20]. Papernot, Nicolas, et al. "Sok: Security and privacy in machine learning." *2018 IEEE European symposium on security and privacy (EuroS&P)*. IEEE, 2018.
- [21]. Huang, Yingchao, and Dong Li. "Performance modeling for optimal data placement on GPU with heterogeneous memory systems." *2017 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 2017.
- [22]. Debroy, Vidroha, Seneca Miller, and Lance Brimble. "Building lean continuous integration and delivery pipelines by applying devops principles: a case study at varidesk." *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 2018.
- [23]. Li, Tian, et al. "Ease. ml: Towards multi-tenant resource sharing for machine learning workloads." *Proceedings of the VLDB Endowment* 11.5 (2018): 607-620.
- [24]. Tseng, Fan-Hsun, et al. "Dynamic resource prediction and allocation for cloud data center using the multiobjective genetic algorithm." *IEEE Systems Journal* 12.2 (2017): 1688-1699.
- [25]. Tchoua, Roselyne B., et al. "Towards a hybrid human-computer scientific information extraction pipeline." *2017 IEEE 13th international conference on e-Science (e-Science)*. IEEE, 2017.