UDC 004.8
DOI https://doi.org/10.24144/2616-7700.2024.45(2).238-248

M. V. Robotyshyn[1], M. M. Malyar[2]

[1] Uzhhorod National University,
PhD student
mykola.robotyshyn@uzhnu.edu.ua
ORCID: https://orcid.org/0000-0001-6567-6974

[2] Uzhhorod National University,
Professor of the Department of Cybernetics and Applied Mathematics,
Doctor of Technical Scienes
mykola.malyar@uzhnu.edu.ua
ORCID: https://orcid.org/0000-0002-2544-1959

THE USE OF MATHEMATIC MODELLING IN FINOPS DOMAIN

The advent of cloud computing has revolutionized the way organizations manage their IT resources, offering scalable and flexible solutions for various operational needs. However, with the increased adoption of cloud services, financial operations (FinOps) have become more complex, presenting challenges in cost management, resource allocation, and financial forecasting. Traditional methods often fail to address these complexities effectively, leading to inefficiencies and suboptimal decision-making. This paper explores the application of mathematical modeling in the FinOps domain as a robust solution to these challenges. By proposing a classification of FinOps problems into distinct classes and suggesting a mathematical formulation for each class, the research aims to enhance the efficacy of FinOps practices. The integration of mathematical models improves accuracy and efficiency while providing a systematic approach to managing financial operations in the cloud. Through case studies and real-world examples, this paper demonstrates the transformative potential of mathematical modelling in driving innovation and operational excellence in FinOps.

Keywords: mathematic modelling, cloud management, data analysis, forecasting, FinOps, optimization processes.

1. Introduction. The rapid advancement and adoption of cloud computing have significantly changed how organizations handle their IT infrastructure. By providing scalable and flexible solutions, cloud technology helps businesses manage their costs more effectively. As a result, companies across various industries have increasingly migrated their operations to the cloud, taking advantage of its cost-saving benefits. However, this shift has also introduced new complexities in financial management. The dynamic and often unpredictable nature of cloud costs has necessitated the development of specialized practices to optimize financial operations in this new environment. This need has given rise to the field of Financial Operations, or FinOps, which focuses on managing cloud spending and maximizing the value derived from cloud investments. FinOps combines financial management principles with the agile, data-driven methodologies inherent to cloud computing, ensuring that organizations can maintain control over their cloud expenses while achieving their strategic objectives.

A central focus of FinOps is to ensure that organizations can effectively manage and optimize their cloud spending. The primary goals of FinOps include cost optimization, financial accountability, and operational efficiency. By fostering a culture of financial responsibility, FinOps encourages collaboration between finance and engineering teams, ensuring that both sides work together to achieve cost-effective cloud usage. This collaborative approach is guided by core principles such as data-driven decision-making and continuous improvement. FinOps practitioners rely on real-time data and analytics to monitor cloud expenses and make informed decisions that align with the organization's financial objectives. Without a structured FinOps approach, organizations often struggle with overspending, lack of transparency, and inefficiencies in their cloud operations.

These challenges underscore the necessity of FinOps in maintaining control over cloud costs and maximizing the value derived from cloud investments. However, implementing FinOps practices is not without its difficulties, especially when it comes to data analysis. The large amount of data generated by cloud environments can be hard to manage and analyze quickly. This can lead to delays in decision-making and overspending. Another major challenge is predicting and controlling cloud expenses, which can change a lot based on usage patterns and pricing models. This makes it hard for organizations to forecast costs accurately and stay within their budget. Additionally, combining financial data with operational metrics adds another layer of complexity. Effective FinOps requires a complete view that includes both financial information and operational data, like resource usage and performance metrics. However, this is often difficult due to data silos and inconsistent data formats. Mathematical concepts and techniques can help solve these problems. For example, statistical models can forecast cloud costs by looking at past spending patterns. Optimization algorithms can help allocate resources more effectively, ensuring that organizations get the most value from their cloud investments. Machine learning models can also identify unusual patterns and inefficiencies in cloud usage, helping to manage costs proactively and improve operations. By using these mathematical tools, FinOps practitioners can overcome data analysis challenges and manage financial operations in the cloud more effectively.

The relevance of this problem is best demonstrated by the amount of money businesses are willing to invest in its research. The global cloud FinOps market was valued at approximately \$832.2 million in 2023 and is expected to grow to \$2,750.5 million by 2028, reflecting a compound annual growth rate (CAGR) of 18.8% [1]. This significant growth underscores the increasing importance of FinOps as organizations seek to optimize their cloud spending and enhance financial accountability. As businesses continue to migrate to cloud environments, the demand for effective FinOps strategies will only increase. This trend highlights the need for advanced solutions, such as mathematical modeling, to address the complex challenges associated with cloud financial management. Furthermore, there are numerous startups leveraging data analysis, AI, and ML to enhance FinOps practices. These startups are driving innovation by providing tools and platforms that improve cost visibility, optimization, and forecasting. The growing number of these tech-driven startups indicates a strong market trend towards integrating advanced analytics and automation in FinOps [2].

Implementing advanced solutions grounded in mathematical modeling can significantly streamline FinOps processes, allowing companies to shift their focus to other important strategic problems. However, it's important to recognize that such automation isn't a cure-all. While mathematical models and algorithms can automate complex tasks like cost forecasting, resource optimization, and anomaly detection, their effectiveness largely depends on the quality of data and the robustness of the models used. Automating these processes can indeed reduce manual effort and improve efficiency, freeing up valuable time and resources. This enables teams to address other critical issues, such as innovation, customer experience, and market expansion. However, the initial setup and ongoing maintenance of these advanced solutions require significant investment in both time and expertise. Organizations must ensure they have the right talent and infrastructure to support these technologies.

For example, advanced mathematical models and algorithms can continuously monitor and optimize cloud spending, providing real-time insights and proactive recommendations. This allows FinOps practitioners to move from a reactive approach to a more strategic and proactive one, focusing on long-term financial health and operational excellence. Nevertheless, reliance on automated solutions also requires vigilance; without proper oversight, there's a risk of over-reliance on these tools, potentially missing out on nuanced, human insights.

In general, by leveraging these advanced solutions, companies can achieve a higher level of financial management maturity, ensuring sustainable growth and competitive advantage in the rapidly evolving cloud landscape. However, it's crucial to balance automation with human oversight to maximize the benefits and mitigate potential risks.

This work focuses on providing a comprehensive understanding of the FinOps domain, highlighting the main challenges it faces, and exploring how mathematical modeling can offer solutions. It includes an overview of existing optimization techniques that leverage mathematical modeling and various algorithms to address these challenges. Additionally, this paper contributes a unique classification of FinOps problems based on a mathematical perspective, considering the different types of data used to solve these problems. In the author's opinion, the application of mathematical modeling in FinOps is highly relevant in today's context and continues to trend upward. The insights gained from this research can guide organizations in implementing more effective FinOps practices, ensuring better financial management and operational efficiency in the cloud environment.

2. Different classes of problems in FinOps. The FinOps domain is still evolving, and there is no definitive way to classify the various problems it addresses. Different organizations may prioritize different aspects based on their unique needs and perspectives. For instance, some might focus on operational efficiency and cost control, while others might emphasize the development of robust solutions using advanced mathematical models. Additionally, the type and quality of data available for analysis can also influence how these problems are classified.

The classes defined by the authors in this paper are based on a mathematical perspective, also taking into account the different types of data used to solve these problems. This approach aims to provide a structured method for understanding and addressing FinOps challenges. These classifications are not absolute, and other approaches can be equally valid and useful. Furthermore, the authors' perspective is dynamic and evolves with new innovations and insights in the field. The key is to apply the most relevant methods to effectively manage and optimize financial operations in the cloud.

2.1. Resource utilization. Resource utilization is a common problem for companies that heavily rely on cloud services. The goal is to use exactly as much as you need at any given moment and not more, because in the cloud, you typically pay for what you use. Depending on a company's maturity, this problem can vary in scale. Efficient resource utilization is crucial for maintaining cost-effective and high-performance cloud operations. Underutilization and overprovisioning are common issues that can lead to unnecessary expenses and reduced operational efficiency.

One major challenge is resources running idle or using significantly less than their capacity. This often happens when resources such as EC2 instances are provisioned based on peak demand but remain underutilized during off-peak times, leading to wasted expenditures. Other common scenarios include forgetting to turn off instances after use, over-provisioning storage volumes, and leaving unused services running. These situations result from not accurately aligning resource provisioning with actual workload requirements, causing financial inefficiencies and wasted cloud spending. To address these issues, data on compute instances such as CPU usage, RAM utilization, and disk I/O can be collected. This time-series data provides a continuous stream of performance metrics over time, enabling detailed analysis. Various methods, from simple statistical techniques to complex machine learning models, can be applied to this data to optimize resource utilization.

Detecting patterns in utilization data is essential for identifying inefficiencies and opportunities for optimization. By analyzing these patterns, organizations can better understand usage trends and predict future demands. This allows for more accurate provisioning of resources, ensuring they are used efficiently. Predictive analytics uses historical data and machine learning algorithms to forecast future resource demands. By predicting usage patterns, organizations can better plan and allocate resources to match demand, reducing idle time and avoiding over-provisioning.

To formalize this approach, let's define an optimization metric and an algorithm. For a single resource, the optimization metric could be the cost over a given day C_d:

$$C_d = \sum_{i=1}^{24} Cost\,(R, t_i),$$

where $Cost\,(R, t_i)$ represents the cost of resource R at hour t_i.

For multiple resources, or a workload, the cost can be extended to consider all resources over a 30-day period C_{30}:

$$C_{30} = \sum_{d=1}^{30} C_d.$$

Then we can define f as an algorithm that takes various data sources D (such as CPU usage, RAM utilization, disk I/O, historical data, and forecasted demand) and outputs optimized resource allocations. The cost function can be expressed as:

$$minimize\ f(C_{30}, D),$$

where the objective is to minimize 30-day cost given the data D. Now, with this mathematical problem statement, we can use different algorithms as f and different data sources or even different time granularities for D. The reason cost was selected as the metric is that we usually focus on reducing cost. Even though other metrics could be applied, this one is suitable for quantifying and explaining because cost is a key concept in FinOps.

2.2. Cloud cost forecasting. In mature companies, planning cloud spending in advance is a standard practice, which creates a necessity for accurate cost forecasting. It is essential to ensure that actual spending aligns with these forecasts. Accurate forecasting helps companies to manage their budgets, allocate resources effectively, and avoid unexpected expenses. Forecasting is a well-known problem in the mathematical world, with many established methods available to tackle it. These methods include time series analysis, regression models, machine learning algorithms, and stochastic processes. However, cloud cost forecasting is particularly challenging due to the multiple dimensions involved. These dimensions include the variability in cloud usage patterns, different pricing models offered by cloud providers, the impact of reserved instances and spot instances, seasonal trends, and sudden changes in demand or application performance. The complexity of these factors makes it difficult to create accurate forecasts, requiring sophisticated models that can account for the multifaceted nature of cloud cost data.

In paper [3], the authors investigated various methods such as ARIMA time series models and machine learning algorithms to forecast cloud costs. They found that machine learning models, particularly those using neural networks, provided better accuracy in dynamic environments. Another study [4] explored the use of regression models and ensemble methods for cloud cost forecasting, highlighting that combining multiple models often leads to more reliable predictions. The main finding was that hybrid models combining statistical methods with machine learning techniques can handle the complex patterns in cloud cost data more effectively. Additionally, a study by Hofmann and Rutschmann [5] emphasized the integration of different data sources in demand forecasting. They showed that big data analytics could significantly enhance forecast accuracy by utilizing diverse data inputs, including historical data, performance metrics, and market trends. Another relevant work [6] reviewed various AI-based forecasting methods in financial accounting, finding that hybrid models combining support vector machines, fuzzy logic, and genetic algorithms provided reliable forecasts. Access to diverse and comprehensive data is crucial for accurate forecasting. Several open data sources can be utilized for cloud cost forecasting. The FOCUS (FinOps Open Cost and Usage Specification) initiative is an open-source specification that aims to standardize cost and usage billing data across different cloud vendors. This initiative helps reduce complexity for FinOps practitioners by providing consistent datasets for analysis [7]. Additionally, platforms like Kaggle offer various datasets that can be used for machine learning projects, including cloud cost forecasting [8]. Amazon Forecast also provides tools and templates for deploying time-series forecasting models, which can utilize historical usage data stored in Amazon S3 [9]. Google BigQuery public datasets offer another resource for accessing diverse datasets that can be integrated into forecasting applications [10]. These open data sources provide the necessary historical data and contextual information needed to develop robust forecasting models that can accurately predict future cloud costs.

From a mathematical standpoint, forecasting is a well-defined problem. It involves predicting future values based on historical data using various statistical and machine learning methods. The core idea is to build a model that can understand

and extrapolate the patterns in the historical data to make accurate future predictions. Common methods include time series analysis, regression models, and machine learning algorithms like neural networks. However, when you apply forecasting algorithms to different domains, the mathematical formulation can be adjusted to better suit the specific characteristics and requirements of that domain. Each domain may have unique factors that influence the forecasting process, necessitating modifications to the standard forecasting models.

FinOps is no exception, and because companies plan their budgets, we need to consider these unique factors in our forecasting models. Cloud cost forecasting must account for various dimensions such as seasonality (e.g., summer, winter), company growth plans, cost dynamics, team growth, business domain, and other variables depending on the company. These factors are often difficult to quantify numerically, adding complexity to the forecasting models. This makes it necessary to adapt traditional forecasting techniques to the specific needs of FinOps.

To formalize the problem, we consider the following mathematical formulation. Let Y_t represent the cloud cost at time t. The goal is to predict future costs $Y_{t+1}, Y_{t+2}, \ldots, Y_{t+h}$ for a forecast horizon h. The feature set X_t includes variables such as season, company growth plans, cost dynamics, team growth, business domain, and other variables depending on the company. The forecasting model f

$$Y^{\wedge}_{t+h} = f(Y_t, Y_{t-1}, \ldots, Y_{t-p}, X_t, X_{t-1}, \ldots, X_{t-p}),$$

where p is the number of lag observations considered in the model. The objective is to minimize the forecast error, which can be measured using metrics such as Mean Absolute Error (MAE), Mean Squared Error (MSE), or Root Mean Squared Error (RMSE):

$$\sum_{i=1}^{n} L\left(Y_{t+i}, Y^{\wedge}_{t+i}\right) \to \min,$$

where L is the loss function, and n is the number of forecasted time points. By defining the problem this way, we can leverage existing forecasting techniques while incorporating the unique aspects of cloud cost management. This approach ensures that our models are tailored to the specific needs of FinOps, providing more accurate and reliable cost forecasts.

This forecasting definition is applicable not only to the overall cost but also to costs for specific services, cloud accounts, compute types, and other granular levels of cloud expenditure.

2.3. Anomaly Detection. In mature companies, monitoring cloud spending in real-time is essential to ensure financial efficiency and prevent unexpected costs. Despite careful planning and forecasting, actual cloud expenses can sometimes deviate significantly from expectations due to various reasons such as sudden changes in usage patterns, misconfigurations, or even security breaches. Detecting these anomalies quickly is crucial to mitigate potential financial losses and operational disruptions. Anomaly detection mechanisms play a vital role in identifying unusual spending patterns that deviate from the norm, allowing companies to take immediate corrective actions.

Anomaly detection is closely tied to cloud cost forecasting. Effective forecasting provides a baseline of expected costs and usage patterns, which anomaly detection systems use to identify deviations. While forecasting aims to predict future costs based on historical data, anomaly detection focuses on identifying unexpected variations in real-time. Although they are interrelated, we consider them distinct classes of problems because forecasting primarily addresses planning and budgeting, whereas anomaly detection deals with operational monitoring and immediate response to deviations.

As discussed in the cloud cost forecasting chapter, the vast amount of data generated in real-time from cloud services, combined with the variability in usage patterns, presents a significant challenge for anomaly detection. Additionally, integrating anomaly detection systems with existing monitoring and management tools is crucial. These systems need to provide real-time alerts and insights while fitting seamlessly into the organization's current infrastructure. Ensuring compatibility and smooth integration can be complex but is essential for effective anomaly detection and timely responses to potential issues. However, the implementation details of such integrations are a separate topic of discussion and are beyond the scope of this paper.

To better explain the anomaly detection problem in the FinOps domain, let's focus on the cost of compute instances (EC2) rather than the whole cloud cost. This does not mean we lose any information or limit our scope. The principles applied to EC2 can be extended to other services and resources, preserving the original problem statement's generality. Anomaly detection in this context revolves around three key variables: Baseline, Threshold, and Time-Granularity.

The *baseline* represents the expected cost, which can be defined using statistical measures (such as average or median) or based on values provided by domain experts. This helps establish a reference point against which current costs can be compared. The *threshold* defines the acceptable deviation from the baseline before an anomaly is flagged. It can also be determined using statistical measures (such as standard deviation or the 75th percentile) or based on expert judgment. This value helps to distinguish between normal fluctuations and actual anomalies.

Anomaly detection can be applied at different *time granularities*: hourly, daily, weekly, etc. It's important to understand the chosen granularity because an anomaly might be detected on an hourly level but not on a daily level due to the dynamic nature of workloads. For example, a high cost at a particular hour might be offset by lower costs in other hours, making it important to consider the context in which anomalies are detected.

Define the baseline $Y_t^\wedge$ using an appropriate statistical measure or expert input:

$$Y_t^\wedge = StatisticalMeasure\,(Y_t, Y_{t-1}, \ldots, Y_{t-n}).$$

The anomaly score S_t can be calculated as the difference between the current cost Y_t and the baseline $Y_t^\wedge$:

$$S_t = |Y_t - Y_t^\wedge|.$$

Set or calculate a threshold θ to determine when a deviation is considered an anomaly:

$$Anomaly\ if\ S_t > \theta.$$

This method is straightforward and computationally efficient. It provides a clear and immediate indication of anomalies based on recent cost patterns without the need for complex forecasting models. The simplicity of this approach makes it well-suited for real-time monitoring and quick response to deviations in cloud spending. Even though we defined the anomaly score as an absolute difference, in practice, we will likely care only about cases where costs are higher and not lower.

A cascade architecture can be used to enhance anomaly detection. In this approach, anomalies are evaluated at multiple levels of granularity. For example, an anomaly detected at an hourly level (Anomaly Level 1) may prompt immediate investigation. However, if this anomaly does not affect the daily cost (Anomaly Level 2), it might be considered a transient fluctuation rather than a true anomaly. This method ensures that short-term anomalies are validated over longer periods, reducing false positives and providing a clearer understanding of cost deviations. The authors believe that due to the unique nature of FinOps, where costs are visible at different levels and need constant monitoring, the cascade architecture is particularly well-suited.

Implementing anomaly detection in cloud cost management involves several practical considerations, such as ensuring data quality and consistency through pre-processing steps like normalization, handling missing values, and outlier removal. Deploying the detection mechanism to operate in real-time is crucial, providing immediate alerts and insights. Seamless integration with existing cloud monitoring and management tools allows for automated responses and comprehensive visibility.

Defining any abstract problem using mathematical concepts is crucial to ensure that the problem statement is clear and precise. Once the mathematical formulation of the problem is established, various parameters such as the algorithm and data sources can be adjusted without altering the core definition of the problem. This approach ensures consistency and accuracy in addressing the problem, allowing for effective monitoring and management of cloud spending, ensuring financial control and operational efficiency.

2.4. Cost management and optimization. As cloud technology continues to evolve, the variety of pricing and commitment options available to organizations has increased significantly. Cloud providers offer numerous pricing models such as on-demand, reserved instances, spot instances, and various discount plans. Additionally, there are multiple ways to commit to these resources, each with its own set of terms and potential cost savings. With these options, there are more combinations and strategies for managing cloud costs based on an organization's specific needs and usage patterns. However, this increased flexibility also adds complexity, making it challenging to determine the most cost-effective approach.

When an organization grows large enough, managing cloud costs at scale for a vast number of resources becomes a significant challenge. The dynamic nature of cloud pricing, combined with the need to monitor and optimize resource usage continuously, requires sophisticated tools and strategies.

One major challenge is understanding and selecting the most cost-effective pricing options while considering usage patterns. It is difficult to estimate how many commitments to buy and to understand if, at any given moment, you can manipulate instances to avoid driving usage to unused commitments, thus wasting money on resources that are not being fully utilized. Cloud providers frequently update their pricing structures, which necessitates continuous adaptation and strategy updates. This complexity, together with the variability in cloud usage, further complicates cost management. These challenges necessitate the use of advanced mathematical models and optimization techniques to manage and optimize cloud costs effectively. FinOps practitioners must leverage these tools to navigate the complexity of cloud pricing, manage resource usage efficiently, and ensure that spending aligns with business objectives.

To address the challenges of cost management and optimization in the FinOps domain, we will focus on the most dynamic and complex resource type, which is EC2 (Elastic Compute Cloud) in AWS. The variability and cost structure of EC2 instances present unique challenges that require sophisticated mathematical models and optimization techniques. Let's define a more detailed cost function that accounts for different types of EC2 instances (Box, Spot) and the impact of purchasing commitments (Reserved Instances, Compute Savings Plans, EC2 Savings Plans). The total EC2 cost is the sum of the costs for both Box and Spot instances. For simplicity, we define an EC2 instance that is not Spot as Box. Each cloud provider has its own terminology for these instances.

Cost of Spot Usage can be defined like this:

$$Spot\ Cost\,(N_{spot}) = \sum_{i=1}^{N_{spot}} (SpotPrice_i \times Usage_i),$$

where:
- N_{spot} is the total number of spot instances;
- $SpotPrice_i$ is the price of the i-th spot instance;
- $Usage_i$ describes amount of usage of spot instance (normalized units or amount of minutes/hours).

Cost of "Box" Usage can be defined like this:

$$Box\ Cost\,(N_{box},\ commitments) = \sum_{i=1}^{N_{box}} (PricingOption_i \times Usage_i),$$

where:
- N_{box} is the total number of spot instances;
- $commitments$ is a variable that defines amount of commitments customer has;
- $PricingOption_i$ is the cost of instance depends on $commitments$ it can be different even for same instance;

Then the total EC2 Cost is given by:

$$EC2\ Cost = Box\ Cost\,(N_{box},\ commitments) + Spot\ Cost\,(N_{spot}) + CW,$$

where new variable CW means commitments waste or amount of cost customer wasting because of not proper use of commitments and more exacly because of commitments underutilization.

To minimize the total $EC2\ Cost$, we need to find the optimal values for $(N_{spot},\ N_{box},\ commitments)$ for cost function:

$$EC2\ Cost\ (N_{spot},\ N_{box},\ commitments) \to \min.$$

Intentionally for simplicity time granularity dimension was ignored but as in any other FinOps problem this cost function can be viewed on hourly, daily, weekly and monthly level, etc.

In a more general view current approach can be applied to any cloud services or to the whole cloud account at once. However, different compute types and services could have different ways of optimization. In this paper we focused on the most complex case, which is compute instances.

3. Conclusion. FinOps faces numerous challenges that arise from the dynamic and complex nature of cloud environments. The increasing variety of pricing models, the need for real-time data analysis, and other complexities create many abstract problems. Mathematical modeling provides a way to address these abstract problems effectively, enhancing decision-making and operational efficiency.

For any abstract problem to be effectively addressed, it must be defined mathematically. Mathematical formulations provide a structured way to solve problems, allowing for the application of advanced techniques like optimization algorithms and machine learning. This approach not only clarifies the problem but also offers systematic methods to address it.

This paper overviews the main FinOps problems and contributes by classifying these problems into distinct categories. By proposing mathematical formulations for each class, we think that integrating statistical models, optimization algorithms, and machine learning techniques can improve the accuracy and efficiency of FinOps practices. This structured approach guides organizations toward better financial management and operational excellence in the cloud.

References

1. State of FinOps. Retrieved from: https://data.finops.org
2. AI and FinOps Predicted to Lead Observability Innovation in 2024. Retrieved from https://www.infoq.com/news/2024/02/observability-2024-predictions
3. Gaur, R. R., Budruk, S. C., & Patil, N. V. (2022). Survey on Cloud Load Prediction using Machine Learning Techniques. *IJIRT,* 9(1), 792–796.
4. Archana, Y., Kushwaha, Sh., Gupta, J., Saxena, D., & Singh, A. K. (2021). A survey of the workload forecasting methods in cloud computing. In *Proceedings of 3rd International Conference on Machine Learning, Advances in Computing, Renewable Energy and Communication.* MARC. Singapore: Springer Nature Singapore.
5. Hofmann, E., & Rutschmann, E. (2018). Big data analytics and demand forecasting in supply chains: a conceptual analysis. *The international journal of logistics management,* 29(2), 739–766.
6. Kureljusic, M., & Karger, E.(2023). Forecasting in financial accounting with artificial intelligence — A systematic literature review and future research agenda. *Journal of Applied Accounting Research,* 25(1), 81–104.
7. FinOps Open Cost and Usage Specification. Retrieved from https://focus.finops.org
8. Find Open Datasets and Machine Learning Projects. Retrieved from https://www.kaggle.com/datasets
9. Automate the deployment of an Amazon Forecast time-series forecasting model. Retrieved from https://aws.amazon.com/blogs/machine-learning/automate-the-deployment-of-an-amazon-forecast-time-series-forecasting-model/
10. BigQuery public datasets. Retrieved from https://cloud.google.com/bigquery/public-data

Роботишин М. В., Маляр М. М. Застосування математичного моделювання у домейні FinOps.

Розвиток хмарних обчислень революціонізував підхід організацій до управління IT-ресурсами, пропонуючи масштабовані та гнучкі рішення для різноманітних операційних потреб. Проте, з підвищеним впровадженням хмарних сервісів, фінансові операції (FinOps) стали більш складними, створюючи виклики в управлінні витратами, розподілі ресурсів та фінансовому прогнозуванні. Традиційні методи часто не здатні ефективно вирішувати ці складнощі, що призводить до неефективності та субоптимального прийняття рішень. У цій роботі досліджується застосування математичного моделювання в сфері FinOps як потужного рішення цих викликів. Пропонується класифікація проблем FinOps на окремі класи та математична постановка задачі для кожного класу, що має на меті підвищити ефективність FinOps-практик. Інтеграція математичних моделей покращує точність та ефективність, забезпечуючи систематичний підхід до управління фінансовими операціями у хмарному середовищі. Через вивчення кейсів та прикладів з реального світу ця робота демонструє трансформаційний потенціал математичного моделювання у стимулюванні інновацій та досягненні операційної досконалості в FinOps.

Ключові слова: математичне моделювання, управління хмірними обчисленнями, аналіз даних, прогнозування, FinOps, оптимізаційні процеси.

Список використаної літератури

1. State of FinOps. URL: https://data.finops.org (date of access: 25.08.2024).
2. AI and FinOps Predicted to Lead Observability Innovation in 2024. URL: https://www.infoq.com/news/2024/02/observability-2024-predictions (date of access: 25.08.2024).
3. Gaur R. R., Budruk S. C., Patil N. V. Survey on Cloud Load Prediction using Machine Learning Techniques. *IJIRT*. 2022. Vol. 9, No. 1. P. 792–796.
4. Archana Y., Kushwaha Sh., Gupta J., Saxena D., Singh A. K. A survey of the workload forecasting methods in cloud computing. *Proceedings of 3rd International Conference on Machine Learning, Advances in Computing, Renewable Energy and Communication* : MARC. Singapore : Springer Nature Singapore, 2021. pp. 539–547.
5. Hofmann E. and Rutschmann E. Big data analytics and demand forecasting in supply chains: a conceptual analysis. *The international journal of logistics management.* 2018. Vol. 29, No. 2. P. 739–766.
6. Kureljusic M. and Karger E. Forecasting in financial accounting with artificial intelligence — A systematic literature review and future research agenda. *Journal of Applied Accounting Research.* 2023. Vol. 25, No. 1. P. 81–104.
7. FinOps Open Cost and Usage Specification. URL: https://focus.finops.org (date of access: 25.08.2024).
8. Find Open Datasets and Machine Learning Projects. URL: https://www.kaggle.com/datasets (date of access: 27.08.2024).
9. Automate the deployment of an Amazon Forecast time-series forecasting model. URL: https://aws.amazon.com/blogs/machine-learning/automate-the-deployment-of-an-amazon-forecast-time-series-forecasting-model/ (date of access: 27.08.2024).
10. BigQuery public datasets. URL: https://cloud.google.com/bigquery/public-data (date of access: 25.08.2024).

Recived 11.10.2024