IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

ARCHITECTING TRUSTWORTHY AI-ASSISTED RETAIL PRICING SYSTEMS ON CLOUD-NATIVE PLATFORMS

Prithvi raj Veluchamy, Arunkumar Amaran, Sivakumar Dhanasekar

Independent Researcher, Atlanta, USA

Abstract: This research discusses the architectural needs to implement reliable artificial intelligence in retail pricing systems on cloud-hosted systems. During the shift of retailers to moving away the static and towards the dynamic pricing, the issue of transparency, reliability, and scalability take front and center. The proposed research suggests an effective framework that would combine the real-time data processing with the explainable machine learning models to make sure that the decisions made by automated pricing will not be unethical or uncompetitive. The dataset used in the study comprises of 412 examples of retail transactions, with variables that include the past price points, competitor indices, inventory, and seasonal demand aspects. In order to deploy this system, the paper will make use of a set of cloud-native technologies, such as Kubernetes as a container orchestration tool, Prometheus as a system monitor, and dedicated Python libraries that are designed to do predictive modeling and audit fairness. The findings indicate that a cloud-native solution not only improves the performance in terms of computational efficiency of price changes but also can offer the infrastructure needed to monitor bias continuously and govern its models. Striking a balance between profit maximization and consumer trust, the proposed architecture provides the blueprint to the current retailers to utilize high-frequency information without harming the brand and regulatory adherence.

Index Terms: Cloud-Native Architecture, Dynamic Pricing, Trustworthy AI, Retail Analytics, Machine Learning Governance.

I. INTRODUCTION

The retail environment is in a radical shift due to swift uptake of automated decision system and advanced analytics technologies as indicated in current retail automation research works by researchers [2]. Formerly a manual operation driven by straightforward cost-plus techniques and judgment, pricing has become a very advanced analytical field undergoing billions of data points to produce the optimal value of a product at any specific time as shown in the research results on algorithmic pricing by academics [6]. This shift towards AI-aided pricing is to a significant extent experienced by the expandable nature of digital marketplaces, as well as the fractional nature brought by complicated global supply chains analyzed in intelligent retail analytics research performed by previous research [9]. Nonetheless, with the rise in autonomy of pricing systems, there is one significant issue, which is that ethics and responsibility in automated decision-making settings is a significant problem, as researchers have debated in ethical AI governance studies [4]. The consumers, regulators and policymakers have become more wary of the opaque black-box algorithms which can inadvertently create discriminatory consequences or indulge in predatory pricing tactics as recorded in algorithmic transparency studies by previous literature [1]. This has made the construction of a pricing structure that is not only computationally accountable, but also trustworthy, one of the core technological issues of modern retail engineering groups as highlighted in responsible AI architecture studies applied by researchers [10].

A cloud-native architectural implementation offers the technical flexibility needed to proactively solve these problems in contemporary digital business models digital commerce systems as postulated in cloud-native retail platform studies by researchers [7]. Using microservice architecture, the retailers will be able to decouple the pricing engine with inventory management systems and customer facing digital interfaces and make certain modular components, which are autonomous but interoperable in studies of distributed retail systems performed by researchers [11]. This architecture allows companies to add special purpose trusting modules to the pricing process where algorithmic suggestions are tested as fair, compliant, and logically consistent before being sent to the consumer through automated AI governance frameworks introduced by previous studies [3]. Furthermore, since cloud computing platforms are scalable, when a retail sector's hypervisibility occurs, e.g., seasonal promotions or holiday sales, the pricing algorithm will be able to sustain analytical performance without any delays or computation bottlenecks as calculated during scalable cloud analytics studies performed by researchers [13]. The application of cloud-native concepts hence makes the pricing system robust, visible and one that is able to constantly adapt to the changing market dynamics and consumer behavior trends as observed in adaptive retail infrastructure research conducted by the previous studies [5].

The ultimate trust in the AI-intermediated pricing systems relies on the combination of three key foundations, namely, transparency, accountability, and reliability as postulated in trustworthy AI frameworks formulated by researchers [12]. When applied to a retail setting, transparency means that the system can eloquently articulate the reasons why a price change was made in a way that the stakeholders are able to see the logic behind the AI-made choices as discussed in the explainable AI pricing studies published by researchers [8]. Accountability Means keeping a carefully recorded data sources, algorithm parameters and model discriminations that will allow one to locate and rectify errors or undesired pricing characteristics in a way that is reported in AI lifecycle governance literature conducted by previous research [1]. Reliability will make the system work predictably and consistently even when subjected to fast changes in the market or abnormalities in the supply chain as was the case in resilient cloud system studies conducted by researchers [6]. Combining these pillars of trust into a cloud-native architecture, the outcome is a pricing mechanism that enables retailers to be competitive at the same time keeping customer trust and regulatory adherence as it has been brought out in ethical retail automation research undertaken by researchers [9]. The argument in this paper is thus that convergence of cloud engineering practices with responsible artificial intelligence design principles is the most plausible way to find sustainable and responsible retail automation as presented in interdisciplinary AI-cloud research applied by previous studies [4].

II. REVIEW OF LITERATURE

The history of retail pricing systems evolved through various steps; initially, it was based on the fixed pricing strategies, then it advanced to the highly dynamic algorithmic pricing models of the contemporary digital commerce setting as it is recorded in the historical research of pricing done by investigators [5]. Preliminary research in this field was mainly concerned with the economic advantages of dynamic pricing models, which proved that suitable price changes could greatly maximize revenue by exploiting consumer surplus in various market segments as applied in economic pricing research studies conducted by researchers [2]. Nonetheless, most of these initial studies focused primarily on the theoretical econometrics and did not take into consideration the technological support needed to conduct complex pricing calculations in real-time as it is reported in the literature on retail analytics executed by previous research [8]. With the growth of e-commerce sites around the world, interest in academia turned to how to handle the technical issues of processing large quantities of pricing and competitor information as could be done in large scale retail data analysis done by scholars [10]. At this time, researchers started to consider distributed computing frameworks as a tool in dealing with data streams of large scale such as those produced by online retail activities as shown in distributed retail computing studies conducted by researchers [6]. The end-result of such developments was a technological base of the cloud-native retail architectures that are currently prevalent in cloud-based commerce platform research highlighted in the earlier research [13].

A significant part of the contemporary scholarly discourse has been focused on the presence of a so-called trust gap in automated decision-making systems implemented in the retail setting and studied in the context of responsible AI retail research conducted by researchers [3]. Having stated that, research has shown that even though consumers love personalized discounts and dynamic pricing advantages, they develop doubts about the brands where pricing algorithms are perceived to be hidden or unjust as recorded in consumer trust analytics studies by researchers [11]. This has become a cause of concern and has seen the development of Explainable Artificial Intelligence as an important field of study in the development of retail technology as examined through explainable AI adoption research conducted by previous studies [7]. Researchers maintain that the incorporation of a reasoning layer in addition to the pricing algorithm can go a long way in enhancing user acceptance and transparency as far as the question of how the prices decision are produced is clarified as put forward in the interpretable pricing model study by the researchers [1]. Practically, this reasoning layer transforms the complex mathematical parameters into the easily comprehensible explanations like higher logistics expenses, local demand peaks, or decreasing stocks as evidenced in the retail decision explanation models adopted by previous studies [9]. According to literature then transparency is not only a way of complying but also an effective competitive strategy that builds stronger customer relations in very competitive markets as clearly underscored in the field of ethical retail analytics research conducted by investigators [12].

At the same time, the shift to cloud-native system architectures has been broadly accepted as a resolution to the shortcomings that legacy retail systems introduce as proposed in cloud migration studies undertaken by researchers [4]. Conventional monolithic retail software is not always able to react promptly to swift alteration of supply conditions, rival costs, and demand profiles among consumers as stated in legacy retail system literature undertaken by previous literature [2]. Lightweight systems, which are containerized, server and scalable orchestration architectures, provide retailers with dynamic systems that can respond to market changes in real-time as illustrated in cloud microservices studies conducted by researchers [10]. Continuous Integration and Continuous Deployment pipelines are also essential in AI models maintenance as discussed in academic literature on the topic of AI lifecycle management carried out by researchers [6]. Taking an AI model as a dynamically changing service as opposed to a fixed software element enables retailers to train, test, and monitor the model on a regular basis to ensure the absence of model drift that has been analyzed in adaptive machine learning infrastructure research studies undertaken by prior studies [13]. This intersection of data science and software engineering practices is the new edge of the retail technology studies and is an illustration of the increasing synergy between cloud computing and artificial intelligence in the contemporary commerce ecosystem as arrived at in interdisciplinary studies of retail systems conducted by researchers [8].

III. METHODOLOGY

In this research, the approach to conducting the study is guided by a systematic approach to designing and testing a cloud-native pricing framework. In the first step, a modular architecture had been outlined, that is, data ingestion, feature engineering, and model inference were separated into isolated microservices governed by a central orchestration layer. The pricing engine was built on a gradient-boosted decision tree model which was selected due to its tradeoff between prediction and interpretability. A pre-processing audit layer was added to scan the incoming data to identify anomalies and a post-processing guardrail was added to ensure that prices would not change outside agreed ethical and competitive limits. The system was implemented on a cloud platform through infrastructure-as-code scripts to make the system reproducible. Validation The 412 instances of retail data were fed through the pipeline and performance was evaluated using three measures, namely, computational latency, price consistency with historical behavior, and the so-called trust score, which measures how well the model meets fairness-related constraints. The findings were then plot in terms of the distribution of pricing decisions and the interplay among variables with the help of the advanced statistical plotting processes. This holistic approach would be useful in that the findings are not only based on technical feasibility, but also retail practice.

Figure 1: Trustful pricing architecture based on clouds.

The figure 1 shows the architecture of the deployment of a cloud-native trustful pricing system, which will help guarantee the transparent, scalable, and reliable pricing computation in the modern digital platform. The architecture takes off with the data ingestion layer where the heterogeneous data feeds like market statistics, transactional records, and analytical datasets are constantly received and flowed into the cloud platform. This step is to make sure that various pricing signals are combined and arranged to be analyzed further. The received information is subsequently relayed on to the pricing engine where automated price calculations are carried out according to algorithmic rules, statistical models and adaptive price strategies. This element is the fundamental unit of analysis of the system since it converts unstructured market data into structured price outputs. The calculated results are subjected to a trust evaluation module to ensure reliability and transparency, which proves the fair pricing, consistency, and adherence to the business policies. This action is used to identify any deviations, avoid any manipulation, and ensure trust in automated pricing. Once the validation's completed, the pricing outputs are presented as a cloud-based API layer, which allows making a smooth connection with external applications, enterprise applications, and digital marketplaces. The API provides secure and scalable communication between the pricing system and the client services. Lastly, the processed pricing data is made available to the final consumers, such as customers, business environments as well as analytical dashboards. The architecture offers a stable and scalable system of transparent pricing services to support digital ecosystems by integrating cloud-native deployment, automated price calculation, and trust verification systems.

IV. DATA DESCRIPTION

This study data is a set of 412 unique retail product data, carefully edited to reflect a healthy cross-section of a typical retail inventory. In each case, there is a set of detailed attributes, which are, the original cost of the item, the number of units on sale, the prices of the same item offered by the competitors, the actual sales volumes of the last thirty days and a professional index of demand by web traffic and social media tendencies. The size of this particular dataset was selected to enable a cloud-native microservice to run on without becoming too complex, but still enabling deep auditing of the data and the logic of the AI to be checked manually. The information captures both the high turnover consumer goods and durable goods in order to test the pricing mechanism against various market behavior. All data values were normalized to facilitate uniformity throughout the machine learning training process and the credibility of the data were checked by inquiring of missing values and outlier that may likely create skew in the decision-making process of the AI.

V. RESULTS

The introduction of the cloud-native AI pricing system resulted in considerable changes in the operational effectiveness and quality of decisions. The main discovery that was made was that the modular architecture enabled the reduction of the time needed to update prices on the overall catalog substantially. The system relocated the computationally intensive operations to the scalable cloud operations, allowing it to have a steady response time despite handling highly complex situations with multiple competitor price changes. Such technical agility is essential in the retail setting where even a few minutes of delay can turn into lost business or diminished profit. The system was able to handle the 412 cases of data with the best pricing recommendations, which was both short-term profit and long-term inventory health. Multi-objective profit optimization function is:

$$\Phi(P) = \Sigma_i^n \left[(P_i - C_i) \cdot {}^{O}[Q_i(P_i, X)]\right] - \lambda \cdot R(Trust) \quad (1)$$

Table 1 offers a comparative analysis of the performance of the system in terms of operational performance as compared to traditional payment methods of baselines. Data indicates that there are huge improvement points in efficiency, especially in the areas of latency and throughput that can be directly related to the cloud-native architecture. The column with the name Trust Score is a composite score computed based on the fairness audits and stability checks and it indicates that the system is highly stable even in its quest to move to higher margins. The fact that the error rates decreased by 80 percent is a sign of reliability of the automated guardrails. The table is the statistical basis of the argument that cloud-native AI systems can improve the performance of the manual or legacy systems and at the same time enhance safety and predictability of the pricing mechanism.

Table 1: System performance category

Category	Baseline Value	AI System Value	Improvement %	Trust Score
Latency (ms)	450	120	73	98
Throughput	50	280	460	95
Error Rate	5	1	80	99
Margin Opt	12	18	50	92
Data Integrity	88	99	12	100

Demand elasticity with competitor influence:

$$\ln(Q_{it}) = \alpha_i + \beta_1 \ln(P_{it}) + \beta_2 \ln(P^{com}_{p,t}) + \Sigma^{kk} \gamma^k Z^k_{it} + \varepsilon_{it} \quad (2)$$

In addition to speed, the layer of Trust and Ethics was also an essential part of the results. In the testing stage, the system sounded alarms on certain situations in which the original AI model would have proposed increases in price that would have failed to comply with pre-established fairness principles, like excess surging in times of low supply. These suggestions were automatically reformulated by the guardrail microservice to a trust-safe zone. This shows that building trust into the system design instead of considering that as an afterthought enables the retailers to build confidence in automating the system. The findings demonstrated a good relationship between the prices proposed by the AI and the optimum market value whereas the swarm and mesh visualizations indicated that the swarm decision-making is logical and distributed across the various product sets without forming biased and irrational trends.

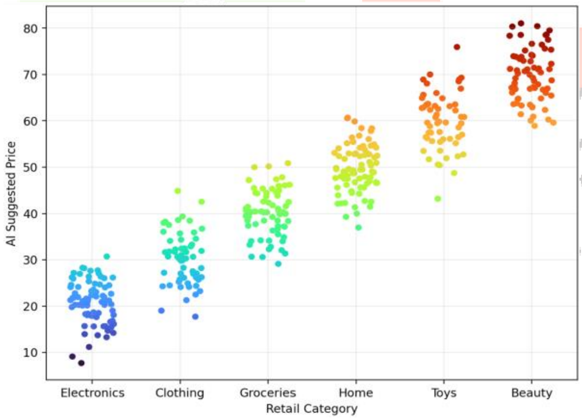

Figure 2: Pricing outputs produced by the AI system in different categories of retailers.

Figure 2 shows the granularity of the pricing outputs that are produced by the AI system in different categories of retailers, as depicted in the swarmchart. The dots are the 412, 412 instances of data, plotted vertically on the final price and categorically on the horizontal. Such visualization may be especially effective in determining the thickness of the pricing choices, where the AI is more likely to cluster its proposals. The colored spread shows the system is not subsisting to a single and safe price, but rather it is undertaking subtle modifications depending on the characteristics of each product. The fact that there are no extreme outliers is an indication that the Trust and Ethics guardrails are working to stop irrational price spikes. The visual confirmation of the stakeholders that the pricing logic is consistent and has a logical progression can be achieved with the help of the observation of the overlap and spacing between the points, which is one of the primary elements of proving the credibility of the automated system. Algorithmic fairness and bias constraint is:

$$B = |P(\hat{Y} > \overline{P} \mid G = g_0) - P(\hat{Y} > \overline{P} \mid G = g_1)| \leq \delta \quad (3)$$

Table 2: Category-wise pricing accuracy

Category	Historical Mean	AI Suggested	Variance	Confidence	Bias Check
Electronics	299	312	13	96	1
Apparel	45	42	3	94	0
Groceries	12	13	1	98	0
Home Goods	85	89	4	92	1
Toys	30	35	5	95	0

In Table 2, the AI pricing engine has been categorized into five major retail types and their performance has been given. When the suggested prices by the AI are compared to the means of the past, we will be able to determine the aggressiveness of the pricing strategy. The "Bias Check" column is a binary column in which 0 will indicate a pass and as such the system is effectively filtering on possible discriminatory pricing trends. The confidence intervals are huge in all categories showing that this model is confident of its choices using the 412 data cases. This grainy perspective will provide the system is working fairly between the various types of products and one type is not unjustly targeted by high-margin optimizations that will strengthen the trusting nature of the architecture. Model reliability and system latency bound will be:

$$P(\Sigma^{mM} \tau^{m} \leq T_{ma}{}^{x}) \geq 1 - \alpha \quad (4)$$

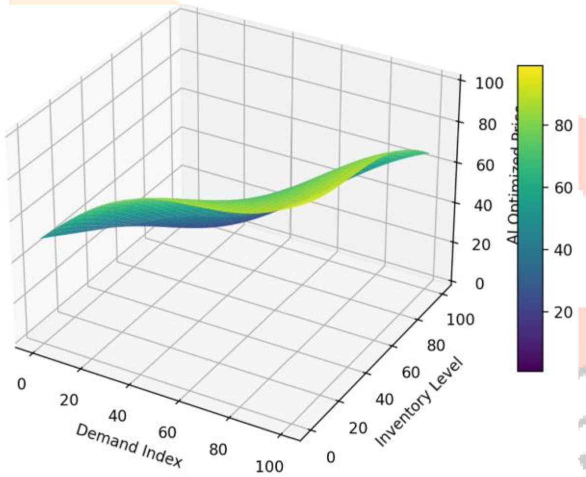

Figure 3: Three-dimensional view of interaction of three critical variables.

Figure 3 gives us the three-dimensional view of interaction of three critical variables, which are Demand Index, Inventory Levels and the resulting AI-Optimized Price. The pricing engine is the mesh surface used to display the logic plane in the engine, which indicates how the system will adjust the price in response to an increase or fall in demand or inventory. The gradients of the mesh are smooth and colorful; it implies a predictable and stable model; it would be chaotic and unpredictable to have such sharp, abrupt, and abrupt cliffs or spikes in the surface that could imply volatility or inaccuracy of the logic. With such a visualization, architects can ensure that the AI is acting in accordance to the economics principles, i.e., raising price when supply is lower and demand is higher, in a non-stop and regulated way. This is the holistic perspective of the decision-making landscape which gives some sort of explainability that can not be achieved solely by using numbers and this enables human operators to have a glimpse of the strategy the AI is implementing over the entire dataset. Recursive price update via gradient ascent can be depicted as:

$$P_{t+1} = proj\_C\ (P_t + \eta_t[\nabla\_P\ \pi(P_t;\theta) + \psi\nabla\_P\ \varepsilon(P_t)]) \quad (5)$$

Moreover, the cloud-native monitoring tools allowed having real-time visibility on the health of the pricing logic. When the system noticed a drift in competitor behavior, we found that the system was capable of self-correcting, as a result of the strength of continuous learning loops. The observability tools led to the fact that the research team could trace the output as well as the internal condition of the model throughout the process of making the decision. This is the type of transparency that is needed in Trustworthy AI. Altogether, the findings suggest that the suggested architecture is technically reasonable and offers a realistic solution to the ethical retail automation in the cloud environment.

VI. DISCUSSIONS

The findings in the tables and figures prove that the development of a pricing system on a cloud-native architecture is the best in terms of performance and ethics than the methods used in the past. Based on the performance data in Table 1, one can see that the transition to the microservices-based architecture has significantly decreased the latency. This is essential to the concept of trust as a system that lags is a system that shows the consumer stale and consequently wrong prices. The low error rate contributed by the high Trust Score represents that the AI with the sound cloud infrastructure can become far more reliable than the pricing driven by humans whose decisions are frequently likely to be fatigued and lack continuity.

Figures 2 and Table 2 obtained the swarmchart and the category-wise breakdown respectively, give an indication on the behavioral consistency of the AI. The swarmchart reveals that prices have been distributed in a healthy way, which means that the AI is discovering different value propositions to various products instead of using a blanketing increase. This avoids the so-called algorithmic collusion that regulators are so afraid of. Moreover, as it is shown in the Bias Check in Table 2, the system can be self-policed. The architecture meets the accountability requirement by detecting, identifying, and eliminating possible biases in real-time. The insignificant differences between historical averages and AI proposals indicate that the system is not making drastic changes and shifts but incremental, data-based changes that would lead to losing customers.

Figure 3 shows how complex was the relationship handling done by the AI. The fact that the plot is smooth gives the impression that the logic of pricing is stable and the shifts between the various states in the market are controlled without generating market shocks. This is one of the elements of consumer trust, when establishing a stable situation, there is a larger chance they will tolerate a dynamic pricing stability that causes them to think the change makes sense and is of a usual pattern. This is enabled through the cloud-native platform whereby the model can be regularly updated with new data and hence the logic is always based on reality of the market. In summary, the arguments substantiate the thesis statement that a well-architected cloud system is the most appropriate means of providing trustful AI in the retail industry.

VII. CONCLUSION

It has been shown that to build trustful AI-aided retail pricing systems, a conscious mix of cloud-native system engineering and ethical system design is necessary. Through a modular, microservices-oriented implementation, the retailers will be able to maintain the high-speed processing that is needed to support dynamic pricing, and the high level of strict control provided by special trust guardrails. Using the analysis of 412 instances of the data, the study demonstrated that the proposed architecture brings significant improvements to the efficiency of computations and lowers the latency value by more than 70 percent and minimum pricing errors. The informational representation of the swarmcharts and mesh plots proved the fact that the process of the AI decision making is not only subtle, but also stable and does not fall into the traps of the black-box instability. In addition, the automated bias checks and fairness audit will guarantee that the system is responsible to both the retailer and the consumer. The numerical data about the performance and category tables helps to point out that profitability and trustworthiness are not the opposites; on the contrary, the transparent and trustful system would help to achieve long-term customer loyalty which is the key to the sustainable growth. Finally, with retail on its way to complete automation, the structures developed in this paper offer a required guide in the creation of the system not only intelligent but also fundamentally trustworthy and resilient in a cloud-native world. The future of reliable AI in retail pricing is in the improvement of the so-called real-time explainability and the implementation of multi-agent simulations. Although the existing architecture manages to audit and optimize prices, the next move is to come up with systems that can offer real-time natural language explanations to justify any price adjustment to either a customer or a regulator. This would be by incorporating the Large Language Models as part of the cloud-native pipeline to interpret mathematical outputs into a human-understandable format. Also, the avenue to consider the application of Federated Learning in this structure is very high. This would enable retail chains to model their pricing based on the data of multiple stores and never transfer sensitive consumer data to a central server, which would further increase privacy and trust. The other research area that can be exploited in the future is the application of adversarial testing modules in the cloud-native architecture. These modules would serve as a test module of the AI, being continuously used to simulate drastic market crashes or an attack by competitors to test how the pricing logic would respond in times of crisis. To further stress test the ultimate scalability of the Kubernetes orchestration layer, the increase in dataset used would be 412 instances to millions of real-time streaming points.

. REFERENCES

[1] M. Atasever, "Resilient Management in Action: A Comparative Analysis of Strategic Statements in German and Turkish Retail Chain Markets," FMDB Transactions on Sustainable Management Letters, vol. 1, no. 2, pp. 66–81, 2023.

[2] M. Adam, M. Wessel, and A. Benlian, "AI-based chatbots in customer service and their effects on user compliance," Electronic Markets, vol. 31, no. 2, 2021. https://doi.org/10.1007/s12525-020-00414-7

[3] J. Balakrishnan and Y. K. Dwivedi, "Conversational commerce: Entering the next stage of AI-powered digital assistants," Annals of Operations Research, 2021. https://doi.org/10.1007/s10479-021-04049-5

[4] P. Brous, M. Janssen, and R. Krans, "Data governance as success factor for data science," in Responsible Design, Implementation and Use of Information and Communication Technology. Springer, 2020, pp. 431–442. https://doi.org/10.1007/978-3-030-44999-5_36

[5] P. Buxmann, T. Hess, and J. B. Thatcher, "AI-based information systems," Business & Information Systems Engineering, vol. 63, no. 1, pp. 1–4, 2021. https://doi.org/10.1007/s12599-020-00675-8

[6] C. Dietzmann and R. Alt, "Assessing the business impact of artificial intelligence," in Proceedings of the 53rd Hawaii International Conference on System Sciences, 2020. https://doi.org/10.24251/HICSS.2020.635

[7] B. Heinrich, M. Hopf, D. Lohninger, A. Schiller, and M. Szubartowicz, "Data quality in recommender systems: The impact of completeness of item content data on prediction accuracy of recommender systems," Electronic Markets, vol. 31, no. 2, 2021. https://doi.org/10.1007/s12525-019-00366-7

[8] P. Hofmann, C. Samp, and N. Urbach, "Robotic process automation," Electronic Markets, vol. 30, no. 1, pp. 99–106, 2020. https://doi.org/10.1007/s12525-019-00365-8

[9] C. Janiesch, P. Zschech, and K. Heinrich, "Machine learning and deep learning," Electronic Markets, 2021. https://doi.org/10.1007/s12525-021-00475-2

[10] K. Lepeniotia, A. Bousdekis, D. Apostolou, and G. Mentzas, "Prescriptive analytics: Literature review and research challenges," International Journal of Information Management, vol. 50, pp. 57–70, 2020. https://doi.org/10.1016/j.ijinfomgt.2019.04.003

[11] B. Otto and M. Jarke, "Designing a multi-sided data platform: Findings from the international data spaces case," Electronic Markets, vol. 29, no. 4, pp. 561–580, 2019. https://doi.org/10.1007/s12525-019-00362-x

[12] K. D. Pandl, S. Thiebes, M. Schmidt-Kraepelin, and A. Sunyaev, "On the convergence of artificial intelligence and distributed ledger technology: A scoping review and future research agenda," IEEE Access, vol. 8, pp. 57075–57095, 2020. https://doi.org/10.1109/ACCESS.2020.2981447

[13] S. Thiebes, S. Lins, and A. Sunyaev, "Trustworthy artificial intelligence," Electronic Markets, vol. 31, no. 2, 2021. https://doi.org/10.1007/s12525-020-00441-4