A Cross-Disciplinary Knowledge Management Framework for Artificial Intelligence in Business Projects

Andrzej Wodecki
Warsaw University of Technology, Poland

andrzej.wodecki@pw.edu.pl

Abstract: This paper presents a knowledge management framework for designing, monitoring, and optimizing the business value of intelligent services using best practices from both management and machine learning engineering areas. The increasing interest in artificial intelligence has highlighted the difficulty of designing an appropriate business case and monitoring the real business value generated by such services due to the complexity of machine learning methods and techniques. Managers have proven frameworks for constructing key performance indicators, machine learning engineers have developed advanced methods and techniques for evaluating and monitoring Machine Learning (ML) models, and scientists have tested methods for empirically evaluating the added value of innovations. By leveraging the strengths of each discipline, the proposed framework enables managers, machine learning engineers, and scientists to efficiently work together and optimize the value of intelligent services during their lifecycle. We propose processes and procedures for creating, capturing, organizing, storing, sharing, and using knowledge in advanced machine learning projects. He also presents key machine learning methods and techniques for monitoring and optimizing the value of intelligent services over their lifecycle, including the adaptation of MLOps methodology for continuous monitoring, reinforcement learning for continuous improvement, and CausalML methods for identifying the root-causes of changes in the business value. These methods and techniques support knowledge management activities and help formulate a competency framework for team members and project stakeholders. We point out the potential of academic researchers and external advisors as catalyzers in such projects based on real-life implementation. He also proposes a method for designing the ML knowledge flywheel to ensure continuous knowledge transfer and improvement in the business-engineer-academy triangle. The approach is illustrated by a case study of the implementation of a marketing communication optimization system in a large, multinational financial company for more than 20 thousand customers in two European countries. Managers and machine learning engineers can implement the proposed knowledge management framework in various organizations for the efficient design, monitoring, and optimization of the business value of intelligent services.

Keywords: Knowledge management, Project team competencies, Machine learning projects

1. Introduction

The relationship between artificial intelligence and knowledge management is becoming increasingly complex. On the one hand, AI solutions can effectively support KM processes (Wang et al., 2022), on the other hand, machine learning projects require teams with diverse competencies (e.g., in the marketing industry (Blomster and Koivumäki, 2022)). Additionally, the rapid development of Large Language Models (LLMs) and, more broadly, generative models, create a new bridge between machine learning and knowledge management, enabling, for example, the improvement of machine learning model quality by enriching features through LLM models.

The rapid development of machine learning methods and techniques makes it difficult for many organizations to keep track of the latest solutions that could improve the quality of their models. The shortage of data science specialists leads to their overload in projects and hinders the development of professional competencies. A solution to this problem may be projects carried out in multidisciplinary teams, where specialists from not only different fields but also organizations and environments meet.

The development of machine learning takes place in the areas of data preparation methods, techniques, and tools, training, evaluation, production deployment, and business value assessment after production launch. It is worth emphasizing that these are not only "quantitative" changes (e.g., new machine learning algorithms or new software libraries supporting the implementation process) but also, in a sense, qualitative ones (e.g., the development of experimental methods enabling a real assessment of the value of a machine learning model compared to previous methods of operation).

Narrow specialization within data science teams, IT engineers, and business representatives makes it difficult to develop and implement "horizontal" solutions: those that optimally integrate the competencies and needs of each of these groups. For example, data scientists measure the efficiency of their models using statistical measures, while business representatives using models are more interested in using business measures (various Key Performance Indicators (KPIs) measuring, for example, process efficiency). Data scientists usually evaluate models before production deployment (although model monitoring and observability systems in production are becoming increasingly popular), while the business is more interested in managing the value of business models over time.

In the following sections, we will present various insights and observations, largely derived from a case study of a financial institution, which illustrate different aspects of the interplay between machine learning and knowledge management. These insights will provide an understanding of the challenges and opportunities that arise when integrating AI solutions into business processes and the value of multidisciplinary teams in addressing these challenges. We will discuss the role of communication and collaboration between various stakeholders, the importance of aligning business and data science objectives, and the impact of emerging technologies on the development of machine learning models. Finally, we will highlight the lessons learned from these observations and their implications for both practitioners and researchers in the fields of AI, machine learning, and knowledge management.

2. Business Value of Machine Learning Models

In this section of the article, we will focus on three key aspects of evaluating machine learning models: statistical methods, business methods, and experimental methods. We will explore their respective roles in assessing model quality, their relevance to different stakeholders, and their application in various stages of the project lifecycle.

2.1 Evaluation

Evaluation methods for machine learning models can be divided into three groups: statistical methods, business methods, and experimental methods. Among Data Scientists, statistical methods are the most popular. For example, to assess the quality of classification models, measures such as ROC AUC, accuracy, F1/F_beta, precision, recall, or confusion matrix can be applied. For regression models, measures like R2, RMSE, or MAPE can be used, while the Silhouette coefficient is used for clustering.

While statistical measures are good from a Data Science perspective, a manager formulating a business justification for a project using machine learning in practice is more interested in the real impact of the solution on processes or products. To assess business value, measures such as the confusion matrix, lift (see e.g. (Provost and Fawcett, 2013), or simply KPIs used to evaluate process performance or measure product quality can be applied.

The methods described above, assuming the availability of appropriate training sets (real historical data), allow for the evaluation of a model's quality before its production deployment. In practice, however, the actual business value may significantly deviate from the forecasted value. Factors influencing the real value generated by information technology can be divided into organization-dependent and independent factors. An organization can influence, for example, the level of adoption of a solution by employees or customers, but it will be more challenging to control market factors affecting sales dynamics.

Therefore, the optimal solution is to use the statistical and business methods mentioned earlier at the stage of constructing the business case (as a basis for, e.g., an investment proposal) and to use experimental methods (e.g., A/B/n tests) for the actual assessment of business value. Such experiments should be planned already at the project development, especially by identifying KPIs as the basis for measuring effectiveness, research hypotheses to be verified within the experiment, the sizes of the control and test samples, and the architecture of the IT solutions that enable the research.

2.2 Operationalization, Model Monitoring and Observability

The production deployment of machine learning models presents a considerable challenge, with costs often significantly exceeding those of training (as suggested in the literature). During the operationalization stage, the model is made available on the production infrastructure, an application layer is created, intelligent solutions are integrated with existing systems, and users are appropriately trained. Additionally, it is necessary to implement the previously designed systems that enable the execution of experiments assessing the actual business value of the models.

The situation is further complicated by the fact that models trained on historical data tend to lose their quality over time. This phenomenon is commonly attributed to data drift and model drift. Data drift refers to changes in the data distributions characterizing the objects being predicted (e.g., customers or products), typically resulting from natural quantitative changes in the economic environment (e.g., inflation, competitors' behaviors, market shifts, etc.). In contrast, model drift usually arises from qualitative changes, such as the emergence of new features characterizing the predicted objects (see Figure 1).

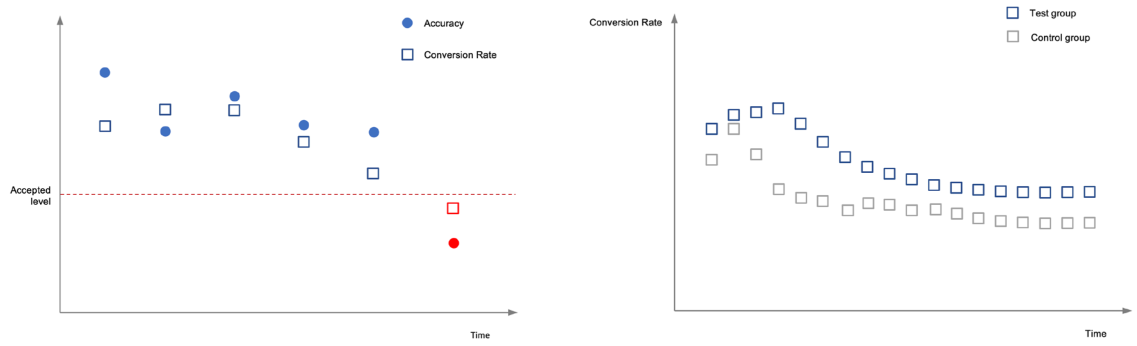

Fig. 1a Fig. 1b

Source: Own elaboration.

Figure 1: Model Drift

On the left (Fig. 1a) exemplary statistical (Accuracy) and business (Convertion rate) drift are presented. On the right (Fig. 1b) we present experiment results, comparing the test group (here: marketing communication basing on Machine Learning model) with the test group (communication basing on current rules). One of the key projects goals is to manage the business measure drift rather than the statistical measure. The data presented in this figure are illustrative and not based on actual values.

Consequently, the experimental verification of an intelligent solution's business value (immediately following production deployment) should be complemented with solutions that enable the monitoring of this value over time.

As demonstrated, the processes of designing, implementing, and maintaining production-ready intelligent services are highly complex. Consequently, it is no surprise that numerous systems have been developed to facilitate their execution. Specifically, solutions aiding the observation and diagnostics of models can be categorized into purely analytical systems (e.g., www.arize.com) and those utilizing fundamental machine learning techniques (e.g., Bayesian inference in solutions like www.nanny.ml). Additionally, the dynamically evolving CausalML methods and techniques of the CausalML (refer to libraries like DoWhy (Sharma and Kıcıman, 2020), CausalNex (Kumar and Ravi, 2022), or gCastle (Zhang et al., 2021) hold great potential, as they enable modeling causal relationships and ultimately identifying factors with the most significant impact on model decay. Reinforcement Learning, which allows designing self-improving systems, is also noteworthy in this area, both for experiment optimization. At the management level, MLDevOps methodologies, supported by dedicated IT solutions (e.g., Google Vertex AI, Amazon SageMaker, or Microsoft Azure), play a crucial role in automating model maintenance.

A considerable challenge at present is that most of the above-described solutions support the detection and diagnostics of model drift based on statistical measures. While selected AutoML or model observability tools already offer the ability to optimize training and monitor model quality using custom measures (e.g., business-related), the habits of Data Scientists who are accustomed to employing statistical measures represent a significant barrier. As of now (April 2023), we are not aware of any solutions that integrate model monitoring with experiments utilizing, for instance, A/B/n tests (although separate solutions exist to support the execution of such experiments, they are not integrated with model monitoring systems or, more broadly, AutoML or MLDevOps).

3. Competency Framework for Team Members and Project Stakeholders in Machine Learning Projects

As engineering projects increasingly necessitate the intersection of diverse disciplines, it is crucial to understand the role of knowledge management (KM) in fostering cross-disciplinary competencies. Fakhar et al. (Fakhar Manesh et al., 2021) examined the intellectual structure and trends of knowledge management (KM) in the context of the fourth industrial revolution, Industry 4.0, uncovering six major themes and proposing future research avenues. Bjornson and Dingsøyr (Bjørnson and Dingsøyr, 2008) conducted a systematic review of empirical studies on KM initiatives in the domain of software engineering, highlighting a significant focus on explicit knowledge and the necessity of considering tacit knowledge. Carrillo and Chinowsky (Carrillo and Chinowsky, 2006) investigated KM initiatives among major U.S. engineering design and construction firms, revealing different KM practices in design firms and construction companies. Kamara et al (Kamara et al., 2002) assessed KM strategies in the Architecture, Engineering and Construction (AEC) industry, concluding that effective KM requires an integrated approach addressing both technological and organizational/cultural issues. Chung et al (Chung et al., 2003) presented a collaborative project that aimed to enhance current workflow systems to better manage dynamic and collaborative processes, particularly focusing on new product development in the chemical industries. Finally, Ardito et al (Ardito et al., 2019) investigated the specific roles that universities play in managing knowledge for smart city projects, identifying universities' roles as knowledge intermediaries, gatekeepers, providers, and evaluators.

While the study of knowledge management in engineering projects has been extensively covered in the literature, there remains a noticeable research gap in understanding its implementation and impact specifically within machine learning projects in the business context. Below we address the competency framework key for successul implementation of such initiatives.

A typical intelligent service implementation project can be divided into stages of formulating a business case, planning, Data Science, operationalization, and maintenance.

Business representatives, Data Science specialists, and IT professionals usually participate in such project, with the entire process coordinated by a project manager. Each of these roles requires a specific set of skills. For the purposes of this study, we will categorize these skills as key competencies and complementary competencies, which are not directly related to the role but are essential for the project's success, focusing on the latter. The reflections presented below are based on the implementation of numerous machine learning projects for medium and large organizations - an example project is briefly described at the end of the paper.

T-shaped Competencies

Business representatives are often the initiators of the project, well-versed in market conditions, and capable of measuring the effects of actions. To effectively plan business cases, monitor project execution, and manage its business value as comprehensively as possible, they should also possess in-depth competencies in designing and conducting experiments, supplemented with skills in Data Science and IT. Using the metaphor employed in competency management, these individuals should have T-shaped competencies, where the vertical bar represents in-depth business skills (from a given area, like internet marketing), and the horizontal bar encompasses skills in general management, experiment design, data science, data engineering, and MLOps.

Among the IT professionals involved in implementing such projects, we can distinguish data engineers, machine learning engineers, integrators, application developers, and those responsible for deployment and maintenance (Ops). Naturally, each of these roles requires competencies in their respective IT/MLDevOps fields (vertical bar) and horizontal competencies, particularly in the areas of Data Science and business(see Figure 2).

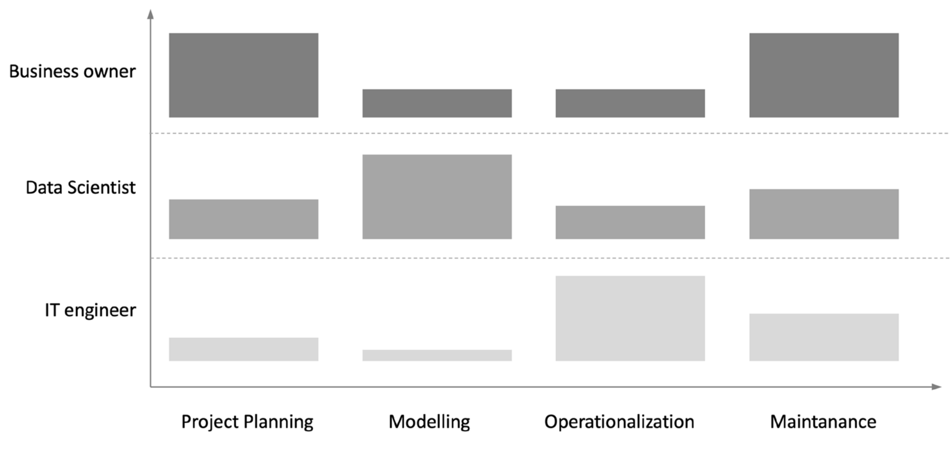

Source: Own elaboration.

Figure 2: Estimate Level of Involvement of Individuals in Various Roles in a Machine Learning Project During its Different Stages

The high involvement of the business representative during the maintenance phase is due to the necessity of monitoring the business value of the solution and recommending subsequent experiment iterations.

Pi-shaped Competencies

Data Science specialists should naturally possess in-depth skills in the Data Science field (including data acquisition, analysis, and preparation, modeling, evaluation, and model quality monitoring) complemented by "horizontal" skills in business and software engineering. However, in the context of managing the value of an intelligent service, the skills related to designing and conducting experiments are so crucial that they can be treated as a second vertical bar in the competency profile, indicating the concept of Pi-shaped competencies.

M-shaped Competencies

Managing a project of this type poses a unique challenge. It requires possessing numerous horizontal competencies (e.g., in the areas of general project management, change management, or business) and at least three vertical pillars: Data Science, IT/MLDevOps, and business experiment design and planning. Consequently, individuals fulfilling this role should have a competency profile known as M-shaped.

4. Knowledge Management Framework for Designing, Monitoring, and Optimizing the Business Value Of Intelligent Services

In the section dedicated to competency profiles, we identified four key competency groups: business, data science, experiment design, and information technology. Individuals fulfilling various roles in the project (business owners, data scientists, IT engineers, and project managers) should possess an appropriate mix of these competencies: T-shaped, Pi-shaped, or M-shaped.

The success of the undertaking is contingent upon the proper flow of knowledge from these areas, both within the scope of the executed project and by leveraging experiences acquired from similar endeavors in the past. Inappropriate "proportions" between these skills may result in errors in assessing the business value of an intelligent service and suboptimal management of that value.

To measure the value (evaluation) of a machine learning model, one can employ business, statistical, and experimental approaches. Each of these approaches has its advantages and disadvantages. A predominance of competencies in one of these areas may naturally lead to an excessive focus on measures of one type, neglecting measures from another category. This may have the following consequences:

- In the case of a predominance of statistical approaches, the model evaluation may either not meet business requirements or be incomprehensible to the business recipient.
- In the case of a predominance of business approaches: an excessive amount of business metrics may generate difficulties both during training (evaluation and validation of models) and in experiments (too many tests to be conducted, difficulties in comparing results).
- In the case of a predominance of experimental approaches, difficulties arise during the formulation of the business case (assessment of potential based on historical data) and during training (evaluation and validation of models).

As a result, it is crucial for the project to select team members in such a way that their competency profiles comprehensively cover not only the skills described above for model evaluation but also business, experimental, data science, and IT competencies. Additionally, ensuring effective knowledge flow between teams and leveraging already accumulated experiences is of great importance.

This is the primary motivation for developing a knowledge management framework, particularly the processes and procedures for creating, acquiring, organizing, storing, and disseminating knowledge using appropriate technologies.

4.1 Key Knowledge Management Processes

Below, we explore various aspects of knowledge management in machine learning projects, from knowledge creation to sharing and application. We discuss the unique nature of ML projects and their potential as valuable sources of knowledge. By examining traditional and contemporary methods for capturing, organizing, storing, and sharing knowledge, we aim to provide insights into effective strategies for managing knowledge in the context of ML projects. Furthermore, we highlight the potential of conversational interfaces and generative technologies in transforming the way knowledge is shared and utilized within organizations.

Knowledge creation. Practically every machine learning project is unique, addressing various business problems using different methods and technologies and carried out by diverse teams. As a result, a well-managed ML project can be a source of valuable knowledge.

Knowledge capturing. This knowledge can be generated at various stages of the project in different areas outlined in the section dedicated to competency profiles. Typical situations conducive to knowledge creation include recurring or ad hoc project meetings, milestone summaries, documentation creation, or interpretation of experimental results.

To capture knowledge, both "traditional" methods (electronic documents, logs from project management systems, or various electronic communication channels such as email or instant messaging) and modern tools utilizing machine learning (e.g., automatically generated meeting transcriptions) can be employed.

Knowledge organization. Similarly, knowledge organization can involve traditional methods such as organizing documents into folders or assigning appropriate tags (facilitating multidimensional classification). Contemporary machine learning methods may additionally enable semantic analysis of collected content (Maulud et al., 2021), automatic tagging and annotations (Lin, 2022), topic modeling (Churchill and Singh, 2022), knowledge graph creation (Chen et al., 2020), or semantic embeddings (Sezerer and Tekir, 2021).

Knowledge storing. Organized knowledge can be stored in folders with electronic documentation, databases, or project archives. However, the rapidly evolving generative technologies encourage considering more advanced methods such as prompt retrieval (Cheng et al., 2023), document embeddings (Sinoara et al., 2019), or generative model weights pretrained on a proprietary knowledge corpus.

Knowledge sharing and application. Properly organized and stored knowledge can be utilized both within the project and in subsequent ventures. The efficiency of these processes will significantly depend on the convenience of access interfaces. Scattering documents across various locations (spatial, technical, or organizational dimensions) hinders knowledge utilization in new projects. A promising group of technologies that can solve this problem is conversational interfaces (text, audio, and graphical modes) based on the aforementioned generative technologies.

It is worth emphasizing that, as of April 2023, there are two dominant methods for using publicly available models to manage internal organization knowledge: prompt-engineering-based approaches (where knowledge is part of a query to a general generative model) and slightly more challenging implementations involving fine-tuning open models for specific tasks on proprietary knowledge resources.

Both approaches are currently implementable, so an increase in popularity of such solutions can be expected in knowledge sharing for machine learning projects as well as projects in other areas. Nevertheless, projects that create intelligent services, due to the profile of the participants, can naturally become the testing ground for developing and testing such modern knowledge transfer solutions.

4.2 Academic Researchers and External Advisors as Catalyzers in ML Projects

As mentioned earlier, in addition to business competencies, crucial skills for implementing machine learning projects and managing the value of such solutions involve Data Science, IT, and experimental design abilities.

An imbalance between these competencies can hinder the assessment of a solution's potential, the evaluation of model quality, and the verification of its business value, ultimately reducing an organization's ability to manage this value in the future.

Our observations indicate that while Data Science and IT competencies are gradually becoming commonplace in mature organizations, knowledge in the area of experimental design and implementation remains a challenge for many organizations.

Moreover, carrying out projects within the same human teams (whether internal or with consistent subcontractors) fosters the development of habitual patterns and conceptual schemas. While this has many advantages (e.g., process streamlining and standardization or improved communication), it can also result in overlooking errors or new problem-solving approaches.

For these reasons, as part of the Knowledge Management Framework, it is worth considering inviting external individuals to the team, such as freelancers or representatives from the academic world. Such individuals can be a source of new ideas, methods, and techniques, as well as bring fresh perspectives to the problem, ultimately improving the solution's quality. Additionally, mature experimental design and implementation methods used in the scientific community can help fill competency gaps in this area.

4.3 Business-Engineer-Academy Knowledge Flywheel

A well-designed knowledge flow within the Business-Engineer-Academy triangle (for simplicity, we combine the roles of Data Scientist and IT Engineer here) can facilitate the so-called flywheel effect: a situation in which successive iterations of knowledge flow enhance the processes of managing the value of intelligent services.

For example:

- The business recommends KPIs that are expected to improve upon implementing an intelligent solution.
- The engineering team (Data Science and IT) technically implements the solution.
- Scientists conduct experiments, analyze the results, verify hypotheses, and present key findings to business representatives.
- Informed by the new knowledge, the business proposes further experiment iterations and refines success metrics.

A well-designed Knowledge Management (KM) flywheel allows for continuous improvement of the process of managing the value of intelligent services, ultimately transcending the paradigm of model monitoring and diagnostics for maintaining their value. This is possible thanks to:

- Systematic processes of generating, capturing, organizing, storing, and sharing knowledge throughout the entire cycle.
- Machine learning methods and techniques employed, particularly CausalML for modeling causal relationships and, as a result, identifying factors that genuinely impact the level of business metrics, and Reinforcement Learning enabling optimization of conducting business experiments (see book on Experiments).

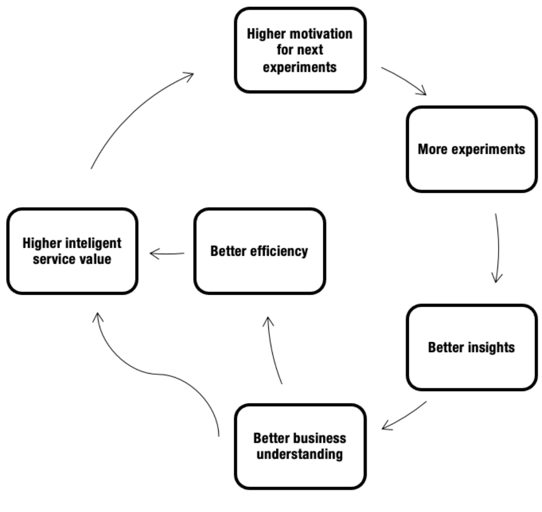

Source: Own elaboration.

Figure 3: An Exemplary Knowledge Management Strategic Flywheel for Machine Learning Projects

Consequently, an exemplary KM flywheel may follow the following scheme:

- New experiments ...
- Generate new insights ...
- ... enabling a better understanding of the business ...
- which improves
    - *business efficiency (measured by KPIs)*
    - *and increases the value of intelligent services*
- which encourages the organization to conduct further experiments.

Thus, a well-designed knowledge flow within the Business-Engineer-Academy triangle can lead to a flywheel effect, where continuous iterations of knowledge exchange improve the management of intelligent services' value. This is achieved through systematic processes of knowledge generation, capture, organization, storage, and sharing, along with the utilization of machine learning methods like CausalML and Reinforcement Learning. As a result, the knowledge management flywheel enables better understanding of the business, improved efficiency, and increased value of intelligent services, encouraging organizations to conduct further experiments.

5. Case Study

ABC (a pseudonym of a real company, which didn’t permit disclousure of their actual name) is an international company operating in the financial industry, one of the larger European institutions providing loans through digital channels, with over 2 mln customers in 7 countries. A key area of operational activity is product promotion, carried out in digital channels, particularly through personal contact (phone, SMS, email) and online advertising.

The costs of conducting promotional campaigns are high, both for ABC (Contact Center services, media purchases) and for customers (too frequent contacts are irritating). Therefore, the company is developing technologies to optimize customer contact.

The primary objective of the project, which serves as a case study in this research, is to increase profits by reducing communication costs and modifying discount policies. To achieve this goal, a model predicting customer interest in a financial loan is developed and utilized as a key method. The model is run weekly and forecasts whether each customer will purchase a financial product in the following week. Based on this prediction, contact is established with the customer, the form of which depends on the customer's purchase history and the predicted probability of purchase.

The project was initiated by a mid-level manager coordinating marketing activities in multiple countries. Based on experience from one market, a decision was made to design a similar solution for two more, for over 100,000 active customers. The project is, however, carried out by a newly formed team, which includes representatives from business, IT, and academia. This approach was motivated by the heavy workload of the existing Data Science team, and the need to look at the problem from a fresh, new perspective. The project management methodology is based on the proven CRISP-DM methodology, but its specifics required some changes in this case.

Due to allowing representatives from the academic world to access sensitive data, it was necessary to implement appropriate legal procedures and grant IT permissions. This procedure significantly extended the initiation of the project but laid the foundation for further ventures of this type.

At the design stage, a decision was made to focus on the experimental verification of the business value of machine learning models. As a result, it was necessary to develop an experiment plan and a dedicated IT architecture.

Planning the experiment, especially selecting the business measure (KPI) to help determine the business value of the model and the historical values and expected test reliability, required the alignment of many concepts and was the first situation allowing for intensive knowledge exchange between company representatives and academics. In turn, the plan to integrate production machine learning systems with systems automating marketing communication forced changes in many communication rules. In summary, the planning and implementation of the experiment required not only the development of new concepts but also significant changes in the architecture of marketing communication systems.

At the modeling stage, one challenge was determining the subject of the forecast. As already mentioned, the primary goal was to optimize marketing communication, focused on "just-in-time" communication: sending the right message to a person with a specific level of product purchase readiness. From a Data Science perspective, this could be achieved either by classification methods (predicting whether (YES/NO) the customer will buy the product in the next week) or regression (forecasting how many days from the current date the customer will contact to purchase). Business representatives, on the other hand, having experience and certain habits from previous similar projects, expected a probability of purchase prediction (which is not the goal itself, but an artifact of the predictive process). Achieving conceptual consensus in this area was somewhat challenging, but by the end of this stage, it was achieved. Moreover, the approach of basing communication not so much on the binary prediction of the model but on the probability value of purchase propensity, which was the starting point for complex procedural rules, was a new experience for representatives from the academic world.

The production launch, using the Google Vertex AI platform, was an opportunity for ABC Company to test the latest MLDevOps tools in production and a basis for assessing their potential in subsequent implementations.

In summary, the experiences described above indicate 1. The significant value of involving individuals from outside the organization in the process of developing intelligent services, 2. The substantial impact of implementing experimental measures on the processes of training, operationalizing, and evaluating the business value of machine learning models, and 3. The necessity of having diverse (T-shaped, Pi-shaped, and M-shaped) competency profiles among team members.

6. Conclusions

This article explores the complex relationship between artificial intelligence and knowledge management, focusing on the practical application of machine learning models in a business context. It discusses various evaluation methods, operationalization challenges, and monitoring and observability techniques for such models.

Drawing from the case study of a financial loan interest prediction model several key lessons were learnt. First, the initial phase of the project (data acquisition) can be significantly shortened by implementing a data governance system and introducing a Data Engineer into the team. Notably, this approach may also simplify procedures concerning data access. Next, in interdisciplinary projects, aligning terminologies in the early stages is of vital importance. It has been observed that the same term may be interpreted differently by different individuals, based on their prior experiences, which can lead to misunderstandings. Therefore, investing time and resources in creating a shared vocabulary can be a significant factor in preventing such communication pitfalls. Finally, one of the crucial business competencies that can influence the pace and quality of a project utilizing machine learning is the practical skill of planning and executing business experiments, and cultivating a culture of their use for empirical validation of the value of new initiatives. This culture can enhance the practicality and relevance of the project, ensuring that initiatives are not only theoretically sound but also empirically validated.

For future research, it is suggested to investigate the effectiveness of cutting-edge machine learning techniques in knowledge management and to explore the potential of advanced methods like CausalML for managing the value of intelligent services, ultimately leading to a better understanding of causal relationships within business processes.

References

Ardito, L., Ferraris, A., Messeni Petruzzelli, A., Bresciani, S., Del Giudice, M., 2019. The role of universities in the knowledge management of smart city projects. Technol. Forecast. Soc. Change, Understanding Smart Cities: Innovation ecosystems, technological advancements, and societal challenges 142, 312–321. https://doi.org/10.1016/j.techfore.2018.07.030

Bjørnson, F.O., Dingsøyr, T., 2008. Knowledge management in software engineering: A systematic review of studied concepts, findings and research methods used. Inf. Softw. Technol. 50, 1055–1068. https://doi.org/10.1016/j.infsof.2008.03.006

Blomster, M., Koivumäki, T., 2022. Exploring the resources, competencies, and capabilities needed for successful machine learning projects in digital marketing. Inf. Syst. E-Bus. Manag. 20, 123–169. https://doi.org/10.1007/s10257-021-00547-y

Carrillo, P., Chinowsky, P., 2006. Exploiting Knowledge Management: The Engineering and Construction Perspective. J. Manag. Eng. 22, 2–10. https://doi.org/10.1061/(ASCE)0742-597X(2006)22:1(2)

Chen, X., Jia, S., Xiang, Y., 2020. A review: Knowledge reasoning over knowledge graph. Expert Syst. Appl. 141, 112948. https://doi.org/10.1016/j.eswa.2019.112948

Cheng, D., Huang, S., Bi, J., Zhan, Y., Liu, J., Wang, Y., Sun, H., Wei, F., Deng, D., Zhang, Q., 2023. UPRISE: Universal Prompt Retrieval for Improving Zero-Shot Evaluation. https://doi.org/10.48550/arXiv.2303.08518

Chung, P.W.H., Cheung, L., Stader, J., Jarvis, P., Moore, J., Macintosh, A., 2003. Knowledge-based process management—an approach to handling adaptive workflow. Knowl.-Based Syst. 16, 149–160. https://doi.org/10.1016/S0950-7051(02)00080-1

Churchill, R., Singh, L., 2022. The Evolution of Topic Modeling. ACM Comput. Surv. 54, 215:1-215:35. https://doi.org/10.1145/3507900

Fakhar Manesh, M., Pellegrini, M.M., Marzi, G., Dabic, M., 2021. Knowledge Management in the Fourth Industrial Revolution: Mapping the Literature and Scoping Future Avenues. IEEE Trans. Eng. Manag. 68, 289–300. https://doi.org/10.1109/TEM.2019.2963489

Kamara, J.M., Augenbroe, G., Anumba, C.J., Carrillo, P.M., 2002. Knowledge management in the architecture, engineering and construction industry. Constr. Innov. 2, 53–67. https://doi.org/10.1108/14714170210814685

Kumar, S., Ravi, V., 2022. Application of Causal Inference to Analytical Customer Relationship Management in Banking and Insurance. https://doi.org/10.48550/arXiv.2208.10916

Lin, B., 2022. Knowledge Management System with NLP-Assisted Annotations: A Brief Survey and Outlook. https://doi.org/10.48550/arXiv.2206.07304

Maulud, D.H., Zeebaree, S.R.M., Jacksi, K., Sadeeq, M.A.M., Sharif, K.H., 2021. State of Art for Semantic Analysis of Natural Language Processing. Qubahan Acad. J. 1, 21–28. https://doi.org/10.48161/qaj.v1n2a44

Provost, F., Fawcett, T., 2013. Data science for business: what you need to know about data mining and data-analytic thinking, 1. ed., 2. release. ed. O'Reilly, Beijing.

Sezerer, E., Tekir, S., 2021. A Survey On Neural Word Embeddings. https://doi.org/10.48550/arXiv.2110.01804

Sharma, A., Kıcıman, E., 2020. Causal inference and counterfactual reasoning. Presented at the ACM International Conference Proceeding Series, pp. 369–370. https://doi.org/10.1145/3371158.3371231

Sinoara, R.A., Camacho-Collados, J., Rossi, R.G., Navigli, R., Rezende, S.O., 2019. Knowledge-enhanced document embeddings for text classification. Knowl.-Based Syst. 163, 955–971. https://doi.org/10.1016/j.knosys.2018.10.026

Wang, X., Zhang, X., Cheng, Y., Tian, F., Chen, K., Ordóñez de Pablos, P., 2022. Artificial intelligence-enabled knowledge management, in: The Routledge Companion to Knowledge Management. pp. 153–168. https://doi.org/10.4324/9781003112150-13

Zhang, K., Zhu, S., Kalander, M., Ng, I., Ye, J., Chen, Z., Pan, L., 2021. gCastle: A Python Toolbox for Causal Discovery. https://doi.org/10.48550/arXiv.2111.15155