Zarządzanie Innowacyjne w Gospodarce i Biznesie nr 2(43)/2026
ISSN 2391-5129

Damian Kociemba https://orcid.org/0000-0002-6604-858X
Independent scholar
e-mail: kociemba.damian.kd@gmail.com

Practical implementation of an AI model (Ministral 3B) in the logistics of dangerous goods

Praktyczne wdrożenie modelu sztucznej inteligencji Ministral 3B w logistyce towarów niebezpiecznych

https://doi.org/10.25312/ziwgib.888

Abstract

This article analyses the practical implementation of a local large language model (Ministral 3B) supported by Retrieval-Augmented Generation (RAG) in dangerous goods logistics, with a focus on verifying ADR exemptions, reducing AI hallucinations, and maintaining data privacy. The study uses a critical literature review and an empirical case study. The model was tested locally under stringent parameter constraints, including temperature 0.1 and strict system instructions. The analysis covered one comprehensive query concerning the transport of goods packed in limited quantities (LQ). Model responses were assessed for substantive compliance with ADR regulations and the commentary of a Dangerous Goods Safety Adviser (DGSA), comparing outputs generated with and without RAG. The results indicate that the RAG-supported model synthesized legal information more accurately and substantially reduced hallucinations in the analysed case compared with the raw model, which fabricated non-existent legal structures. However, conservative parameter settings may cause the omission of minor technical nuances. The findings suggest that local LLMs combined with RAG can support DGSA work, but only as auxiliary tools subject to continuous expert verification.

Keywords: artificial intelligence, Large Language Models (LLM), Retrieval-Augmented Generation (RAG), ADR Agreement, dangerous goods logistics, AI hallucinations

Streszczenie

Celem artykułu jest analiza praktycznego wdrożenia lokalnego dużego modelu językowego (Ministral 3B) wspieranego architekturą RAG (Retrieval-Augmented Generation) w logistyce towarów niebezpiecznych, z naciskiem na weryfikację wyłączeń z przepisów umowy ADR przy jednoczesnym ograniczeniu zjawiska halucynacji AI i zachowaniu pełnej prywatności danych. W pracy zastosowano krytyczny przegląd literatury oraz metodę empirycznego studium przypadku. Model przetestowano w środowisku lokalnym przy rygorystycznych ograniczeniach parametrów (temperatura 0.1, ścisłe instrukcje systemowe). W ramach badania przeanalizowano jedno kompleksowe zapytanie dotyczące przewozu towarów pakowanych w ilościach ograniczonych (LQ). Odpowiedzi modelu oceniano pod kątem zgodności merytorycznej z przepisami ADR i komentarzem doradcy ds. bezpieczeństwa ADR (DGSA), porównując warianty z RAG i bez niego. Eksperyment wykazał, że model wspierany przez RAG skutecznie syntetyzuje informacje prawne, niemal całkowicie eliminując zjawisko halucynacji, któremu uległ model „surowy" (fabrykujący nieistniejące struktury prawne). Zauważono jednak, że bardzo zachowawcze ustawienia parametrów mogą prowadzić do pomijania drobnych niuansów technicznych przez algorytm. Integracja lokalnych modeli LLM z architekturą RAG stanowi bezpieczne, tanie i wysoce użyteczne narzędzie wspierające pracę doradców ds. bezpieczeństwa. Niemniej, ze względu na rygorystyczny charakter przepisów transportowych, sztuczna inteligencja musi pozostać jedynie narzędziem pomocniczym, podlegającym stałej, merytorycznej weryfikacji przez wykwalifikowany personel.

Słowa kluczowe: sztuczna inteligencja, duże modele językowe (LLM), Retrieval-Augmented Generation (RAG), umowa ADR, logistyka towarów niebezpiecznych, halucynacje AI

Introduction

In an era of widespread fascination with Large Language Models (LLMs) such as GPT-4 or Claude, researchers often overlook the fact that in rigorously regulated areas – to which the logistics of dangerous goods undoubtedly belongs – the key factor for success is not the model's "creativity", but its deterministic precision. For dangerous goods safety advisers (DGSAs) and employees involved in operations related to the movement of dangerous goods, the priority is not generating multi-page essays, but obtaining an instantaneous and error-free answer to a specific question regarding the regulations contained in the ADR Agreement.

This article analyses the practical implementation of the Ministral 3B model as an assistant in the process of verifying conditions for exemptions from regulations. The choice of a model with a relatively small scale (3 billion parameters) was dictated by pragmatism – the desire to maintain full locality of the process on standard office hardware. This solution is dedicated to logistics companies that focus primarily on administrative and

operational tasks and do not possess advanced computing infrastructure, while simultaneously prioritizing data security.

The following research hypothesis is formulated in this article: the use of the Ministral 3B language model in combination with RAG (Retrieval-Augmented Generation) architecture and rigorous parameter configuration allows for a tool with high substantive utility that effectively limits the phenomenon of AI hallucinations in the interpretation of fragments of ADR Agreement regulations, while maintaining full privacy of user data.

To verify the formulated hypothesis, a set of tools was used to run the model in a local environment. A key aspect was the configuration of the model's "decision variables" to maximally reduce its tendency to confabulate in favour of a faithful synthesis of the provided material. These variables, along with a description of the tools, are included in the final section of the publication.

To verify the research hypothesis and provide a reliable approach to the issues discussed, this work utilizes a critical analysis of the literature and the empirical case study method. The structure of the article is designed to clearly guide the reader from theoretical foundations to conclusions from practical implementation. First, the evolution of artificial intelligence and its inherent weaknesses (such as the phenomenon of hallucinations) are outlined. Next, contemporary applications of AI in the logistics industry are discussed, serving as a starting point for the main research part of the study, which presents the test environment and conducts a comparative analysis of the results generated by the model supported by RAG architecture versus the model deprived of this context. The whole concludes with an evaluation of the obtained results.

The evolution of artificial intelligence

The fascination with machines mimicking human thought dates back to ancient times; however, it was Aristotle who laid the foundations for modern logic by introducing the concept of syllogisms (Rai, 2024: 2). The formal birth of the field of artificial intelligence (AI) is dated to 1956, when during the Dartmouth Conference, American computer scientist John McCarthy first used this term (Alalaq, 2024: 3; Rai, 2024: 5–6). Over the following decades, the development of AI was primarily based on symbolic systems and so-called expert systems, such as the chemical DENDRAL or medical MYCIN from the turn of the 1960s and 1970s (Singh, Banerjee, Gosh, 2025: 4–5). These required rigorous, manual coding of thousands of decision rules, which made them useful in strictly defined areas, yet extremely rigid and incapable of independently understanding the nuances of natural language.

To realize the dynamics of contemporary technological progress, it is worth looking back just a few years. In an article published in 2019, the author of this article analysed the implementation of artificial neural networks in the process of creating autonomous vehicles. At that time, commercial applications of AI in transport and logistics were based on so-called Narrow Artificial Intelligence (Weak AI). These systems focused on mathematical machine learning – the laborious selection of weights for artificial

neurons to process sensory signals. Their task was to trigger pre-defined reactions, such as maintaining the vehicle in its lane or activating emergency brake assist (EBA). The research problems of that period centred on technical limitations of algorithms and the lack of legal regulations for machines making decisions on the road (Kociemba, 2019: 171, 176, 178–179, 185). These tools, although revolutionary, remained "blind" to complex context – the machine was able to stop a truck in front of an obstacle, but depending on the training of the algorithm (i.e., the presentation of a given event to the algorithm during the learning stage), independent and safe avoidance of said obstacle was not always possible.

The turning point that moved artificial intelligence from the era of simple image analysis and rigid algorithms to the era of fluent text understanding was the introduction of the Transformer architecture. The "self-attention" mechanism used within it allowed models to analyse vast portions of text in parallel and understand the context of words regardless of their position in a sentence. This discovery led to an unprecedented explosion of so-called Large Language Models (LLMs). In 2020, the GPT-3 model was presented with a scale of 175 billion parameters, which proved the ability for "zero-shot" learning – it was able to generate text, translate, and solve logical problems without the necessity of previously programming specific rules for it. The release of ChatGPT at the end of 2022 proved, in turn, that these models are capable of fluently conducting multi-turn conversations and reasoning in the face of ambiguity (Alalaq, 2024: 6; Rai, 2024: 18; Singh, Banerjee, Gosh, 2025: 6, 17–18, 21–22, 25).

Equally important, the years 2023–2024 brought a flood of powerful open-source models, such as the LLaMA or Mistral families, to which the Ministral 3B model analysed in this study belongs. They offered performance close to technological giants, but with drastically lower demand for computing power and the possibility of full control over data privacy (Singh, Banerjee, Gosh, 2025: 26, 29–30).

Thanks to this very evolution – from narrow algorithms physically controlling a vehicle in 2019 to lightweight, open (and constantly developed) generative models – it became possible to realize the implementation described in this article. Personnel involved in operations related to the transport of dangerous goods can gain an innovative supporting tool that optimizes the manual document search process. However, as further analysis will show, this technology – although it reduces cognitive debt and drastically shortens regulation verification time – does not eliminate the human factor, remaining a supporting tool that requires constant substantive verification by trained personnel.

Contemporary problems related to the development of AI

Entering the world of artificial intelligence, it is easy to succumb to initial fascination. This phenomenon is perfectly illustrated by what researchers describe as an adapted Dunning-Kruger effect in the context of using AI. Initially, the user is full of enthusiasm, believing that thanks to AI, they will achieve immediate and error-free results. Over time, however, they hit a wall – wasting time on verification, receiving fabricated data,

and beginning to understand the technology's limitations. Only overcoming this "trough of disillusionment" leads to a conscious and mature treatment of the model simply as another tool (Salvagno, Taccone, Gerli, 2023: 1–2).

The most serious problem of generative artificial intelligence (LLMs) is so-called hallucinations. These occur when the model generates false, misleading, or entirely fabricated information, yet presents it with immense confidence and impeccable linguistic fluency (Joshi, 2025: 38). These are not mere typos. Models can create coherent-sounding but completely non-existent bibliographic references, fabricating well-known authors and publication titles (Salvagno, Taccone, Gerli, 2023: 1). In specialized fields such as law or medicine, the rates of such hallucinations can reach double digits. This problem stems from the architecture of the models themselves. LLMs are, by design, powerful statistical engines that generate text token-by-token (word-by-word), attempting to predict what should come next without true, human "understanding". They lack grounding in reality. When faced with uncertainty, the model prefers to fill knowledge gaps with "plausible-sounding" information, prioritizing linguistic fluency over factual truth (Joshi, 2025: 39).

One must keep in mind that AI models are only as good as the data on which they were trained. In industrial and logistics applications, insufficient, noisy, or incorrect data fed into the model can lead to erroneous predictions, which in turn risks operational disruptions (Sinha, Lee, 2024: 4). At this point, researchers draw attention to an extremely intriguing aspect: AI hallucinations are largely an extension of "human hallucinations". LLMs learn from massive datasets from the Internet, simultaneously absorbing human biases, stereotypes, cognitive errors, and even deliberate misinformation. An AI model simply serves as a digital mirror of human imperfection. If we provide it with a contaminated knowledge base, it will perpetuate and automate that imperfection (Tlili, Burgos, 2025: 68–69).

Additionally, artificial intelligence is exposed to so-called data drift and concept drift. Regulations, the working environment, and technologies are constantly changing. A model trained on static, historical data loses its precision over time because statistical relationships between variables no longer correspond to the new reality (Sinha, Lee, 2024: 7). The absence of mechanisms for continuous learning and adaptation means that an expert can quickly become a dilettante.

Given the above limitations – from the tendency to confabulate and the "inheritance" of human errors to the susceptibility to obsolescence – one might conclude that artificial intelligence is a tool too unpredictable for the business environment. Meanwhile, as researchers note, a complete rejection of this technology in the face of growing market pressure would be a mistake (Sinha, Lee, 2024: 1). The transport and logistics industry, operating on massive volumes of data and under constant pressure to optimize time and costs, now treats AI not as an optional curiosity but as a strategic necessity (Sadiku, Ajayi, Sadiku, 2026: 171, 172). The key to success is therefore not expecting absolute infallibility from algorithms, but building systems that mitigate the risks described above. Awareness of these technical "traps" constitutes the foundation for the proper design and implementation of solutions based on artificial intelligence.

Using AI in logistics

Although AI still struggles with problems, the benefits of its implementation in logistics are invaluable. Skilful risk management allows translating theory into measurable optimization.

Contemporary logistics relies on complex information systems. AI models are becoming a key element of these architectures, taking them to a higher level of efficiency (Dyczkowska, Olkiewicz, Królikowski, 2025: 5199–5202, 5204; Sadiku, Ajayi, Sadiku, 2026: 173–174):

- Supply Chain Management (SCM): these systems are evolving from reactive tools into intelligent, predictive networks. Machine learning (ML) allows for more accurate demand forecasting and delivery time optimization by analysing historical data while taking into account external factors, such as weather or the geopolitical situation. As a result, logistics managers can make more proactive decisions;
- Warehouse Management System (WMS): AI and ML drastically reduce operational errors and optimize resource utilization in the warehouse. Advanced algorithms can continuously modify the warehouse structure and the roadmap for equipment (e.g., forklifts or autonomous robots) to minimize so-called "empty runs". The system can intelligently group products that are frequently ordered together and place them in the same zone, which significantly accelerates the order picking process;
- Enterprise Resource Planning (ERP): implementing AI modules in ERP systems automates, among other things, accounting processes and procurement management, reducing the risk of human error. Algorithms also monitor user behavior, warning of unauthorized access attempts.

Another major application of artificial intelligence in transport is route optimization. AI analyses gigantic volumes of data from satellite navigation and road sensors, determining the most cost-effective (in terms of time and money) and ecological routes. This streamlines the "last mile" – the most expensive segment of the logistics chain (Dyczkowska, Olkiewicz, Królikowski, 2025: 5199; Sadiku, Ajayi, Sadiku, 2026: 173).

AI also plays a massive role in predictive maintenance. Instead of reacting to a failure or relying solely on a servicing schedule, algorithms – based on data flowing directly from machines and trucks – predict the moment when a part is most likely to fail. Such a strategy significantly lowers repair costs and minimizes the risk of unexpected downtime (Sadiku, Ajayi, Sadiku, 2026: 174).

Artificial intelligence is changing the way clients and employees communicate with companies. Fast-reacting chatbots – based on natural language processing (NLP) – not only provide information about a shipment or carrier, but are also capable of solving simple problems, thereby relieving office personnel (Dyczkowska, Olkiewicz, Królikowski, 2025: 5200; Sadiku, Ajayi, Sadiku, 2026: 174).

Since models based on natural language processing (NLP) successfully automate processes and support communication, the natural next step is adapting them to one of

the most demanding operational areas – the interpretation of rigorous legal regulations. The implementation of such a "digital expert" in the specific environment of dangerous goods transport (ADR) does not, however, have to involve multi-million investments in cloud infrastructure or the risk of leaking sensitive operational data. The experiment presented in the following section serves as a bridge between the general capabilities of AI and pragmatic practice, proving that with appropriate configuration, precise analysis of legal texts can be conducted entirely locally.

Description and conclusions from the experiment

In accordance with the assumption of using standard office equipment, the tests were conducted on a computer equipped with 16 GB of RAM, which, as mentioned in the introduction, is intended to prove the accessibility of AI technology for an average transport enterprise.

The software layer consisted of: (1) LM Studio: a runtime environment for the Ministral 3B model, providing control over the processes of loading model layers into memory, and (2) LLMStudio: a program responsible for indexing the content of Volume 1 of the ADR Agreement and dynamically providing relevant text fragments to the model at the time of asking a question. This program contains a RAG mechanism, the task of which is to provide the model running in the LM Studio environment with ADR Agreement regulations, from which the model selects appropriate paragraphs (most often containing appropriate keywords), and then, based on them, synthesizes the answer.

In the process of implementing the AI assistant, rigorous system instructions were applied. The model was instructed on its function as a "specialist in the field of the ADR Agreement" with the following restrictions:

- Source exclusivity: the model received a categorical ban on using general (pre-training) knowledge if the answer was found in the provided document;
- Temperature (0.1): this parameter was set to the minimum level to "freeze" the model's creativity. In the logistics of dangerous goods, where an error in interpretation can result in massive losses or financial sanctions, predictability of the answer is a paramount value;
- Missing data management: the model was instructed that in the absence of precise information in the source text, it should provide the answer: "No legal basis found in the provided fragment of the ADR Agreement", instead of attempting extrapolation.

The above settings are, in fact, technical "muzzles" imposed to protect the verification process against the algorithm's innate creativity and hallucinations. Understanding the weaknesses and challenges outlined earlier – the tendency to fabricate facts, sensitivity to data quality, and unpredictability resulting from the "black box" problem – fully justifies the use of such strict rigors. The extremely low temperature and forcing the use exclusively of the provided PDF files via the RAG architecture were a necessary condition

for artificial intelligence to even be considered as useful support for a DGSA, rather than merely a stochastic text generator.

An important measure was the deliberate selection of data. The model was provided only with a truncated version of Volume 1 of the ADR Agreement. Extensive tables were removed from it (including: tables in part one relating to the list of high consequence dangerous goods, tables in classes: 1 (explosives), 2 (gases), 3 (flammable liquids), 4.1 (self-reactive substances), 5.2 (organic peroxides), and 7 (radioactive material), as well as the dangerous goods list under 3.2). This decision stemmed from the desire to optimize the model's work – providing an excessive amount of text (including Volume 2, which does not contain regulations crucial for this study) could negatively affect the precision of the answers. In the RAG architecture, it is crucial to provide only relevant fragments so as not to "distract" the algorithm's attention. The author's experience gained during the work on this publication indicates that small language models can exhibit difficulties in interpreting data from columns and rows, which poses a risk of incorrectly citing weight limits or UN numbers. Focusing on descriptive (textual) fragments allowed for a higher quality of synthesis. However, it is worth noting that when using more powerful computing units and larger models, including tables in the knowledge base would be desirable, provided it is preceded by a process of optimizing their loading format (e.g., conversion to an appropriate format or applying an Optical Character Recognition (OCR) layer, since some of them are low-quality screenshots).

To verify the effectiveness of these settings, the following section presents a comparison of the results generated by the "raw" model and the RAG-fed model, using the example of regulations regarding exemptions for limited quantities (LQ).

To verify the formulated hypothesis, the model was asked the question: "Based on the provisions of the ADR Agreement, provide all technical conditions for the LQ exemption". The table below compares the model's output with the commentary of an ADR adviser (DGSA).

Tab. 1. AI model response and author's comment

AI model response (removed excess formatting, original spelling preserved)	**Author's comment**
Based on Chapter 3.4 of the ADR Agreement (Volume 1) the technical conditions for the exemption of goods in limited quantities (LQ) are as follows: --- 1. Basic verification – Check Column (7a) of Table A – a value of "0" indicates a prohibition of the LQ exemption for a given substance. ---	1. The model correctly addressed the logic of action in the case of transport of dangerous goods (paraphrased 3.4.1).

AI model response (removed excess formatting, original spelling preserved)	Author's comment
2. Gross mass limits of a package – Packages (inner packagings within an outer packaging): Maximum 30 kg. – Shrink-wrapped or stretch-wrapped trays: Maximum 20 kg. ---	2. The model correctly addressed the maximum limits (paraphrased from 3.4.2)
3. Packaging requirements – Class 8 liquid substances (packing group II): – In glass/porcelain packagings, they must be enclosed in rigid intermediate packagings. – Fragile inner packagings (glass, ceramics, stoneware): – Must be secured against breakage within the outer packaging, in accordance with construction requirements (e.g., 6.1.4 ADR). ---	3. The model correctly provided packaging requirements by referring to the provisions in 3.4.4, 3.4.3 and 3.4.1, but omitted the information that for aerosols and small gas receptacles the use of inner packaging is not necessary.
4. Marking of parcels – Diamond mark (square shape set at a 45° angle): – Top and bottom portions and the border line – black. – Central area – white or contrasting. – Minimum dimentions – 100 × 100 mm, minimum width of the border line = 2 mm. – Possibility of reduction to 50 × 50 mm (1 mm line) if the mark remains clearly visible. ---	4. The model correctly described the marking of packages, referring to 3.4.7 (the model uses the terminology "parcel" as a synonym for "package", which is semantically acceptable, although it does not appear explicitly in 3.4). Note that the model used the word "dimentions" instead of "dimensions".
5. Overpacks – Additional marking: – The inscription "OVERPACK" (letter height ≥ 12 mm). – If the marks on the packages inside are not visible, the LQ diamond mark must be placed on the overpack. ---	5. The model correctly described the marking of overpacks (i.e. packaging in which multiple items are placed – e.g. on a pallet) using the information provided in 3.4.11, but omitted the requirement to use the official language of the country of origin and (if required) English, French or German.

AI model response (removed excess formatting, original spelling preserved)	Author's comment
6. Fees and transport obligations – The consignor must provide the carrier with documented information on the total gross mass of LQ goods before the start of the transport. – Transport units >12 t: – Marking on the front and rear (250 × 250 mm diamonds) only when the mass of LQ goods exceeds 8 tonnes gross per unit.	6. The model introduced an artifact – the use of the word "fees" is a hallucination – the ADR Agreement does not define them. The information provided below is generally correct; the model correctly used the information provided in 3.4.12, but did not supplement it with container markings (two fronts and two sides).

Source: own study based on the AI model and the Government Statement of 13 March 2023 on the entry into force of amendments to Annexes A and B to the Agreement concerning the International Carriage of Dangerous Goods by Road (ADR), done at Geneva on 30 September 1957.

It should be noted that the above result was the effect of multiple iterations and parameter tuning.

To definitively confirm the effectiveness of the RAG architecture, the author conducted a control test. After obtaining a satisfactory response, the knowledge base in the form of ADR Agreement regulations was "detached", leaving only the same general system instructions. The identical question was asked, to which the model replied: "There are no unambiguous, general technical conditions for exemption for the (LQ) category in the ADR Agreement […]. According to ANNEX III and ANNEX IV, there are general rules for hazardous liquids". This result is a hallucination (and partially a confabulation). The model fabricated a structure of annexes that the ADR Agreement does not use in this context. This proves that without an "external memory" in the form of RAG, an advanced language model is not useful in operational practice.

The conducted experiment yields promising conclusions for entities involved in operations handling dangerous goods:

– The use of RAG is a "ready-to-use" solution. In contrast to time-consuming and costly model fine-tuning, RAG allows for an instantaneous update of the knowledge base with every change in the ADR regulations (which occurs cyclically every 2 years);
– Obtaining a reliable assistant does not require large financial outlays or specialized programming knowledge.

Despite promising results, the implementation of the AI model demonstrated a tendency to omit detailed requirements, such as linguistic nuances in marking or specific technical details regarding transport units. It should be noted, however, that the observed deficits may not result from the inherent limitations of the Ministral 3B model, but from the conservative configuration strategy adopted in this study. The use of technical "safeguards" in the form of minimum temperature and restrictive instructions intended to protect against hallucinations may have paradoxically limited the algorithm's ability to synthesize

regulations scattered throughout the text. This technological compromise between precision and the ability to connect facts further justifies the necessity for constant supervision by an ADR adviser and qualified personnel over the process of generating and verifying answers. In their current form, LLM models should be treated solely as a supplementary tool, the main task of which is to reduce the time spent manually searching the extensive documentation of the ADR Agreement.

Directions for future research and development perspectives

The case study presented in this paper is merely an introduction to a broader discussion on the implementation of local language models in the logistics of dangerous goods. The promising results of the experiment pave the way for further stages of research and development, among which the following should be prioritized:

- Optimization of structured data processing: further research should focus on testing methods for integrating extensive tables (e.g., Table A from Chapter 3.2 of the ADR Agreement) with the RAG knowledge base. This requires investigating the effectiveness of advanced OCR (Optical Character Recognition) techniques and converting tables into formats semantically understandable to LLM models;
- Improvement of vectorization and chunking techniques: to eliminate the problem of omitting information scattered throughout the text (such as language requirements for overpacks), it is necessary to test other text splitting algorithms (so-called chunking) and apply hybrid search (combining semantic search with keyword search);
- Evolution from an assistant to an agent: a forward-looking direction is the transition from a passive model (answering questions) to agentic systems that will be capable of proactive action – for example, the automatic verification of the correctness of a generated transport document based on an uploaded scan of a waybill.

The implementation of artificial intelligence in logistics (not only of dangerous goods) is no longer a thing of the distant future, but a technological reality. The democratization of access to AI tools means that the barrier to entry for these innovations has drastically decreased. The real challenge for the industry today is not access to technology, but the skilful and responsible management of it. AI will not replace experts – it will only replace those experts who do not know how to use AI.

References

Alalaq A. (2024), *The History of the Artificial Intelligence Revolution and the Nature of Generative AI Work*, "DS Journal of Artificial Intelligence and Robotics", vol. 2(4), pp. 1–24.

Dyczkowska J., Olkiewicz M., Królikowski T. (2025), *The use of artificial intelligence in logistics information systems*, "Procedia Computer Science", no. 270, pp. 5198–5206.

Government Statement of 13 March 2023 on the entry into force of amendments to Annexes A and B to the Agreement concerning the International Carriage of Dangerous Goods by Road (ADR), done at Geneva on 30 September 1957 (Oświadczenie rządowe z dnia 13 marca 2023 r. w sprawie wejścia w życie zmian do załączników A i B do Umowy dotyczącej międzynarodowego przewozu drogowego towarów niebezpiecznych (ADR), sporządzonej w Genewie dnia 30 września 1957 r.).

Joshi S. (2025), *Comprehensive Review of AI Hallucinations: Impacts and Mitigation Strategies for Financial and Business Applications*, "International Journal of Computer Applications Technology and Research", vol. 14(6), pp. 38–50.

Kociemba D. (2019), *Wykorzystanie metod sztucznej inteligencji w tworzeniu pojazdów autonomicznych*, "Zarządzanie Innowacyjne w Gospodarce i Biznesie", vol. 28(1), pp. 171–188.

Rai D. (2024), *Artificial Intelligence Through Time: A Comprehensive Historical Review*, Tribhuvan University, Institute of Science and Technology, https://www.researchgate.net/profile/Dilli-Hang-Rai/publication/385939923_Artificial_Intelligence_Through_Time_A_Comprehensive_Historical_Review/links/676fb15600aa3770e0c2516f/Artificial-Intelligence-Through-Time-A-Comprehensive-Historical-Review.pdf [accessed: 13.10.2026].

Sadiku M., Ajayi S., Sadiku J. (2026), *Artificial Intelligence in Logistics*, "International Journal of Trend in Scientific Research and Development (IJTSRD)", vol. 10(1), pp. 171–178.

Salvagno M., Taccone F., Gerli A. (2023), *Artificial intelligence hallucinations*, "Critical Care" vol. 27, articele no. 180.

Singh G., Banerjee T., Ghosh N. (2025), *Tracing the evolution of artificial intelligence: A review of tools, frameworks, and technologies (1950–2025)*, preprint, https://doi.org/10.31224/5802

Sinha S., Lee Y. (2024), *Challenges with developing and deploying AI models and applications in industrial systems*, "Discover Artificial Intelligence", vol. 4, article no. 55.

Tlili A., Burgos D. (2025), *AI Hallucinations? What About Human Hallucination?! Addressing Human Imperfection Is Needed for an Ethical AI*, "International Journal of Interactive Multimedia and Artificial Intelligence", vol. 9(2), pp. 68–71.

About the Author

Damian Kociemba – graduate of engineering studies in transport and master's studies in management. He is the co-author and author of several academic and specialist publications in the fields of logistics, safety, and optimization. He holds a safety advisor certificate for the transport of dangerous goods by road (ADR).

O autorze

Damian Kociemba – absolwent studiów inżynierskich na kierunku transport oraz magisterskich na kierunku zarządzanie. Współautor i autor kilkunastu publikacji naukowych i specjalistycznych z zakresu logistyki, bezpieczeństwa i optymalizacji. Posiadacz świadectwa doradcy ds. bezpieczeństwa przewozu towarów niebezpiecznych w transporcie drogowym (ADR).

Ten utwór jest dostępny na licencji Creative Commons Uznanie autorstwa-Na tych samych warunkach 4.0 Międzynarodowe.

Open Access