Unlocking the Potential. An Introduction to Artificial Intelligence and its Potential Use in Logistics

May 2024

CIFFA

Introduction

What is AI?

Artificial intelligence (AI) refers to the simulation of human intelligence in machines that are programmed to think and learn like humans. This technology encompasses a broad range of capabilities, from basic algorithms capable of performing simple tasks to highly sophisticated systems that can process and interpret vast amounts of data. At its core, AI is about creating systems that can operate autonomously, make decisions and perform tasks without constant human guidance. This is achieved through various subfields, such as machine learning, where computers are trained to learn from and adapt to new data without being explicitly programmed for every task.

Types of AI

Artificial intelligence (AI) encompasses a variety of types and subfields, each with unique capabilities and applications. At a high level, AI can be categorized into two main types: narrow AI and general AI.

Narrow artificial intelligence

Narrow artificial intelligence, also known as weak AI, is the only type of AI that exists today. It can be trained to perform a single or narrow task, often far faster and better than a human mind can. However, it can't perform outside of its defined task. Instead, it targets a single subset of cognitive abilities and advances in that spectrum.

Artificial general intelligence (AGI)

Artificial general intelligence (AGI) is a branch of theoretical artificial intelligence research working to develop AI with a human level of cognitive function. AGI would allow machines to comprehend, learn, and perform intellectual tasks much like humans, with a self-aware consciousness that has the ability to solve problems and plan for the future.

Subfields

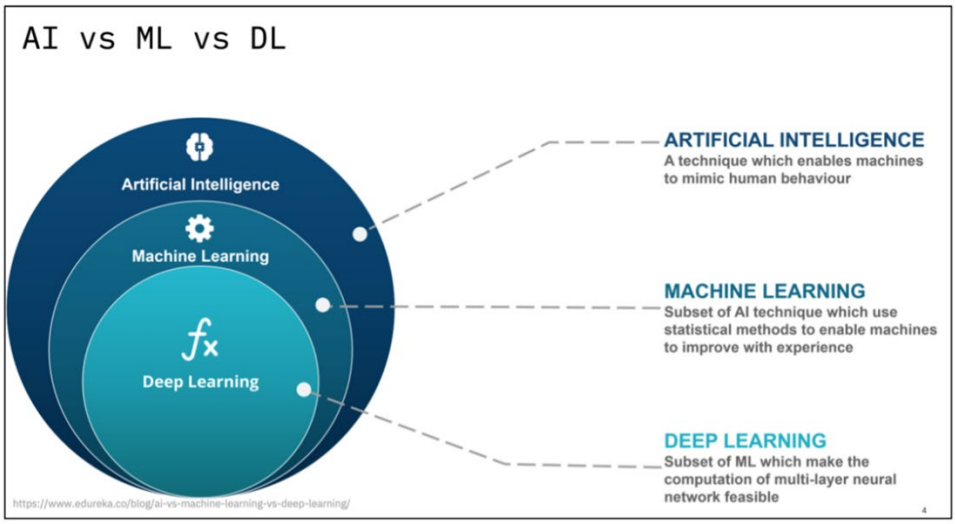

Machine learning

Machine learning (ML) and AI are closely related fields, but they are not synonymous. AI is a broader concept that refers to machines designed to act intelligently, mimicking human cognitive functions like learning, problem-solving, and decision-making. It's an umbrella term that encompasses various technologies, including ML, where machines can learn from data, identify patterns, and make decisions with minimal human intervention.

Machine learning, a pivotal subfield of AI, is specifically focused on developing algorithms that enable computers to learn from and make predictions or decisions based on data. Unlike traditional programming, where a machine follows pre-defined rules, ML systems improve automatically through experience.

Deep learning

Deep learning is a subset of machine learning that uses multi-layered neural networks, called deep neural networks, to simulate the complex decision-making power of the human brain. Some form of deep learning powers most of the artificial intelligence in our lives today.

Deep learning drives many applications and services that improve automation, performing analytical and physical tasks without human intervention. It lies behind everyday products

and services—e.g., digital assistants, voice-enabled TV remotes, credit card fraud detection, as well as still-emerging technologies such as self-driving cars.

Natural language processing

Natural language processing (NLP), a critical subfield of artificial intelligence, enables machines to understand, interpret, and respond to human language. NLP combines computational linguistics—rule-based modeling of human language—with statistical, machine-learning, and deep-learning models. These technologies allow computers to process human language in the form of text or voice data and 'understand' its full meaning, complete with the speaker's or writer's intent and sentiment.

The applications of NLP are vast and varied. In customer service, AI-powered chatbots and virtual assistants use NLP to interact with customers in natural language, providing responses that are contextually relevant and personalized. In the realm of data analytics, NLP is used to sift through large volumes of textual data—such as social media posts, customer reviews, and news articles—to extract insights, monitor brand sentiment, and identify trends. This ability to analyze unstructured data can be invaluable for businesses and organizations seeking to understand the vast landscape of public opinion and behaviour.

Furthermore, NLP enables translation services that can instantly convert text or speech from one language to another. Additionally, NLP is used in voice-recognition systems, allowing for hands-free control of devices and accessibility features for those with physical or visual impairments.

Additional subfields

Another significant subfield is robotics, which involves designing intelligent machines capable of performing tasks in the physical world. Computer vision, on the other hand, enables machines to interpret and make decisions based on visual data from the surrounding environment. Finally, there's predictive analytics, which uses data, statistical algorithms, and ML techniques to identify the likelihood of future outcomes based on historical data. This is particularly useful in logistics for forecasting demand, optimizing routes, and managing inventory. Each of these subfields contributes to the overall landscape of AI, offering tools and technologies that are increasingly relevant to the logistics industry.

When to use AI

Large data sets

Artificial intelligence is particularly adept at handling and extracting value from large data sets through machine learning. Systems can be trained to analyze data, learn from it, and make predictions or decisions based on their learning. This process is dynamic and self-improving; the more data the AI system processes, the more accurate it becomes.

Pattern recognition

At its core, pattern recognition involves the identification of regularities and anomalies within data, a task for which AI, especially machine learning algorithms, is exceptionally suited. These algorithms can sift through vast and complex data sets, discerning patterns that are not immediately apparent to human analysts. This capability is crucial in environments where understanding trends, behaviours, and relationships within the data can lead to significant insights and strategic decisions.

AI's strength in pattern recognition is particularly evident in its ability to learn and improve over time. With each new data point, AI systems refine their algorithms and can manage the sheer scale of data much more efficiently than traditional methods, enabling the analysis of larger data sets without compromising speed or accuracy.

Repetitive tasks

Global trade is a complex and ever-changing industry. Yet many of the day-to-day functions within organizations could be categorized as routine and repetitive.

By identifying the repetitive tasks in our organizations that consume the greatest amount of resources, we can work to choose the right AI product to reshape our manual chores into automated processes, and refocus our people to a higher level of work and contribution. Examples include:

- Email and communication responses (“smart replies”, email classification)
- Tracking updates and data scraping (transferring data between application, reducing manual intervention and data entry)
- Report generation (using AI reporting tools to increase productivity)

Predictive analytics

At its essence, predictive analytics involves using historical data, statistical algorithms, and machine learning techniques to identify the likelihood of future outcomes.

The strength of AI in predictive analytics lies in its ability to handle large and diverse data sets, extracting meaningful patterns and insights that are not easily discernible. Machine learning models are particularly adept at identifying correlations and trends within data, revealing hidden patterns, uncovering causal relationships, and even suggesting actionable strategies based on predictive outcomes. This level of insight is invaluable for strategic planning, risk management, and optimizing operational efficiencies.

Real-time decision making

In scenarios where time is of the essence and the window for decision-making is narrow, AI systems can process and analyze data at unprecedented speeds, delivering insights and recommendations in a fraction of the time it would take for human analysis. This rapid processing capability is crucial in environments where conditions change rapidly and decisions need to be both timely and data-driven.

AI continuously analyzes streaming data from various sources, such as sensors, user interactions, and live market feeds. By employing advanced algorithms and machine learning techniques, AI can identify patterns, anomalies, and trends in this data as it's being generated.

Moreover, AI's role in real-time decision making extends to predictive and prescriptive analytics. AI systems can predict future scenarios based on current data and can suggest the best course of action to achieve desired outcomes, particularly valuable in strategic planning and operational management.

Customization and personalization

In a world where one-size-fits-all solutions are increasingly inadequate, AI can analyze individual behaviours, preferences, and patterns, allowing for the creation of highly personalized experiences. This level of customization is achieved through the analysis of vast amounts of data.

Guidelines for AI projects

Starting AI projects effectively requires first identifying the specific problem to be resolved. By conducting audits of existing operational processes and engaging in stakeholder consultations, companies can effectively identify areas of challenge.

Operational process audit

This type of audit generates a clear description that reviews the intricacies and interdependencies of operational processes. Once mapped, the focus extends to understanding how these processes interact with one another. This stage involves a thorough examination of the organization's current workflows, spanning from initial quotations and bookings, to documentation and invoicing, where inefficiencies can ripple through the entire operational chain.

Stakeholder consultation

Stakeholder consultation for needs analysis forms the base for understanding the problem from multiple perspectives. This phase involves engaging with a broad range of stakeholders, including employees, customers, and vendors, to gather insights into their experiences, expectations and frustrations.

By cross-referencing the challenges identified by different stakeholder groups, paired with in-depth operational process auditing, companies can prioritize AI initiatives based on their potential impact across multiple business dimensions.

Choose the right AI product

Once you know the problem you are trying to solve, it's time to look at solutions. Implementing a new system can involve considerable expense, along with time and resources. You want to ensure not only that the product fits your needs but also that the company that is providing it will be around to provide updates and support.

Some important things to consider:

- Are there products that will meet your specific needs?
  - There are AI products created for specific industries and subjects.
- Do you plan for additional AI products in the future?
  - Some AI options may have additional features you do not need now but may need later. Having one product that solves multiple needs can save time by requiring staff training on only one system.
- Pricing
  - Some products will charge based on how much you use them and if they use an API connection.
- Features
  - Along with ensuring the software will have the features you need, look at everything else included. It can be good to include less robust software if it provides all the features you need now and for future uses.

Data availability

Relevant and accurate data is needed to get the best results from AI. Many products use generic data for various purposes, which may not be your best solution. If your use requires more specific data, ensure that the product you are looking at either already uses it or has the ability for you to add it.

System integration

If your needs include integration with your existing systems, you will need to ensure the product has the ability to integrate. Some products have this ability and may have prebuilt connections to popular software. Alternatively, if the software on both ends has API capabilities, you may be able to use an API connection service.

Plan user training and adoption

Developing a comprehensive training program before the phase of user adoption is essential to ensure a smooth transition and successful integration of AI technologies. Training activities should be tailored to different roles within the organization.

Internal user adoption

By streamlining the progression from training initiatives to adoption, organizations can ensure that their workforce is well-prepared and supported as they transition into an AI-enhanced operational setting. It is essential to recognize that training and adoption are interconnected processes that need to be cohesively integrated to optimize the success of these projects.

Early adoption is the initial phase that focuses on AI tool implementation into daily operations. Pilot initiatives can evaluate the efficacy of the AI tools, test the full functionalities, assess their impact on operations and provide valuable insights into the AI system's practicality.

Transitioning from the early adoption phase, the full adoption stage represents a pivotal expansion of integration across the entire organizational spectrum, with AI tools integrated thoughtfully into existing workflows, complementing and augmenting current processes to enhance efficiency without causing disruption.

Furthermore, dedicated help desks or support teams must be made readily accessible to provide users with the necessary guidance to navigate the AI systems. Technical support is instrumental in boosting the confidence and trust of employees, thereby helping to ensure a smooth transition to AI-enabled operations.

Continuous improvement and process iteration

Continuous improvement and process iteration are fundamental to the success of AI applications.

Data collection and iteration

Adopting a data-centric strategy is fundamental to understanding the effectiveness of AI systems in real-world scenarios. Through this iterative process, logistics organizations can harness the full potential of AI to optimize their operations, enhance customer experiences, and navigate the complexities of the global supply chain environment.

Feedback loops

Establishing feedback loops that involve users, AI systems, and developers is essential for continuous improvement and adaptation. User feedback such as error rates and performance metrics will be relayed back to the development team, which can provide insights indicating areas where the AI projects meet expectations and where further developments in technical refinements are required.

Agile development practices

Implementing agile development practices that allow for incremental improvements and frequent updates to AI systems enables companies to adapt quickly to changes in the industry and incorporate new findings and technologies into their AI solutions.

Cross-functional teams

Create cross-functional teams that include business domain experts, data scientists, AI engineers, and a technical support team to ensure a complete approach to continuous improvement. This collaboration can help in identifying practical challenges and solutions that are aligned with business objectives.

Monitor performance

It is essential to emphasize the significance of continuous performance monitoring. This involves the establishment of a real-time analytics framework that captures and evaluates data related to key performance indicators (KPIs). For example, in an AI project aimed at automating documentation handling, KPIs could include metrics such as the average shipping instruction processing time and the accuracy rate of data entry. Project leaders have the option to implement a variety of strategies that incorporate real-time data monitoring, stakeholder feedback, and regular formal evaluations. The following are some practices that may be considered:

- Set up a dashboard that integrates with the organization’s operating system to display real-time KPI data. Highlight the deviations from expected performance benchmarks.

- Implement AI-driven anomaly detection to identify unusual patterns that could indicate problems.
- Use graphs or other visual tools to highlight trends and issues at a glance.
- Use a centralized platform to collect feedback from various stakeholders.
- Conduct correlation analysis between feedback data and quantitative performance metrics.
- Keep detailed records of changes made to the AI system, including the rationale, implementation details, and the observed impacts.
- Determine a schedule for formal performance evaluations and reports.

Return on investment

When evaluating the return on investment (ROI) for AI implementations within the organization, the calculation must factor in the immediate and long-term costs associated with AI deployment. These expenditures form the baseline against the financial returns and efficiency gains from AI applications.

On the cost side, the financial investment in AI may include:

- Upfront expenses of acquiring and setting up the technology
- Recurring costs related to software updates and system maintenance
- The need for additional hardware or infrastructure upgrades
- Staff training costs and process re-engineering to ensure seamless adoption

On the other hand, the benefits of integrating AI into operations can extend beyond mere cost savings. In practical terms, AI can help optimize cargo loading, automate customer interactions, refine demand forecasting, and streamline warehouse inventory management leading to significant operational efficiencies. This translates to reduced manual tasks, minimized errors, and enhanced throughput, which directly contributes to the business bottom line.

Potential challenges

Not achieving the desired results

Here are five things to consider if your AI project is not delivering to your expectations.

1. Investigate why the project is not performing as expected. Use this checklist to validate the project implementation:
   - ☑ Is there a problem with the data?
   - ☑ Is the underlying algorithm correct?
   - ☑ Is there a design flaw?
   - ☑ Was it implemented according to your plan?
   - ☑ Was there sufficient testing?
   - ☑ Was it deployed according to specification?
   - ☑ Are your expectations realistic?
   - ☑ Were your objectives clear and well communicated?
   - ☑ Did you assign sufficient resources?
   - ☑ Are there technical issues?
   - ☑ Did you encounter unexpected ethical or legal issues?
2. Share the results with key stakeholders such as staff, clients, and partners. Let them know what went wrong, key takeaways, and your plans for moving forward.
3. Assess the significance of missing expectations. Is it having a negative impact on finances or your relationships with stakeholders? Has your credibility or brand image been tarnished? Ask for input from your stakeholders.
4. Decide whether you should continue with your existing plan or make changes to your direction based on what you've learned.
5. Once you make the decision, move forward. You have gained valuable insight. Use it to help get the results you were originally expecting.

Bias and fairness

AI systems can inadvertently perpetuate biases present in their training data, leading to discriminatory outcomes. Legal challenges arise when AI is used in hiring, lending, and criminal justice decisions, as it can disproportionately impact certain groups. Establishing

guidelines and regulations to address AI bias and discrimination is vital to ensure fairness and prevent legal disputes.

Transparency and explainability

Transparency and explainability are crucial for building trust in AI systems. As AI algorithms become more complex and opaque, it becomes challenging to understand how they arrive at their decisions. This lack of transparency can lead to concerns about bias, discrimination, and unfair outcomes.

To address these concerns, researchers and policymakers are focusing on developing explainable AI (XAI) techniques. XAI aims to provide insights into how AI systems make decisions, enabling users to understand the underlying logic and factors influencing the outcomes. By enhancing transparency and explainability, XAI can help identify and mitigate biases, ensure accountability, and build trust in AI systems.

Data privacy

AI systems rely on vast amounts of data to learn and make informed decisions. However, the collection, storage, and processing of personal data raises significant privacy concerns.

Global regulations, such as the European Union's General Data Protection Regulation (GDPR), have been enacted to protect individuals' privacy rights. These regulations require organizations to obtain explicit consent for data collection, provide transparency in data processing, and implement measures to ensure data security. Organizations must also adopt privacy-by-design principles, which involve incorporating privacy considerations into the design and development of AI systems from the outset.

Security

As AI systems become more sophisticated and autonomous, ensuring their security becomes paramount. The potential risks associated with AI security breaches include data breaches, adversarial attacks, and the manipulation of AI systems to produce biased or malicious outcomes.

To mitigate these risks, organizations must implement robust security measures throughout the AI lifecycle. This includes secure data storage and transmission,

authentication and access controls, encryption, and regular vulnerability assessments. Additionally, organizations should adopt ethical guidelines and best practices to ensure the responsible development and deployment of AI systems.

Ethical considerations

Questions about the use of AI in surveillance, hiring decisions, loan approvals, and deepfake technologies highlight the need for ethical considerations to be enshrined in the legal framework. Addressing these concerns will remain one of the defining challenges of AI technologies, and the answers will impact public trust in AI systems.

The issue of accountability in AI systems also presents an ethical challenge. When AI makes a decision, it can be difficult to determine who is responsible for the outcome, especially if it results in harm or loss. This "black box" problem, where the decision-making process of AI is not transparent or understandable, complicates the assignment of responsibility. To address this, there's a growing demand for explainable AI.

As AI systems become more capable, there's a risk that they will replace human workers in various industries, hence a need for policies that support workforce transition, including retraining and education programs.

Another area of ethical concern is that, as AI systems become more integrated into daily life, they can influence human behaviour, decision-making, and interactions. Societal values, cultural norms, and human dignity should be considered.

Another emerging concern is around AI and the energy consumption required for training and running large AI models. It's important to consider the environmental footprint of AI systems and strive for more energy-efficient technologies.

Finally, the pace of AI development poses an ethical challenge in itself. The rapid advancement of AI technologies can outstrip the ability of societies and regulations to keep up, leading to gaps in oversight and protection. A collaborative approach among policymakers, technologists, and ethicists is needed to ensure that AI development is guided by ethical principles and societal needs.

Regulatory compliance

A regulatory framework for the use of AI has been so far lacking worldwide. Only as of late 2023 have regulations been partially implemented in the EU, China and the USA.

The Canadian federal government has proposed legislation, Bill C-27, the Digital Charter Implementation Act, 2022, to modernize Canada's privacy laws with the introduction of the Artificial Intelligence and Data Act. If passed, the new law would regulate the use of AI systems for both the companies that use these systems, and providers of such systems, as well as data processors who deploy AI systems in the course of data processing.

In September 2023, the Minister of Innovation, Science and Industry announced the Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems. This code temporarily provides Canadian companies with common standards and enables them to demonstrate, voluntarily, that they are developing and using generative AI systems responsibly until formal regulation is in effect. Based on feedback received during a consultation process on the development of a Canadian code of practice for generative AI systems, this code aims to help strengthen Canadians' confidence in these systems.

Under the Artificial Intelligence and Data Act, businesses will be held responsible for the AI activities under their control. They will be required to implement new governance mechanisms and policies that will consider and address the risks of their AI system and give users enough information to make informed decisions.

Legal and liability issues

Legal and liability issues are among the most significant challenges when implementing AI technologies. One of the primary legal concerns is the determination of liability in cases where AI systems cause harm or make erroneous decisions. Traditional legal frameworks are based on human decision-making, and the autonomous nature of AI complicates the assignment of responsibility. Whether it's a malfunctioning autonomous vehicle or a flawed decision from an AI-powered medical diagnostic tool, establishing who is at fault—the developer, the user, or the AI itself—is a complex issue that requires careful consideration and possibly new legal frameworks. Insurance coverage should also be reviewed to determine how it applies when losses are caused by the decisions made by AI systems.

Intellectual property (IP) rights also present a legal challenge in the realm of AI. The question of who owns the output of an AI system, such as a piece of art, a written work, or a piece of software code, is still largely unresolved. Current IP laws are not fully equipped to address the nuances of creations made by AI, leading to potential disputes and uncertainties about rights and ownership.

Common use cases for AI in logistics

Email automation

Email automation streamlines communication by automating routine email interactions, freeing up valuable time for customer service representatives to focus on more-complex queries. AI-driven email automation systems are designed to handle a wide range of customer interactions, from order confirmations and shipping updates to feedback requests and support inquiries. By utilizing natural language processing (NLP) capabilities, these systems can understand and respond to customer emails in a way that is both accurate and contextually relevant.

The effectiveness of AI in email automation lies in its ability to analyze past interactions to improve their understanding of customer inquiries and enhance response accuracy.

Another significant advantage is scalability. As the volume of customer interactions grows, AI systems can easily handle the increased load without compromising response quality or speed.

In addition to handling routine queries, AI-driven email automation can aid in gathering valuable customer insights. By analyzing the content and patterns of customer interactions, AI can identify common issues, preferences, and trends. This information can be used to improve service offerings, tailor communications, and inform business strategies.

Lastly, the integration of AI in email automation contributes to a more cohesive and consistent brand experience.

Chatbots

AI-driven chatbots offer instantaneous, round-the-clock interaction with customers,

providing answers to queries, facilitating transactions, and offering support. Unlike traditional customer service systems, AI chatbots are capable of handling a vast number of queries simultaneously, ensuring that customer inquiries are addressed promptly, thereby enhancing the overall customer experience.

The sophistication of AI chatbots lies in their ability to understand and process NLP. This technology enables chatbots to comprehend customer inquiries in their natural, conversational form, and respond in a manner that is both accurate and contextually appropriate.

AI chatbots in logistics can assist customers with tracking shipments, providing estimated delivery times, and updating order information. Moreover, AI chatbots can handle routine tasks such as scheduling pickups, processing orders, and providing quotes, thereby streamlining operations and reducing the workload on human customer service agents.

They can also personalize interactions by accessing and analyzing customer data, providing tailored recommendations and solutions, and enhancing the customer's experience with the brand. For example, if a customer frequently ships items to a particular destination, the chatbot can remember this preference and expedite future transactions. This personalized approach not only saves time for the customer but also builds a sense of loyalty and satisfaction.

Document management

One of the challenges facing the logistics industry is transcribing data from forms into the various systems that we use. Examples of these forms include:

- AP/AR invoices
- Commercial invoices
- Customs declarations
- House/master air waybills
- House/master ocean bills of lading
- Shipper's letters of instruction

Employees manually re-keying data from these types of documents is a waste of resources. Not only is it unfulfilling for the employees doing the work, but it also adds little

value, and can be error-prone. The goal should be to free up valuable employees from this type of work so they can focus on more value-added, customer-focused activities.

A solution to this challenge is optical character recognition (OCR). Figure 1 depicts a simplified OCR flow.

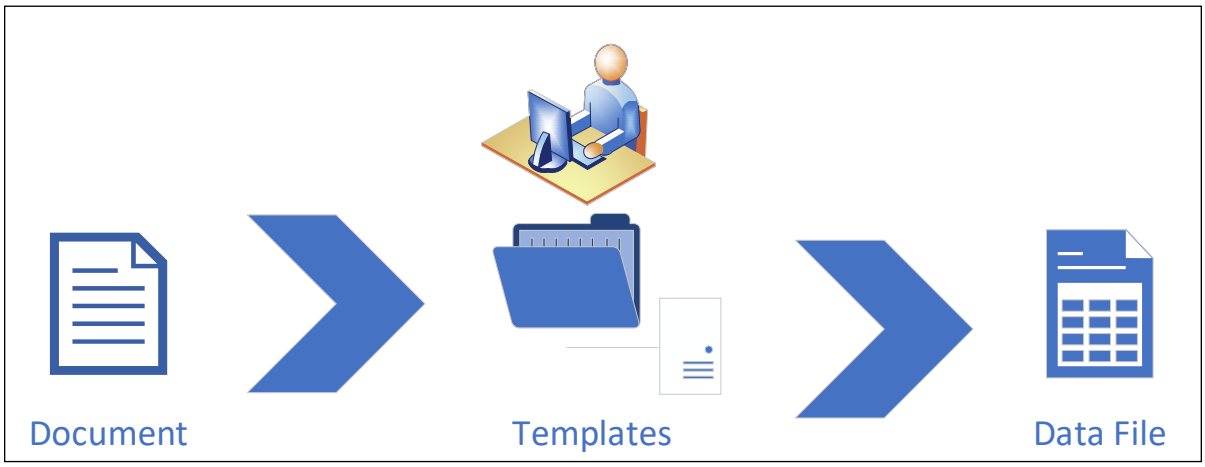

Figure 1 – Simplified OCR flow

Documents

Documents can be received via e-mail, from FTP sites, or direct uploads and can be in various formats such as PDF, Word, Excel, and CSV files. They may also be in different languages. In some cases, they may even be handwritten.

Machine learning is used to decipher the information in the forms. For example, determining whether a character is the letter 'O' or the number '0'.

Templates

Templates are used the tell the OCR where to find information within the form. Human intervention is required to select the appropriate template based on the form being processed. If there is no existing template, the operator must create a new template. The operator may have to adjust the template for a specific instance of a form.

Machine learning is used to accurately process forms as well as suggest data elements to speed up the creation of new templates.

The creation of templates can be a make or break component of a successful OCR implementation since the time to create and manage the template library can be labour intensive. When considering the implementation of OCR, make sure that this process is simple and speedy.

Data file

This is the output of the OCR. It is usually a CSV or TXT file that can be read by the system that will be accepting the data.

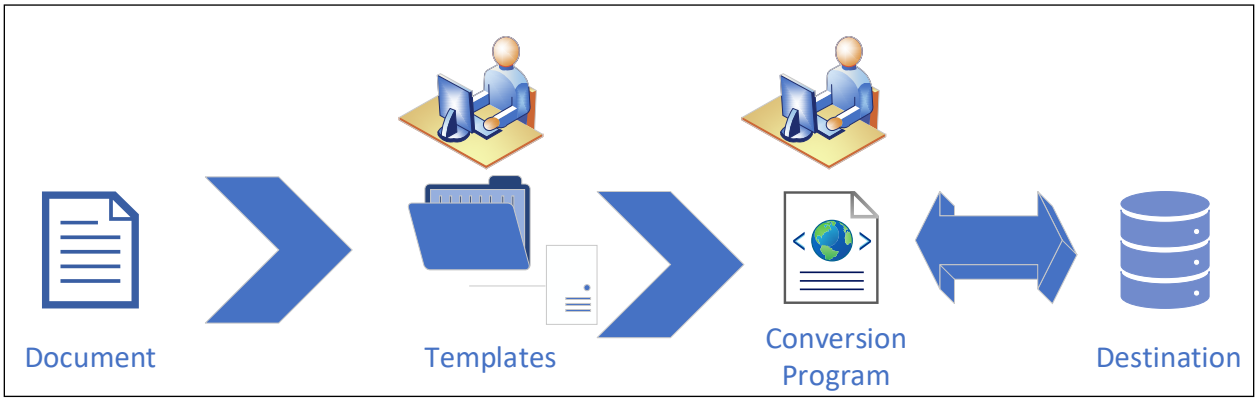

Figure 2 – Advanced OCR

Figure 2 shows an advanced OCR flow where the data file has been replaced by a conversion program and a connection to the destination. This connection could be direct access to the destination database or an API provided by the vendor of the destination system.

Conversion program

The conversion program will take the data extracted from the document and write it to the destination system, eliminating the step of processing an OCR data file. It also allows for data validation. For example, it can check to ensure that part numbers on a commercial invoice exist in the destination system and, if not, create them.

Machine learning is used to enhance this process and suggest part numbers based on descriptions. This reduces the creation of duplicates.

Destination

This is the receiving system. It is essential to understand the implications of an OCR implementation for the destination. For example, there is danger that the OCR will create duplicate parts in the destination system, which can grow into a database maintenance challenge.

The future of document management

AI will advance OCR further, thereby reducing the labour required to extract data from forms. Here is what to look for:

- Automatic selection of the correct template for the document being processed
- Automatic creation of templates
- Automatic determination of tariff codes

In all cases, it is essential to understand the accuracy rates. Vendors must provide documented evidence of their success rates in applications similar to the one you are considering implementing.

Shipment visibility

AI-supported predicted estimated time of arrival (ETA) is a significant application of AI in the logistics industry, enhancing the efficiency and accuracy of shipment visibility. In a sector where timely delivery is crucial, AI's ability to predict ETAs with high precision is transformative. Traditional methods of estimating arrival times often rely on static variables and historical data, which can be unreliable due to the dynamic nature of transportation and logistics. AI, on the other hand, incorporates real-time data analysis, offering a more-accurate and up-to-date estimation.

The power of AI in predicting ETAs lies in its ability to process and analyze vast amounts of data from various sources simultaneously. This includes traffic conditions, weather reports, vehicle speed, and shipping routes. By considering these dynamic factors, AI can make real-time adjustments to the estimated arrival times, providing more-accurate predictions. This level of accuracy is crucial for logistics companies in managing their operations, planning for contingencies, and communicating with customers.

AI-supported predicted ETAs also enhance supply chain efficiency. With more-accurate arrival times, companies can optimize their inventory management, reducing the need for excess stock and minimizing the risks of stockouts. This optimization leads to cost savings and improved operational efficiency, making AI an invaluable tool in supply chain management.

Furthermore, AI-driven ETA predictions improve customer satisfaction. With AI, logistics companies can offer their customers real-time updates on their shipments, increasing transparency and trust. This capability is especially important in a competitive market, where customer experience can be a key differentiator.

Lastly, the use of AI in predicting ETAs paves the way for more advanced logistics operations. The data gathered and analyzed by AI systems can be used to identify patterns and inefficiencies in the shipping process, enabling continuous improvement. For instance, if AI identifies recurrent delays on certain routes, logistics companies can investigate and address the underlying issues.

Information extraction from unstructured data sources

The ability of AI to extract information from unstructured data sources has opened new avenues in enhancing shipment visibility within the logistics sector. Unstructured data, which includes emails, documents, social media posts, and images, contains a wealth of information that is often underutilized due to the complexity of processing and analyzing it. AI, with its advanced algorithms and machine learning capabilities, is adept at sifting through this unstructured data, extracting relevant information, and converting it into actionable insights.

One of the primary strengths of AI in this context is its ability to process natural language. Through natural language processing, AI can understand and interpret the text in emails, documents, and social media, extracting critical information such as shipment dates, locations, and status updates.

Additionally, AI can analyze images and videos, a feature increasingly relevant in logistics for monitoring cargo conditions and ensuring security. For instance, AI can scan images and videos from surveillance cameras to identify damaged goods or security breaches in warehouses and cargo holds. This real-time analysis and reporting significantly enhance the ability to respond quickly to potential issues, ensuring the integrity and security of the shipments.

Conclusion

In conclusion, this paper has shed light on the myriad benefits that AI offers across various sectors. From revolutionizing businesses with predictive analytics to enhancing productivity in industries through automation, AI's potential is boundless. Moreover, its ability to analyze vast amounts of data quickly and accurately enables informed decision-making and drives innovation.

Furthermore, AI is not just a tool for efficiency but also for inclusivity, as it can help bridge gaps in accessibility and provide solutions to societal challenges. However, it's crucial to navigate the ethical considerations surrounding AI deployment, ensuring transparency, accountability, and fairness.

As AI continues to evolve, collaboration between researchers, policymakers, and industry leaders will be vital to harness its full potential while mitigating potential risks. Embracing AI responsibly promises to usher in a future marked by unprecedented advancements, improved quality of life, and equitable opportunities for all.

Appendix A: AI-powered sales

There are AI applications that have been developed specifically to assist in a training environment. AI-powered sales interactions are like a flight simulator for the sales team. Just as you wouldn't let a pilot fly a passenger jet without lots of practice, why would you put a salesperson in front of your customers without lots of practice?

With sales-focused AI, sales staff interact with avatars who have been developed to support proper sales techniques, such as probing for needs, understanding the budget, and so on.

Their characters range from enthusiastic to very reluctant to buy. The sales reps will encounter unpredictable reactions, which keep them engaged and help them prepare for the real world. Simulations can be quickly configured to reflect the industry and products/services being sold by entering product/service information, meeting information, and by picking an avatar and their personality type.

In addition to simulating genuine customer conversations, the AI provides actionable feedback on the student's performance. Key metrics that are specific to the sales process are measured and tracked including:

- What went well
- Areas to improve
- Volume
- Pace
- Use of filler words
- Pain point discovery
- Confirmation of decision makers
- Solution requirements
- Budgets

Depending on the avatar/personality chosen, closing the sale can be surprisingly difficult. Salespeople learn how to be at the top of their game to earn more business with the opportunity for virtually unlimited practice.

Appendix B: AI-assisted software development

The integration of large language models (LLMs) in software development, exemplified by tools like GitHub Copilot, represents a groundbreaking application of AI in the logistics sector. These AI-driven coding assistants are transforming the software-development process, making it more efficient and accessible. By understanding and generating code, LLMs like GitHub Copilot assist developers in writing, reviewing, and debugging code more quickly and accurately. This advancement is particularly beneficial in logistics, where custom software solutions are often needed to manage complex supply chains and optimize operations.

LLMs in software development work by analyzing existing codebases and providing suggestions or completions based on the context of the code being written. They are trained on vast libraries of code, enabling them to understand a wide range of programming languages and coding styles. This capability allows them to offer relevant and syntactically correct code suggestions, reducing the time and effort required for coding tasks. For logistics companies, this means faster development of custom software solutions, whether it's for tracking shipments, managing inventory, or analyzing operational data. The speed and efficiency brought by LLMs can significantly accelerate the deployment of new software features, enhancing the agility of logistics operations.

Beyond just code generation, the use of LLMs in software development fosters a more collaborative and learning-oriented coding environment. These AI tools can serve as educational resources for developers, offering insights into best practices and new approaches to problem-solving. They can also help identify potential bugs or inefficiencies in the code, contributing to higher-quality software. For logistics companies, where reliable and efficient software is crucial for smooth operations, the quality improvements and learning opportunities provided by LLMs are invaluable. These AI assistants not only optimize the software development process but also contribute to the continuous improvement and innovation in logistics technology solutions.

Definitions

Artificial intelligence (AI) – refers to the simulation of human intelligence processes by machines, especially computer systems. These processes include learning (the acquisition of information and rules for using the information), reasoning (using rules to reach approximate or definite conclusions), and self-correction. AI systems can perform tasks such as understanding natural language, recognizing patterns in data, making decisions, and solving problems, often with a level of efficiency and accuracy surpassing that of humans.

Artificial general intelligence (AGI) – refers to the hypothetical ability of an artificial intelligence system to understand, learn, and apply knowledge across a broad range of tasks and domains, similar to the cognitive capabilities of a human being. AGI aims to exhibit general intelligence comparable to or exceeding that of humans.

CSV file – A CSV (comma-separated values) file is a plain text file format used to store tabular data, typically in a spreadsheet or database format. In a CSV file, each line represents a single row of data, and the values within each row are separated by commas (or other delimiters like semicolons or tabs).

Explainable artificial intelligence (XAI) – refers to the design and development of AI systems that can provide understandable explanations for their decisions and actions. The goal of XAI is to increase transparency and interpretability in AI systems, enabling users to comprehend the reasoning behind AI-generated outcomes.

Intellectual property (IP) – refers to creations of the mind, such as inventions, literary and artistic works, designs, symbols, names, and images used in commerce. Intellectual property is protected by law through patents, copyrights, trademarks, and trade secrets, which give the creators or owners exclusive rights to their creations for a specified period.

Large language models (LLMs) – are sophisticated artificial intelligence systems designed to process and generate human-like text on a vast scale. These models, often built using deep learning techniques such as transformers, are trained on enormous data sets containing text from diverse sources like books, articles, websites, and other written material.

Machine learning (ML) – is a subset of AI that focuses on the development of algorithms and statistical models that enable computers to perform tasks without explicit programming instructions. Instead, ML systems learn from data, identifying patterns and making predictions or decisions based on that data.

Narrow artificial intelligence – also known as weak AI, refers to artificial intelligence systems that are designed and trained for specific tasks or applications. Narrow AI systems are limited in scope and functionality but tailored to perform well-defined tasks within a predefined set of parameters and constraints.

Natural language processing (NLP) – is a branch of AI that focuses on the interaction between computers and human language. NLP enables computers to understand, interpret, and generate human language in a way that is both meaningful and contextually relevant. The primary goal of NLP is to bridge the gap between human communication and computer understanding by enabling machines to process and analyze natural language data.

Optical character recognition (OCR) – is a technology used to convert scanned images of text into machine-readable text data. The primary goal of OCR is to enable computers to interpret and extract text from scanned documents, images, or printed materials, making the text searchable, editable, and analyzable by computer systems.

Return on investment (ROI) – is a financial metric used to evaluate the profitability or efficiency of an investment relative to its cost. It measures the ratio of the net profit or benefit gained from an investment to the initial investment cost.

Contributors

This report was planned and written by the members of CIFFA's Technology Committee.

Committee Chair

Marc Bibeau	President & CEO	OEC Group

Committee Members

John Berry	Senior VP of Information Technology	OEC Group
Paul Glionna	Vice President – Systems Dev & Ops	Universal Logistics Inc.
Karl-Heinz Legler	General Manager	Rutherford Global Logistics
Jillian Li	Tariff & System Manager Canada	ECU WORLDWIDE (CANADA)
Kris Runge	IT Specialist	Complete Shipping Solutions
Drew Simons	Principal Advisor	Roxville Technology Inc.
Gianclaudio Steinberg	VP, N.A. Marine Marketing Leader	Marsh Canada Limited