

COGNITIVE DOCUMENT PROCESSING: *How* DID WE GET HERE AND *What* IS COMING NEXT

By Douglas Heintzman, Hugh Frost, Jet Theuerkauf, and Daren Hanson

2025

Table of Content

Whitepaper

Key Messages	02
Introduction	03
The Evolution of Document Processing	04
Generation 1: OCR	07
Generation 2: Robotic Process Automation Document Processing	10
Generation 3: Intelligent Document Processing	13
Generation 4: LLM/GPT Document Processing	16
Generation 5: Cognitive Document Processing	19
Generation 6: Ecosystem Information Management	23
Generations Summary	27
Conclusion	28

Key Messages

- Physical and digital documents are used everywhere, and they play a key role in connecting many business processes.
- Document processing technology has been advancing through six generations—from optical character recognition to ecosystem information management.
- The latest generation, cognitive document processing, uses human-like cognition to understand context and data element relationships to extract information from documents.
- Understanding the evolution of document processing technologies provides a useful foundation to develop document processing strategies.

Introduction

Documents are fundamental to almost all business and governmental operations. They are used to encapsulate information, convey instructions, and record transactions. Document processing involves extracting their content, understanding their meaning and significance, and transferring their information to business systems.

Documents can be physical and digital—and there are a lot of them. For example, on the physical side, the global trade sector alone creates approximately 4 billion paper documents a day. On the digital side, it is estimated that there are more than 2.5 trillion PDF documents in the world.

In 2025 it is forecasted that we will generate 181 zettabytes of new data. Some 80% to 90% of that will be unstructured. Considering the massive volume of documents and their critical importance to the functioning of business processes, it is no surprise that businesses and governments have invested huge effort to digitize and automate document handling and processing.

The Evolution of Document Processing

Until the late 1980s, documents were processed by people. And people are very good at processing documents. Their proficiency improves with experience and they don't need extensive training with a document type because they can—for the most part—figure things out. People are adaptable and can tolerate significant variability. Unfortunately, they are also an expensive resource. Additionally, document volumes have long since eclipsed a manual workforce's ability to keep up with demand. As computers and the internet emerged, we looked to new technology to digitize documents so they could be managed at scale affordably.

Document processing has undergone many technology-driven improvements over the last thirty years, such as advancements in optical character recognition (OCR), the introduction of machine learning for better classification, and the integration of natural language processing (NLP) to enhance data extraction, each generation building on the advancements of the previous. As a result, document processing has evolved from being a keyboard-oriented digital transcription activity to a fully automated integrated part of complex business processes.

Document processing is now a core component of digital transformation because it enables the efficient handling of critical business information, driving automation, reducing costs, and enhancing decision-making capabilities. Organizations understand that documents—whether paper or electronic and whether structured, semi-structured,

or unstructured—hold critical business information that, when captured, organized, and understood, represents a significant business asset. Advancements in technology have made the return on investment of document processing more compelling than ever.

This paper will examine the evolution of the five most recent generations of document processing technologies: OCR, robotic process automation (RPA) document processing, intelligent document processing, large language model (LLM)/ and generative pretrained transformer (GPT) document processing, and cognitive document processing (CDP). It will also take a brief look forward at how document management might evolve into ecosystem information management (EIM). We will explore the technologies behind each generation, examine how those technologies integrate and complement process automation systems, discuss their strengths and weaknesses, and highlight real-world use cases to illustrate their value.

THE RAPIDLY INCREASING VOLUME AND COMPLEXITY OF UNSTRUCTURED DATA ARE DRIVING DEMAND FOR COGNITIVE DOCUMENT PROCESSING

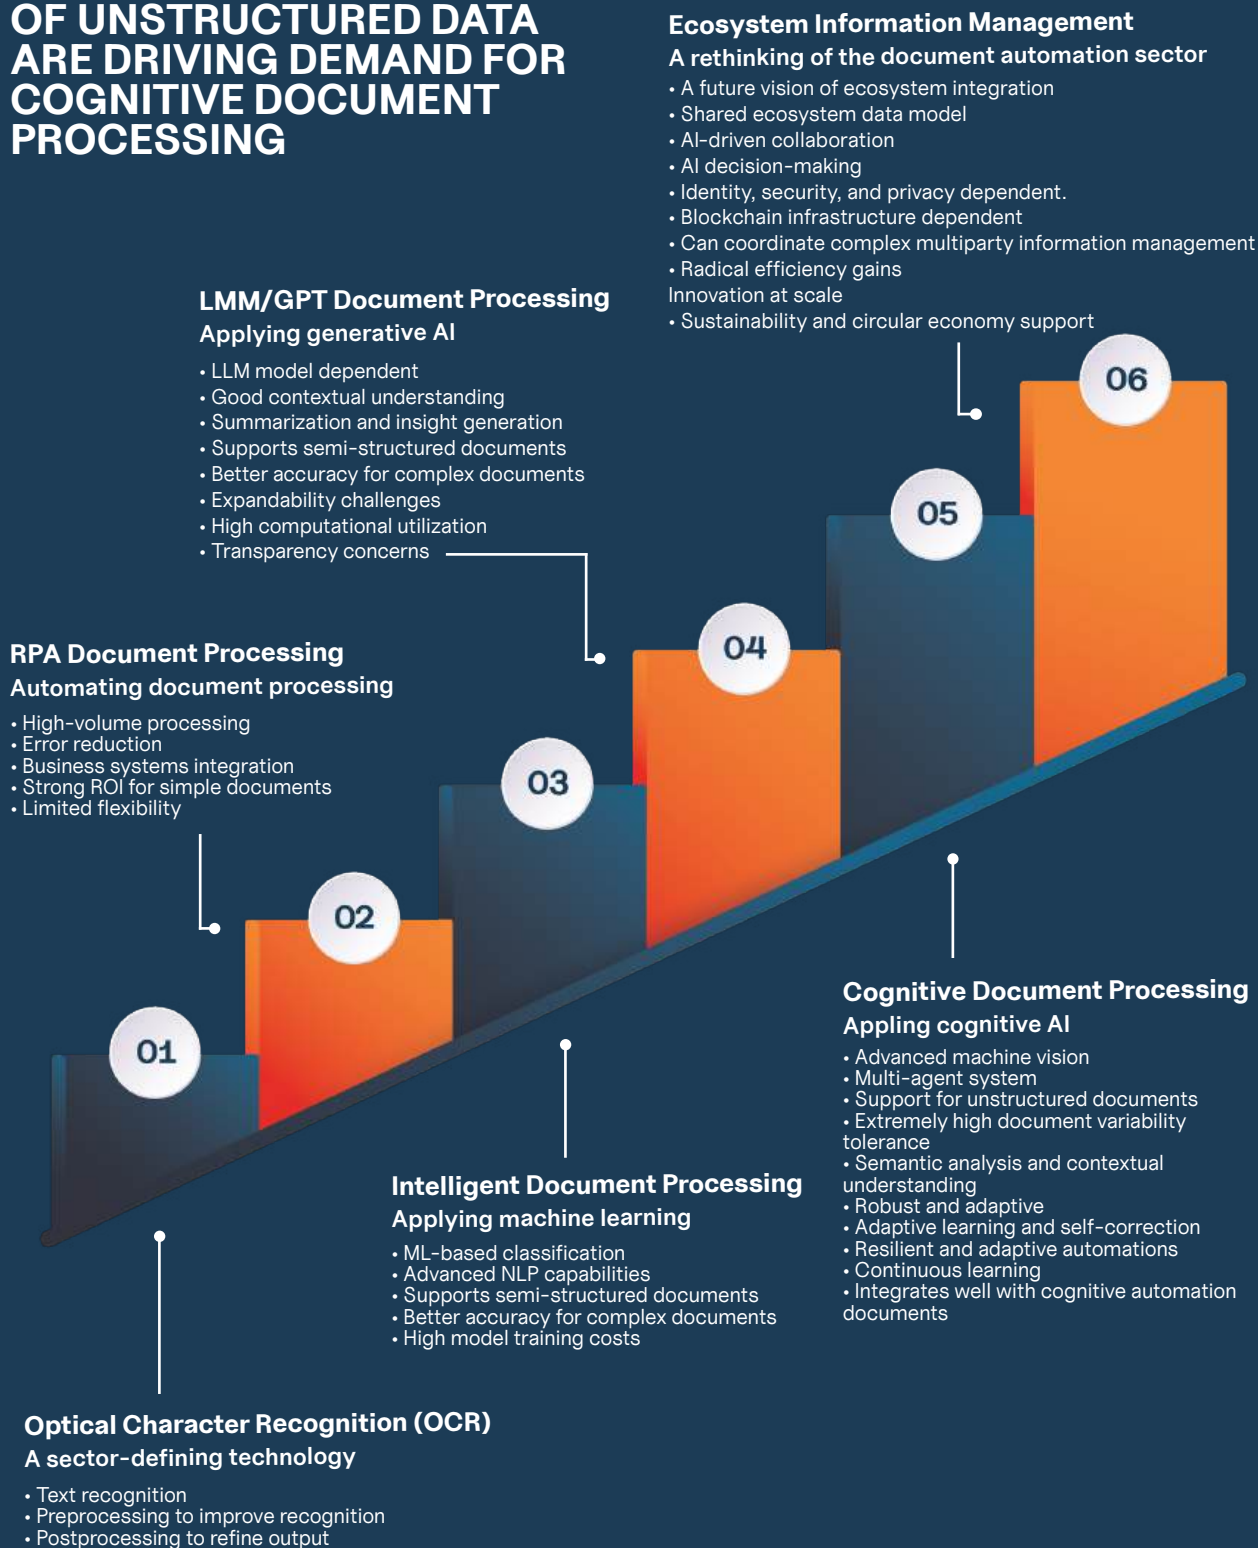


Figure1

OCR



Generation

01

The idea behind OCR is to optically detect characters in an image of a document so that a machine-readable format of its contents can be used in digital systems.

Preprocessing:

Once text is captured, various preprocessing techniques, such as noise removal, de-skewing (straightening tilted text), and binarization (converting to black-and-white), are applied to make the text clearer and more distinct.

Post processing:

After character capture, postprocessing techniques, such as spell-checking, dictionary-based corrections, or NLP, are applied to refine the output and fix errors.

Strengths:

Traditional OCR worked very well on clean printed text in standard fonts. Despite its limitations, it dramatically improved document processing efficiency and accuracy, while improving data searchability and reducing costs.

Limitations:

OCR was much less effective in digitizing unstructured or handwritten text. This meant that it was less effective for the majority of document use cases.

Key Characteristics of OCR

- Text recognition
- Image preprocessing
- Pattern and feature recognition
- Multiple language and font support
- Multiple output formats
- High scalability
- Pre and Postprocessing

Example Use Cases for OCR:



DOCUMENT DIGITIZATION FOR ARCHIVING

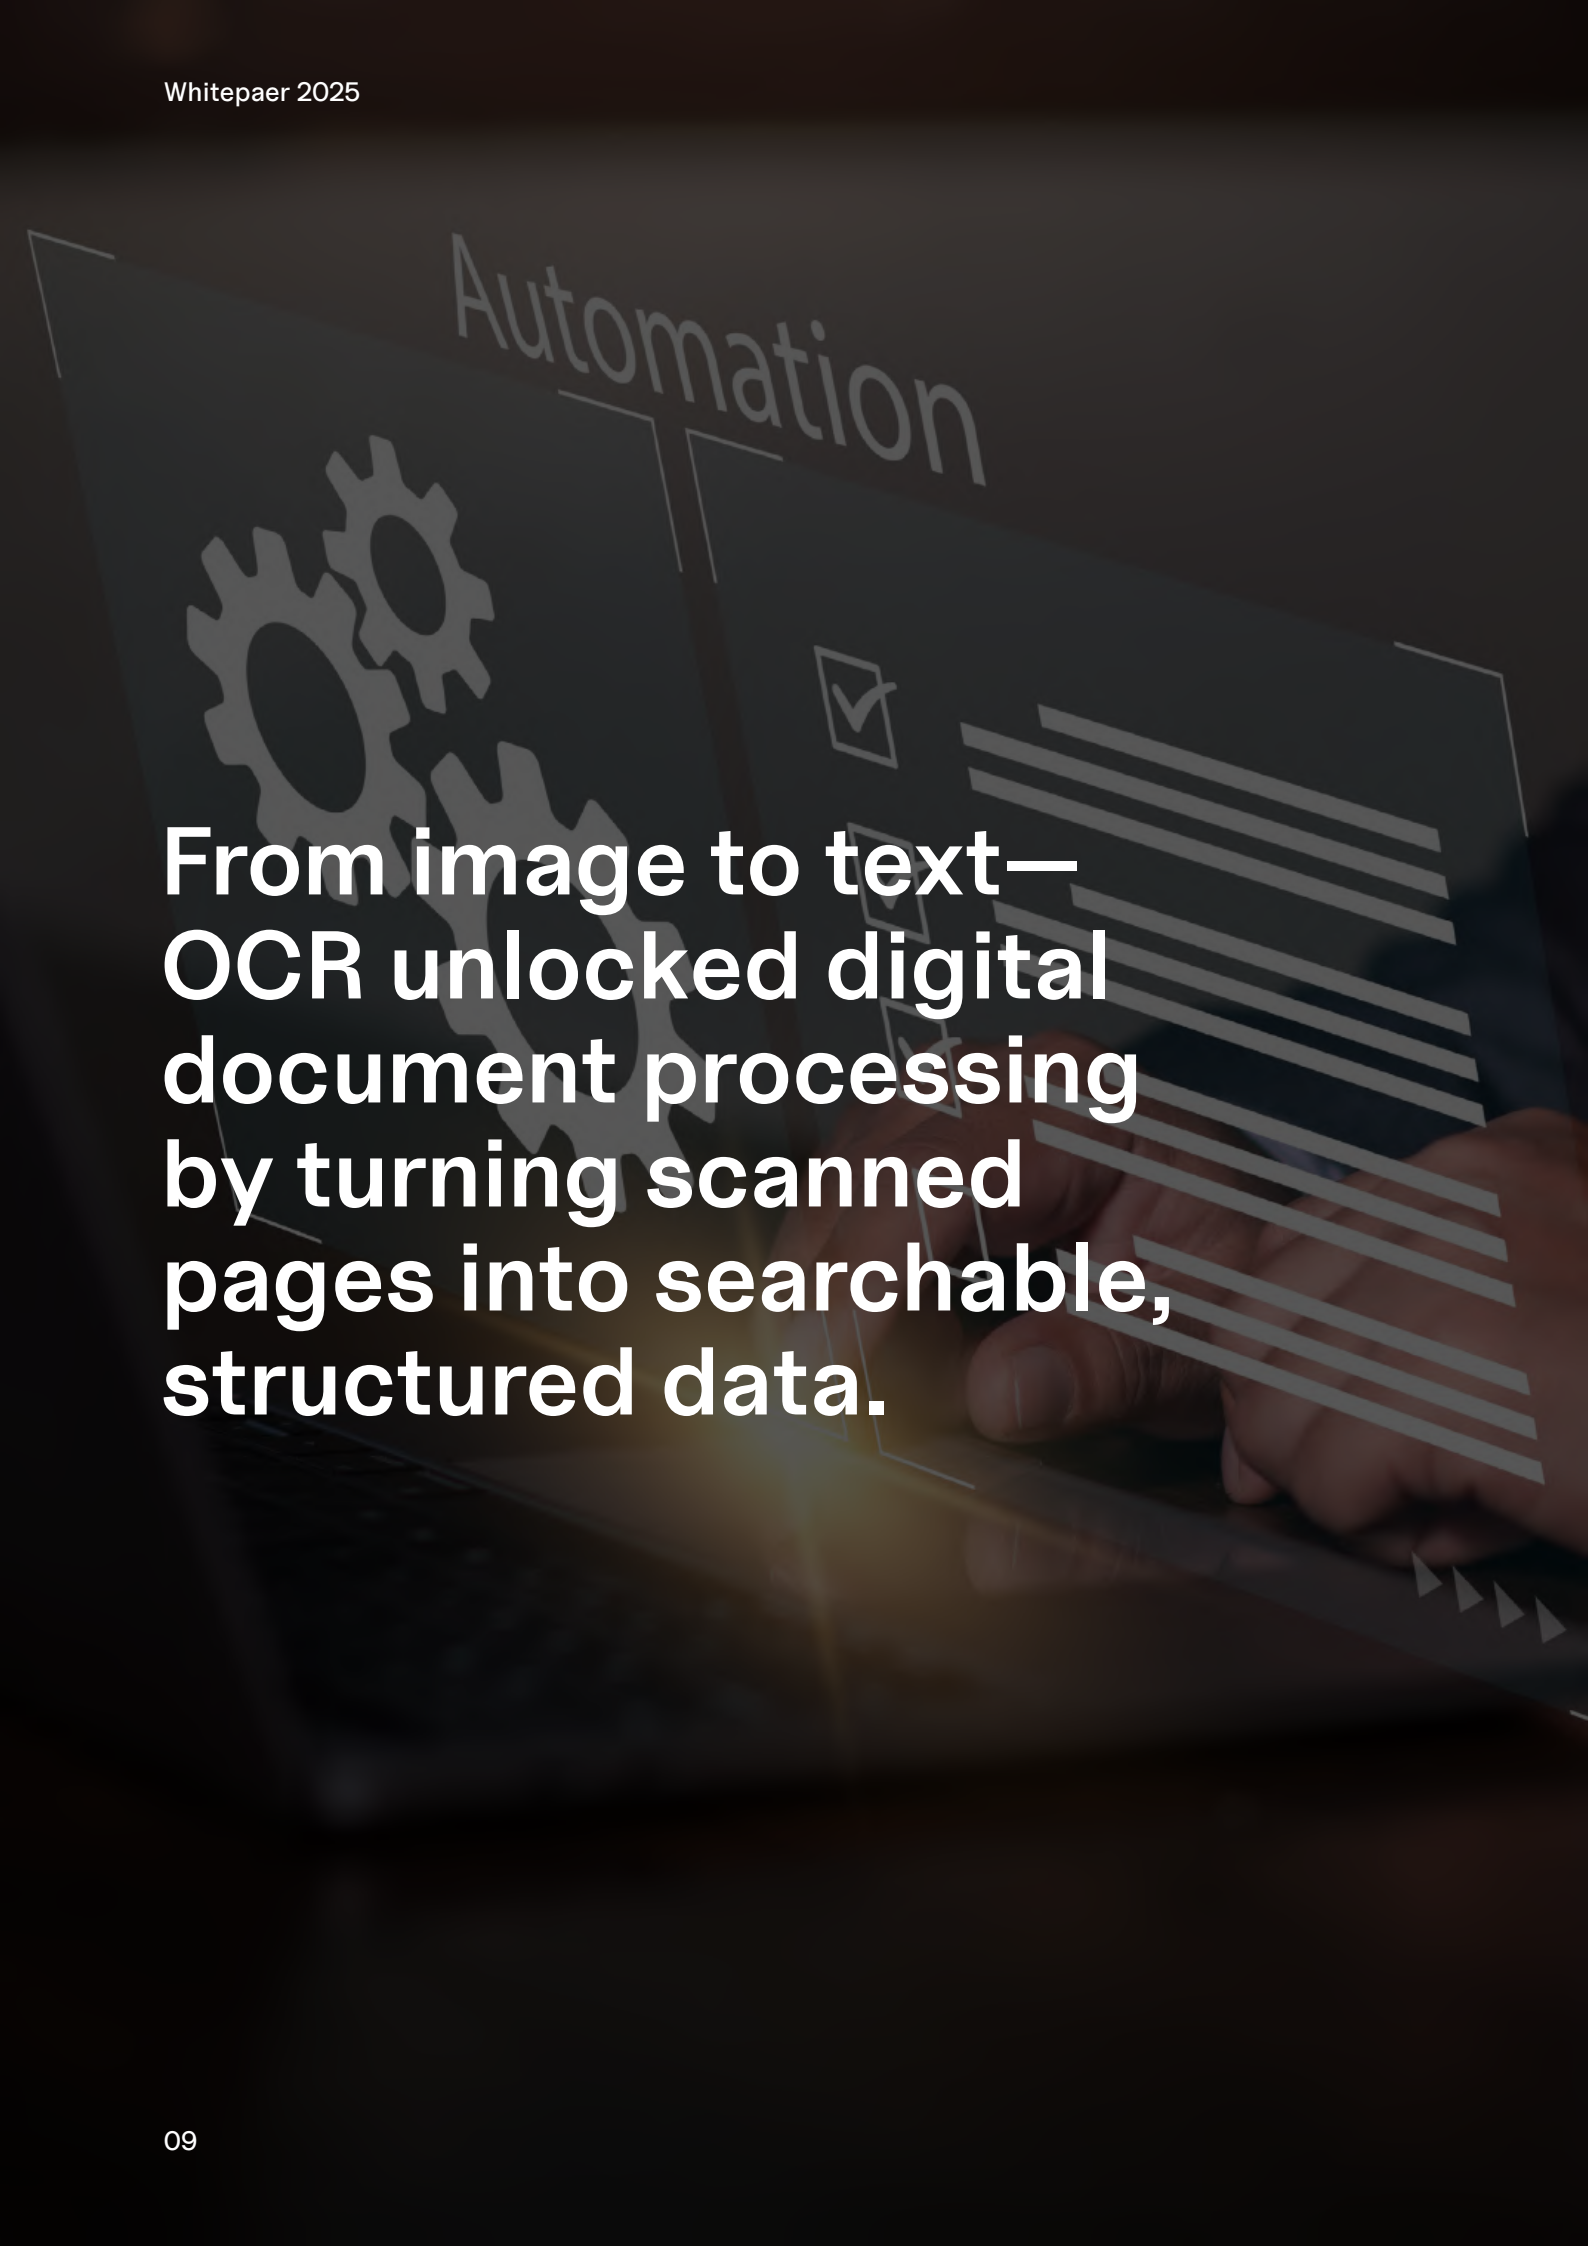
An insurance company uses OCR to convert physical documents, such as manuals, contracts, and records, into digital formats. This allows the organization to reduce its reliance on physical storage and make information more accessible.



AUTOMATED DATA ENTRY FROM FORMS

A shipping company uses OCR to automate the extraction of text from structured forms, such as invoices, receipts, and application forms. This eliminates manual data entry, saving time and reducing errors.

As revolutionary as OCR was, it was limited. As the technology advanced, it would be enhanced by many other technologies to improve its accuracy and usefulness. The next big step forward was its integration into process automation.



**From image to text—
OCR unlocked digital
document processing
by turning scanned
pages into searchable,
structured data.**

Robotic Process Automation Document Processing



Generation

02

The digitization of documents is a process itself, and documents are an integral part of many business processes. As such, OCR and RPA have been intimately linked. RPA provided an efficient way to streamline document processing workflows.

Repetitive task automation:

RPA document processing excels at automating repetitive tasks, such as invoice processing and patient data management. It uses OCR to read, interpret, transform, and classify documents, and it uses code called automations or “bots” to load that data into databases and enterprise resource planning systems to use the extracted information in business processes. RPA document processing relies on rule-based logic and predefined templates, making it an ideal solution for handling large volumes of well-structured documents in predictable formats.

Strengths:

RPA document processing is largely limited to digitizing relatively simple well-structured documents. It has been successful in processing documents at scale with low costs and high quality. It can also reduce human error, and it allows businesses to redeploy workers to higher-value tasks. RPA has a lower initial investment and lower overall cost compared with more complex artificial intelligence (AI)–driven systems, making it a good choice for simple processes.

Limitations:

RPA document processing is well-suited to simple, well-structured, and non-variable documents conforming to a fixed format. However, most organizations have large volumes of complex, unstructured, and variable documents. RPA document

processing also lacks adaptability because it relies on static rules and templates that require constant maintenance and updating if document formats change.

Some businesses have been successful at processing moderately complex documents through extensive coding, customization, and exception handling, but these systems still lack adaptability and are relatively fragile. RPA document processing is further challenged by its lack of ability to interpret the meaning and relationships of data elements and understand document context. This limits its utility to relatively simple data extraction and the business value of the extracted information.

Key Characteristics of RPA Document Processing:

- Extends the value of OCR
- Well-suited for high-volume structured document processing
- Has a strong return on investment for simple automations
- Reduces human transcription errors
- Integrates with many business systems
- Has limited flexibility
- Lacks contextual understanding

Example Use Cases for RPA Document Processing:



INVOICE PROCESSING

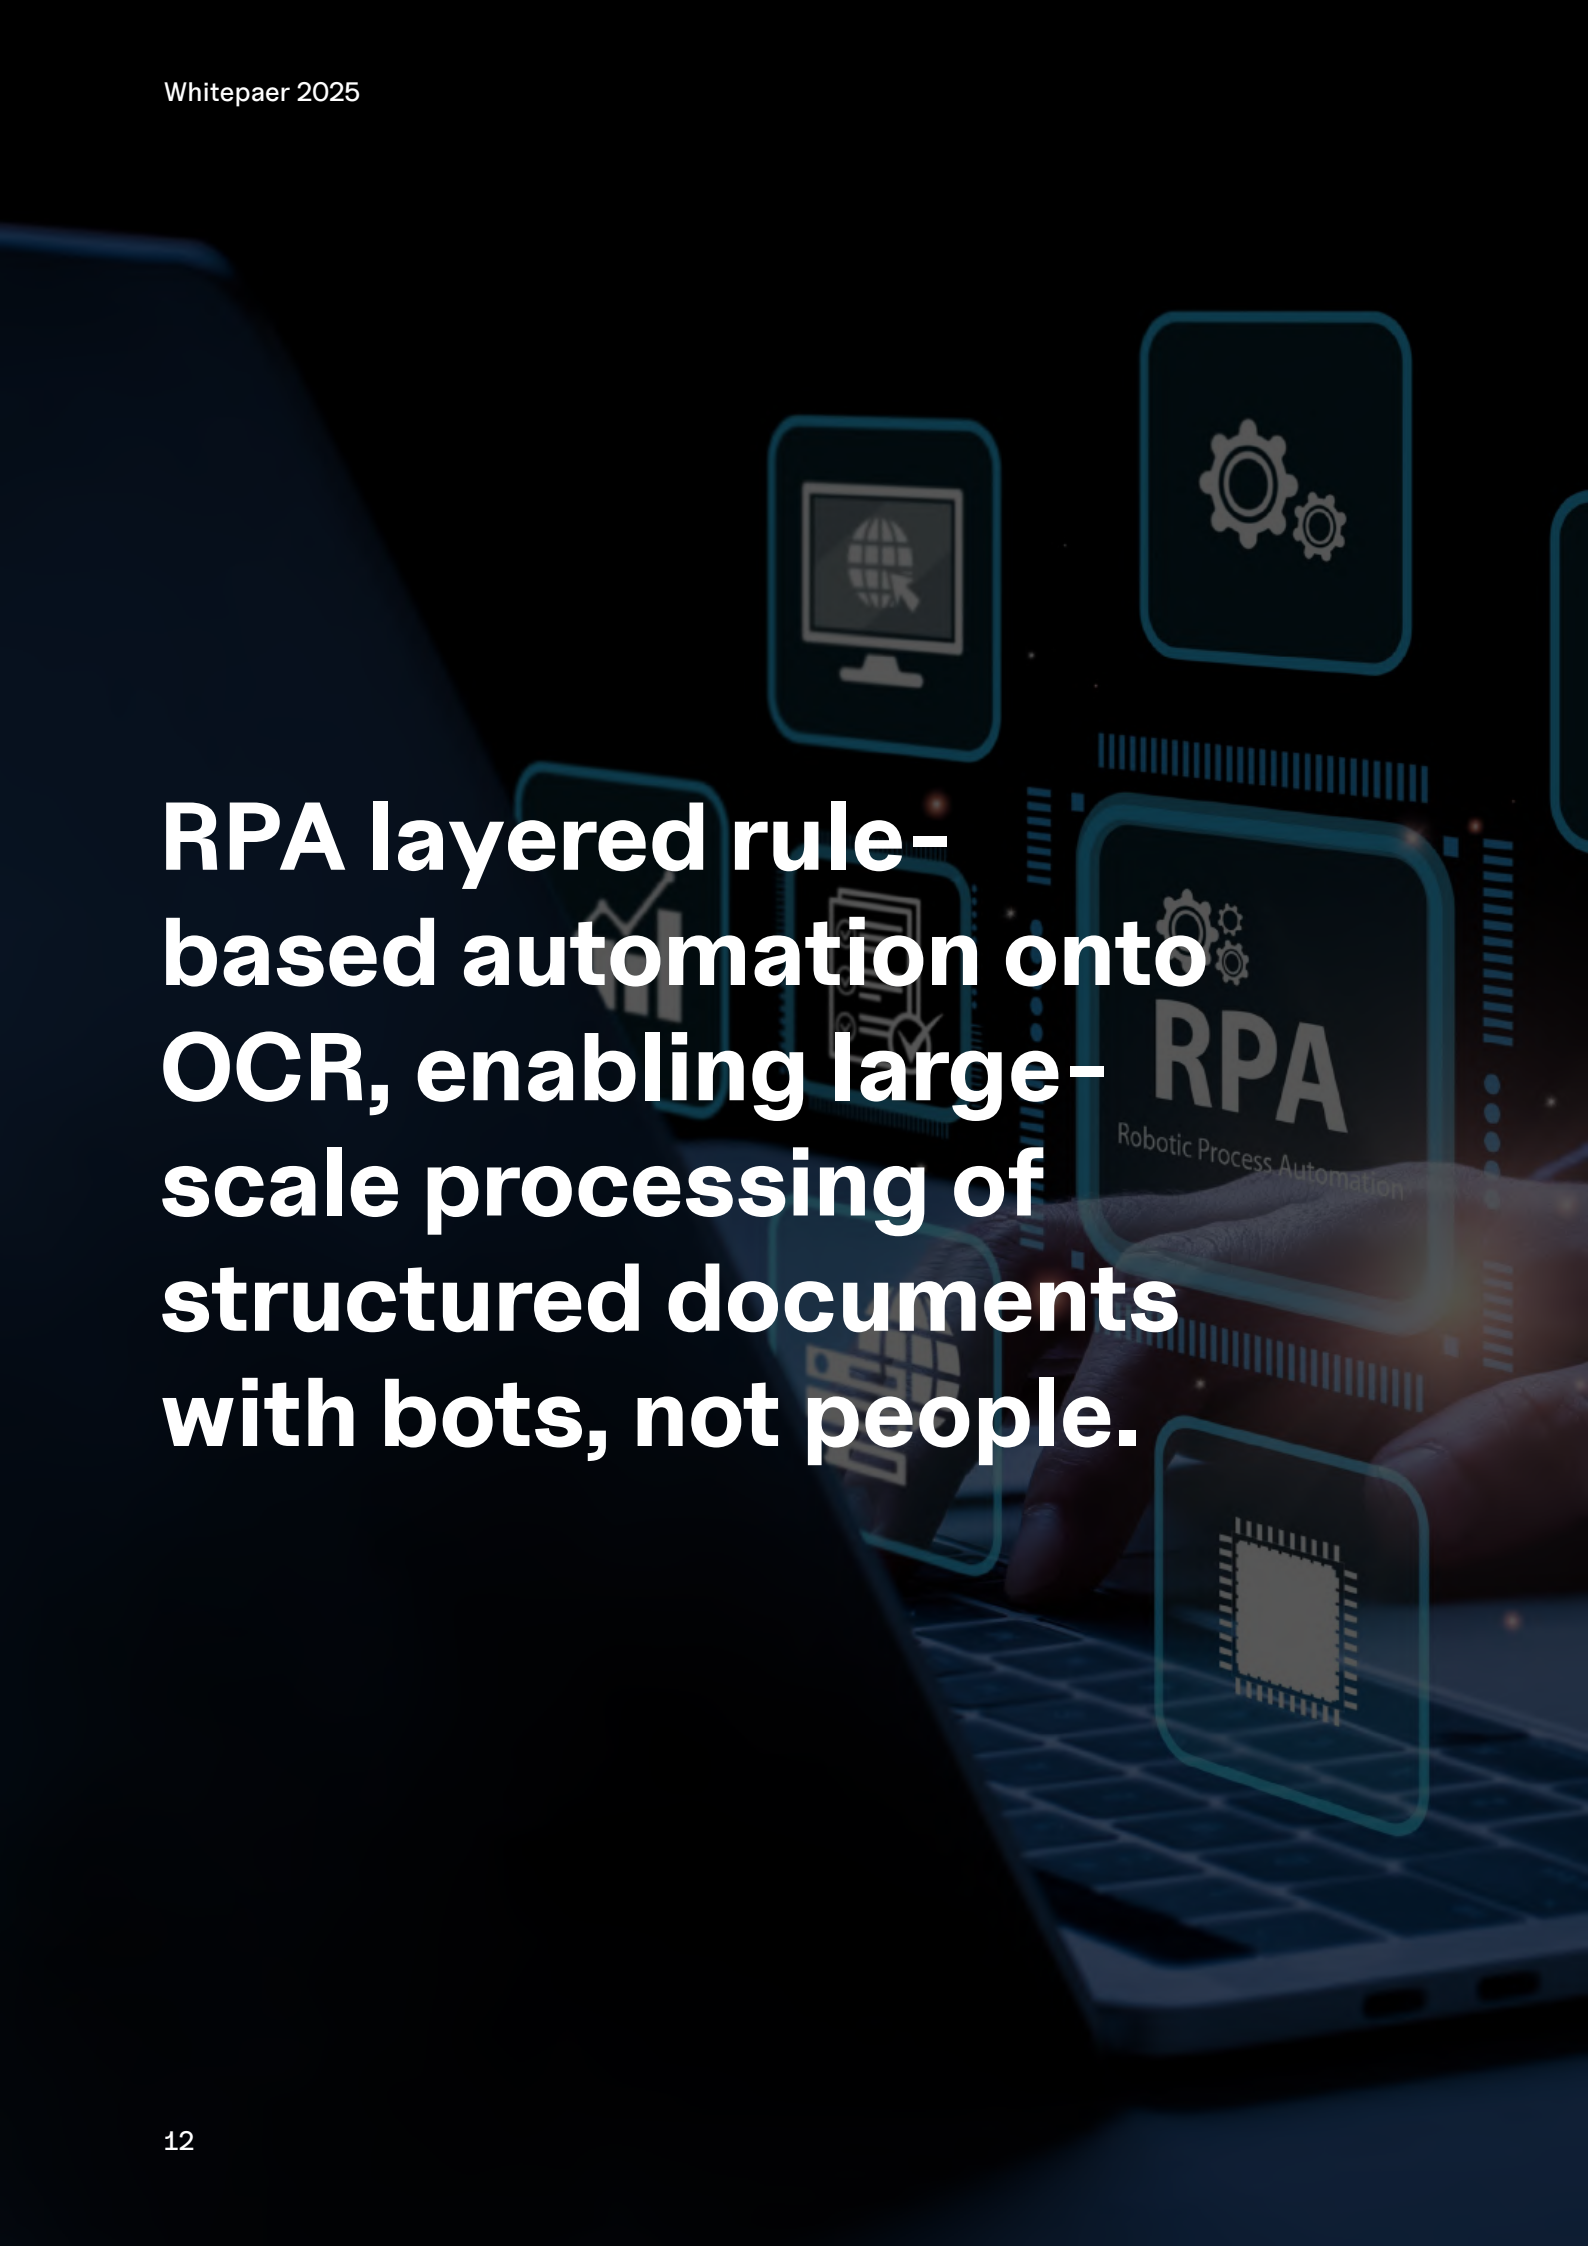
A finance department uses RPA document processing to extract invoice data, including totals, vendor details, and payment terms. An RPA bot transfers the extracted data into the accounting system for approval. This improves efficiency by reducing manual effort and accelerating payment cycles.



FORMS PROCESSING

A health care provider uses RPA document processing to scan and extract patient information from structured admission forms. An RPA bot transfers the information to a patient management system. Automating form processing helps improve patient experience by reducing wait times and minimizing administrative errors.

RPA document processing was the norm until the mid-2010s when AI technologies started to add comprehension capabilities to document processing.



RPA layered rule-based automation onto OCR, enabling large-scale processing of structured documents with bots, not people.

Intelligent Document Processing



Generation

03

IDP extends RPA document processing with a variety of AI technologies, including machine learning, deep learning, neural networks, machine vision, and NLP capabilities. These capabilities revolutionized how digital text, handwriting, and image-based data were processed, and allowed IDP systems to classify documents and extract key information from both structured and semi-structured formats. IDP is significantly more capable of processing different formats of the same document type, including complex invoices, purchase orders, and emails. It also makes document processing much more valuable to organizations.

Classification:

Machine learning's ability to find patterns gives IDP a much-improved classification capability over previous technologies. It is much easier to route information using IDP, which significantly reduces the need for human input. For example, IDP can automatically categorize incoming documents like invoices or contracts, allowing them to be forwarded to the appropriate departments without manual intervention. As a result, document content can be more easily integrated into sophisticated business processes.

Data Extraction and Interpretation:

IDP isn't only better at extracting content from documents, it can use machine learning and NLP to analyze, extract, and interpret key data elements, like names, dates, places, and amounts, with great accuracy.

Limitations:

IDP added many new capabilities, but it still struggles with highly variable, unstructured documents, which impacts its effectiveness in more complex scenarios. This is due in part to its reliance on

machine learning models trained on large volumes of similar documents. To develop these models carries significant upfront costs and complexity because its machine learning models must be trained on large amounts of representative labeled documents. IDP is better than RPA document processing at understanding documents, but its ability to understand the context of data and the relationships between data elements is basic. This limits its usefulness when making informed decisions and ensuring data accuracy for when used business processes.

Applicability:

The addition of machine learning and advanced NLP to document processing expanded the variety of documents that could be captured successfully to about 70%. The remainder of documents still required human intervention.

Key Characteristics of IDP:

- Machine learning–driven classification and extraction
- Higher accuracy OCR
- Advanced NLP capabilities
- Semi-structured document capability
- Human-in-the-loop validation is required for fewer documents
- Dependence on templates
- Expensive model development and maintenance
- Limited scalability for complexity

Example Use Cases for IDP:



INSURANCE CLAIM PROCESSING


An insurance company use IDP to process claim forms, classifying documents, extracting key data such as policy numbers and claim amounts, and passing that information to a business process. This improves response times and reduces costs.



EMAIL PROCESSING

A customer service department uses IDP to classify incoming emails based on content and automatically routes them to the correct team (e.g., billing inquiries, support requests). This reduces response times and improves customer satisfaction by ensuring emails are directed appropriately and efficiently.

IDP's use of machine and deep learning technologies allows it to address a much greater variety of documents. It's ability to understand the nature of different data types significantly improves classification and makes information much more useful to businesses.



IDP brought machine learning into the mix, enabling smarter classification and data extraction from semi-structured content.

LLM/GPT Document Processing



Generation

04

Document processing took another leap forward with the advent of transformer-based AI. The new LLMs powered by GPTs are far better at understanding conversation. As a result, it has brought better contextual understanding to document processing, making it easier to extract insights from unstructured text in documents.

Low code:


Some tools in this emerging era are using GPT's generation capability to build automation code without coding. This approach has some significant shortcomings for complex workflows. However, the user-friendliness of no-code and low-code tools lowers barriers to adoption.

Understanding:

LLM/GPT document processing can read documents, and it also has some ability to understand them. For example, GPT-powered systems can analyze complex contracts, automatically identify key clauses, and provide summaries in seconds—a task that typically requires significant manual effort. The ability of LLMs to understand context across multiple languages and domains makes them versatile tools for a wide range of industries.

Limitations:

The context GPTs can intuit and the insights that they can come up with can be very valuable in a business workflow. However, the way GPTs come to their conclusions can be problematic for some use cases. Unlike traditional models that can trace outputs to specific rules or weights, GPT



models operate as black boxes, making it hard to verify their reasoning. This is particularly concerning in sensitive applications, such as legal or health care document processing, where errors can have severe consequences.

Costs:

LLMs usually require fine-tuning of specific tasks or for different industries. This can be very expensive and time consuming because it requires domain expertise and a significant amount of labeled data. LLMs also require substantial computational resources, which can make large-scale deployments costly. As a result, LLM/GPT document processing is not a one-size-fits-all solution.

Key Characteristics of LLM/GPT Document Processing:

- Contextual understanding
- Multilingual capability
- Adaptability across domains
- Summarization and insight generation
- Excels at processing unstructured or semi-structured documents
- Challenges with explainability
- High computational resource demand

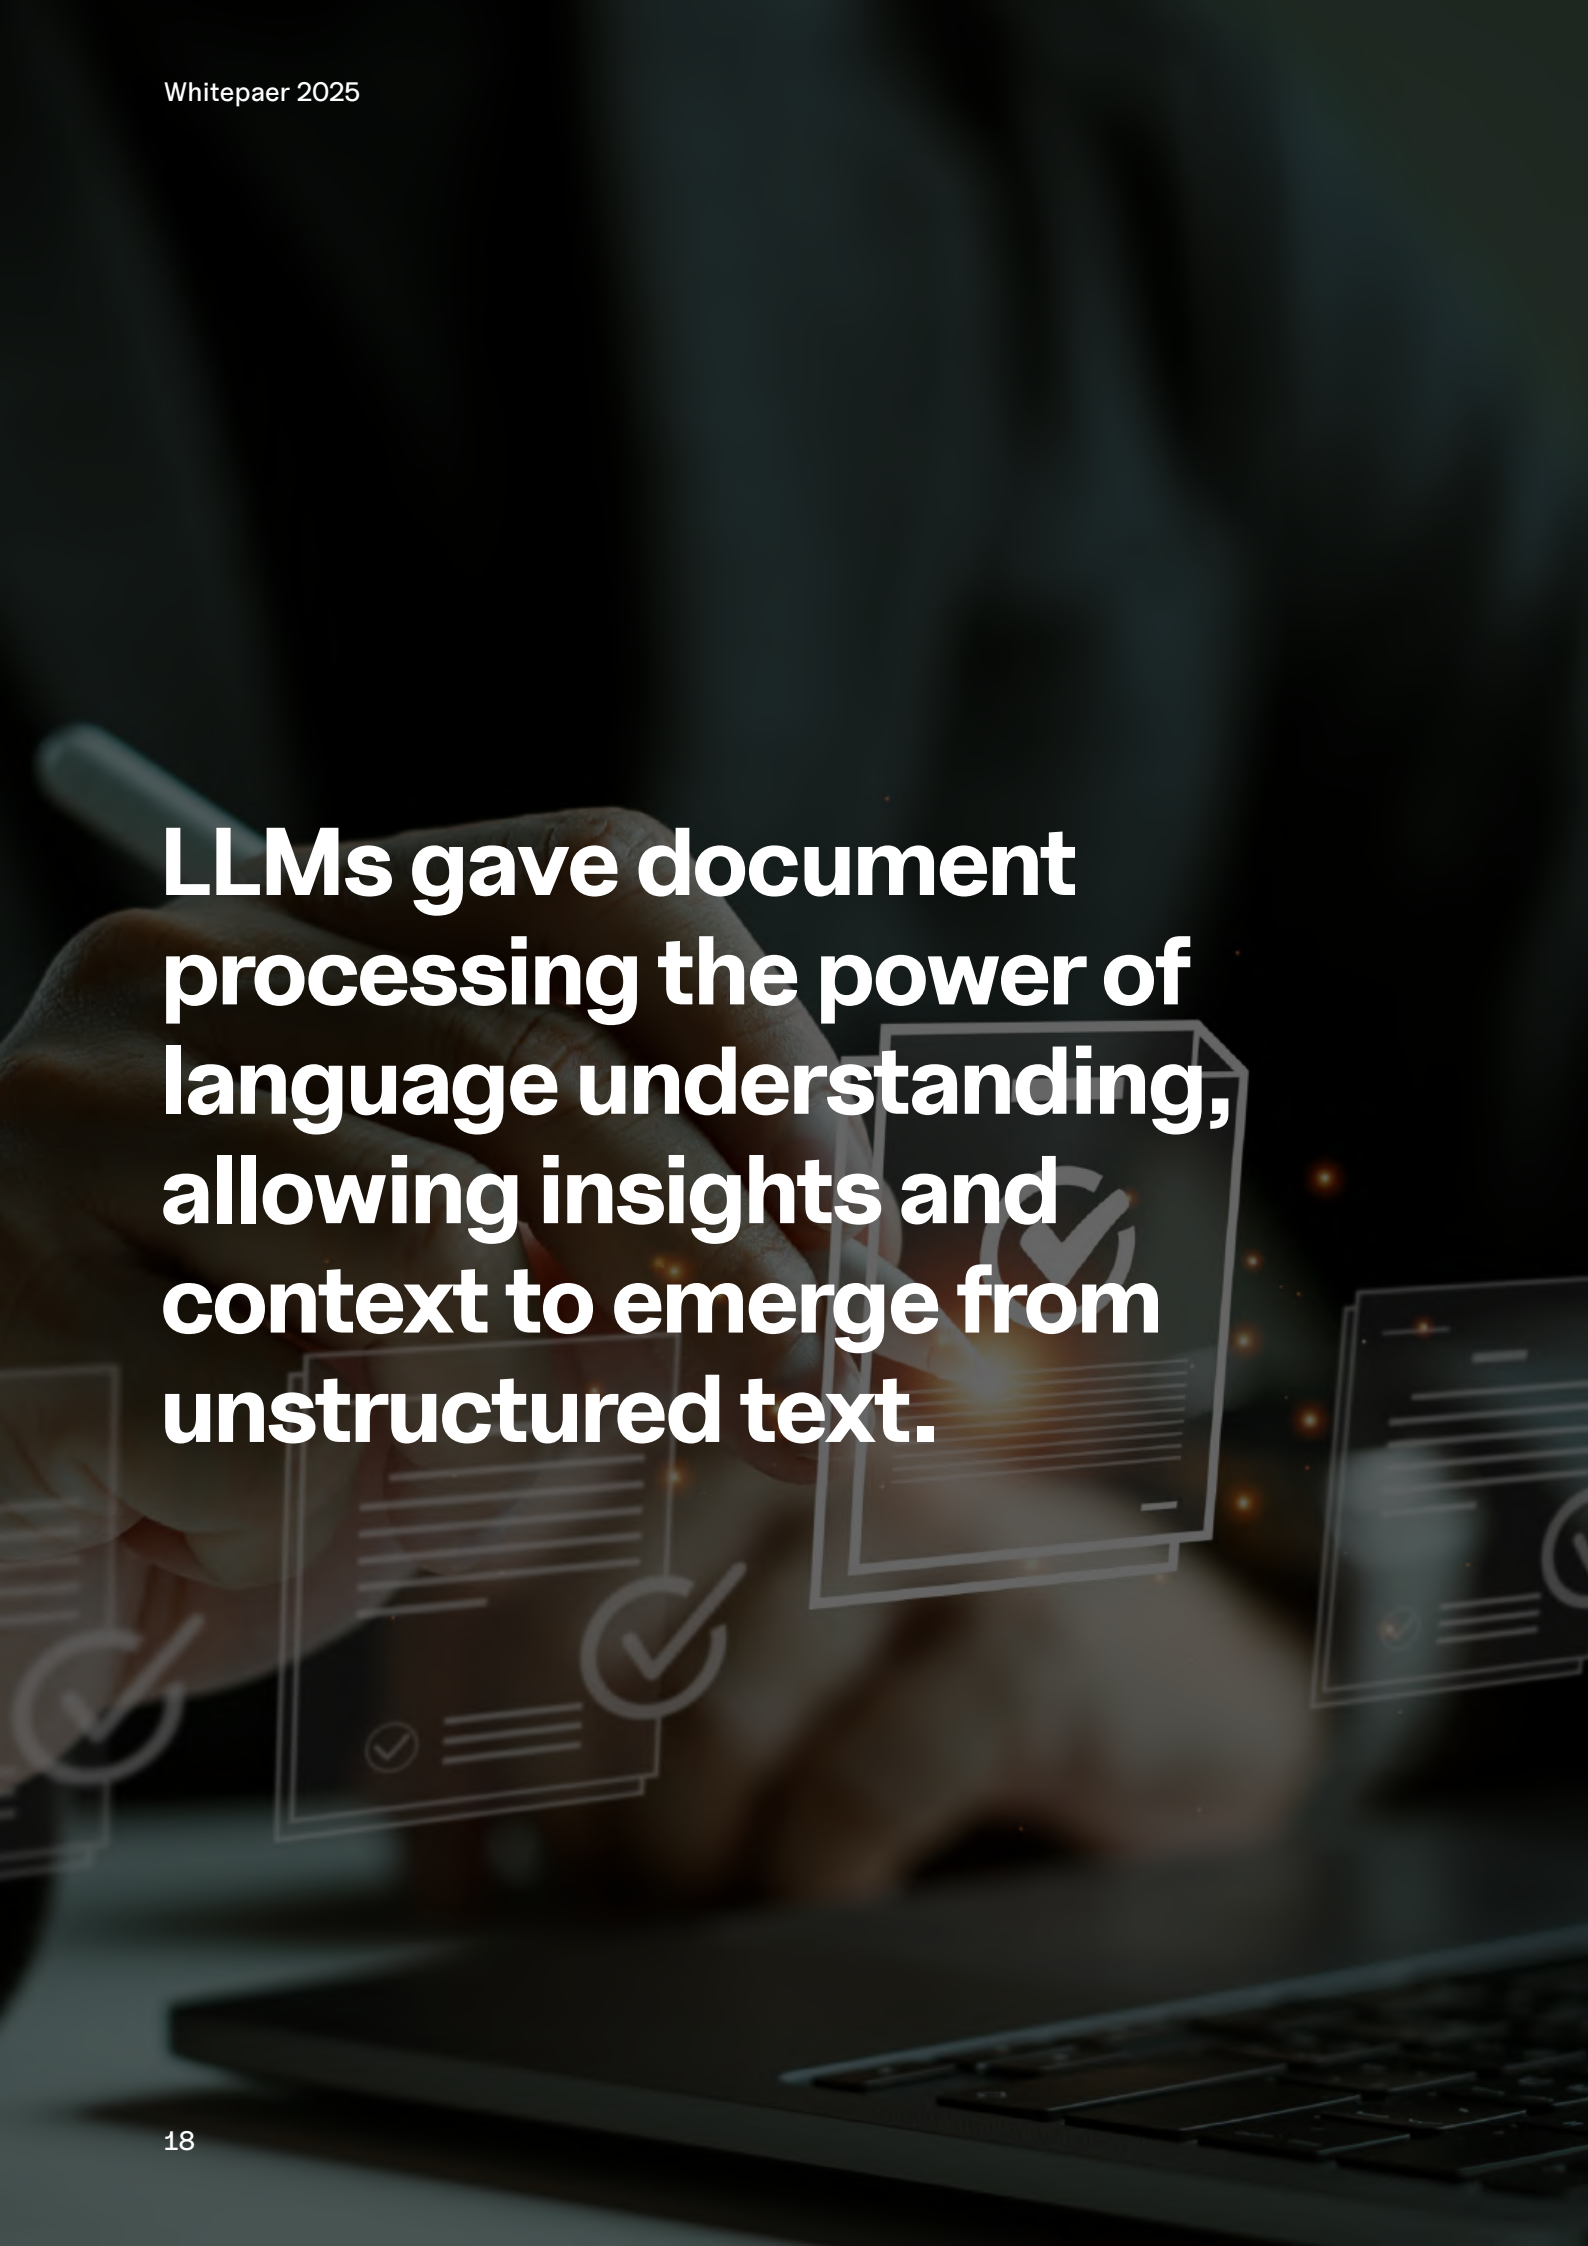
Example Use Cases for LLM/GPT Document Processing:**MORTGAGE APPLICATIONS**

A financial services company uses LLM/GPT document processing to process mortgage applications by extracting relevant information from lengthy documents, such as income, assets, and liabilities. This streamlines a process that can be a bottleneck in home buying.

**IMPROVED PATIENT OUTCOMES**

A hospital uses LLM/GPT document processing to process clinical notes, identify key patient data, and even suggest preliminary diagnoses based on patient records. This leads to improved utilization of hospital resources and patients getting better care faster.

While LLM/GPT document processing can be expensive and is not a one-size-fits-all solution, their ability to combine contextual understanding with adaptability makes them powerful tools for businesses looking to automate and enhance document processing workflows.



LLMs gave document processing the power of language understanding, allowing insights and context to emerge from unstructured text.

Cognitive Document Processing

Generation 05

CDP is a very new technology that expands capabilities and can address a broader range of documents than previous generations of document processing technology. It also delivers better results and is less computationally intensive than LLM/GPT systems. It uses an approach that functions much more like a human brain. It uses advanced machine vision as a cognitive preprocessor to mimic perception and filtering. It then uses element grouping, pattern recognition, and synthesis to better integrate diverse pieces of information, ideas, or perspectives into a unified understanding.

Zero-shot:

Because of its use of cognitive preprocessing, high-quality data can be extracted from documents zero-shot, without the need to pretrain a model. This saves development and runtime costs and allows the system to be much more tolerant of variability.

Multiagent analysis:

CDP uses a multiagent system that involves machine vision, advanced OCR, deep learning, advanced NLP, and semantic analysis. These technologies work together to provide comprehensive document analysis and understanding. To improve results and minimize the dangers of model bias, multiple agents can generate output and perform a best result analysis. The resulting enhanced contextual understanding allows CDP systems to better extract information from unstructured documents, such as contracts, legal briefs, and research papers, with very high accuracy.

Contextual metadata:

The captured context and relationship metadata makes the

information encapsulated in a document even more valuable. It enhances decision-making capabilities and improves downstream business processes by providing deeper insights into document content.

Self-correcting:

The multiagent approach and the improved understanding of context and the relationship between data elements improves the quality of understanding and provides a basis for continuous learning. As a result, CDP systems can improve over time and self-correct. This allow CDP systems to be much more flexible, adaptive, and maintainable, and much more adept at handling document variability and complexity.

Lower costs:

Once data is preprocessed, advanced machine vision and the cognitive techniques, fuzzy logic or LLMs, can be used to extract information. The cognitive system approach allows CDP to process documents of greater complexity at higher quality and lower cost than IDP and LLM/GPT solutions.

Key Characteristics of CDP:

- Multiagent system
- Advanced NLP and machine vision
- Context-aware data extraction and analysis
- Semantic understanding of document content
- Unstructured document handling
- Adaptive learning from feedback
- Document variability tolerance and application versatility
- Flexibility in document handling
- Robust and resilient
- Faster deployment and reduced costs
- Improved data accuracy and quality

Example Use Cases for Cognitive Process Automation:



CONTRACT ANALYSIS


A legal department uses CDP to analyze complex contracts, automatically identifying key clauses, obligations, and risks. The system highlights areas that require review or negotiation, reducing the workload of legal teams and enabling them to focus on more strategic tasks.



REGULATORY POLICY ENFORCEMENT

An energy company uses CDP to ensure compliance with environmental and safety regulations by automating the extraction and analysis of critical data from inspection reports, permits, drilling permits, environmental impact assessments, safety inspection reports, and emissions logs. The system extracts key data points, cross-reference the data with local and international regulatory standards and flags non-compliance issues or anomalies in real time.

CDP processes document much more like humans do. This gives CPA systems human-like flexibility in workflows. The cost and flexibility advantages of this approach are likely to accelerate its adoption in processes and workflows that involve complexity and variability of documents.



CDP mimics human cognition, enabling zero-shot, context-rich understanding of complex documents with unmatched flexibility and accuracy.

Ecosystem Information Management



Generation

06

A possible future is ecosystem information automation (EIM). This process allows AI-driven document processing to be seamlessly integrated into broader business ecosystems. Unlike previous generations that focused on individual document handling and automation, EIM aims to create a fully interconnected environment that enhances collaboration and real-time decision-making across entire business networks.

Cross ecosystem workflows:

EIM will be able to process documents in real time as part of end-to-end automated workflows that span multiple systems across multiple companies. For example, an EIM system could integrate customer relationship management and compliance systems to ensure that customer data is correctly processed and verified across departments in multiple companies without manual intervention. These systems go beyond extracting data—they view documents as part of an interconnected information ecosystem not just isolated data sources. This enables a more comprehensive ecosystem-wide approach to decision-making.

Data fabric:

EIM relies on advanced cognitive and reasoning AI systems, real-time data sharing, agentic workflow, and interoperability between systems using common data models residing on a blockchain-based infrastructure. This will enable organizations to operate in highly efficient, coordinated, and collaborative configurations. Blockchain services will provide the essential identification and reputation mechanisms. These mechanisms will give

ecosystem participants confidence in the veracity of the information contained in documents and the metadata associated with that information.

Efficiency:

EIM will reduce the need for manual input and duplication for workflows such as data entry and document verification. It will manage documents and their encapsulated information as part of a broader, integrated process. It will also be able to process and analyze documents in real time, providing essential connectivity between multiple businesses, eliminating ecosystem silos and improving overall operational efficiency.

Challenges:

EIM could potentially lead to radical efficiency gains and innovation at scale. However, it will require a significant investment in technological infrastructure and process redesign as well as the development of standards and advanced decentralized governance systems. Integrating the many different systems in most ecosystems will also be technically challenging and require careful planning and execution. Challenges may include ensuring data compatibility across platforms, managing data privacy concerns, and coordinating between multiple stakeholders with differing technological capabilities. EIM will also require a cultural shift because businesses must adapt to fully automated, real-time, multiparty operations.

Key Characteristics of EIM:

- Holistic integration across ecosystem partners
- Shared data environment
- Real-time document processing integrated with business systems
- Real-time analytics and decision support
- AI-driven decision-making
- AI collaboration
- Interoperability with multiple backend systems
- Improved ecosystem efficiency
- Security and privacy management
- Dependence on complex digital blockchain-based infrastructure

Example Use Cases for EIM:



PROCUREMENT AUTOMATION


A multinational corporation with complex multiparty supply chains might use EIM to automate an entire procurement process, from receiving purchase orders to negotiating contracts, managing suppliers, and processing payments. This can significantly reduce procurement cycle times and minimize human errors, enhancing overall efficiency.




LEGAL COMPLIANCE

A financial institution facilitating trade finance might use EIM to automatically analyze and process regulatory documents across multiple jurisdictions, ensuring real-time compliance with changing legal requirements. This helps avoid penalties and improves audit readiness, providing greater operational assurance.

EIM and the associated ecosystem economy implies both a technical and a behavior complexity that will be very difficult to overcome. Like all technology and business change it won't just appear it will evolve. We will likely see the technology, change management, and business cultural changes that are required to manifest inside of large companies and then inside of highly concentrated ecosystems. Eventually highly composable ecosystems will evolve that will leverage digital infrastructure that includes EIM functionality.





**ELM envisions documents
as part of a connected
data fabric—powering
real-time, AI-driven
decision-making across
entire ecosystems.**

Generations Summary

Table 1: Comparison of Generations of Document Processing Technologies

Generation	key Features	Strengths	Weaknesses	Example Use Cases
Optical character recognition	<ul style="list-style-type: none">•Text recognition•Image preprocessing•Pattern recognition•Postprocessing	<ul style="list-style-type: none">•Improved efficiency•Searchability•Cost-effective for basic digitization	<ul style="list-style-type: none">•Struggles with unstructured text•Limited to standard fonts and clean documents	<ul style="list-style-type: none">•Digitizing archives•Automating data entry
RPA document processing	<ul style="list-style-type: none">•Rule-based•Works with RPA•Limited to structured data	<ul style="list-style-type: none">•Cost-effective for highly structured documents•Integration with business systems	<ul style="list-style-type: none">•Limited to structured data•Inflexible	<ul style="list-style-type: none">•High volume data entry•Invoice reconciliation
Intelligent document processing	<ul style="list-style-type: none">•OCR integration•Machine learning•Template-based processing	<ul style="list-style-type: none">•Handles semi-structured data•Reduces human errors	<ul style="list-style-type: none">•Relies on templates•Needs training and human intervention	<ul style="list-style-type: none">•Invoice scanning•Compliance checking
LLM/GPT document processing	<ul style="list-style-type: none">•Transformer-based AI•Contextual understanding•Multilingual•Advanced summarization	<ul style="list-style-type: none">•Ease of use•Excels with unstructured data•Adaptable across industries	<ul style="list-style-type: none">•High computational demand•Expensive fine-tuning•Lack of explainability	<ul style="list-style-type: none">•Mortgage application processing•Clinical notes analysis
Cognitive document processing	<ul style="list-style-type: none">•Machine vision•Semantic analysis•Zero-shot•Self-learning	<ul style="list-style-type: none">•Contextual understanding•Handles variability•Reduced setup time•Low processing costs	<ul style="list-style-type: none">•Limited explainability when using LLMs for postprocessing	<ul style="list-style-type: none">•Insurance claims processing•Policy enforcement
Ecosystem information management	<ul style="list-style-type: none">•Holistic ecosystem automation•AI-driven collaboration	<ul style="list-style-type: none">•Cross-organization automation•Real-time analytics	<ul style="list-style-type: none">•Complex to implement•Requires collaboration standards	<ul style="list-style-type: none">•Supply chain automation•Smart cities

Conclusion


The evolution of process automation era document processing technologies—OCR, RPA document processing, IDP, LLM/GPT document processing, CDP, and eventually EIM—has been driven by the need for greater efficiency, accuracy, and scalability in handling ever-increasing volumes of business-critical information. OCR and RPA document processing provided the initial automation of simple, rule-based tasks, while IDP introduced machine learning to handle semi-structured data. LLM/GPT document processing added a much greater ability to understand and contextualize language. CDP improved adaptability without extensive training requirements and added continuous learning and self-correcting mechanisms. EIM represents a future state where all these advancements are combined with next-generation digital data infrastructure, creating a fully interconnected and collaborative environment.

As organizations implement more advanced document processing technology, the capability to process more complex and unstructured documents improves, human intervention decreases, and the value of document-derived information is increased.

Businesses adopting advanced technologies, such as CDP, will unlock significant operational efficiencies and competitive advantages, including faster decision-making and reduced processing costs. However, these advancements also require careful planning, investment, and organizational readiness, such as training employees, updating internal processes, and ensuring effective change management. The implementation of these technologies can deliver significant return on investment and improve business agility and productivity. It is also an essential step toward a future of document processing when fully integrated, AI-driven ecosystems will fuel business processes and shape real-time decision-making.

The real question for most businesses is not whether to automate document processing, but how far they should go toward that goal and how quickly they need to embrace these technologies to stay competitive.

About the Authors:



Douglas Heinzman
CEO
Co-Founder
Syncura



Hugh Frost
Head of Product & Services
Co-Founder
Syncura



Jon "Jet" Theuerkauf
Head of Client & Partner
Co-Founder
Syncura



Daren Hanson
Head of Sales & Marketing

About Syncura:

Syncura is a leading provider of advanced cognitive document processing and cognitive process automation solutions. Its technology enables businesses to automate complex processes and handle dynamic document workflows with unmatched accuracy and efficiency. Its groundbreaking observational cognitive AI technology mimics human-like reasoning, allowing organizations to streamline operations, reduce costs, and improve productivity without extensive manual setup or retraining.

Syncura's Cognitive Document Processor and Cognitive Automation Platform continuously adapt and learn, ensuring flexibility and scalability across industries such as finance, legal, healthcare, and others. With Syncura, businesses can confidently automate their most challenging processes and tackle complex document management tasks.

Syncura is committed to driving innovation in the automation space by combining cutting-edge AI technologies with deep industry expertise. By providing solutions that evolve with changing business needs, Syncura empowers organizations to stay ahead in an increasingly automated and interconnected world.

For more information, visit www.syncura.ai or send us an email at info@syncura.ai



The Road to Cognitive Automation and Beyond

www.syncura.ai