



Really Intelligent Document Processing

by Sarah Burnett



CMR+

is a market leading Intelligent Document Processing (IDP) platform that combines advanced techniques to streamline and automate document processing. The advanced techniques include computer vision, character recognition (OCR/ ICR/ OMR), machine learning (ML), deep learning (DL), and natural language processing (NLP). With CMR+, enterprises can automate their document-intensive processes to significantly reduce manual work and increase process efficiency. CMR+ has been built for handling enterprise concerns around integration, configurability, security and scalability and has been deployed across Fortune 500 companies including in regulated sectors.

How it works

CMR+ leverages a mix of technologies to process documents starting with computer vision to take in the visual components and layout of the documents. It then uses a variety of techniques to recognize characters to extract text and data from the documents. The techniques include Optical Character Recognition (OCR), Intelligent Character Recognition (ICR), and Optical Mark Recognition (OMR). These are described in more detail further on in this document. CMR+ takes advantage of ML algorithms for training purposes. The algorithms help it to learn to recognize patterns and understand the context of the extracted information. Deep Learning (DL) models using neural networks are applied to further enhance the accuracy and reliability of CMR+'s document processing. NLP techniques enable the system to comprehend and analyze the textual content of the documents.

CMR+ operates in four stages: enhancement, classification, extraction and inference, and validation.

Enhancement

The first stage is enhancement, when the extracted data is enhanced and optimized to improve its quality. Enhancement involves cleansing the captured data, normalizing, and augmenting it. These steps ensure the accuracy, consistency, and completeness of the information that CMR+ has captured from the documents. It can also apply noise reduction algorithms to enhance any images in the document to improve their quality so that they can be processed with greater effectiveness.

Classification

Stage two is classification, to efficiently process documents. CMR+ utilizes various classification techniques to understand the nature and purpose of each document to categorize them. There are predefined classes in CMR+ but custom-defined classes can be added to it. Classification allows CMR+ to use the most appropriate techniques and logic to process many different types of documents, such as invoices and insurance claim forms.

It learns to classify new types of documents by analyzing folders of documents to build the initial models. Then it takes human feedback to check its output and learn from the corrections.

Extraction

During the extraction phase, using pre-defined rules and learned patterns, CMR+ identifies and extracts specific data items from documents that it is processing. Using ML for confidence scoring, it generates scores as the data is being extracted. It leverages parameters around how the data can be improved and presents the score to the user. The data can be of the structured variety such as names, addresses, invoice amounts, or order numbers, or unstructured data like text in a paragraph or clause.

Inference involves analyzing a paragraph of text to understand the context, expected value or classification.

Another feature of CMR+ is sentiment analysis, that it applies to understand the tone of the text in a document. This is particularly useful in customer interactions such as reviews, feedback, and complaints.

CMR+ has prebuilt solutions and accelerators for Policy Intake and Policy Review, Slips, and Loss Run for Insurance brokers and carriers; In addition to Invoices and Pegasystems Accelerators. Other models can be built either by AntWorks, its partners, by citizen developers or automation professionals.

Validation & Business Rules

Validation ensures the accuracy and reliability of the extracted information by checking it against predefined rules and business logic. The extracted data is also validated against external reference sources and databases. The techniques used for validation include cross-referencing, data matching, data type validation and data reconciliation. Consequently, errors can be identified, and discrepancies resolved to reduce the risk of processing errors.

Business Rules are a set of out-of-the box rules which can be applied to the extracted data encompassing formatting, mathematical operations, string manipulation, and reconciliation rules.

AntWorks provides training, expert services support and customer success resources to enable customers to maximize their investments.

Trusted by



CMR+ ENABLED BY GENERATIVE AI


Generative AI technologies, and specifically Large Learning Models (LLM) have driven product evolution significantly. In the context of documents, one major aspect of Generative AI is different to IDP, and that is its ability to generate text, based on the users' requirements. IDP captures information from scanned and digitized documents, but it does not generate text. However, LLMs can help enhance IDP in many ways.

Streamlining access to data for business users

Utilizing a natural language interface for system interaction. This facilitates seamless navigation between structured and unstructured data, enabling the retrieval of information to address business inquiries (e.g. reviewing specific clauses in contracts for the top five borrowers in the subprime segment).

Empowering Enterprises with Integrated Document and Data Processing

Enabling new capabilities for enterprises in processing documents and data through an integrated approach. This encompasses tasks demanding language comprehension, contextual integration, and the capacity for planning and logical reasoning to provide precise responses.



CMR+ platform leverages the following guiding principles in its integration with GenAI:

- Establishing a seamless fusion of conventional methodologies and LLM-based capabilities to optimize cost, control, and performance.
- Prioritizing customer concerns regarding data security and control by allowing content regionalization to minimize external LLM data transfer. Additionally, the platform facilitates redaction and the application of data guardrails, thereby emphasizing security and control measures.

CMR+ implements standardized data pipelines with plug-and-play interfaces across both SaaS and private LLMs, alongside its storage infrastructure. The platform's development entails the creation of domain-specific, fine-tuned models tailored for distinct business domains, with further customization and contextualization options available.



AI-Generated images. Prompt used: "Visualize IDP"

Empowered by GenAI capabilities, CMR+ platform supports the following functions:

- Offering an alternative approach to extraction with minimal reliance on prior training or limited fine-tuning processes (e.g. interpreting tables, deducing paragraph meanings, etc.).
- Synthesizing multiple data inputs to generate output, including:
 - a. Summarization across documents of varying sizes through direct prompts (for concise excerpts) and overlapping embeddings (for extensive documents)
 - b. Language translation across diverse languages and complete documents or identified sections
 - c. Classification into predefined categories based on prior training
 - d. Analysis of sentiment and determination of customer mood to facilitate effective actions and downstream queue allocation
- Enabling comprehensive search and comparison of clauses while considering business context:
 - a. Engaging in context-retaining conversations with documents based on semantic meaning
 - b. Verification against a datastore listing similar clauses for the closest match
 - c. Comparison of two clauses within the same document or across separate documents
 - d. Evaluation of a clause against the enterprise's standard clause to flag potential risks
- Facilitating reasoning through multi-step decision flows involving logical steps derived from structured and unstructured sources. For example, determining customer pricing based on volume slabs and special incentives, or evaluating inputs and guidelines for loan underwriting.

UNDER THE HOOD

Under the hood, CMR+ is powered by advanced document processing capabilities including:

- Computer vision that enables it to interpret visual elements and their context within documents.
- OCR, ICR, and OMR technologies that enable it to extract different types of text, including handwriting, handwritten comments, and editorial marks. These combine OCR with artificial intelligence to improve the accuracy of information that CMR+ captures.
- ML algorithms are key to CMR+ as they enable it to learn patterns and identify document layouts and structures as well as improve its processing accuracy.
- DL models are utilized to improve CMR+'s capability to process more complex document layouts as well as unstructured data.
- NLP techniques are fundamental to CMR+'s ability to understand and analyze the content of documents, conduct sentiment analysis, and understand the language to process textual content.
- Confidence engine to power confidence scoring which can enable enhanced productivity by ensuring the user focusing their time on the right information.

Automation and Straight Through Processing

CMR+ automates document processing leading to highly efficient straight-through processing (STP) for the enterprise seamlessly without the need for manual intervention. Accordingly, organizations can eliminate error-prone and lengthy manual document-based tasks and reduce document processing cycle times.

CMR+ can be an integral part of an end-to-end automated process that includes many other automation tools, such as Robotic Process Automation (RPA), e-mail agents and conversational interfaces with chatbots.

Human in the Loop (HITL)

Human-in-the-loop combines the power of AI and humans to ensure the accuracy and quality of IDP outputs. It allows humans to handle exceptions and ambiguities, to check for errors and validate information. Humans typically interact with the automated IDP process at the validation stage, reviewing and correcting the results of the automated extraction. Additionally, by providing feedback to the IDP system, humans help to teach, update and improve the machine learning algorithms. The feedback boosts the IDP's accuracy and with experience improves its ability to process different document types. Accordingly, over time, the scope of the automation can increase, and exceptions and errors minimized.

When IDP is combined with Generative AI capabilities, many aspects of HITL can be improved with features such as conversational interfaces to the IDP system, better search, extraction, summarization, and translation of information. These powerful features make the work easier for humans helping them complete tasks faster, leading to higher productivity. As a result, the humans' sense of achievement can increase, leading to better job satisfaction and engagement.

On an average, in high value processes, Human-in-the-loop reviews of extracted data delivers 80% productivity gains compared to manual extractions and review.

USE CASES

CMR+ has applications in many different industries and business functions, including insurance, banking, manufacturing, purchase-to-pay and contracting processes. Examples include:



Insurance



- **Policy Review:** In Marsh, CMR+ extracts and compares multiple complex data items from expiring policies, binders, quotes like limits, deductibles and endorsements to ensure that the policy is consistent and compliant with agreed source documentation. It speeds process execution with very high accuracy in addition to helping improve customer satisfaction.
- **Policy Administration:** CMR+ excels at the extraction and validation of policy information. This encompasses coverage specifics and premium data. By doing so, it spearheads the modernization of policy administration, resulting in heightened service quality and elevated customer satisfaction levels and reduce costs.



Banking



- **Compliance and Regulatory Reporting:** In US Bank, CMR+ extracts relevant information from different types of financial documents, including records of transactions, purchase orders, invoices, and receipts, to minimize risks of regulatory failures, maintain good audit trails and accurate reporting.
- **Account Opening:** CMR+ speeds up the Know Your Customer (KYC) and on-boarding process. Firstly, it automatically extracts and validates the customer information from account application forms. Secondly, it can check the customer's ID against internal and external reference databases. Finally, as part of a joined-up process automation workflow, it can send the information to a human, a banking system, or a robot to set the account up.
- **Trade Finance:** CMR+ revolutionizes trade finance with intelligent document processing, automating verification, Letter of Credit processing, and compliance checks. It digitizes documents, streamlines payments, and offers valuable trade analytics & insights, ensuring a seamless and secure international trade ecosystem.



Manufacturing



In Unilever, CMR+ automatically ingests and analyses the results of lab tests. It does so by analyzing data in Certificates of Analysis (COA) that provide details of laboratory tests that are conducted on manufactured products. CMR+ automates the extraction of the test parameters, the measurements, and quality control data, enabling automated analysis and comparison.



ESG (Environmental, Social, and Governance)

CMR+ aids companies in collecting and analyzing ESG-related data from various sources. This includes extracting information on environmental sustainability practices, social impact metrics, and governance policies from corporate reports, disclosures, and other relevant documents. By automating this process, CMR+ facilitates efficient ESG reporting, helps companies meet their sustainability goals, and fosters transparency in their ESG efforts. As companies increasingly prioritize their ESG performance, CMR+ serves as a valuable tool to promote responsible and sustainable business practices across industries.

These are just a few examples of CMR+ applications in various industries. The technology offers the advantage of automating the extraction and processing of information from a wide range of documents, reducing manual effort, improving data accuracy, and increasing operational efficiency.

THE BENEFITS

By automating document-related processes, CMR+ drives a step-jump in productivity while reducing processing times and data capture errors. It empowers enterprises to scale up rapidly while ensuring regulatory compliance, manual work, faster processing and significantly increased productivity.

In our customer implementations, clients report benefits including productivity jumps of up to 80% and automated accuracy rates exceeding 95% in complex processes.

100%


Data extraction accuracy with human-in-the-loop

95%+

Classification accuracy

80%

Increase in productivity



Drive straight-through-processing (STP) with business driven confidence engine



Speed to value through citizen development



Enterprise scalability



Document, process and industry agnostic

HOW TO GET STARTED

- Choose the first use case
- Set the objectives for implementing CMR+
- Work out if any preparation is needed in terms of the basic quality of the documents and the way that they are organized
- Conduct a pilot implementation to gain experience and to use the lessons learnt when scaling up
- Training CMR+, you will require 100 of each document type, for example, purchase orders, invoices, insurance slips
- Provide training to the relevant staff and manage change effectively, with on-going monitoring to ensure continued success

By following this approach, organizations have effectively adopted CMR+ to streamline their document processing and achieve the desired business outcomes.



Sarah Burnett

**Technology Evangelist and Board Advisor, Author,
Analyst, AI Accelerator founder | FBCS**

Sarah Burnett FBCS is a renowned analyst, advisor, and author in intelligent automation. Her acclaimed book, "The Autonomous Enterprise – Powered by AI," was published by BCS, the British Chartered Institute for IT, in 2022. Her current clients include AntWorks, and KYP.ai where she serves as Chief Technology Evangelist. Previously, she was executive vice president and distinguished analyst at Everest Group. Earlier in her career, Sarah worked in software development in financial services and e-Government programme management.

A passionate advocate for gender diversity in tech, Burnett chaired BCS Women and founded and ran AI Accelerator from 2016 to 2021, promoting AI industry diversity. She's been acknowledged by Computer Weekly as a top influencer in UK IT and entered their Influential Women in Technology Hall of Fame in 2021.

For more information, please visit www.ant.works