

Cognitive Invoice Processing

Unlocking value from a deluge of documents

wipro holmes



intel
AI Builders

Introduction

Document management issues continue to be a major pain point for businesses even in the era of digitization. Digest this—a staggering 46% of workers waste time every day on paper-intensive workflows. There is also a steep rise in the volume of digital documents such as application or claim forms, checks, passports, bills, and many others, in various formats. Processing these items, particularly the semi-structured and unstructured documents, to unlock insights from document data at high speed, precision, and accuracy is a task prone to errors and inefficiencies. However, business leaders often ignore the value at stake from transforming document processing. Thankfully, as Artificial Intelligence (AI) infiltrates every aspect of business, innovative organizations are realizing the potential of AI-driven Intelligent Document Processing (IDP).



Intelligent Document Processing is rapidly emerging as a category of solutions that help extract data from documents that drive business – like insurance claims, mortgage applications, contracts, purchase orders, invoices, engineering drawings, and so on.

This paper discusses the key drivers of IDP adoption in enterprises and introduces Wipro HOLMES® Invoice Processing solution, powered by Intel® Xeon® Scalable processors, for cognitive invoice processing capabilities.

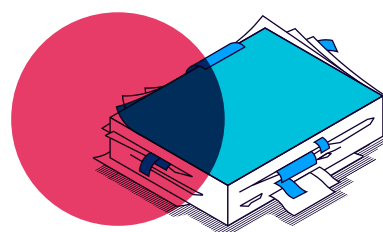
From navigating challenges to driving opportunity: Moving up the curve

Although many high-performing organizations during the last decade have been investing heavily in Robotic Process Automation (RPA), they have reached a saturation point – exhausting the low-value opportunities available through such simple task-based automation. For instance, traditional RPA bots fail to classify the documents by type and nature, or work with bad quality of scanned documents. Most OCR engines fail to recognize the data from scanned documents accurately, making it challenging to achieve at least 80%-90% of automation. Access rights management capabilities with security are also lacking in traditional RPA solutions, besides heavy dependency on experts in specific domains. Integrating IDP technologies such as AI, machine learning (ML), computer vision and others with RPA helps enterprises automate document-driven use cases. This drives re-imagination of business processes end-to-end, rather than getting stuck in an endless loop of task-based automation. IDP has the potential to elevate the level of document processing to a strategic enabler of business

success. The IDP market is expected to grow to USD 3,855 million by 2024, at a CAGR of 54.5%.

IDP solutions increase productivity by assisting experts to process more documents per day, thereby reducing turnaround time, while also improving the quality of reporting through automation.

Besides the ability to define, extract, and validate relevant data, self-learning IDP algorithms improve accuracy and reduce errors over time, even when processing complex semi-structured and unstructured documents. Other than typical financial documents, IDP solutions can also be extended to perform Know Your Customer (KYC) checks from ID cards, reports (both unstructured and semi-structured), forms and bill processing, and research papers, engineering drawings, etc.



Driving intelligent automation with Wipro HOLMES®

Wipro HOLMES®, Wipro's AI and hyper automation platform, connects enterprise assets, data, analytics, and processes to drive automation in an enterprise and successfully address challenges in the building, deployment, and usage of AI solutions. The platform is powered by machine learning and strengthened by years of cumulative experience of Wipro HOLMES® experts. It provides vendor-neutral advisory services and prebuilt solution assets on ecosystem partner marketplaces. It helps to

handhold customers in their digital journey through identification of right business problems, defining success criteria, evaluating right technologies, smooth implementation and adoption, risk and change management, and governance.

The offerings of HOLMES have been successfully deployed across industry verticals and functions driving the 3Es – Efficiency, Economics, and Experience. Figure 1 illustrates how Wipro HOLMES® Automation Studio (HAS) works.

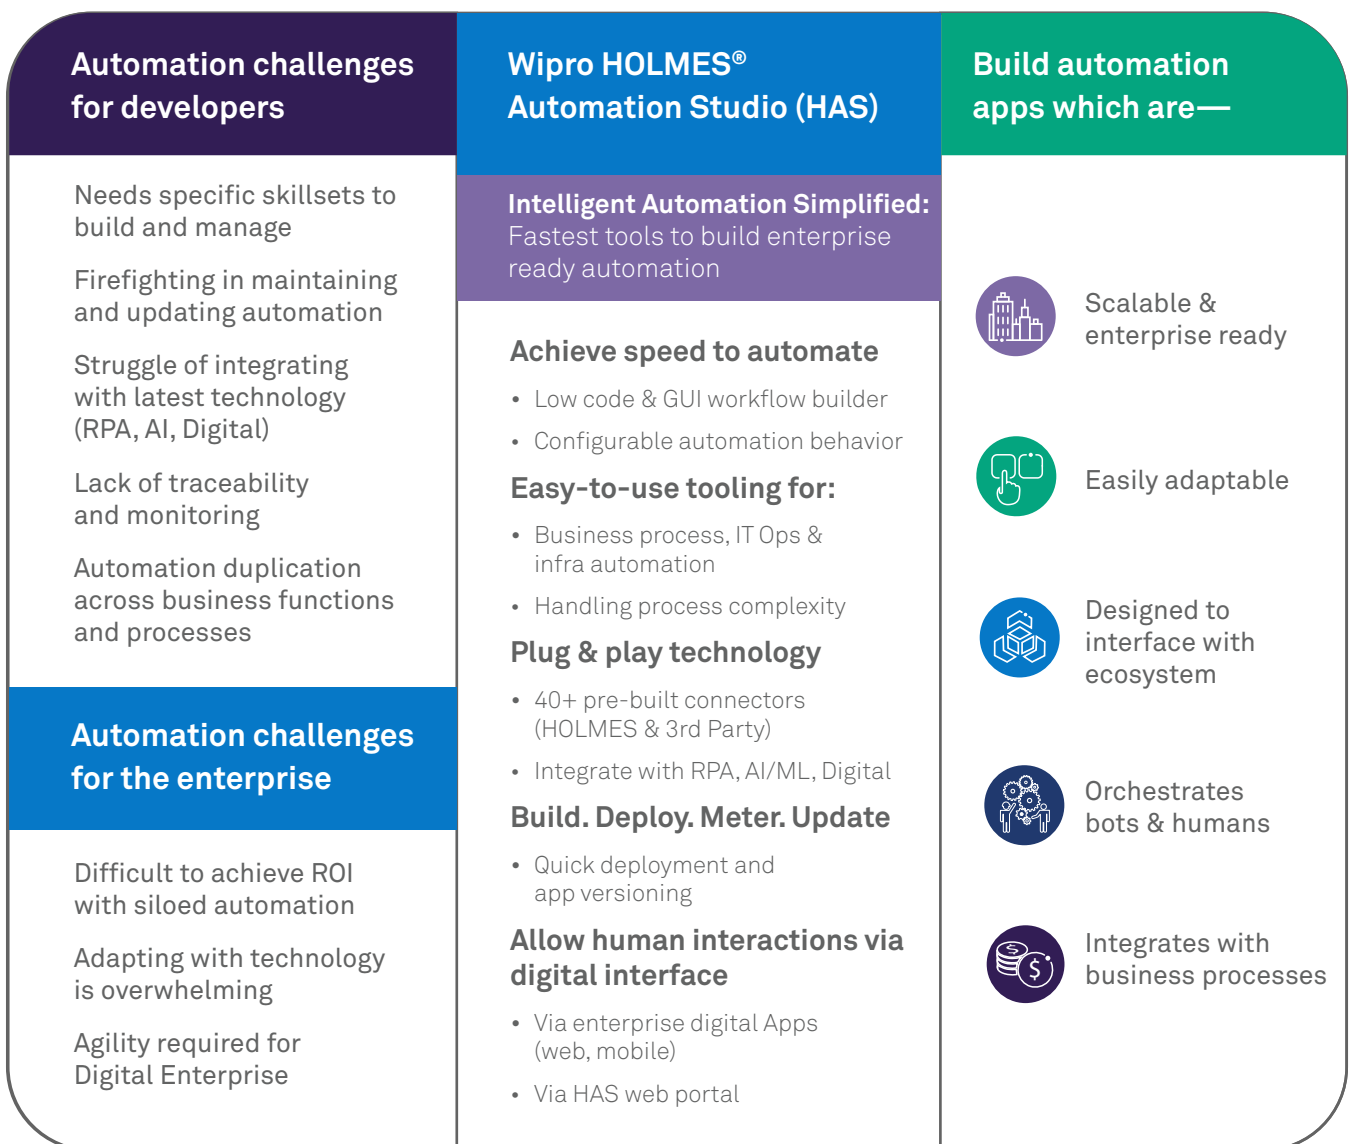


Figure 1: Wipro HOLMES® Automation Studio (HAS)

Introducing Wipro HOLMES® AI-based invoice processing

Wipro HOLMES® Intelligent Invoice processing solution is an AI application that leverages advanced ML and deep learning (DL) algorithms for image processing, table detection, and required value extraction. This is extremely useful when the number of invoices increases, making the process of manual data entry, data validation and approval costly. In such cases, a rules-based approach to

invoice processing, though helpful, is more computational and needs continuous changes in extraction rules to accommodate new invoice types. An automated intelligent invoice processing system is therefore a more scalable option that can help enterprises adopt a standardized approach to invoice processing, without the need to be customized for different types of invoices. Figure 2 illustrates the key steps in Wipro's AI-based Invoice Processing solution.

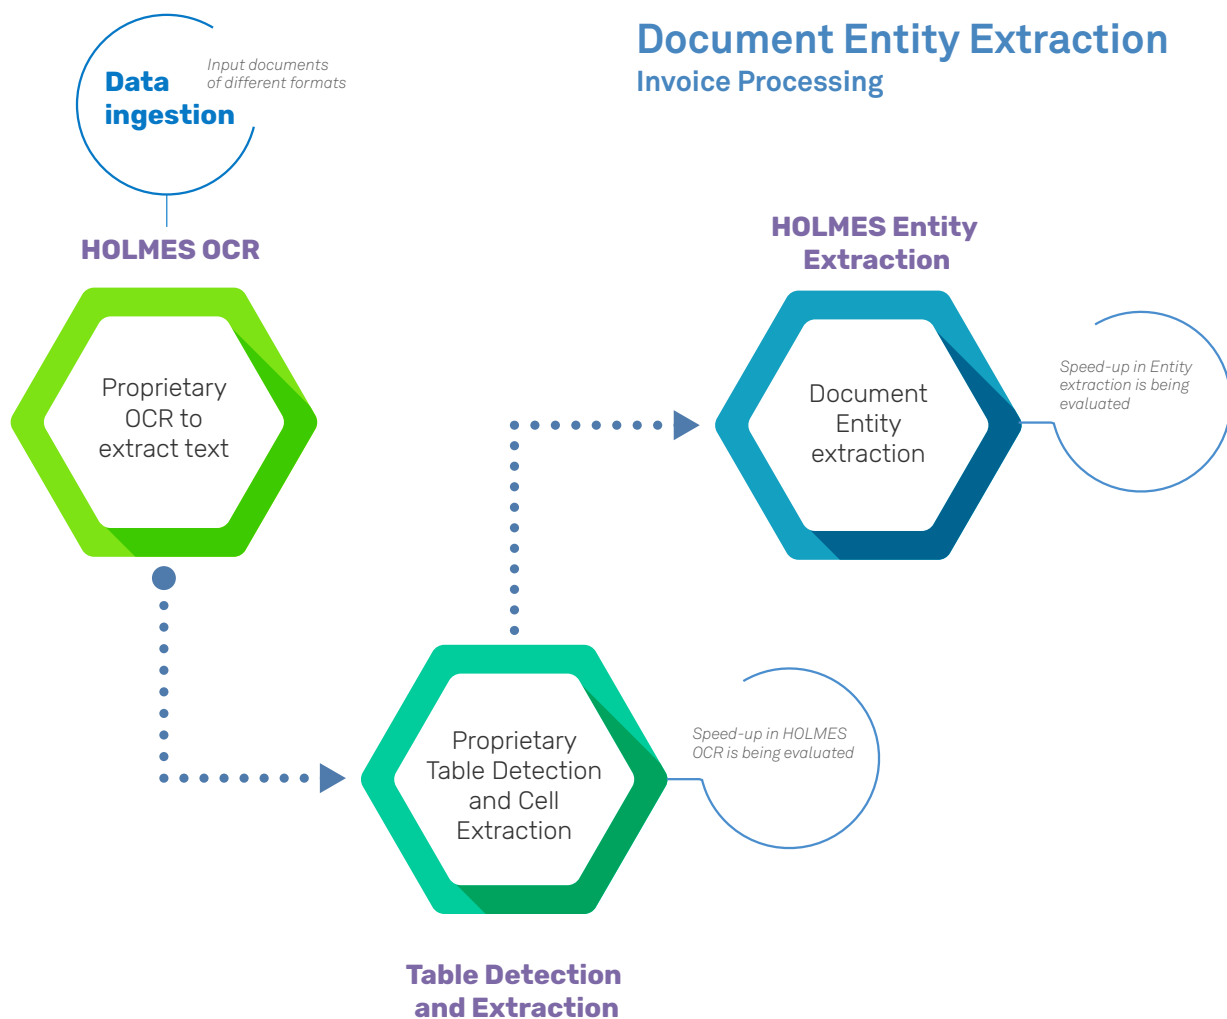


Figure 2: How Wipro's AI-based Invoice Processing solution works

The Invoice Processing solution offers multiple benefits for modern enterprises facing a deluge of documents, including:

- Seamless handling of multi-page invoices and documents in multiple languages.
- Position independent entity extraction from

invoices as opposed to rule-based extraction in most other solutions.

- Competency in handling various quality images including noisy, broken characters, blur, and rotated and multiple image formats.
- Optimization of training time and faster execution of inference based on multi-core CPU

Wipro HOLMES® Table Detection and Entity Extraction

Wipro HOLMES® has designed and developed its own proprietary OCR, table detection and entity extraction to extract data with high accuracy. Designed using advanced DL techniques, it is customized to ensure maximum extraction accuracy from invoices. The table detection is trained on scanned annual reports documents to ensure maximum accuracy. With cognitive validation of table detection and entity extraction results, the model is continuously updated and customized. This improved accuracy of the OCR, table detection, and entity extraction increases extraction rate of the entity extraction module, thus reducing overall digitization time.

Overcoming AI's anticipated risks: How Wipro's human-first approach helps

While the benefits of AI-driven intelligent document processing abound, risks associated with lack of human intervention and control are among the top concerns of any enterprise adopting AI and automation. Any business process being automated needs to be considered from the perspective of criticality and the impact of decisions made by the AI model. Therefore, defining the degree of human involvement is critical during implementation of AI and automation solutions.

Wipro HOLMES® can help in negating these risks through the human-first approach. It is defined based on the level of human involvement and control required in the AI model decision-making process. This includes pre-processing of data before the AI model takes a decision, features considered during decision making, and the after impact of the decision. Humans need to validate the decisions for their correctness, consistency, dependability, etc.. Key metrics of the human-first approach include human interaction type, control over model, and exceptions.

Human interaction type is defined as the kind of human interaction required in handling the output of an AI model which can be classified as:

1. Human in the loop: A human is always part of the AI model, enabling quick identification of problems and requirements that may not be easily identified by other means of simulation.

2. Human out of the loop: AI systems are capable of selecting targets and delivering output without any human input or interaction.

3. Human over the loop: Human is involved wherever model confidence score or similar metric is less than the threshold. Here, a human will aid in checking correctness of only those decisions made by the AI system.

Human-first dimension is based on the level of explainability, transparency, and interpretability of the AI model and its decision-making capability. This helps in checking the human involvement needed in validating and controlling decisions taken by the AI model wherever required.

Success Speaks

The Challenge

The client's manual document processing system led to time wastage, errors, and dependency on domain experts.

Benefits delivered by Wipro's Automated Invoice Processing Solution

- Enabled 80% zero touch - straight through processing
- Enabled 20% Human-in-loop instances
- Reduced assignment time to less than 1 minute

Accelerating invoice processing with Intel® AI technologies

Today, Intel® technologies power some of the most promising AI use-cases in business, society, and research—from medical discoveries to autonomous vehicle intelligence, from tiny devices to massive cloud, turning the promise of a transformative AI model into a global-scale reality. However, their biggest advantage is their democratized enabling of AI leveraging the Intel® Xeon® Scalable based platform that gives an incredible ROI or TCO advantage. Especially in scenarios like RPA, where the datasets available to train the models today are still growing.

Intel® worked with the Wipro HOLMES® team to optimize the HOLMES Table Detection and Extraction framework and Wipro HOLMES® Deep Extraction framework for Text (DEFT) pipeline for Invoice Entity Extraction on Intel® Xeon® Platinum 8124M processor-based platform @ 3.00GHz and the solution is optimized further with Intel® Optimization for TensorFlow, which includes the Intel® Math Kernel Library for Deep Neural Networks (Intel® MKL-DNN).

Based on Fully Convolutional Neural Networks (FCNN), the OpenVINO™ toolkit enables enhanced performance in networks used for tasks of inferencing of table data:

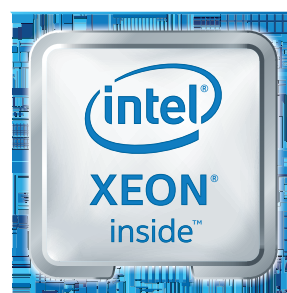
Evaluation of table detection and extraction pipeline was performed on in-house datasets having an average of 100 pages per document. These pipelines were evaluated with default TensorFlow, the Intel® Optimization for TensorFlow and the Intel® Distribution of OpenVINO™ toolkit v2018.5.455 on an Intel® Xeon® Scalable processor-based AWS instance (c5.9xlarge).

In the enterprise landscape, building architecture cognizance for the DL model runtime environment can make all the difference to performance. The lowest hanging fruit, of course, is to leverage Intel®

optimized frameworks and libraries. Building in cognizance of NUMA memory accesses, multi-threading, software thread affinization to cores, etc. are the next level of easy-to-build cognizance for both training and inference. We additionally used distributed training framework called Horovod developed by Uber over MPI (any flavour optimized for x86) for RPC to effectively scale our deep learning model training across the powerful CPU cores. It helped reap maximum performance benefit from the platform. Further, higher CPU utilization for productive computation adds to the ROI and TCO benefit. The same can also be extended to multi-node training.

Transforming document processing

Through framework optimizations and development of software libraries and toolkits, Intel® has enabled significant performance enhancements for deep learning training and inference on Intel®-based platforms using Intel® Optimizations for TensorFlow and the Intel® Distribution of OpenVINO™ toolkit. Working with Intel®, Wipro has successfully run its AI Framework Wipro HOLMES® powered Invoice Extraction pipeline for on-premise systems. Wipro has achieved breakthroughs in technology and outstanding results for Cognitive Automation for Wipro HOLMES® Document Processing System, which has helped launch our solution to work for multiple customers. Together, we are pushing the boundaries and leading change in the way organizations manage their documents.



About the author

Raghavendra Hosabettu has worked extensively in the Machine Learning and Artificial Intelligence, Computer Vision, Mobility, and embedded systems domains. He has ideated, architected, and implemented innovative AI solutions across industry verticals. Currently, Raghavendra is Chief Platform Architect for the Wipro HOLMES® Platform and a Senior Member of Wipro DMTS. He plays the role of a Product owner for Wipro HOLMES® Cognitive Image Processing and Wipro HOLMES® Deep Extraction Framework for Text framework. For more information, contact him at raghavendra.hosabettu@wipro.com

References

- i Adobe, The State of AI in Document Management, <https://adobe.ly/39K76Ea>
- ii Research and Markets, Global Document Analysis Market Report 2019-2024 - Increasing Adoption of Cloud-based Intelligent Document Processing Solutions Drive Market Growth, <https://bit.ly/2L4gZCz>





Wipro Limited

Doddakannelli,
Sarjapur Road,
Bangalore-560 035,
India

Tel: +91 (80) 2844 0011

Fax: +91 (80) 2844 0256

wipro.com

Wipro Limited (NYSE: WIT, BSE: 507685, NSE: WIPRO) is a leading global information technology, consulting and business process services company. We harness the power of cognitive computing, hyper-automation, robotics, cloud, analytics and emerging technologies to help our clients adapt to the digital world and make them successful. A company recognized globally for its comprehensive portfolio of services,

strong commitment to sustainability and good corporate citizenship, we have over 180,000 dedicated employees serving clients across six continents. Together, we discover ideas and connect the dots to build a better and a bold new future.

For more information, please write to us at **info@wipro.com**