Using AI to Accelerate Document Processing for a Major Freight Carrier

CORECOMPETE

DELIVERING VALUE FROM BIG DATA

01. The Mainstreaming of Artificial Intelligence

Most of us are beginning to see how Artificial Intelligence (AI) can be applied to improve our everyday lives. Mobile phones obey voice commands. Simply say “Ok Google” to your Android phone or “Hey Siri” to an iPhone to make phone calls, send text messages, search the web or find directions. With Amazon Echo, say “Alexa” and she can play music, give a news update, read a book, check the weather, or start a voice or video call.

Shazam can listen to a short sample and identify the names of the songs, movies, TV shows. Image recognition from Google, Microsoft and Amazon has improved to the point that it not only can tell the difference between a photo of dog and cat based, but it can even identify the breed of your dog. In the example to the right, Google Lens is an AI-powered technology that uses your smartphone and deep machine learning to scan a picture of your dog, detect the dog, and compare it to an expansive database of dog breeds until it finds a match. If the AI-system finds a strong match, it presents the specific breed and if the probability of a match is lower, for example if you have a mixed breed dog, it will present several possible options.

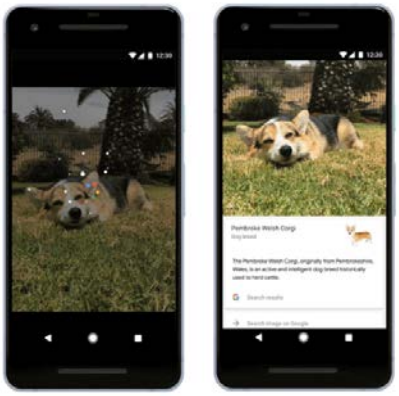

Breed recognition through AI.

With AI becoming more pervasive and usable in the consumer space, the natural question in the business world is “Where can AI and image recognition be applied to address real-time business problems?” One practical use case scenario already adding value is in accelerating the workflow for businesses with high volumes of paper documents such as invoices and bills of lading.

02. AI Adding Value to Document-Centric Processes

While many businesses have a lot of paperwork that require manual processing, the challenges in the freight industry are impressive by any measure. Core Compete worked with a Top 10 LTL (Less Than Truckload) trucking company to implement an Intelligent Process Automation solution designed to improve ROI by reducing errors, lowering costs and speeding up response times. These improvements mean the company can accelerate dispatching, scheduling and moving their fleet assets.

Question:

Where can AI and Image Recognition be applied to address real-time business problems?

The system is able to learn the thousands of different forms customers use that enable it to automatically improve scan rates over time as the system learns and evolves to the constant influx of new customers, forms and ways of specifying the key data on bills of lading.

03. The Challenge: Processing Thousands of BOL's Every Day

The challenge facing the freight company is that they receive over 70,000 bills of lading (BOLs) every day. That's over 2 million per month and each document needs to be processed by a person who reads the BOL and manually keyboards the data into the company's information system. Two additional factors compound the challenge. First, a logjam is created because 60,000 of the BOLs arrive between 4 PM and 1 AM after the freight is delivered. And second, all the work needs to be completed before the next day begins, requiring people to work 3rd shift and process all the BOL's.

Having a large team of people manually keystroking data for long shifts is slow, costly and prone to error.

From an operational effectiveness standpoint, having a large team of people manually keystroking data for long shifts is slow, costly and prone to human error; while also creating a challenging work environment.

According to an Aberdeen Group study (Sept 2015) even the Top Performing 35% of companies take an average of 6 days to process an invoice and 21% have either incorrect or incomplete information. For poorer performing companies processing time averages 18 days. Despite these inefficiencies the majority of invoice processing involves manually keying data from paper.

04. The Promise of Intelligent Process Automation

There is great promise in applying the same type of image recognition and AI we see used in identifying the breed of a dog and applying it to recognize the information contained in documents such as BOL's and invoices. In fact, many companies are already seeing benefit. A McKinsey study on Robotic Process Automation (RPA) showed ROI improvements between 30 and 200 percent in the first year.

Robotic Process Automation showed ROI improvements between 30 and 200 percent in the first year

McKinsey Study, RPA December 2016

Importantly, the study highlights that automation provides more benefits than just faster processing of BOL's, invoices, claims or policies. Workers can be shifted from data entry to more value-added work such as working with customers or dealing with more complex business issues. Furthermore, workers welcome automation because it reduces the time spent on tedious, repetitive tasks.

05. A Look into Intelligent Process Automation for Bills of Lading

Intelligent Process Automation refers to a solution that can be programmed to perform basic tasks in a manner typically done by employees. A practical example is a typical BOL process, shown below.

Trucks deliver their freight (1) and provide a bill of lading (2) that is scanned (3). In step (4), Optical Character Recognition (OCR) and Artificial Intelligence identifies the information in the scanned image and converts it into digitized data that is input into the company's information system (5) and stored in a database (6). The information is routed through the workflow to the team for review and processing (7), delivering productivity improvements (8).

To run the business effectively, there are many critical pieces of data that need to be extracted:

- Shipper's & Receiver's Names and Address
- PO / Account Numbers
- Date of Shipment
- Number of Shipping Units
- Type of Packaging, including Cartons, Pallets, Skids and Drums
- Special Instructions
- Description of Items
- NMFC Freight Classification
- Weight of Shipment
- Declared Value

06. The Importance of AI in Delivering High Accuracy


Bills of Lading

At first glance, it might be easy to underestimate the challenges facing the intelligent automation of the BOL process. After all, if Google Vison can tell a Welsh Corgi from a Welsh Terrier, how hard could it be to extract the information from a BOL? Not so fast. While image recognition and optical character recognition are powerful technologies, a massive amount of intelligence and processing power and dozens, if not hundreds, of sophisticated machine learning algorithms are required to produce optimal results.

Why? Because BOL's in the real world are not consistent or simple. Trucking companies deal with BOL's in many different formats and flavors. The same information such as "Ship To" will be located in different locations on a form and blocks of content can have different titles. One form might have the label "Ship To" and another will refer to the same information as "To: Consignee".

Some forms will prepopulate lines with the words "Name:", "Address:", "City/State/Zip:" and others will leave the fields blank. Some forms might print the "Ship To" label vertically. While reading vertically is intuitive for a person, it creates significant challenges for OCR.

Additionally, forms can have stamps applied over top of key information or contain handwritten information. And, scanners can create lines or spots which make character recognition difficult or a trucker might fold the BOL, creating a black line on the scanned document.

Another example of an inconsistency is when a document is scanned at an angle instead of being vertical. While it is easy for the human mind intuitively to adjust to this irregularity, it is not so simple for the AI-system which uses the "x-y coordinates" of each character to determine which characters and words should be assigned to a particular line in a content block. In these situations where ambiguity is introduced, a higher level of Artificial Intelligence is required in order to make inferences and judgments that will accurately produce the correct outcome.

07. A Look at How OCR and AI Work in the BOL Process

Now, let's look into how OCR and AI are used in the Bill of Lading process. First, we'll explore how humans perform the task and then we'll dig into how the AI-based solution does it. Simply put, the need is to read content from the blocks of information from the scanned BOL (below left), extract it and input it into the corresponding fields in the company's information system (below right).

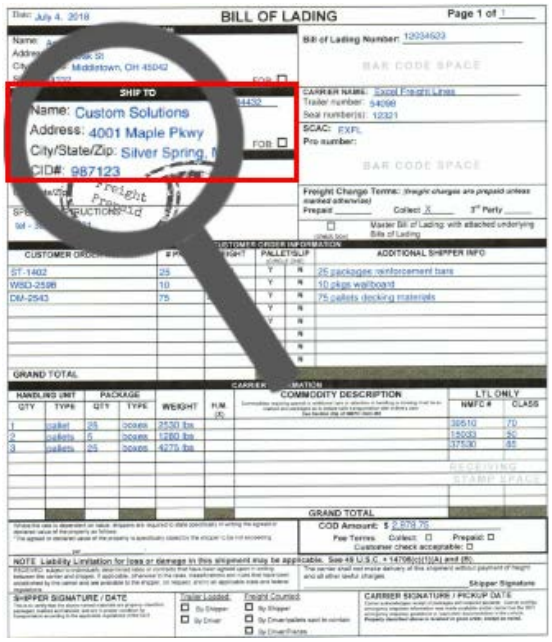

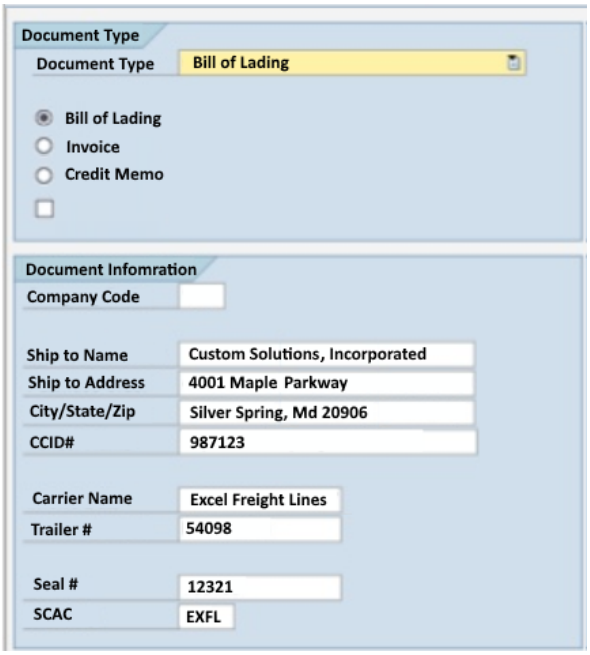

08. Human-Centric Process: A Tops-Down Approach

In the manual, human-centric process, a person can instantly and intuitively use the powerful pattern recognition inherent in the human mind to identify blocks of information. For example, see the "SHIP TO" information block in the Scanned Bill of Lading above. We easily see the "SHIP TO" label and intuitively recognize all the content associated with that block, including the company name, address, city/state/zip and CID#. Then, we enter it into the appropriate fields in the company information system.

If there are any errors, abbreviations or inconsistencies, a skilled data entry person can typically recognize the issue and correct it. For example, the name "Custom Solutions" is easily recognized as an alternative name for "Custom Solutions, Incorporated" in the Master Database and the term "Pkwy" is intuitively understood as an abbreviation for "Parkway". Furthermore, if there are scanning issues or stamps, or handwritten entries that make parts of the document difficult to read, the person can often use reasoning and judgment to determine the proper information. In addition, a company might have a look up table so that as soon as the company name is entered (e.g. Custom Solutions) the rest of the information (e.g. 4001 Maple Pkwy, Silver Spring, Md 20906, 987123) can be auto-populated. After completing the "SHIP TO" information, the data entry person will continue to the next blocks of information and repeat the process.

09. AI-Centric Process: A Bottoms-Up Approach

In our BOL case study, the automation starts by using Google Cloud Vision and OCR which recognizes each individual character in the document at an "atomistic" level. In contrast to the human-centric approach, the AI-centric method is more bottoms-up where we start by identifying characters and build up to words, paragraphs and finally into blocks of information, as shown to the right. The Google Vision character recognition and AI is quite accurate, making this a reliable methodology.

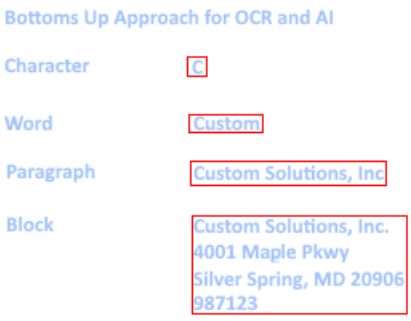

10. A Deeper Dive into Intelligent Process Automation & Named Entity Recognition

First a bit of terminology. Named Entity Recognition (NER) refers to the process of identifying, extracting, organizing and storing data such as names of people, organizations, addresses, quantities, monetary values and descriptions. As such, NER is a good description for what we are accomplishing in intelligent process automation. While it is beyond the scope of this paper to explain many of the terms associated with NER, and their nuances, it is helpful to explain a few terms such as tokens, entities and named entities.

Named Entity Recognition refers to the process of identifying, extracting, organizing and storing data.

Let's looks at these two strands of information as an example:

Custom Solutions, Inc.
4001 Maple Pkwy

Token: A single grouping of characters, either alpha or numeric. "Custom", "Solutions," "Inc." and "4001" are each referred to as a Token. While some people may refer to two adjacent tokens such as "Custom Solutions" a two-token company name, we will refer to tokens as single entities. It is helpful to note that in certain cases, specific characters such as a semi-color can be identified as a token.

Entity: An Entity may consist of one or more tokens. As such, "Custom", "Custom Solutions" and "Custom Solutions, Inc." are all entities, but consisting of different numbers of tokens. Likewise, "4001 Maple Pkwy" is also an entity and consists of both numeric and alpha tokens. Importantly, at this point, each of these entities is unclassified and is simply a generic word or set of numbers.

Named Entity: When an Entity gets classified as a specific type of information such as a name, company name, address, quantity, or monetary value it becomes a Named Entity. Once the entity is named it is much easier to assign it to the proper information block and associate it with the other information associated with the same block. For example, a 5-digit zip code can be classified as a zip code which would then allow the algorithm to associate it with a "SHIP TO" or "SHIP FROM" information block.

NER is focused on tagging words or word groups that correspond to predefined categories such as the ones found in forms such as invoices or BOLs.

11. Four Technologies that Aid Automating the BOL Process

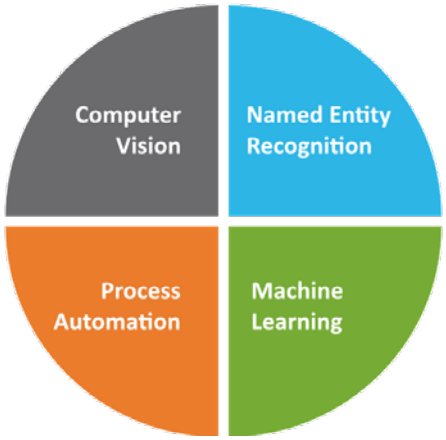

A robust AI-solution is built using multiple elements, each providing important capabilities. Let's take a look at four key elements and the value that each provides in the Automated BOL process:

- Computer Vision
- Named Entity Recognition
- Machine Learning
- Process Automation

12. Computer Vision

Computer vision is the first step in converting the image into digital output. Simply put, a scanned document is just dots of ink on a page and it has no real information. There is no concept of characters, words, lines, spaces, returns, or content blocks.

OCR output recognizes the dots of ink on the page and gives them (x, y) locations. Using this data, Google Vision assembles the dots into individual characters and then, based on proximity, into tokens or entities. When sufficient space is identified between characters, a new token is created. As shown below, the red boxes identify where tokens have been identified by the OCR. Upon close inspection we can see that the OCR does a pretty good job of identifying discreet individual words. However, it's important to appreciate that it still has no concept of lines or blocks. In essence, what we have is a bunch of disconnected entities that have no structure or contextual relationship to one another.

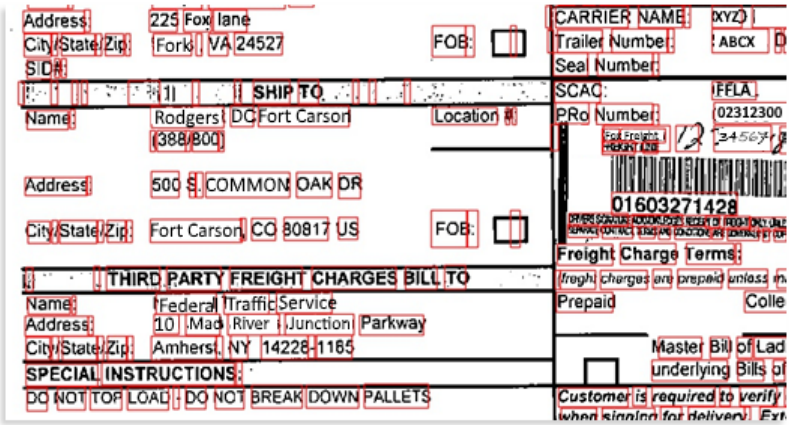

13. Named Entity Recognition & AI

The next step is to take these entities, classify them in to relevant categories and begin building them into blocks of information. This starts by identifying the required information out of the unstructured, non-positional text and using AI and machine learning to "teach" the computer to think like a person.

High value, common tokens are automatically identified from historical bills of lading. For example, the phrases "Bill of Lading", "Bill of Lading Number", "Ship To", "Destination", etc. can all be thought of as common terms that are found in a large number of the historical bills. The AI is trained to identify unknown objects (e.g. 02312300) based on the location to all of the common tokens found on the bill. The machine may learn that when the token is to the right of tokens such as "PRO Number" it's likely the PRO Number.

The machine is able to automatically learn as new shippers introduce new forms.

Rather than the time-consuming process of manually defining these rules for each of the thousands of formats, the machine is able to automatically learn them from examples, and more importantly continue to learn them as new shippers continue to introduce slightly new forms into the system.

14. Process Automation

OCR and AI are powerful enabling technologies; however, they represent only part of a complete AI-based solution. Maximum value is added only when the Intelligent Process Automation fits into company's workflow and information systems. As such, Core Compete works holistically to fully integrate the AI-based solution into the business processes and systems.

Furthermore, Core Compete understands the mission-critical nature of the BOL process and deploys systems that are designed and operated to be highly available and fault tolerant. With our 24x7 operations team, the solution is designed to take full advantage of the Google Cloud fault tolerance, with load balancing and multi-data center failover. This is because the BOL process must be able to scale seamlessly as demand increases and decreases as well as to be resilient enough to withstand the unexpected such as the loss of one or more computing resources.

Maximum value is only added when the Intelligent Process Automation fits into the company's workflow.

15. Cloud Architecture for BOL Scanning

A diagram for the cloud architecture for an in-market BOL solution is shown below. In our freight line case study, the solution used the Google Cloud Platform with Google Vision, but other cloud-based image recognition and classification solutions could be used as well.

Importantly, given the mission-critical nature of the BOL process and fact that during peak demand 60,000 BOL's flooded the system, high availability and auto-scaling were essential to the system architecture. Cloud Load Balancing is able to react instantaneously to changes in users, traffic, network, backend health and other related conditions in order to manage traffic. Seamless Autoscaling ensures that the system can easily handle huge, instantaneous spikes, enabling it to scale from zero to full throttle in a matter of seconds. In summary, this provides peace of mind that the solution is always available and ready to handle even the heaviest level of demand.

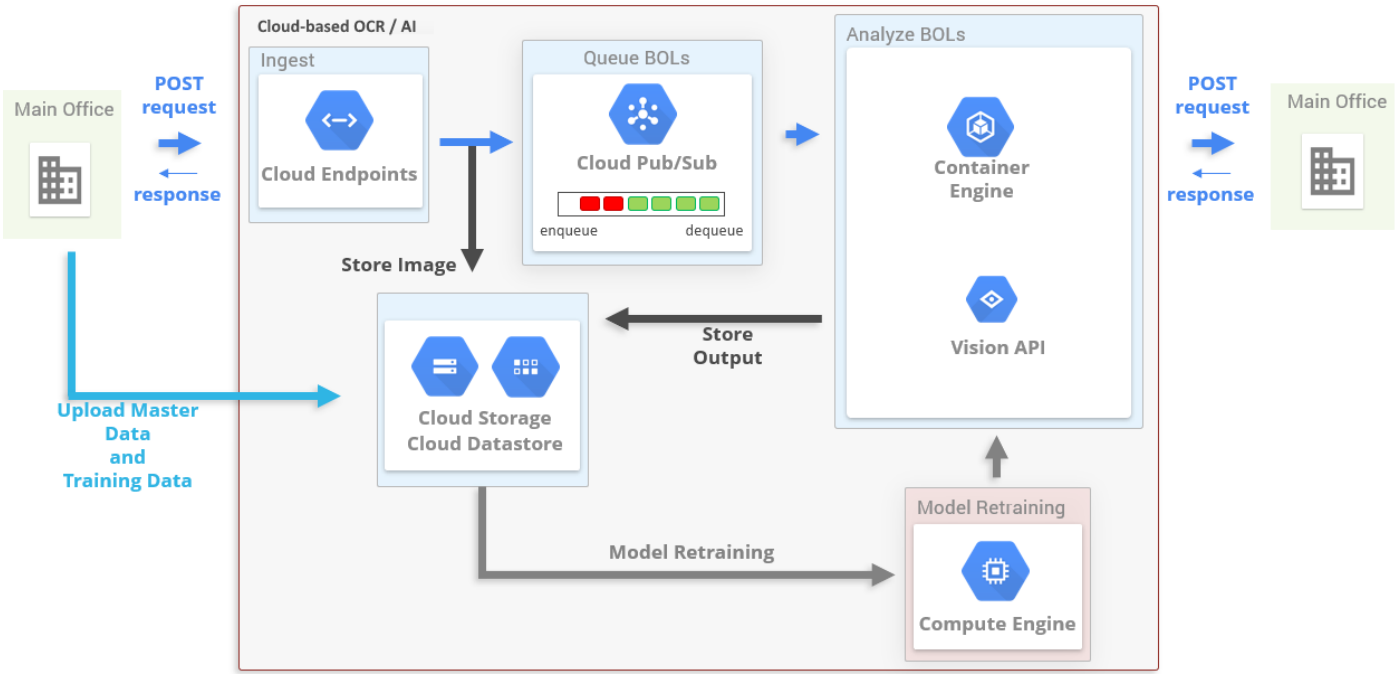

16. Benefits of Intelligent Process Automation

For this BOL process and well as other processes dealing with large quantities of paperwork, significant benefits can be realized using intelligent process automation.

- Reduced data entry headcount.
- Shift headcount to more value-added activities in the organization.
- Automated system scales elastically to support growth
- Lower error rates

17. Final Thoughts and Recommended Approach

Artificial Intelligence is advancing rapidly in both the consumer and commercial space. Many companies are already seeing significant operational benefits when using AI, image recognition and machine learning. As we have reviewed, these technologies can be deployed as part of a fully integrated intelligent process automation system in order to streamline inefficient, cumbersome processes such as bills of lading and invoicing.

The approach that Core Compete recommends is to start by identifying an important process where AI can add value. Then, define a specific project where the business impact can be realized in weeks or quarters instead of years. Test, validate and launch the first initiative. Upon successful delivery, you can take the data and technology assets and deploy them in the next initiative. With this approach, companies need not wait to get started deriving business value from AI solutions. They are able to think big and start small, stand up a new solution, derive value and then scale to the next initiative.

How to Get Started

Define a specific project where the business impact can be realized in weeks...

Contact Us

For more information,
email Core Compete at
info@corecompete.com