

Driving Global AP Efficiency with Smart Capture

Accelerating the Digital Transformation of Invoice
Processing in the Global Economy





Table of Contents

The Dark Ages 1

- Challenge #1 - Paper, Paper, Paper
- Challenge #2 - Digitization but not Transformation
- Challenge #3 - Invoices are like Snowflakes
- Coming into the Light

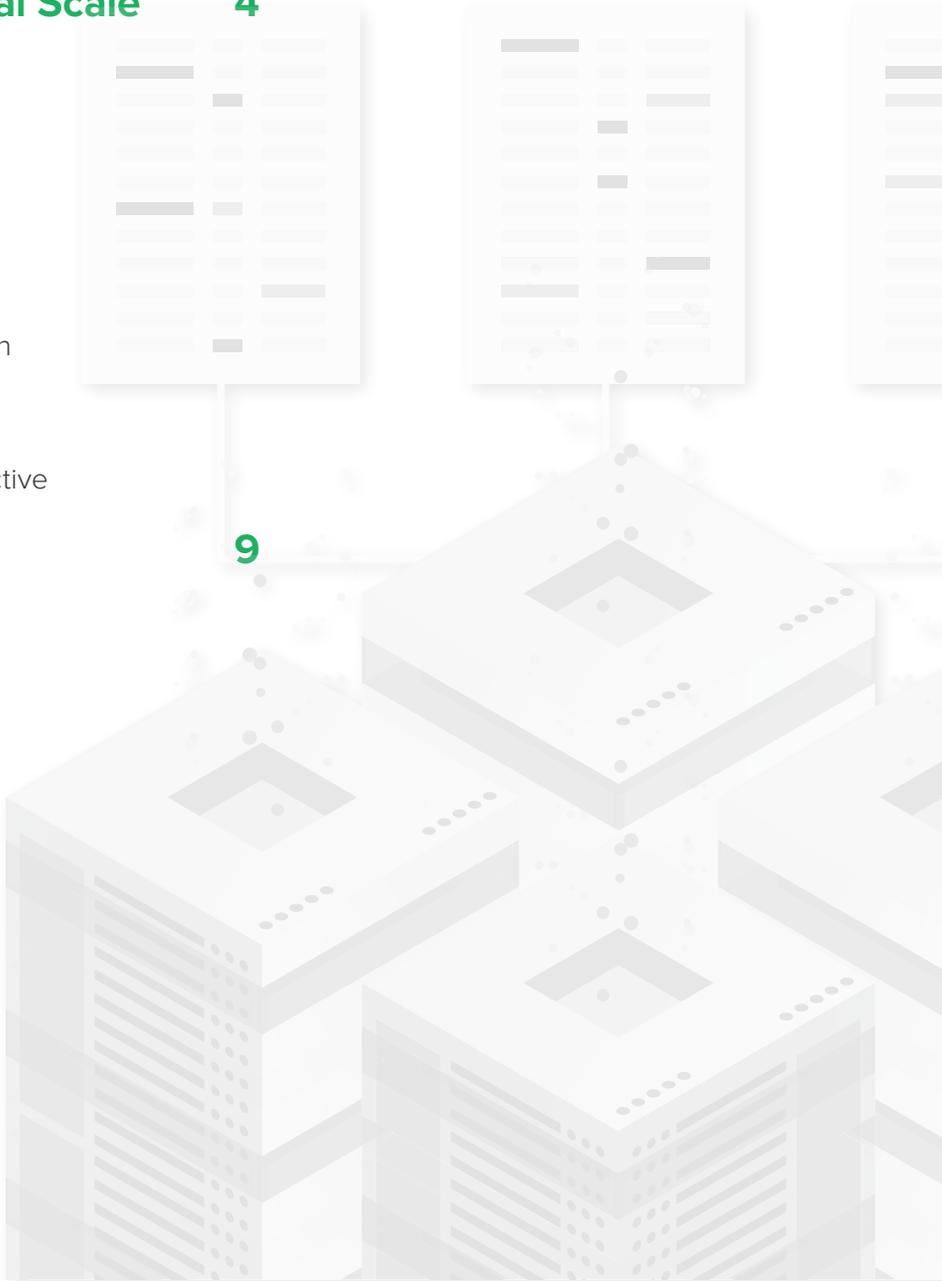
Invoice Capture with Smart Capture® 3

- Capturing the Smart Way
- Pre-configured Intelligence
- Supervised Machine Learning

Invoice Recognition on a Global Scale 4

- There is an app for that... sort of...
- Invoices and Languages
- Language Does Not Equal Country
- Invoice Separation (Classification)
- General Validation
- Validation of Tax Amounts
- Line Item Extraction and Table Recognition
- Invoice Recognition
- Integration into ERP, BPM and RPA
- Invoice Recognition from a GDPR Perspective

What is the ROI? 9



The Dark Ages

Managing the payment of financial obligations is a process as old as business itself. All successful organizations must pay their bills on time and efficiently in order to “keep the lights on” and maintain their financial health.

Despite this fact, Accounts Payable, in common with many other core back office business functions, has not seen any significant transformations until recently. The accounts payable process still fairly antiquated compared to the advances in other line of business, ERP, CRM or RPA systems.

That’s not really surprising – any process where you have a highly diverse and widespread set of constituents represents a challenge. There are a number significant challenges that need to be addressed:

Challenge #1 - Paper, Paper, Paper

As much as we would like it to disappear, paper is not going away. Yes, it’s true, while we don’t send as much snail mail any more, we still have printers, buy paper and invoices continue to exist in paper form.

The problem with paper is that it is an analog format. The information contained within it is not easily accessible or searchable. Paper files are only typically accessible via one or two fields. For example, looking for Invoice #1002 from Acme Widgets Inc., one would possibly look in the Acme Widgets Inc. file and search the enclosed invoices (arranged in Invoice # order) until we come to the invoice we are looking for.

Challenge #2 - Digitization but not Transformation

Enterprise Content Management (ECM) or Content Services can offer some help here. In fact, since the early Document Imaging systems first arrived as far back as the 1980s, there have been improvements. Capturing an electronic image of the invoice and indexing using a few key fields is certainly an improvement.

However, this merely digitizes an image of the document, typically in a PDF format. What it fails to do is to digitize the business information that the vendor invoice contains in such a way as to enable the data to be passed electronically to be acted upon by application responsible for the approval and ultimate payment of the invoices.



Challenge #3 - Invoices are like Snowflakes

Optical Character Recognition (OCR) technology has been around since the beginning of document capture (and earlier) but it is not enough to read each character on the page. It is necessary to understand the context of what is on the page. What does the invoice mean?

OCR does a pretty good job at processing fixed forms, i.e. identical forms where data fields are located in exactly the same location on the form. Creating a rigid template that looks for text in a predetermined location and assigns it to a specific data field. This capability can deliver huge benefits in productivity and cost reduction by automating time consuming and expensive manual data entry.

The problem here is that invoices are not fixed forms. Like snowflakes, while they follow certain structured principles, they are as unique as the vendors they come from. While they all contain the same underlying information, the manner in which it is presented is determined by the vendor and not the purchaser.

For a large company dealing with hundreds or thousands of vendors, it is impractical to creating hundreds or thousands of templates to accurately extract the data from all the invoices received.

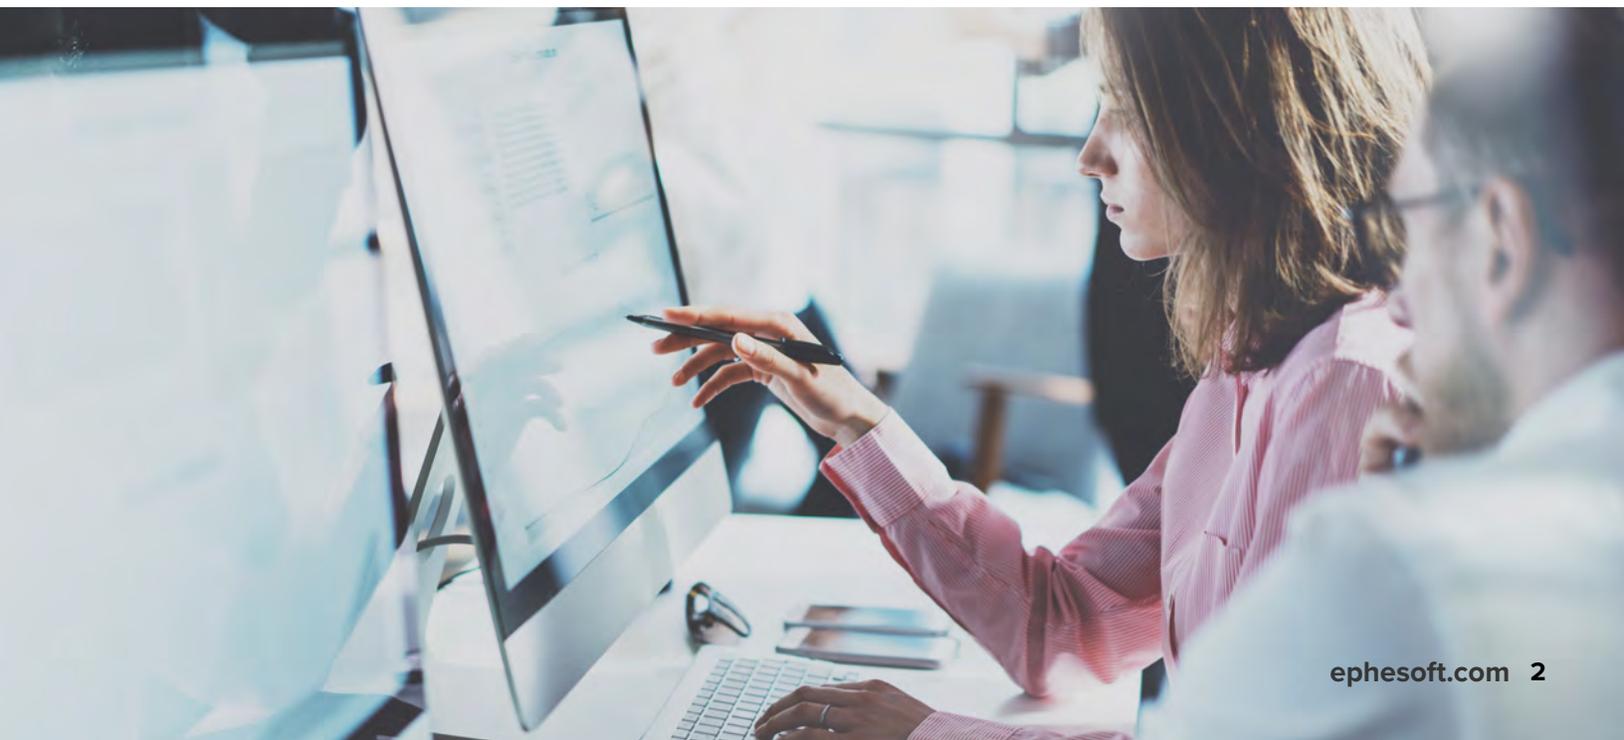
Coming into the Light

What is required is the ability to not only capture documents but to understand them. This understanding is necessary on a number of levels.

First, the system should be able to recognize the type of document that is being processed. Invoices are not the not to the only type of document that could be received. Unless you want to manually sort all the documents by type before processing, it makes sense to have a system that can automatically recognize documents and classify them into different types.

Once a document has been classified, the system knows what information it is looking for. The system should then be able to seek out the desired information regardless of the document layout and where certain information is located in it.

That is where Smart Capture® comes in to help organizations support smart decisions with smart data.



Invoice Capture with Smart Capture®

Processing a document with OCR, gets you a “bucket full of words” and, if you are lucky, maybe some information as to how these words are laid out on the page. What is required is some intelligence in understanding what these words represent in the content of the type of document, such as a vendor invoice, that is being processed.

Capturing the Smart Way

Once the document has been classified as a vendor invoice, a Smart Capture® platform can seek out the desired data from the text that has been extracted.

Artificial Intelligence (AI), and specifically supervised machine learning, is a critical technology in achieving this goal. Supervised machine learning can be used to “train” the system by using sample invoices and “teaching” the system where the critical data resides.

Pre-configured Intelligence

Certain base level intelligence can be pre-configured ahead of time relating to specific solutions such as Accounts Payable. This would include predefined labels for the various fields contained on an invoice. Synonyms for these fields would also be necessary as different vendors may use different labels for the same data, i.e. “Invoice No:” or “Invoice #”.

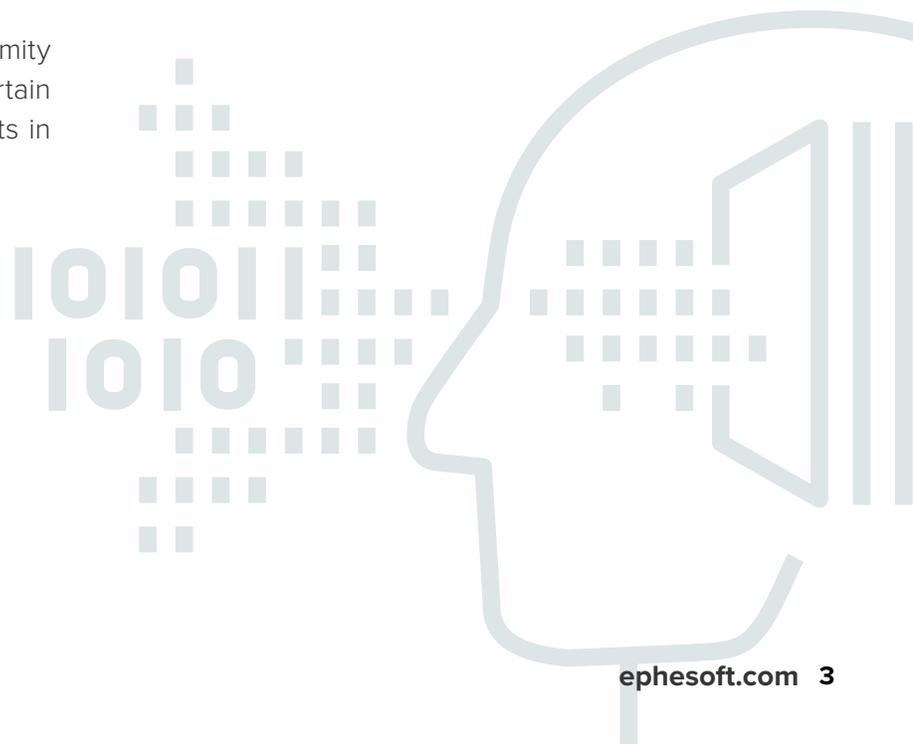
The desired data is located by its relative proximity to these labels or “anchors” and meets certain predefined patterns, such as various date formats in the case of the Invoice Date field.

Supervised Machine Learning

Pre-configured intelligence will only take you so far. Supervised machine learning takes it to the next level. As the system processes vendor invoices, it recognizes the various data fields automatically and forwards documents to users when additional assistance is required.

When users review the document and manually identify and select the necessary data fields, this information is used to train the system further, increasing the level of intelligence.

For example, when a new vendor’s invoice is seen for the first time it may not be processed entirely autonomously. Once referred to the user the information is identified, located and validated. Once the user has validated the information, the system intelligence is updated and subsequent invoices are processed with minimal user intervention.



Invoice Recognition on a Global Scale

As a developer of content capture and data discovery software, vendor invoice recognition is a core use case for many customers around the world. Moving across continents and country lines, a broader international scope of digital transformations necessitates more complex deployments that accommodate multiple languages and financial markets.

This is very common when working with larger companies with international supply chains. These companies are often surprised by how complex such a project can become especially when automating processes for larger companies.

There is an app for that... sort of...

In our app-based economy there seems to be an app for everything and if there is no app, there is a Web Service built by a smart startup providing invoice recognition, which makes the process look like a piece of cake. Just upload the invoice and it will be read ... sort of.

When we look at the recognition of invoices in SME and large global companies the process is a bit more complicated than most of these apps or easy to use Web Services can address. The key consideration customers have when investing in a product like Ephesoft Transact, is the desire to get a quick and accurate return.

This is only warranted if a high degree of automation within their use case can be achieved. This means we go far beyond the classic fields, often addressed by these simple services - beyond Invoice Number, Invoice Date, PO Number and Total.

To achieve a proper return on a vendor invoice automation project, the scope of it goes beyond a few fields, and will address:

- Different languages in a mixed batch without prior sorting
- Purchase Order (PO) matching
- Invoice separation as part of classification and automatic splitting
- Line item extraction
- Tax related fields and their validation
- Validation of an invoice related to its origin and destination (matching the legal entity)
- Table recognition and vendor matching

Beyond these tasks, vendor invoice projects are usually the starting point. Benefits may also be realized from automating the processing of credit notes and customer invoices. Such automation should include the key elements such as:

- Proof of delivery
- Customer invoices
- Delivery notes
- Bills of lading

Finally, for international projects, custom declarations and similar documents need to be included. The recent rise of GDPR in Europe and similar compliance policies elsewhere have resulted in additional process complexity being added to many projects.

Let's examine and discuss many of the above elements and provide an overview of what is possible and potential pitfalls.



Invoices and Languages

In an increasingly globalized market, it is decreasingly likely that documents will only exist in a single language – at least when doing projects in larger companies with international sourcing of products. Therefore, a first key element of the project will be to determine the language of the document, since it is required in order to:

- Identify keywords to look for; Invoice Number vs. Rechnungsnummer (German) vs. Numéro de Facture (French) vs. Numer Faktury (Polish) vs. номер на фактура (Bulgaria) vs. Fakturanummer (Swedish) vs. 发票号码 (Chinese) etc.
- Determine the origin of an invoice, which will impact:
- Splitting (you might have to keep trailing documents linked to the invoice, such as a customs declaration or documentation for hazardous goods)
- Origin (country) will also impact automatic validation of tax amounts (in European VAT)

The first example of the word “invoice number” represents the challenge, since when receiving invoices from different countries the recognition engine needs to be set up to look for variations of the words, such as invoice numbers, inv-number, number, but also reference possible synonyms, ref-nr., ref-no, our number etc.

There are infinite variations of the word invoice number alone and those will have to be handled for each language and field type. Therefore, handling of a language goes beyond the pure capability of the OCR language, but requires careful setup of the system in relation to the use case.

Language Does Not Equal Country

Language not only plays an important role for the extraction of values, but also for the validation. A key topic which has to be handled in international vendor invoice projects is the validation of tax amounts or the determination of “completeness of an invoice”. All these elements require an understanding of the origin of the invoice and which legal entity it needs to be matched against.

There are engines available which can identify the language of a document/text, but is it safe to say that every English invoice originates from the UK or the US or Canada? In fact, every country has different rules and regulations, and, therefore the language, while being important to read an invoice, is not enough.

An invoice recognition project will also have to safely identify the country of origin for an invoice, which can be done, for example, by reading the originating tax ID, finding the vendor address and matching it against the ERP record. However, it is a bit more complex than just using the language. Another good example would be Belgium or Switzerland where a single invoice can hold various languages within a single document.



Invoice Separation (Classification)

If only a few invoices need to be processed, it is easy to scan every invoice on its own or to separate them, with barcodes affixed to the first page. While this provides for stable and solid results when it comes to separation, this will add significant manual work to the process, since those pages will have to be added and removed later on.

Using Ephesoft Transact, we leverage our supervised machine learning based classification and rules to separate invoices which are scanned in a batch, without the need to prepare them. Based upon content, the system is able to identify the first, middle and last page of a document and therefore securely separate invoices.

A key element here is the combination of automatic learning, using supervised machine learning and rules. After several documents to train the system, patented supervised machine learning tools can securely separate documents or recognize that there might be a need to keep certain documents together. An example might be an invoice which has a trailing customs declaration, which outlines deductions of taxes on the invoice itself.

Since a customs declaration isn't an invoice, a supervised machine learning based approach would split the two documents, and by adding rules, we will be able to control splitting based on country or currency of origin and therefore make sure audits can be ran on complete documents.

General Validation

Once data has been extracted obviously validation needs to happen, ensuring the accuracy of data. This validation can range from:

Simple: Quantity x Price = Total

TO

Complex: Compare price of ordered goods using the referenced PO number with the ERP system, to make sure there are no variances

All of these validation steps help drive the accuracy of data and reduce manual labour.

Validation of Tax Amounts

For obvious reasons every country has its own tax regulations for VAT and other commercial taxes. When reading an invoice, the system must have the capability to check if the tax amounts match. This is done by finding:

- a. Total / Gross amounts
- b. Net Amount(s)
- c. The deducted/calculated amounts in relation to the tax rate(s)

The last point (c), is an especially complex one since there is usually not just one tax rate for VAT in a country but various rates. In Germany you have 19% (full), 7% (reduced e.g. for food) and 0% if not taxable, which results into 3 tax rates. While for most companies 19% or 0% will be applicable, but, for example, in the food and beverage industry or retail, 7% plays an important role. So, German invoices have 3 amounts to be handled, while in Poland 4 amounts exist.

These amounts might be listed in a separate table or as part of a list of all the totals. Identifying these amounts with high accuracy is a key element for invoice projects and consumes quite some time in the delivery.

To complicate things, it might be required to read those different amounts as part of table extraction as well as the line items and compare those results against the totals. This is all possible, but will impact the complexity of the project.



Line Item Extraction and Table Recognition

Another key aspect is line item extraction and line item matching. It may sound obvious that when reading an invoice header and footer, it is also important to read the table details. However, you might already realize by now, that header and footer information, for international invoice projects, comes with a certain level of complexity – so do tables.

Before going on to explain what it takes to recognize tables, the question will be to determine if table extraction is really required. Many processes especially that are purchase order based, will focus on the total amounts and ensuring that they match the ordered goods in the ERP system. If this is the case, table extraction can be ignored, which will result in faster and more cost effective delivery.

If, however, line items are required for the process, especially if partial deliveries are common or accuracy is required on an item level, then there will be a need to extract the tables. For this part, there are two drivers for complexity:

- a. Language: Since the table/column header needs to be identified, and, like invoice headers and footers, there is a lot creativity how to name things in each language.
- b. Table and content normalization: This refers to the process needed to make the extracted information available to the ERP system. Table layouts and data differ in many ways, and therefore data needs to be normalized prior to making it available to the ERP system.

While Ephesoft Transact is very flexible in dealing with layouts, there may be a requirement to customize the system towards the provided invoice layouts.

Many people think, that one of the most complex elements would be to handle tables spanning over multiple pages or to deal with more than one table within a document. Either scenario can easily be handled with Ephesoft Transact.

Invoice Recognition

Many people might wonder if there is still a market or case for invoice recognition in 2019, considering that digitalization is a prominent initiative around the world. For the time being, invoices remain image based, while the transmission will shift from paper-based to electronic, via email.

The benefits of having an e-text layer within PDF files can be leveraged, which greatly helps increase the accuracy compared to the accuracy of OCRing an image, ultimately improving extraction results.

Electronic invoice interchange standards are clearly on the rise and are becoming standards. However they suffer from the many “standards” in e.g. in the EU, which are different on a country level. Therefore, for the time being, extracting data from the human readable presentation layer is the easiest approach, while a straight import of the raw data from the data layer of an invoice would be the most efficient.

With the rise of more and more electronic invoice formats, it is also key to understand the requirements of processing them. As an example, in Spain, electronic signatures need to be processed when reading an e-Invoice.

Another benefit of the image-based approach is that it does not stop with the invoice. As mentioned in the beginning, invoices are usually just the starting point for automating documents into an ERP system.



Integration into ERP, BPM and RPA

Invoice processing does not stop with reading the data. Recognition is just the starting point where the raw data is generated and is provided into an ERP, BPM, RPA or similar application platform. This is achieved through:

- Direct connections, such as into Infor ERP or SAP
- Web Services-based integrations (REST or Swagger) for UiPath, Blue Prism, Microsoft Flow, Nintex, K2, etc.
- File-based integrations (XML, CSV) for any other ERP, ECM, etc. process

The importance is that connectivity is provided via a standard method, so the solution is future-proof and can be expanded for additional use cases.

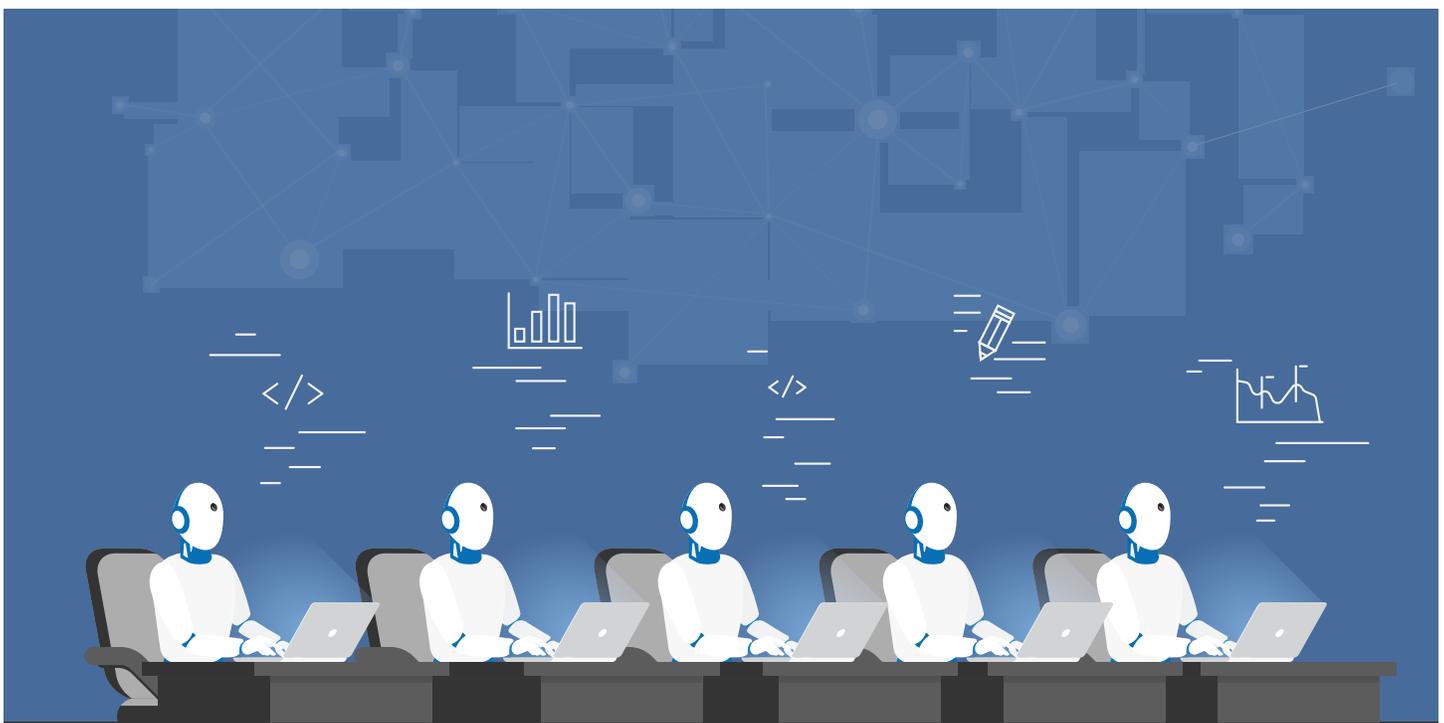
Another element which needs consideration when setting up connectivity to the ERP system, is the bi-directional communication, which is needed for data validation or vendor matching. The system to be used should provide methods via JDBC, ODBC or Web Services (SOAP/REST) allowing for easy integration. In the case of Ephesoft Transact, a fuzzy logic database is populated with data, allowing for vendor and line item matching directly as part of the validation process.

Invoice Recognition from a GDPR Perspective

GDPR and other consumer privacy rights acts have been a hot topic and will continue to be so. It is not limited to just invoices and commercial documents. While invoices might not be a direct subject for GDPR compliance, a delivery note to an end-customer might be. What should we look for?

Consider an individual being referenced in a document in a commercial context (i.e. the name of the person who placed the order, their direct phone extension or email address). While this might count as PII (personal identifiable information), the information that is provided in context to a commercial transaction may be interpreted differently and, therefore, should not be a problem from a GDPR perspective.

However if, for example, a customer is mentioned in relation to sensitive information (i.e. a wheelchair being order for Mr. John Doe in Big Town, High Street 123), this invoice should be flagged since it contains sensitive information. This illustrates that recognition of invoices or of commercial documents might also be extended to other use cases.



What is the ROI?

By now it should be clear that the recognition of vendor invoices can become quite complicated especially when seeing it from an international standpoint. There will always be apps providing for simple single language recognition handling very low volumes. These apps provide for quick results and a faster (yet ultimately smaller) ROI, but they are not meant for high-volume enterprise companies.

To understand the full ROI potential, it needs to be understood how many fields are recognized and checked. In a standard project for invoices, there are 15 fields or more that need to be recognized. Why that many? Think about the classic fields like all the dates and reference numbers, potentially various taxes rates or net and gross amounts.

If you also check validity in context to e.g. German VAT deduction regulations, address, court registration number, VAT number or other fields, the data needs to be read and checked. As mentioned earlier, in Spain, there is also the check of the digital signature. Now

think of a multi-page invoice, where this information needs to be checked over multiple pages and it becomes evident that just entering these 15 fields will consume quite some time. It's proven that automation will lead to a quick ROI, usually in months, not years.

The use case for accurate vendor invoice recognition in an international context is far wider than just reading a few fields. It is about reading fields and variances in different languages, validating them or checking completeness in context to origin and destination and aiding the person with entering this information into an ERP process to do this in an efficient way.

A big driver for Smart Capture® is that with the insurmountable problem of Big Data combined with compliance, now information must be captured, classified and put in an usable format so decision-makers worldwide can make better and faster decisions. Customers who use Smart Capture® solutions see results that boost accuracy, productivity, efficiency and bottom line cost savings.

United States HQ

+1 (949) 335-5335
info@ephesoft.com

Ephesoft, Inc.
8707 Research Dr.
Irvine, CA 92618
United States

United Kingdom

+44 (0) 1184665000
info.eu@ephesoft.com

Australia

+61 2 9056 7490
info.au@ephesoft.com

Germany

+49 6126 5503510
info.eu@ephesoft.com

Italy

+39 (02) 8088 6345
info.it@ephesoft.com

France

+33 1 8288 4002
info.eu@ephesoft.com

Singapore

+65 3163 5499
info.asean@ephesoft.com