


Automating customs declaration processes: the trends, the challenges and how to overcome them

**How AI will be a personal assistant to
the customs declarant and vice versa**



Table of contents

A word on your challenge: automating customs processes	5
Solutions that you may have investigated. Spoiler: they'll let you down	6
Option 1: Request structured customs instructions	6
Option 2: Can OCR offer a solution?	7
Option 3: RPA, the magical add-on for OCR?	8
Option 4: IDP for the win?	9
Rounding up our research (so you don't have to)	12
Customs Declarations Automation done differently	13
How Customaite works to allow 2X faster customs declarations	14
Closing notes: a few guiding principles we use while building our solution	15



In this document, you will learn how companies like yours have attempted to automate (or at least optimize) the process of filing customs declarations.

Technology is always evolving, and we've seen many solutions applied in different sectors, one being more successful than the other.

In the last few years, the Customaite team investigated this domain and looked at the challenges from a totally different angle. Instead of striving for perfection in data extraction using AI algorithms, we start from the observation that these algorithms are inherently imperfect. How can AI and humans work together in symbiosis so that efficiency is maximized? **Can the AI be a personal assistant to the declarant who assists him/her throughout the entire declaration process, while the declarant himself/herself also helps the AI where it has doubts or even fails?**

Executive summary

In this report, we explore the challenges that arise when attempting to automate customs declaration processes. As technology advances, there have been various attempts to automate and optimize the process of filing customs declarations, with varying levels of success. We take a different approach to the problem: we propose that AI and humans can work together in symbiosis to maximize efficiency, with the AI acting as a personal assistant to the customs declarant throughout the entire declaration process, and the declarant helping the AI where it has doubts.

The main challenge faced by those working in logistics and customs is variability – there is no such thing as a standard when it comes to customs requests. Customers often send in their declarations requests in whatever format they want, and they are not always concerned with the quality of the data they provide. Despite this, it is essential to strive for 100% compliance with customs legislation.

In the past, orders mostly came in as paper documents, but today, most ‘document driven’ orders come in via email. These ‘digital’ documents are sometimes just scanned copies of paper documents, making it challenging to copy and paste data to avoid human error. Customs declarants are highly skilled at re-typing information from documents to the user interface of software, but the time required to transfer information can still be significant. Declarants may spend hours processing a single declaration and spend most of their time retrieving and processing information from a provided document set.

The challenge of manual labor could potentially be overcome by adding more people to teams, but international trade is growing so fast that the required workforce cannot be found, and many companies are struggling to hold on to their current workforce. Thus, the sector has been looking for solutions to overcome the burden of manual work that, from a distance, seems easy to automate or optimize. In this document, we’ll explore various options and evaluate which contexts they may be suitable for and where the challenge may soon become too great. Indeed, none of the solutions have proven to provide the magic bullet.

When attempting to find automation solutions in customs processes, the biggest challenge remains in the extraction of information from invoices, particularly invoice lines. We explored several off-the-shelf solutions that promised high quality levels of data extraction from invoices, but we quickly discovered the limitations of this kind of AI. Although some solutions perform well in more common supplier invoices, extracting information from international trade invoices in a customs context is a whole different challenge. Invoice lines, and especially elements specific to the sector, such as tariff codes, weights, and packaging codes, have proven to be the main challenge. Additionally, AI models have a slow learning curve and require numerous examples of similar invoice layouts before they can perform basic data extraction. The approach also requires people to process and label documents in large numbers before the AI can learn from them. Unfortunately, there is no guarantee of success with this method.

Through our research, we discovered that although there are many options available on the market, there is no ‘automagical’ solution to the problem. However, we saw an opportunity to approach the problem from a different perspective, leading to the development of Customaite. We’ll explore our solution in more detail later in the report.

A little word on your challenge: automating customs declaration

Main challenge: variability also known as “the customer is always right”

If you're working in logistics and especially in the customs domain, you'll know there's no such thing as a standard. Customers send orders in whatever form they want and are not very concerned with the data quality. They do expect a perfect outcome and would like to have their declaration filed yesterday.

No matter if you're a customs broker or working for an internal customs department in a company, quality is of the utmost importance, and you are striving to be 100% compliant with customs legislation.

Luckily, we are well behind the times when orders mostly came in as actual paper documents. Today we see that most 'document driven' orders come in via email. Unfortunately, in most cases these 'digital' documents are the result of a 'scan to pdf' from paper documents. A simple copy paste to speed up the process or to avoid human errors (typos) is often not an option. When seeing customs declarants at work, you'll easily acknowledge their immense skills in retyping information from documents to the user interface of some software (e.g. Microsoft Excel or a declaration software). Nevertheless, the time that is required to take over the information can still be significant. We've heard many testimonials of declarants who spend hours to process a single declaration and spend most of their time on retrieving and processing the information from a provided document set.

If manual labor was the only challenge, it could easily be overcome by adding more people to your teams. International trade, however, is growing so fast that the required workforce just can't be found, and many companies are already struggling to hold on to their current workforce.

No surprise the sector has been looking for solutions to overcome the burden of manual work that from a distance seems easy to automate or optimize. In the next chapter, we'll share some of your options and evaluate in which contexts they offer a solution and where the challenge might soon become too big.



Solutions that you may have investigated

Spoiler: they'll let you down.

A few insightful stats from a recent survey among freight forwarders and customs brokers

- 69%** still have **difficulty in finding qualified staff** and only partially or not at all succeed in doing so
- 64%** let their **customers submit their information** in a format of their choice
- 0%** said they were **successful in automating the entire process for all their customers with customer specific templates**
- 91%** have **not yet used AI solutions for automation** while 100% believed it was a promising technology

* Source: customer survey during a Customaite webinar

Option 1: Request structured customs instructions

In the field, we see attempts of structured 'customs instruction forms' in which the customer can provide the information. These forms are then processed by humans or by some other form of automation that we will discuss later.

When automating input that is based on the information that is provided by the customer, we see that it's quite hard to achieve a stable setup as you're still depending on a human that doesn't always do things the way you've intended it.

Compliance remains a risk you'll want to tackle. Even if the information was provided in a structured way, you'll need to check if it's consistent with the documents. The time that was won gets eaten away for a large part by this verification activity.

What to know

We've seen **various levels of success** with companies going this way. One can discuss if shifting the issue to your customer is what you want to do in terms of customer experience and if you do, we see that **customers find it hard to provide qualitative information consistently**.

Option 2: Can OCR offer a solution?

OCR or Optical Character Recognition is a technology that has been around for quite a while already. While this technology has evolved since it was first conceived, it is still not capable of understanding the text and numbers it has created.

When you've applied OCR to a document, you basically get something you can copy-paste in another application and start working with that result. All interpretation of the information itself still needs to be done by a human being. This means you still need to verify addresses of parties, look up tariff codes or use Excel to create pivot tables, do calculations, and copy the resulting tariff lines into your declaration software.

What time do you win? The time of retyping information. Not only is time saved because information does not have to be retyped, but the number of errors will also decrease as retyping information will inevitably lead to human errors.

What to know

Although **OCR** is without any doubt a part of the solution, **there is still a lot of manual work involved**. Not only in the copying and pasting itself, but also in searching through all the documents, verifying, and locating the right information that is needed at a given point in the declaration.

Option 3: RPA, the magical add-on for OCR?

RPA or Robotic Process Automation. This technology also goes back some time but has been maturing and gaining popularity in more recent years. It does exactly what it says. It's a robot that automates human tasks. We've seen this technology being applied with success and the robot even becoming part of the team. Employees would refer to James, Bob, Chives,... (or another name that works well in the team) as their loyal assistant that does the dirty work for them. James doesn't get sick, never takes a holiday and even continues to work at night.

James is very good at doing repetitive tasks, typically what most humans don't like to do that much. James doesn't make any mistakes. Situations in which James, or others of its kind, work very well are when there are standard structures available.

Think about a standard request form like filing a claim to an insurer. The form is quite straightforward and has a fixed structure. In practice, these forms are often implemented as pdf forms where you can type into, and when they are sent to the insurer, all James has to do is copy the information from known fields and paste them in some other software application in the field he was taught to do so. In our example, James would login and copy the information from the form in the software in which the claim must be processed. At the end, he can also send back an email to the customer to inform her/him that the claim was processed correctly.

James however is not as intelligent as his human colleagues. If an error is generated in the software he is working in, he relies on a human colleague to pick up the work from where he got stuck.

When we start talking about pdf's that are not structured and that are just images, James has no clue what to do. This is where the strength of OCR and RPA are combined to, as some will say, offer you the level of automation you're looking for. As before OCR is used to convert the image to text and numbers the computer can understand, James however still has no clue what to do with the text and numbers. What will help him is a template. Think of it as an overlay that you place over the text and numbers, the overlay knows what information is typically present in the various places of a document.



This works extremely well for documents that have a fixed layout. In the world of logistics and customs in particular, you could argue that the typical documents of shipping companies like arrival notices and bills of lading or customs instructions have a fixed layout, and that this technology could be used to optimize the process. However, as soon as we step away from the typical standardized documents, training James becomes a mission impossible. Just think about invoices and the infinite number of formats in which they occur. You'll be happy to find an invoice with a clear table structure, let alone the idea of making templates for every single invoice structure you encounter. Additionally, the slightest deviation in structure (e.g. unexpected whitespaces or newlines, an extra column,...) will make it impossible for James to be sure he's getting the correct information.

During our journey we rapidly learned that an approach with OCR, templates and RPA have led to many disappointments. Unfortunately, there is no such thing as magic...

What to know

RPA works extremely well for documents that have a fixed layout. In the world of logistics and customs in particular, you could argue that the typical documents of shipping companies like arrival notices and bills of lading or customs instructions have a fixed layout, and that this technology could be used to optimize the process. However as soon as we step away from the typical standardized documents, training the robot becomes a mission impossible.

Option 4: IDP for the win?

Now that we are on a roll, let's move to the next promising technology.

Yet another abbreviation: **IDP or Intelligent Document Processing** in full. Intelligent refers to the AI part, more precisely the use of special purpose, more advanced AI algorithms specifically aimed at dealing with text and language in general. While OCR is only able to convert images into text without really giving meaning to it, IDP goes a step further by also understanding and even interpreting the text related to the context in which the text appears. The domain within AI that deals with this is called Natural Language Processing (NLP). It combines linguistic principles with AI and Machine Learning and is therefore at the crossroads of both.

Where template-based solutions suffer from varying structure, IDP solutions offer more intelligent ways to extract 'contextualized' information without the need for a specific template.

Consider the following example.

1. Although 80% of all fish school at some point in their lives, there are a number of fish that live solitary lives.
2. Although 80% of all fish at our school died from starvation during the holidays, we still had plenty left to continue studying them.

NLP algorithms will be able to make the distinction between the meaning of the word school in the first and the second sentence. Without going into details, it will do this based on the context of the words surrounding the word school.

Likewise, NLP algorithms can identify and make the distinction between the name of a consignor, consignee versus the name of a vessel (boat) or the names of goods on an invoice without knowing the exact location of the text on the document, i.e without relying on a template.

In recent months great progress has been made in the field of NLP, one of the most imaginative developments is without a doubt the ChatGPT model (<https://openai.com/api/>) an NLP model that is able to understand very complex language tasks, to execute a wide variety of language-driven instructions and even generate its own language.

Consequently, the market has been flooded with vendors offering NLP services specifically aimed at processing documents. These services are often based on Deep Neural Networks, of which you probably already heard about. The idea is simple but its realization complex: the neural network is trained with a large set of examples and the more examples it sees, the better it gets (by analogy with how the human brain learns). To continue with the example above: when it has seen enough examples of the word school and the meaning of the word school in a certain context, it will eventually be able to make the distinction in meaning on its own. This, of course, also applies to the data that can be found on logistics documents such as names of consignors, consignees or products on invoices, arrival notices and so on...

Some vendors offer off-the-shelf services with pretrained neural networks running in the background, typically trained to do this so-called 'document extraction' on specific document types (e.g. commercial invoices) and specific document layouts (e.g. tables). Examples of such providers include Google's Document AI services and Nanonets' invoice, receipt, and table extraction services. Other vendors provide the ability to train your own neural networks tailored to your own needs (i.e. document types, layouts and information to be extracted). This involves the tedious task of labeling document examples on which the neural network can be trained: drawing boxes on documents and providing them with the labels corresponding to the information the box contains (e.g. consignor names, addresses, VAT numbers, invoice lines, dates,...). Examples of such self-service labeling and training platforms include Metamaze and Rossum. These platforms also offer pretrained models for specific sectors and document types that you can fine tune (this is called transfer learning) to your specific needs.

So, IDP must be without any doubt the holy grail in automated document processing for customs declarations? Unfortunately not, but as it was the case for OCR it is certainly a part of the solution.

What to know

IDP will perform better than trying to template every single invoice lay-out you encounter but it is still quite limited in delivering a high degree of automation when presented with a high variability. In general these solutions work well to identify and extract concepts like consignors, consignees, vessels, invoice dates, etc. from documents, BUT FAIL WHEN CONFRONTED WITH COMPLEX DOCUMENTS LIKE INVOICES AND PACKING LISTS.

The self-service training models offer the ability to add new concepts to be extracted, BUT REQUIRE THE TEDIOUS TASK OF LABELING DOCUMENTS, whereas the pretrained models are limited to what is provided without any possibility to add new concepts to be extracted.

So, IDP **isn't the holy grail in automated document processing** for customs declarations, **but it is certainly a part of the solution - as it was the case for OCR.**

Rounding up our research (so you don't have to)

The biggest challenge when trying to find automation solutions in customs processes remains in the extraction of information from invoices, and more specifically the invoice lines. When we conceived the concept of Customaite, we explored several off-the-shelf solutions that promised high quality levels of data extraction from invoices. Being a startup product, it would make sense to choose a buy over build approach if we could reach the results we were looking for and accelerate our time to market in that way.

Sadly, we were confronted with the boundaries of this kind of AI soon enough. Although there are solutions available that perform pretty well in more common supplier invoices, we learned that extracting information from international trade invoices in a customs context is a whole different cup of tea.

Moreover, the learning curve of the AI models is quite slow, and they need to be trained with many examples of similar invoice layouts before they can start doing some basic data extraction. Invoice lines and especially those elements that are specific for the sector (tariff codes, weights, packaging codes,...) have proven to be the main challenge. In addition, the approach requires that you need people to process and label documents in gigantic numbers before the AI can actually learn from it. And to make things worse there is no guarantee of success...

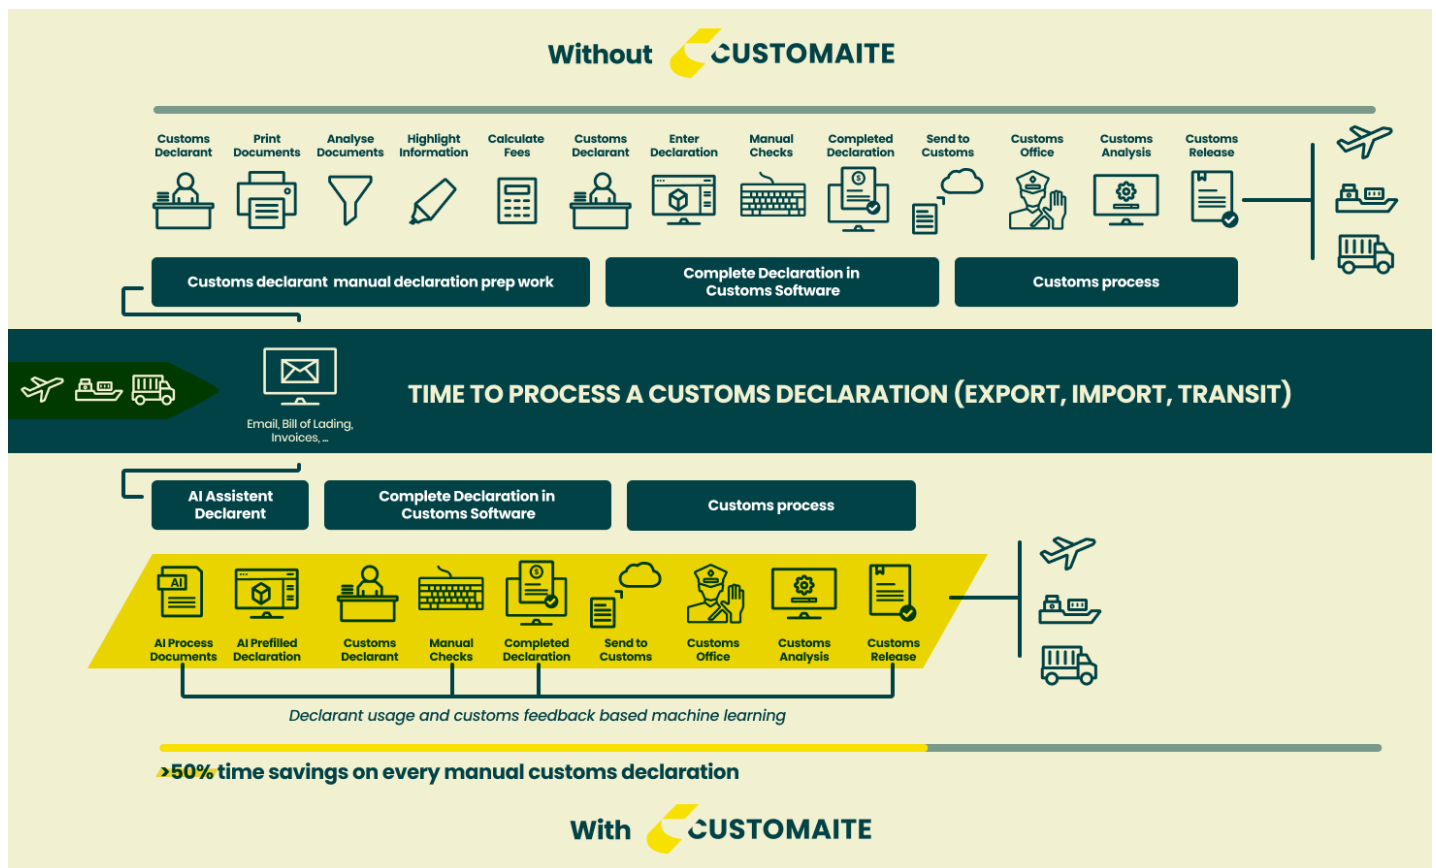
“From our market research into technologies that are available on the market, it became clear that there are quite some options but no ‘automagical’ solution existed as of yet. That’s when we saw an opportunity.”



Customs Declaration Automation Done Differently

At Customaite we take a different approach on the automation of customs declaration processes. The approach is based on our research, on what we have learned from conversations with customs managers, the observation of declarants as well as our own experiences, and in particular the various attempts and failures, in developing a solution to automate customs declaration processes.

How we guarantee 50% time savings for every customs declaration:



Step 1

Simply forward your documents to create a dossier

Managing dossiers from your inbox is a thing of the past. Simply forward documents directly to your virtual Customaite assistant via email or API. The message and its attachments will be analyzed by our artificial intelligence, saving you precious time.

Step 2

We pre-process your customs declarations

After uploading, we will analyze all documents and divide them into the correct types. Bills of lading, invoices, arrival notices and more are processed into clear and editable information.

Step 3

A fully automated proposal is drafted for sign-off

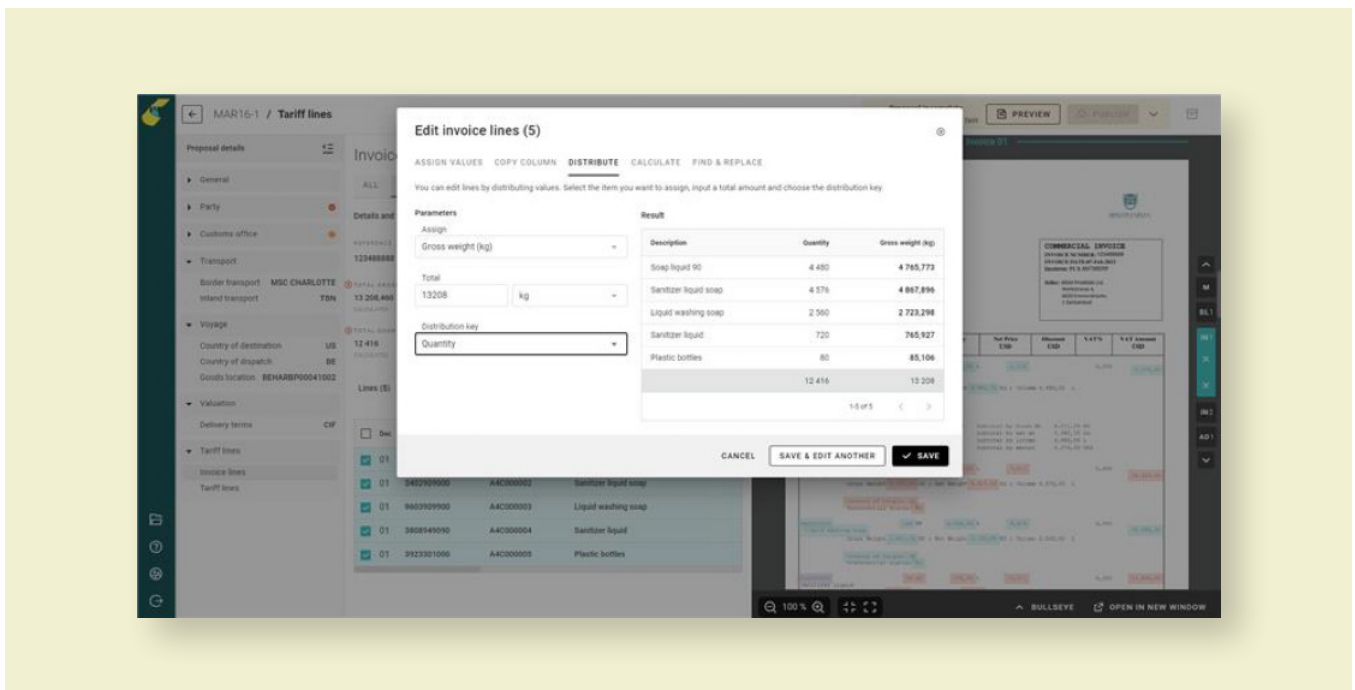
Customaite will draft a declaration proposal based on the uploaded documents. Additionally, it will take into account your historical data and other known sources (Tarbel, Vies, Vessels....) to complete the proposal. The customs agent remains in control of the dossiers at all times. **When everything is green, you're good to go.**

A few guiding principles

In what follows, we provide an overview of the key elements in our approach. We see them as guidelines that we keep in mind with every step we take and every decision we make.

1. Facilitating the declarant's work is the end game, IDP is just a means to that end

Focus on value and don't put all your effort in Intelligent Document Processing alone. Getting the correct information from the documents is of course very important, but there is much more to it when making a customs declaration. Just think of the operations needed to make tariff lines (e.g. distribution of weights, price calculations and aggregation of invoice lines) or the integrated provisioning of derived information, based on what is in the documents, but which must typically be looked up in separate sources (e.g. nationality of vessels or location codes of goods). But also, simple things like auto-completion, drop-down lists and business rules to complete information are part of this. At every step in the declaration process, we think about facilitating the declarant's work with a focus on value, so that declarations can be made faster and with higher quality.

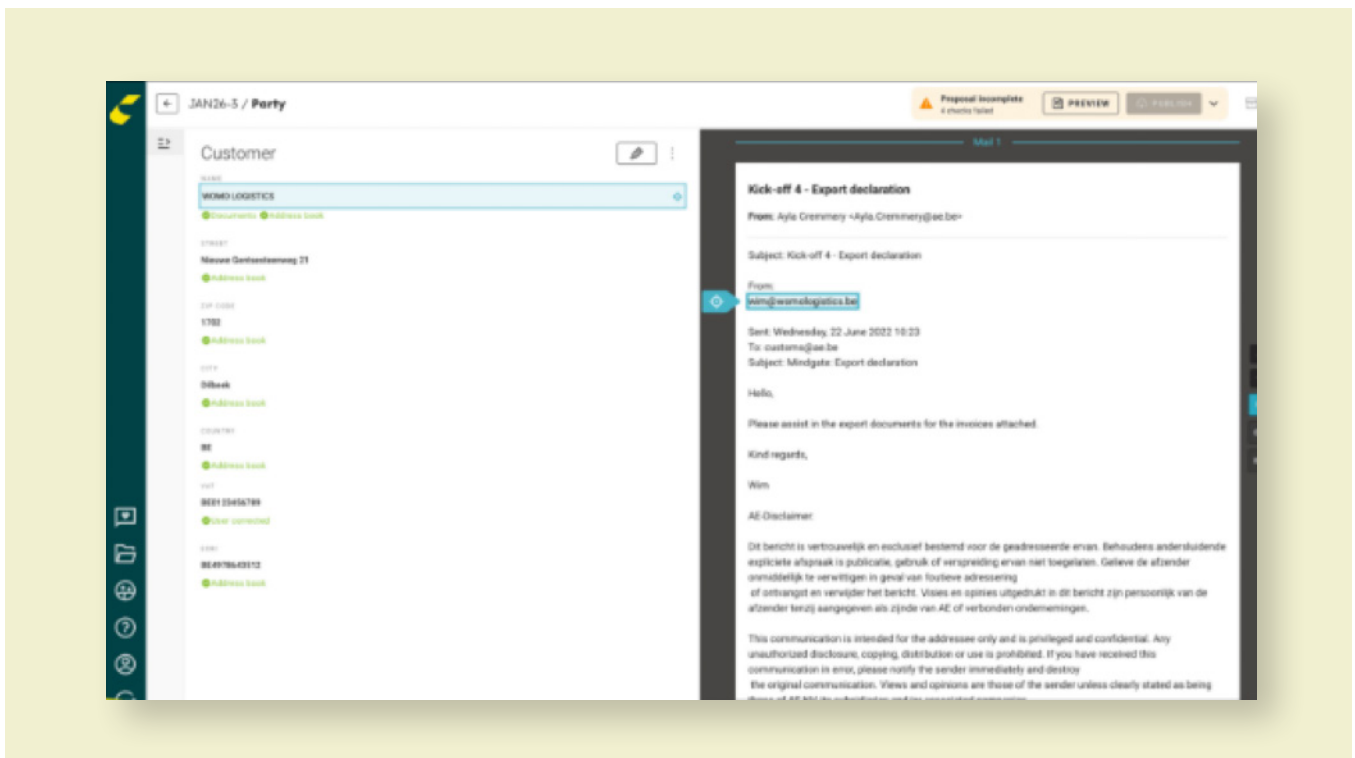


Smart edit functions allow to easily do different transformations and advanced calculations that are typically needed to process invoice lines to tariff lines such as distributing weights, gross to net conversions or find & replace values.

2. Design for imperfection

We're convinced it is better to accept and even embrace the fact that data extraction will never be 100% correct or complete. Striving for perfection in this area would mean putting in untold effort into those few extra percentages to achieve the completeness and quality of the data extraction.

The way to go is making it as easy and efficient as possible for a human to spot the imperfections from the AI algorithms and base his or her work on the preparation that was made by the AI. This can only work if you ensure that human and AI enter into a partnership and work together in symbiosis.



The sources of the information that is provided (on the left side) are always mentioned, making it easy for the declarant to verify it's correctness. In addition, if the information is extracted from a document it can be easily located in the document itself (on the right side) with a simple click of a button, making it even easier to verify it.

3. Enabling a partnership between a human and artificial intelligence

So, bringing the human in the loop is the next ingredient in our approach. We firmly believe in an approach that is inspired by the day-to-day tasks that customs declarants need to do. We believe that the way to lower the learning curve that comes with the adoption of new technologies lies in a custom-made user interface that only contains features and screens that contribute directly to the job to be done, creating a customs declaration as easy, correct and fast as possible.

We partnered with NDN, a reference independent customs broker in Belgium. Not only to enlarge our team's knowledge in customs practices but also in a user interface that makes sense to customs declarants. In our approach, we opted to digitize the full workplace of the customs declarant. The best example we can give is the prominent place that we gave to the documents (invoices, arrival notices,...) in our user interface. The declarant must be able to verify the outcome of the data extraction with the source that was used by the AI in the documents in the blink of an eye.

4. White box over black box

In advanced AI solutions, there is often no transparency about how an AI algorithm came to its results. Exposing the way of thinking of such an AI system is near to impossible and this can be a challenge to build trust from your end user.

At Customaite, we aim to be as transparent as possible about what the AI has done. In our UI, the user is taken to pieces of information in documents by a single mouse click. The information that is highlighted in the document is what was used by the AI to make a prediction for a value in the customs declaration.

Next to that, the user can assist the AI. When the AI is not confident enough about a prediction or has several options to choose from, the user can easily see the alternatives that the AI found. By assisting the AI and making a choice, the AI models learn from the user feedback so that next time a prediction can be made with more confidence and user intervention becomes obsolete.

5. Context is everything

The customs business is a very specific business domain that leaves little to no room for error. High quality in data extraction is therefore of the utmost importance. We are convinced that this can only be done by developing AI algorithms specifically for logistics documents in general and the customs context in particular. However, relying on data extraction solutions, even when they are trained on a customs context alone, is not enough to make the right predictions.

We have built a configurable application context around the AI to ensure that we can predict correct values, not only in a customs context but also a step further. Within the customs domain, different activities can take place like Transit, Export and Import declarations.

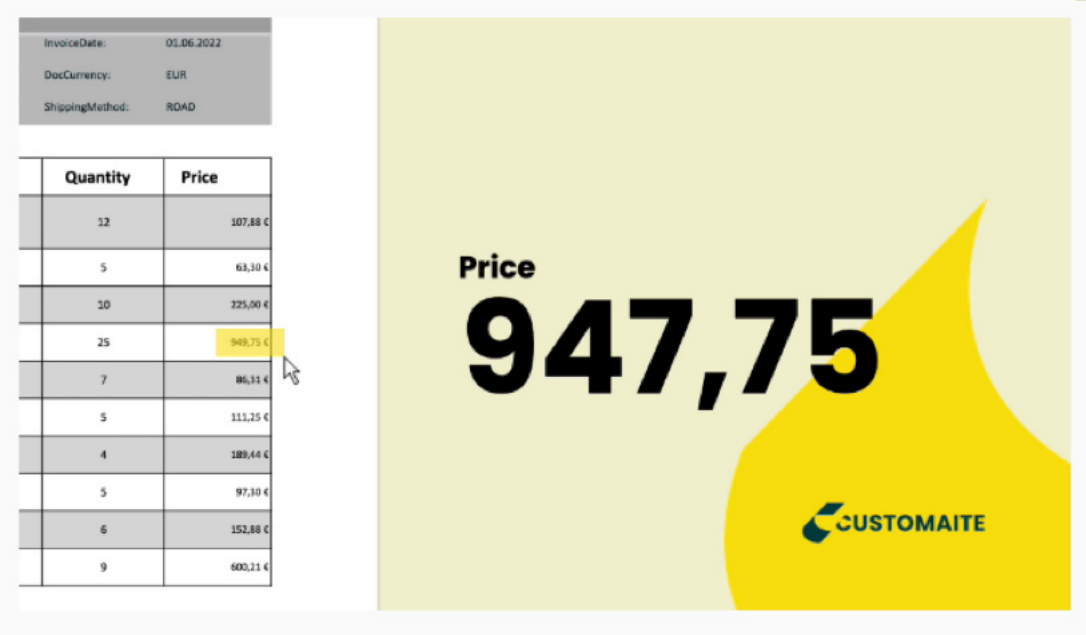
Within each activity, other rules apply, and the context can have a huge impact on the information that needs to be included in the declaration. The same applies to other elements of a declaration such as country of origin and/or country of destination, means of transport as well as the parties involved such as the consignor and the consignee. To realize the time savings we promise, we have chosen to focus on the customs business domain and to include all relevant business knowledge in our software.



6. Easy to implement in your existing way-of-working & business model

From the outset, we designed Customaite to blend into the current way of working of our customers, the customs teams. We are not into revolutionizing the way declarants work. Instead, we focus on efficiency gains with a minimal impact on people and processes. When looking at the process of building a customs declaration from a distance, you can distinguish 2 steps: work preparation and filing a declaration.

It is in the work preparation that humans copy (type) information into an Excel sheet or highlight values on paper documents with different colour text markers. Customaite simply replaces the tools for this work preparation by adding in artificial intelligence to the mix to take care of most of the tedious work to be done. For the declarant, this boils down to simply stop printing documents and start working with a web interface.



The screenshot displays the Customaite web interface. On the left, there is an invoice summary with the following details:

InvoiceDate:	01.06.2022
DocCurrency:	EUR
ShippingMethod:	Road

Below the summary is a table with two columns: Quantity and Price.

Quantity	Price
12	107,88 €
5	63,30 €
10	225,00 €
25	949,75 €
7	86,31 €
5	111,25 €
4	189,44 €
5	97,10 €
6	152,88 €
9	600,11 €

On the right side of the interface, a large yellow graphic displays the word "Price" above the number "947,75". The Customaite logo is visible in the bottom right corner of the graphic.

An easy to use web interface that is like a digital twin of the declarants physical workplace with the same functionalities, but better and faster. And with new functionalities that are even faster and even better, such as copying and automatic processing of information from scanned documents.

Conclusion

Companies that are dealing with large numbers of customs declarations that require significant human effort to create, are all looking for opportunities to increase efficiency. They are not necessarily looking to cut costs but try to deal with increasing market demand.

At Customaite, we have a very clear mission: halve the time needed to make a customs declaration. We do this by taking a fundamentally different approach to automating the customs declaration process. We are committed to the full workspace, to allow him/her to make declarations together with our AI algorithms that have been developed specifically for the customs context.

Sounds interesting?
Book your discovery call now at customaite.ai

Customaite is proudly enabled by AE

