omni:us

A brief guide to ask more.

Our technical FAQ.

Who are we?

We are omni:us. We pave the way for companies in the insurance industry to become digital insurers, and bridge the gap between traditional insurers and modern insurtech companies. We apply machine learning to the insurance space and help companies dealing with insurance be more transparent, affordable and efficient.

Based in Berlin, we are a global team with members hailing from all over the world. On top of these close-knit engineering & scientific teams, we also work together with our research partner, the Computer Vision Centre (CVC) at the Universitat Autònoma de Barcelona. As of the moment, omni:us operates in 6 countries within the EU, as well as, in the US.

Microsoft AI Startup

nvidia Inception Program

Google AI Startup

Red Herring Winner 2017

Plug&Play Insurtech Member

What do we do?

We enable businesses to "outsource" a set of mundane, time-consuming tasks that revolve around specific insurance processes, such as claims handling, to AI systems. These AI systems, fueled by the use of artificial neural networks, can essentially carry out a variety of tasks related to understanding documents and pinpointing directly to information within them. By "understanding" the insurance documents and their semantics and context, our AI system is able to automatically attain the vital information as extracted values - which is also the basis for our claim validation.

Where can you find us online?

You can visit our engineering blog at **omnius.com/engineering**, follow us on twitter **@omniusHQ**, watch video highlights & exclusive interviews from our Machine Intelligence Summits on our YouTube channel **omni:us**, check out our latest job openings at **omnius.com/careers** or drop us a line at **hello@omnius.com**.

Why claims automation?

The claims management process is one of the core insurance activities which determines what, and how much is paid out to insurance clients for each claim. The handling of insurance claims is a meticulous one, as it requires an extensive knowledge of the domain & expertise in a given vertical (automobile, health, household, life etc.)

Every insurance company is different, but the claims value chain is generally streamlined in the following fashion:

1 Intake 2 Triage 3 Assignment 4 Investigation 5 Evaluation 6 Mgmt 7 Resolution 8 Recovery & Compliance

The steps of Intake > Triage > Assignment tend to be rather labor-intensive. The adjunct steps (Investigation, Evaluation, Management, Resolution, R&C) are handled by experts called claims adjusters. They ensure that a claim is valid (non-fraudulent) and appropriate for its coverage. They have the final decision-making in terms of insurance payouts.

omni:us empowers these claims adjusters by providing the vital information digitally, allowing them to work on the more cognitively-inclined tasks, without further need to manually handle documents. By applying our AI system to the highly variable insurance documents, we are able to understand, structure and deliver the necessary vital information. This speeds up the internal insurance processes, allowing for quicker claims handling.

What makes us different?

At omni:us, we are dedicated to transforming the insurance industry by automating and easing each and every step of the claims handling process. Our goal is to ensure that with this digital transformation, the insurance industry attains more efficiency. Throughout the last 3 years, we have developed and specialised our AI to serve the insurance domain and the prevalence of the high degree of variable documents. Our AI system "understands" documents, similar to how humans do so, by first segmenting and distinguishing them based on their visual similarities, ensuring a very profound comprehension of the documents, as well as, their context. Its understanding of document context enables the attainment of better results from the eventual classification and extraction processes. By applying our AI technologies, we are able to extract vital information, even from printed, digital or handwritten texts.

For our AI system to remain a continuous leader in its field, we use a combination of computer vision and NLP to create a state-of-the-art approach that directly addresses and solves the core challenges faced by insurers.

What technologies do we use?

Our relevant approaches are explained below. Our processing approach is related to document handling and information extraction, carried out by the respective technological approach that is most ideal.

Template Alignment

Claims documents very often contain predefined forms, from which information needs to be extracted. In this case, the most reliable approach is to rely on a predefined template which contains the position and the meaning of the relevant fields. Since any misalignment between the scanned page and the template would lead to errors in the subsequent text recognition and information extraction steps, an automatic document alignment module has been developed. It computes a global transformation, in order to align the scanned image to the template image. Since the misalignments are relatively small and the appearance is quite similar, an affine transformation (translation, rotation, scaling) is estimated using a dense pixel based approach (Lucas-Kanade).
However, a similar approach can be used for photographed images by estimating a perspective transformation using a feature based approach (SIFT, SURF, ORB). Depending on the application, the estimated transformation can be applied to the scanned image or inversely applied to the template.

Document Classification

Documents, especially insurance-related ones, tend to be so diverse that it is far more accurate to run multiple classification approaches, either in sequence or in parallel. Thus, we first visually classify the pages, and then take it a step further by textually classifying them to get comparable predictions.

1. Computer Vision

Often, the first step in a processing chain is to assign a certain class or label to a full document or single page. This can be used for a variety of tasks: Detecting rotated pages, checking if a certain page exists in a document, or marking the page for further specific processing downstream, to name a few.

For vision based classification, omni:us uses Convolutional Neural Network (CNNs) architectures, a type of neural network loosely modeled after the visual mammalian cortex. Each layer consists of a set of convolutional filters, which learn to detect more and more abstract features of an image, depending on their place in the network's hierarchy. These high level features are used in the final layer to assign one or multiple classes to an image with a certain confidence. While the filters in an initial layer might pick up on vertical and horizontal lines, later layers might detect a certain block structure of a text or logo. As these features are learned by the network, we can avoid the costly manual feature engineering needed for classical computer vision methods. Big architectures with hundreds of millions of trainable parameters might give a boost of classification accuracy compared to smaller networks, but at the cost of longer training and prediction times and with the possibility of only memorizing seen data. Thus, depending on the task, a diverse set of architectures is used.

2. Natural Language Processing

In some cases, documents of different classes have very similar semantic content, but have visually different layouts. Natural Language Processing (NLP) then becomes a suitable alternative to Computer Vision based approaches. The main idea is that each type of class is associated with certain vocabulary. A document can be represented by the number of occurrences of each word in it (bag of words). Topic modeling or the usage of word embeddings can refine the representation further. Finally, the document is classified using neural networks or classical machine learning approaches like random forests.

Region Detection

For some cases, one would not only be interested if a page contains a certain object, for example a barcode or a signature, but also, where this object precisely occurs. Current state-of-the-art is the usage of object detection neural networks, an extension of the convolutional based classification networks. The network learns more and more complex features at each layer by itself. But instead of predicting a list of labels, it predicts multiple boxes at the locations where an object of interest might be, and assigns labels and confidences to these. This is internally achieved by overlaying a grid of bounding boxes with different aspect ratios on the image and predicting a final vector of parameters, which translate into moving and transforming these boxes to the correct coordinates.

Text Recognition

When processing documents obtained from uncontrolled sources, like in the case of insurance claims, it is very common to receive a mixture of digital and scanned documents. In the case of scanned documents, it is necessary to detect the text present in each page (where the text is) and recognize the text (what is written). When the text is printed, the text recognition process is known as Optical Character Recognition (OCR) and for handwritten text, the process is known as Handwritten Text Recognition (HTR). For HTR or a combination of OCR/HTR, omni:us has internally developed a software that implements state-of-the-art techniques. For text line detection, a fully-convolutional YOLO-like neural network architecture is used. For the text recognition, a convolutional-recurrent (CNN+LSTM) neural network architecture is used, and trained using connectionist temporal classification (CTC).

NLP for Value Extraction

Given the text output from the text recognition step, important information from the documents is gathered. For insurance claims, we might be interested in invoice numbers and prices, for instance. To automatically extract these values, we use word embeddings trained on a large number of claims and feed them into a Bi-LSTM architecture. During the learning process, the network becomes aware of common locations and formats of the relevant information. The approach is not restricted to specific layouts or relies on any hand-crafted rules, which makes it easily scalable to new values of interest and even new insurance sub-domains.

What further research are we working on?

Together with our research partners at CVC we are currently working on the following research areas:

Graph Based Layout Analysis

Claims documents, in particular, administrative ones, often follow specific layouts that consist of regular substructures with semantic meanings, which can be used to ease document classification and information extraction. Due to their representation, graphs are suitable data structures to describe the content and the relationship between these substructures in a document. Generally speaking, the basic elements (words, rules, logos, etc.) can be represented by graph nodes and the spatial relationships between them (distances, orientations, etc.) can be described by the graph edges. Based on these structural models, modern graph matching techniques based on graph kernels, embeddings and graph convolutional networks can be adopted for document classification and information extraction. We are currently working on a graphlet based approach to describe the local substructures and their relationships in a hierarchical way.

Semantic Annotations

The goal is to train a model that extracts information from any kind of semi-structured document images, without separation of the intermediate steps, i.e. the model should perform region-of-interest detection, transcription and semantic annotation in a single step, taking into account all the parts of the process, to get the best possible performance. This is motivated by Sayre's Paradox, that suggests that to attain a high quality text segmentation, it is necessary to understand its content, and to get a good transcription, it is necessary to have good segmentation. This idea can be extrapolated to semantic annotation of the text after transcription. Hence, it makes sense to train a model that performs the whole process in an end-to-end way.

Handwritten Text Recognition

Using a new attention-based sequence-to- sequence model for HTR tasks. Our model has 3 main parts: encoder, attention and decoder. The encoder part consists of CNN and BGRU while the decoder is a one directional GRU to generate the predicted word, character by character. Location-based attention has been applied to connect both encoder and decoder which would segment each character in an unsupervised way. Without segmentation, sliding window, pre-processing, language model and lexicon, our model achieves a character error rate in the test set of 7.25% and a word error rate in the test set of 18.58% in IAM dataset.

What is your current open-sourcing project?

We would like to give back to the community, as we have faced some general problems in handling documents ourselves. Therefore, we are in the midst of open sourcing a couple of document handling microservices and libraries.

These include:

- PDF conversion to images & back
- Ghostscript wrapping as standalone microservice
- Tesseract wrapping as standalone microservice
- Helpers for easy & reliable temporary files creation & disposal
- Helpers for handling & tracking exceptions across microservices
- Helpers for sending files with spring's RestTemplate

Keep an eye out & give us some :hearts: on

https://github.com/omni-us

omni:us

Shape with us an AI-first insurance world.

omnius.com/careers

Go deeper.

Ask the ones who just gave answers.

Marianne@omnius.com
Scientific Engineer Intern

Vithya@omnius.com
Junior Technical Writer

Aleksandar@omnius.com
Senior Developer

Christopher@omnius.com
Junior Scientific Engineer

Lei@omnius.com
PhD Multiwriter

Lutz@omnius.com
Senior Scientific Lead

Manuel@omnius.com
PhD Semantic Annotation

Mauricio@omnius.com
Data Scientist

Nischal@omnius.com
Lead Scientific Engineer

Thomas@omnius.com
Technical Product Lead

Vasanth@omnius.com
Data Scientist