

[Home](#) > [AI/ML](#)

Car Plate Text Extraction: A Comprehensive Analysis of a real-world solution powered by AWS ML services

1 September 2023 - 11 min. read

[Amazon Rekognition](#)

[Amazon Textract](#)

Overview

In today's rapidly advancing technological landscape, the convergence of artificial intelligence and cloud computing has opened up new avenues for solving real-world challenges.

Among these is the **Optical Character Recognition (OCR)** domain, which has unveiled many possibilities across various sectors when applied to car plates.

The ability to accurately extract and recognize text from car plates has transcended beyond mere identification; it has become a cornerstone for seamless automation and operational efficiency.

Consider the multifaceted applications of **car plate OCR**.

For enterprises entrusted with managing bustling car traffic, the integration of this technology enables streamlined processes and swift, automatic payments. Imagine a parking lot where vehicles glide in and out effortlessly, with payments being processed in real-time, eliminating the need for manual transactions. This enhances user experience and optimizes resource allocation and revenue management.

Zooming in on the public sector, the implications of a reliable way to obtain license plate text are equally transformative. Law enforcement agencies can harness the power of car plate OCR to swiftly identify vehicles of interest, aiding in everything from traffic management to locating stolen cars. Moreover, city planners can use this technology to gain valuable insights into traffic patterns, making urban planning more data-driven and responsive.

Delving into the business sphere, industries dealing directly with automobiles can reap substantial benefits. Take a petrol station, for instance. With car plate OCR, fueling transactions can be seamlessly linked to vehicles, and many statistics about traffic and potential customers can be gathered.

Similarly, road maintenance authorities can efficiently manage toll collection and monitor usage on toll roads, enhancing revenue generation and road infrastructure management.

This article is a testament to the dedication and expertise of our team as we embarked on a quest to unravel the complexities of car plate OCR within the AWS environment. Faced with the challenge of identifying the optimal AWS service for this task, we embraced an empirical approach. Our journey took us through various AWS services with unique strengths to determine the most effective solution for accurate car plate text extraction.

Throughout this journey, we encountered common pitfalls and hurdles that often accompany the implementation of advanced technologies. These insights, born from practical experience, serve as valuable guideposts for those navigating the landscape of car plate OCR. Our goal is to equip you with the knowledge of which AWS service to choose and a comprehensive understanding of how to navigate potential obstacles.

Ready? Let's start.

OCR on AWS

Within the comprehensive suite of Amazon Web Services (AWS), a variety of pragmatic Artificial Intelligence (AI) and Machine Learning (ML) services are available to facilitate the development of robust Optical Character Recognition (OCR) solutions. Among these, Amazon Textract, Amazon Rekognition, and AWS's SageMaker play a crucial role.

Anyway, for our specific endeavor, the directive was clear: prioritize the utilization of **managed services** within the AWS ecosystem.

As a result, our attention zeroed in on Amazon Textract and Amazon Rekognition, two instrumental services that align seamlessly with our mission to craft an efficient car plate OCR solution.

Let us formally introduce both of these services to ensure the utmost clarity.

Amazon Textract

Amazon Textract, a standout component of AWS's repertoire, is a sophisticated machine-learning service designed to extract text and data from diverse documents, images, and forms. Its core strength lies in its ability to accurately and efficiently process large volumes of textual information, transforming unstructured content into structured data with

remarkable precision. This service showcases exceptional adaptability, excelling in various scenarios such as invoice processing, content indexing, tables, etc.

The most relevant features include its capacity to distinguish between different types of content within documents, such as tables, forms, and text blocks. This granularity of analysis enhances its ability to extract pertinent information effectively. Its integration with AWS services allows seamless data transfer for further analysis and application. Furthermore, identifying key-value pairs and their context is invaluable, offering a comprehensive view of structured data. As we dissect the offerings within AWS for car plate OCR, Amazon Textract emerges as a compelling solution that holds the potential to streamline our project's objectives significantly.

Amazon Rekognition

At the forefront of AWS' AI and ML offerings, Amazon Rekognition is a formidable image and video analysis service designed to unlock insights from visual content. Its formidable strength lies in its ability to identify objects, faces, and scenes within photos and videos and detect and recognize text within these visual assets. This makes it an invaluable asset in Optical Character Recognition (OCR), particularly for endeavors such as car plate identification.

Amazon Rekognition's most significant features encompass its advanced facial analysis capabilities, enabling the detection of facial attributes, emotions, and even facial recognition. Moreover, its text detection prowess extends beyond mere identification to capturing fine-grained details within images. Integrating customizable labels further augments its applicability in content moderation and safety compliance scenarios.

For our endeavor, the capacity to detect and recognize text within car plate images is paramount, and Amazon Rekognition emerges as a robust contender in this capacity. Its versatility to handle varied use cases and its seamless integration with AWS services substantiate its potential to contribute significantly to our project's objectives. As we meticulously assess AWS offerings for car plate OCR, Amazon Rekognition's ability remains a vital focal point in our exploration.

With the introduction of these services complete, we can now seamlessly proceed with our journey.

Step 1: Image Preprocessing

No matter the chosen service, a preprocessing measure becomes essential. In our case, we opted to harness the capabilities of Amazon Rekognition to facilitate several preliminary stages; notably the extraction of the bounding box encapsulating the license plate.

However, the path ahead was far from straightforward: **real-world images, captured from varying angles, under diverse lighting conditions, and encompassing many vehicles and license plate formats, unveiled anything but trivial challenges**. The endeavor to accurately acquire the bounding box for the license plate proved to be an intricate pursuit.

Submitting the original high-resolution image alone proved insufficient, often resulting in the failure to detect the license plate. This highlighted the complexity of the task, underscoring the intricacies tied to real-world scenarios and the nuanced characteristics of license plates in diverse contexts. The journey to success necessitated leveraging Amazon Rekognition's tools and delving into image preprocessing to enhance the accuracy of subsequent processes.

In light of these challenges, we embarked on a meticulous journey of fine-tuning through a series of intricate **preprocessing steps**.

Our initial focus was on color manipulation.

Recognizing the significance of optimal color representation, we strategically converted the images into grayscale. This strategic choice simplified subsequent analysis and mitigated potential interference from color variations due to diverse lighting conditions.

Subsequently, our attention shifted to contrast enhancement, a pivotal technique in elevating the clarity of visual elements.

With contrast enhancement serving as a cornerstone, we ventured into the domain of posterization—a method that segmented the grayscale image into distinct, visually discernible regions. By doing so, we aimed to accentuate the boundary between the license plate and its surroundings, facilitating its isolation and identification.

Taking the quest for clarity even further, we introduced a **dithering process**.

This technique involved the strategic arrangement of pixels to simulate additional shades and nuances, effectively bridging the gap between grayscale and a richer, more detailed representation. This intricate fusion of visual enhancement techniques aimed at producing a refined image that retained the essence of the original while minimizing the impact of visual noise.

We then proceeded to subject a substantial volume of images to Rekognition, only to discover an intriguing phenomenon: some images in which the license plate was successfully recognized in its original format failed to yield accurate results once processed, and some images in which the plate was not detected at first, were now correctly recognized. This revelation was initially disconcerting, prompting us to delve deeper into the underlying reasons.

Upon closer analysis, an insight emerged: **noise reduction, while seemingly beneficial, did not consistently enhance Rekognition's ability to identify license plates.** This was because Rekognition's underlying model isn't exclusively tailored to anticipate license plates. Instead, the model employs diverse criteria to determine whether a license plate is present within a frame—criteria that our preprocessing inadvertently modified by highlighting only high-contrast letters.

In light of this, we recalibrated our strategy, embracing a nuanced approach. Our revised methodology unfolded as follows:

Submitting the Raw Image: Initially, we presented the unaltered, high-resolution image to Rekognition, eagerly awaiting its response regarding license plate detection and the corresponding bounding box.

Conditional Processing: In instances where Rekognition's initial analysis failed to detect a license plate, we would apply the preprocessing steps detailed earlier. The enhanced version of the image was then resubmitted to Rekognition for a fresh attempt at detection.

Bounding Box Success: Upon Rekognition successfully identifying the bounding box of the license plate, we progressed to the next phase.

Step 2: Extracting text from Car Plates

At this point, we have achieved optimal license plate detection outcomes. With this accomplishment, we proceed to calculate a slightly larger bounding box than the one initially identified by Rekognition. This strategic expansion ensures the comprehensive encapsulation of the license plate, thus affording ample contextual space for other AI-driven services to detect the presence of a license plate. This intentional enlargement safeguards against potential information loss while maintaining a balanced visual frame.

We then proceed to crop the image and apply the preprocessing procedures detailed earlier unless the image has undergone prior treatment.

The resultant image crop presents a uniform composition: a centered and precisely enclosed license plate characterized by minimal visual noise and heightened contrast, and it's ready to be sent to the chosen OCR service.

In our initial exploration, our attention gravitated towards harnessing the capabilities of Amazon Textract. This choice was based on our **previous observations of Amazon Textract's performance in processing various documents** underscored its robust capabilities. This led us to a logical inference: With the implementation of targeted preprocessing steps, the service could be tailored to the specific task of extracting text from car plates.

So, we sent the preprocessed images to Textract for text extraction—a step we anticipated would be straightforward given the clarity of the text. However, reality diverged from our expectations. Our presumption was that presenting Textract with pristine, high-contrast images would yield outcomes akin to its performance with documents.

Unfortunately, the outcome proved to be a cold shower. Regardless of how many images were processed, Textract's performance fell short of delivering usable results. Textract was evidently optimized for document-oriented tasks. Our assumption that preprocessed, contrast-enhanced images would align with Textract's capabilities was flawed. Specifically, when tasked with extracting text from license plates, Textract's performance faltered, at least for our use case.

Consequently, we opted for Rekognition as a versatile tool for text extraction.

In our pursuit, we revisited the same preprocessed images previously submitted to Textract. This time we obtained auspicious outcomes. Our approach encompassed subjecting the solution to a comprehensive evaluation involving images with varying degrees of complexity. Interestingly, while images featuring optimal conditions were accurately detected without the need for preprocessing, more intricate images proved successful only after undergoing preprocessing.

As a result of this thorough examination, we arrived at a decisive conclusion: employing the cropped, high-contrast image constituted the most reliable and effective option. This realization stems from the recognition that preprocessing the image enhances its suitability for Rekognition's analysis and ensures consistent performance across a spectrum of image complexities. Ultimately, this finding fortified our methodology, reinforcing the importance of tailored preprocessing in optimizing the capabilities of AI-driven solutions for text extraction.

Step 3: Output Post-processing

The challenge of post-processing then came to the forefront. Once the text had been successfully extracted, many cases demanded careful handling. Here are a few examples:

One notable aspect that surfaced was Rekognition's inability to differentiate certain characters between the Cyrillic and Latin alphabets precisely. Within the extracted text, visually identical letters were frequently encoded using incorrect UTF8 bit sequences. This resulted in interpreting these letters as Cyrillic characters, consequently generating license plate text that deviated from the original image.

Unfortunately, Rekognition lacks a mechanism to dictate the choice of the alphabet during extraction. Consequently, a post-processing step is necessitated to rectify this discrepancy.

The first post-processing procedure involves the removal of all Cyrillic letters—an operation aligned with the fact that these characters are never present in the license plates under examination in our use case.

Another crucial post-processing task centers around reassembling license plates that span multiple lines, such as those on motorcycles. For this endeavor, Rekognition's capabilities prove valuable yet again. Using the labeling API, Rekognition provides insights into whether an image captures a car or a motorcycle. In cases where a motorcycle is detected, the extracted text needs to be reconstructed and formatted onto a single line.

These post-processing procedures underscore the intricate interplay between AI-driven extraction and the nuances inherent in textual data. The mix of automated analysis and subsequent human-guided refinement ensures that the extracted text aligns accurately with real-world scenarios, resulting in a reliable and robust car plate OCR solution.

Conclusions

In this article, we presented our empirical approach to constructing an optimal car plate OCR solution on the AWS platform, harnessing the capabilities of its managed services. We have compared Amazon Textract and Amazon Rekognition, and evaluated their performance on a dataset of car plate images. We have also discussed the necessary preprocessing steps, such as image resizing, cropping, and enhancement, that can improve the accuracy of the OCR process. We have found that Amazon Rekognition is the best service for car plate text extraction, as it provides the highest accuracy, lowest latency, and most flexibility among the AWS services.

We hope this article was useful for a deeper understanding of OCR. Stay tuned for other content about OCR and ML

About Proud2beCloud

Proud2beCloud is a blog by [beSharp](#), an Italian APN Premier Consulting Partner expert in designing, implementing, and managing complex Cloud infrastructures and advanced services on AWS. Before being writers, we are Cloud Experts working daily with AWS services since 2007. We are hungry readers, innovative builders, and gem-seekers. On Proud2beCloud, we regularly share our best AWS pro tips, configuration insights, in-depth news, tips&tricks, how-tos, and many other resources. Take part in the discussion!



Alessio Gandini

Cloud-native Development Line Manager @ beSharp, DevOps Engineer and AWS expert. Since I was still in the Alfa version, I'm a computer geek, a computer science-addicted, and passionate about electronics all-around. At this moment, I'm hanging out in the IoT world, also exploring the voice user experience, trying to do amazing (lo)Things. Passionate about cinema and great TV series consumer, Sunday videogamer

Copyright © 2011-2023 by beSharp spa - P.IVA IT02415160189