

ABBYY FineReader Engine

Performance Guide

Integrating optical character recognition (OCR) technology will effectively extend the functionality of your application. Excellent performance of the OCR component is one of the key factors for high customer satisfaction.

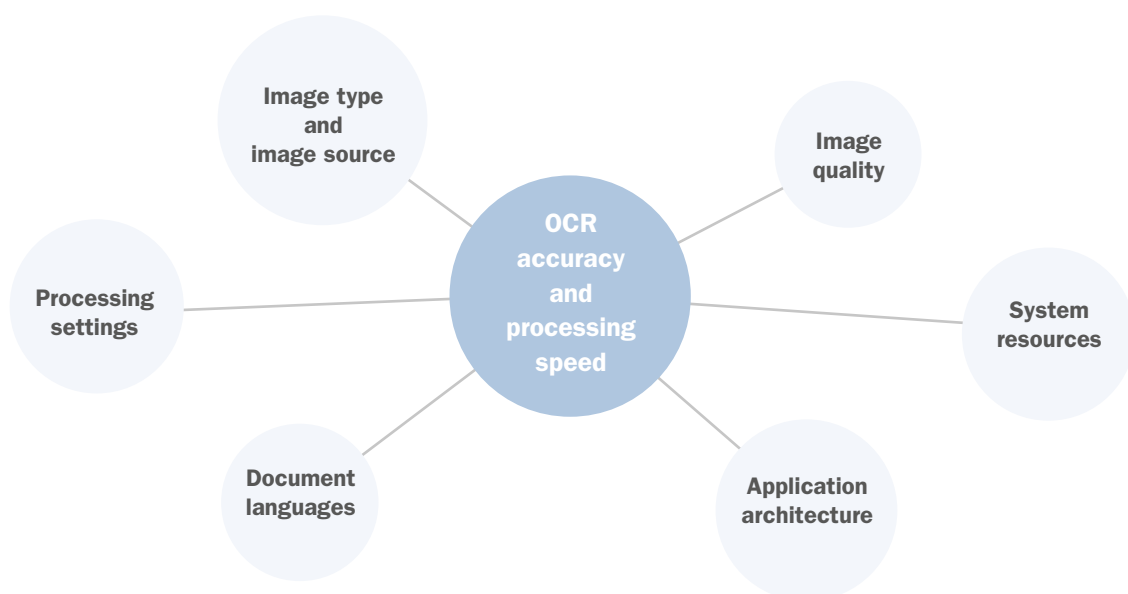
This document provides information on general OCR performance factors and the possibilities to optimize them in the Software Development Kit ABBYY FineReader Engine. By utilizing its advanced capabilities and options, the high OCR performance can be improved even further for optimal customer experience.

When measuring OCR performance, there are two major parameters to consider:

RECOGNITION ACCURACY

PROCESSING SPEED

Which Factors Influence the OCR Accuracy and Processing Speed?



Recognition speed and recognition accuracy can be significantly improved by using the right parameters in ABBYY FineReader Engine.

Image Type and Image Quality

Images can come from different sources. Digitally created PDFs, screenshots of computer and tablet devices, image files created by scanners, fax servers, digital cameras or smartphones – various image sources will lead to different image types with different level of image quality. For example, using the wrong scanner settings can cause “noise” on the image, like random black dots or speckles, blurred and uneven letters, or skewed lines and shifted table borders. In terms of OCR, this is a ‘low-quality image’.

Processing low-quality images requires high computing power, increases the overall processing time and deteriorates the recognition results.

**Image
quality**



**Key factor
for OCR
performance**

On the other hand, processing ‘high-quality images’ without distortions reduces the processing time. Additionally, reading high-quality images leads to higher accuracy results.

Therefore, it is recommended to use high-quality images for the OCR process.

Increase OCR speed and accuracy by enhancing the image quality.

How to Get High-Quality Images

TIPS FOR DOCUMENT SCANNING



Font Size

Documents printed in very small fonts should be scanned at higher resolutions.

Use the following resolution for scanning:

- 300 dpi for typical texts (printed in fonts of size 10 pt or larger)
- 400-600 dpi for texts (fonts size 9 pt or smaller)



Print Quality

Poor quality document, such as old newspapers or books should be scanned in the grayscale mode. This mode retains more information about the letters in the scanned text.

TIPS FOR TAKING PHOTOS



Correct Lighting

- Make sure that lighting is evenly distributed across the page and that there are no dark areas or shadows.
- If possible, use a tripod. Position the lens parallel to the plane of the document and point it toward the center of the text.
- Turn off the flash to avoid glare and sharp shadows on the page.
- If the camera has “White Balance” option, use a white sheet of paper to set white balance. Otherwise, select the white balance mode, which best suits the current lighting conditions.



If There is Not Enough Light ...


- Select a greater aperture value
- Select a greater ISO value for sensitivity
- Use manual focusing if the camera cannot lock the focus automatically

For more details refer to the Developer's Help of the FineReader Engine:

[Guided Tour](#) → [Best Practices](#) → [Source Image Recommendations](#)


Image Quality Enhancement with FineReader Engine

If it is not possible to influence the image quality in advance, it is recommended to enhance it prior to the recognition step. In FineReader Engine, various powerful image preprocessing functions are available:



BASIC
IMAGE PROCESSING FUNCTIONS

Image scaling
Image cropping
Image clipping
Lines straightening
Mirroring and inverting
Noise removal
Local contrast enhancement
Correction of geometric distortions
Adaptive binarisation



ADVANCED
IMAGE PROCESSING FUNCTIONS

ABBYY Camera OCR technology
Auto-splitting of double-pages
Removal of stamps and written notes
Automated image de-skewing
Autodetection of page orientation and rotation
Image despeckling

Predefined Processing Profiles in FineReader Engine

Recognition accuracy and processing speed can be optimized by fine-tuning the individual settings. However, it can be done by applying one of the predefined processing profiles, which are available for the main usage scenarios. The settings provided in these profiles are most suitable in the corresponding situations.

Most of the profiles are available in two forms:

- with settings optimized for the best recognition accuracy
- with settings optimized for the highest speed of processing

If necessary, it is still possible to tune the profiles more precisely by setting other relevant parameters through the API.

SCENARIO	PROFILE NAME
Document conversion for archiving A document requires high processing speed, good visual quality and small size of the resulting PDF. Highest recognition accuracy is not a critical parameter, a satisfactory level is sufficient.	DocumentArchiving_Accuracy DocumentArchiving_Speed BookArchiving_Accuracy BookArchiving_Speed
Document conversion for content reuse In this case, recognition accuracy and document structure retention are the most desirable criteria. As the final document must be error-free, lower recognition rate or wrong layout reconstruction would imply additional work for operators.	DocumentConversion_Accuracy DocumentConversion_Speed
Text extraction on an image, which includes small text areas of low quality. This text information should be later available for search, extracted for further processing or used for document classification.	TextExtraction_Accuracy TextExtraction_Speed
Data capture from exactly specified page areas	FieldLevelRecognition
Barcodes recognition	BarcodeRecognition_Accuracy BarcodeRecognition_Speed
Creating high-compressed PDF files, saved as pictures	HighCompressedImageOnlyPdf
Business cards recognition	BusinessCardsProcessing
Recognition of technical drawings, which usually are of large size and include complex diagrams, as well as different text orientation.	EngineeringDrawingsProcessing



Predefined Recognition Modes

Another possibility to influence the OCR performance is to use recognition modes designed for particular scenarios. FineReader Engine provides following predefined recognition modes:



NORMAL

Using this mode you will achieve the highest recognition accuracy.

This mode is highly recommended when recognizing content is going to be reused in other applications or tasks where high accuracy is critically important.



BALANCED

This mode provides the intermediate values of recognition accuracy and speed.

Generally, this recognition mode provides higher speed than the 'Normal' recognition mode while achieving almost the same level of accuracy.



FAST

Usage of this mode increases processing speed up to 200-250%.

This mode is recommended when processing speed is of the primary importance, such as in high-volume document processing for archiving, content and document management systems.

Document Languages

ABBYY FineReader Engine is capable of recognizing both mono- and multi-lingual (e.g. written in several languages) documents. It is very important to specify the correct recognition language, since an incorrectly specified language can significantly slow down the document processing and decrease the recognition quality.

If the recognition language cannot be specified in advance, it is possible to use automatic language detection.

However, a high number of preselected recognition languages will reduce the processing speed. Therefore, it is not recommended to specify more than five recognition languages.

To increase the recognition accuracy even more, FineReader Engine provides dictionary and morphology support for many languages. When processing documents including subject-specific terms or „structures“ such as product codes, telephone numbers or passport numbers, custom created dictionaries can be imported to ensure high recognition quality.

Document Processing Stages

In terms of OCR, document processing is a multi-step process. Depending on a particular scenario, different functions and parameters can be applied to each OCR stage. In FineReader Engine the process includes following stages:

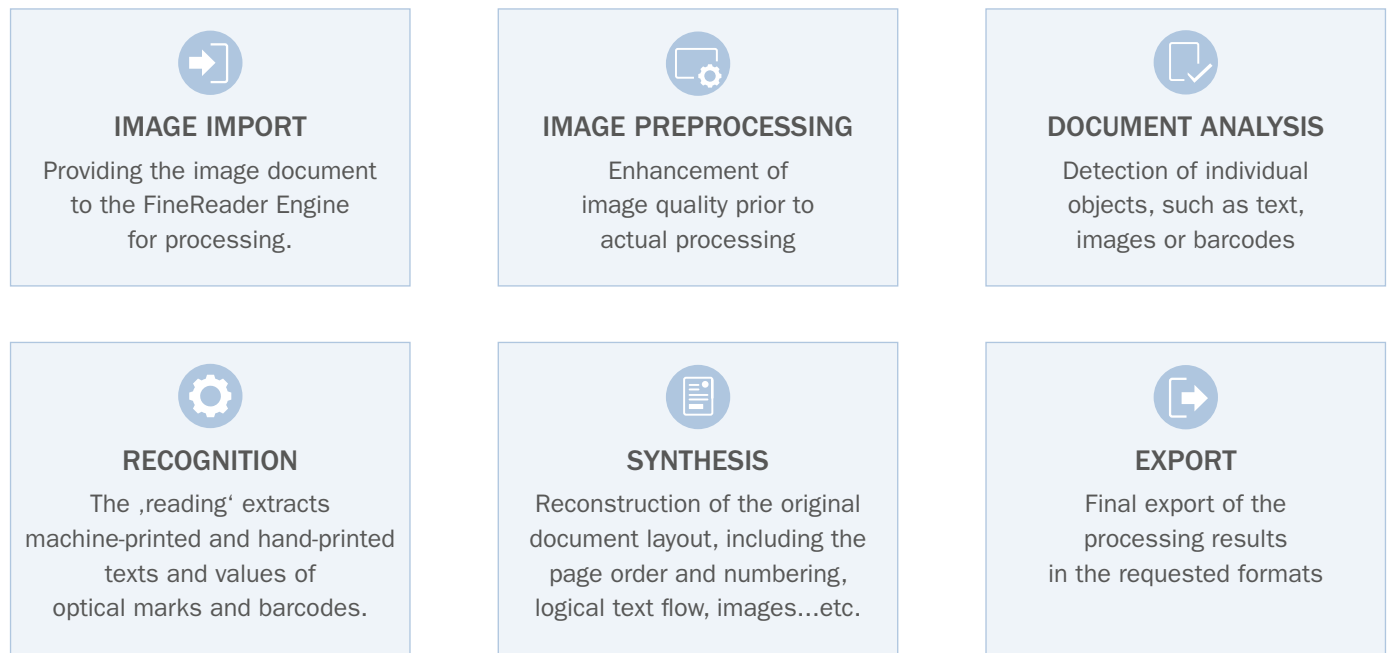


Image Import

Documents can be sent to the FineReader Engine directly from a scanner or imported from the storage system or the memory stream. To obtain images from different sources will require different methods and influence the recognition speed. The image import from memory is generally faster than opening the images from a file storage.

Image Preprocessing

Generally, the OCR process is faster for good-quality images. It is recommended to fine-tune the image preprocessing step accordingly and therefore save time during the actual processing step.


Images can be of different formats and quality. High-quality images, such as digitally created PDFs, typically do not require a lot of preprocessing work. For low-quality images, like scanned documents with incorrect scanner settings or old books, it is necessary to apply advanced image preprocessing functions to improve the recognition results. For the preprocessing of digital photos, the special ABBYY Camera OCR™ technology is applied. Here, the algorithms are optimized specifically for photo enhancement. Usage of different preprocessing functions will individually influence the processing speed.

Document Analysis

Brochures or newspapers often contain text in columns, tables, diagrams and pictures. Technical drawings might be large documents including complex engineering diagrams with different text orientation. For documents with such complicated layouts, the document analysis step will require more processing time. On the other hand, the analysis of simple layout document like letters or contracts is very fast.



The different methods and parameters used for specific processing scenarios will significantly influence the overall processing speed. Discarding unnecessary stages can speed up the entire OCR process. For example, when extracting data from predefined document areas, the document analysis is not required. When exporting the documents to TXT format or to PDF Image Only format, the synthesis stage can be skipped.



Parallel Processing Using Multiple Cores

FineReader Engine can be used to build applications of any scale and complexity – from a client workstation to a server-based solution or a large multi-million page project. When it comes to the OCR performance, it often makes sense to utilize multi-processor or multi-core systems to increase the processing speed.

Built-in multi-core support in FineReader Engine allows different approaches to scale-up the OCR process:

- Utilizing a single Engine instance
- Loading several Engine instances

There are different approaches for processing documents:

Processing of Large Multi-Page Documents

For parallel processing of large documents with many pages the 'FRDocument' object is best suited. In this case, the pages of a multi-page document are processed in parallel on the CPUs available. At the end, the results are combined into one multi-page document. It is the most easy-to-code multiprocessing way. The number of processes needed is detected automatically, depending on factors such as the number of available physical or logical CPU cores, the number of free CPU cores available in the license, and the number of pages in the document. If necessary, the developer can easily change the multiprocessing settings and tune the number of processes to be run.

Processing of Many Single-Page Documents

To process many one-page documents in parallel, which are received from the same source, e.g. a scanner, it is recommended to use the 'BatchProcessor' object. This object is most effective in terms of speed, when document export is not required, like in data capture scenarios with a custom output format.

To perform full processing of many one-page documents in parallel, it is recommended to use a pool of Engine instances. This approach is also best suited for web-service scenarios, when the input document should be processed directly after it was submitted. In this case, the document is passed to an available FineReader Engine instance from the pool and processed immediately.

ABBYY FineReader Engine offers a flexible, modular licensing that allows selecting the best combination of options. For example, the network license allows scaling-up the processing to a very high number of computers (virtual or physical).

FineReader Engine - Speed Testing Results

The table presents the results of internal performance testing. Please be aware that testing results always depend on many factors, such as image quality, used recognition languages and other factors.

	ONE-PAGE DOCUMENTS PAGES/MINUTE	ONE MULTI-PAGE DOCUMENT PAGES/MINUTE	DATA CAPTURE SCENARIO NO EXPORT* PAGES/MINUTE
Sequential processing	60	51	87
Parallel Processing with FRDocument	41	117	57
Parallel Processing with FRDocument while keeping data in memory during processing	55	141	82
Parallel Processing using Batch Processor	99	115	294
Parallel Processing using a pool of Engines	165	10	102

Technical Test Information

Intel® Core™ i5-4440 (3.10 GHz, 4 physical cores), 8 GB RAM, 4 processes running simultaneously.

The performance was tested on 300 documents in English, using the 'DocumentArchiving_Speed' predefined profile. In the scenarios „One-page documents“ and „One multi-page document“ the documents were exported as PDF format.

* The text was extracted from pre-defined areas on one-page documents. No export to any file format was performed.

System Resources

During the OCR process, a range of different algorithms are applied. They depend on image quality, document languages, layout complexity and number of pages in the document. Accordingly, such algorithms might require higher memory resources. It is recommended to set up the system in accordance with the outlined memory requirements to optimize the processing speed by allocating adequate system memory.

MEMORY REQUIREMENTS

Processing one-page documents:

minimum 400 MB RAM
recommended 1 GB RAM

Processing multi-page documents:

minimum 1 GB RAM
recommended 1,5 GB RAM

Parallel processing:

350 MB RAM x number of CPU cores
+ additional 450 MB RAM

Parallel processing of documents in Arabic, Chinese, Japanese or Korean:

850 MB x number of CPU cores
+ 750MB RAM

How to Increase the Overall Processing Speed in FineReader Engine

There are several possibilities to improve the performance of your system:

- Fine-tune the image preprocessing settings to deliver the highest document quality for the processing step.
- During the processing step, use one of the predefined processing profiles optimized for speed and the appropriate recognition mode – balanced or fast.
- Specify the correct recognition languages. Incorrect language can significantly slow down document processing. The more recognition languages are selected, the slower the speed of processing.
- Use the appropriate object (FRDocument or BatchProcessor) and enable parallel processing.
- Specify appropriate parameters of analysis and recognition. For example, disable table detection and page orientation correction if images contain no tables and have correct page orientation.
- Omit the synthesis stage if the processed documents will be exported to TXT format or PDF Image Only format.
- Use the Fast PDF Export Profile, when exporting the documents to the PDF format.
- Use the special object (ExportFileWriter), which is designed for the export of very large multi-page documents into PDF format.

For more information, refer to the FineReader Engine Developer's Help:

[Guided Tour](#) → [Best Practices](#) → [Increasing Processing Speed](#)



How to Improve the Text Recognition Quality in FineReader Engine

FineReader Engine offers high recognition quality. The recognition quality will always depend on factors such as image quality, language and other factors. However, there are several ways to increase the recognition quality:

- Specify the correct text type.
- Specify the appropriate recognition languages.
- Define unique languages and custom dictionaries for the recognition of special characters or documents with specific terminology, e.g. legal or healthcare texts.
- Split the facing pages of scanned books into two separate images.
- Apply the special Camera OCR technology, when digital photos are processed.
- Correct resolution of the image, if it significantly differs from the recommended resolution.

For more information refer to the FineReader Engine Developer's Help:

[Guided Tour](#) → [Best Practices](#) → [Improving Recognition Quality](#)



As you know, in document processing the OCR process can be a very complex task. Depending on the individual document processing scenario, the OCR performance results can significantly vary. The tips for recognition accuracy and processing speed optimisation in ABBYY FineReader Engine should help you to achieve the optimal performance for your business case.

Additional Information Resources

To learn more about the different aspects of OCR performance optimisation and about the SDK ABBYY FineReader Engine, please use following sources of information:

- ABBYY Technology portal: <https://abbyy.technology/>
- ABBYY OCR SDK forum: <https://forum.ocrsdk.com/>
- The help file provided in the ABBYY FineReader Engine distributive
- ABBYY FineReader Engine product pages on www.abbyy.com/ocr-sdk.

Individual Project Support

If you would like to discuss a particular project, please contact us. During the testing period, you can ask standard technical questions to our ABBYY Technical Support or use ABBYY Professional Services for advanced consultancy, in-depth project analysis and individual code review. Using ABBYY's technical resources can shorten your development work and speed up your project. Please contact the ABBYY sales manager, if you wish to further explore these options.



If you have additional questions, contact your local ABBYY representative listed under www.abbyy.com/contacts or use the online contact form www.abbyy.com/ocr-sdk/#request-demo.