

Every industry has vitally important data trapped in PDFs and paper documents. With all of the document variations, it's hard to maintain accuracy. This is the GLYNT annual accuracy report, describing GLYNT how well GLYNT meets these challenges? What is the accuracy? performance. How fast does GLYNT learn? What can the user expect?

GLYNT'S ANNUAL ACCURACY REPORT

AUGUST 2021

GLYNT'S 2021 ACCURACY REPORT

Summary

Welcome to our annual report on GLYNT's accuracy.

As we all suspect or have experienced, AI is not perfect. It does make mistakes. Yet, we all want 100% accurate data. So the bottom line is two key questions: How well does the AI work? How quickly can the user get to 100% accuracy?

This report is GLYNT's answer. This is our third annual report on accuracy and the results just get better and better. From the start, GLYNT is highly accurate. We're averaging 99%. And it takes less and less time to train our AI. In the past year we cut the number of documents needed to get to peak performance in half.

With results like these, you won't be surprised at our goal:

"GLYNT is the fastest path to 100% accurate data from documents"

Take a look at our deep dive into accuracy. Send us your questions. We're constantly scrutinizing our accuracy and worrying about all the details. So you don't have to.

Onwards!

The GLYNT Team



2021: Higher Accuracy

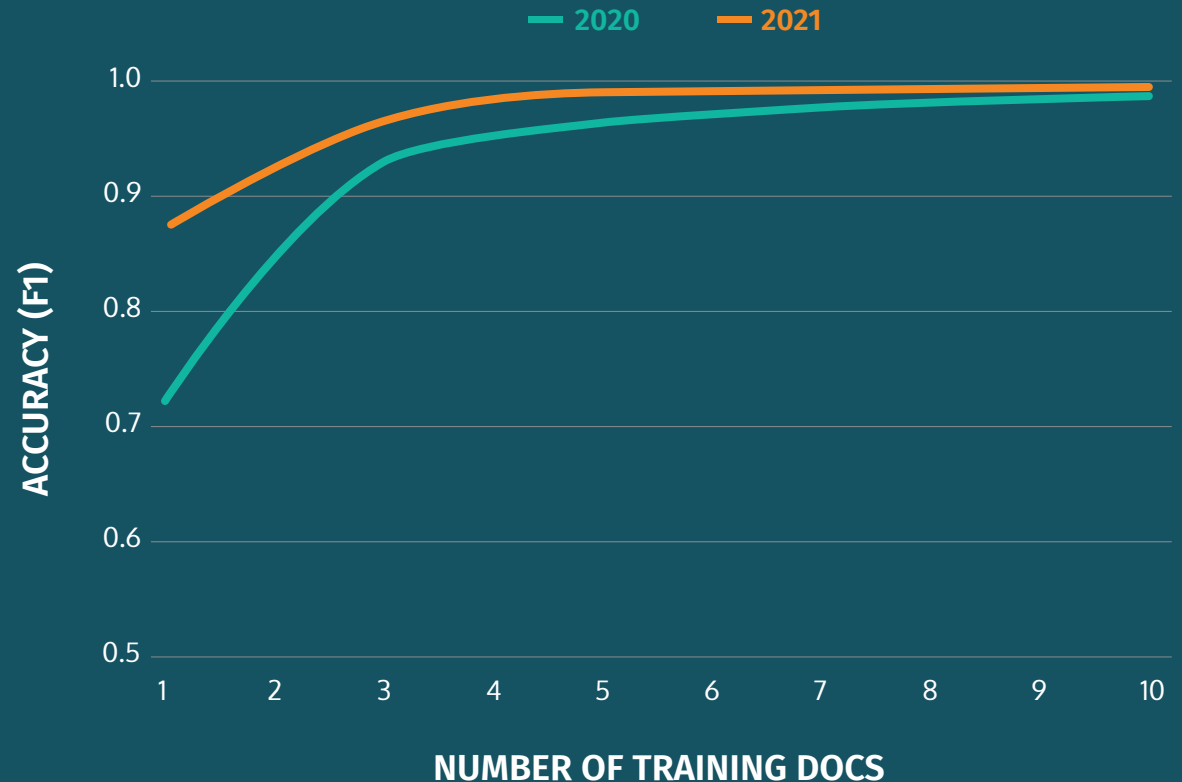
Let's jump to the bottom line: GLYNT is more accurate than ever before. We tested against a variety of document types – utility bills, invoices, healthcare documents – and included scans of documents and original issue PDFs.

The results shows that GLYNT is highly accurate after training on just one document. We went from an F1 score just above 0.7 to nearly 0.9 in the past year.

And we're hitting F1 of 0.99 after training on just 5 documents.

Strong performance, year after year. And always getting better.

Note: F1 is a standard measure of accuracy in machine learning. See page 9 for details on how it is calculated




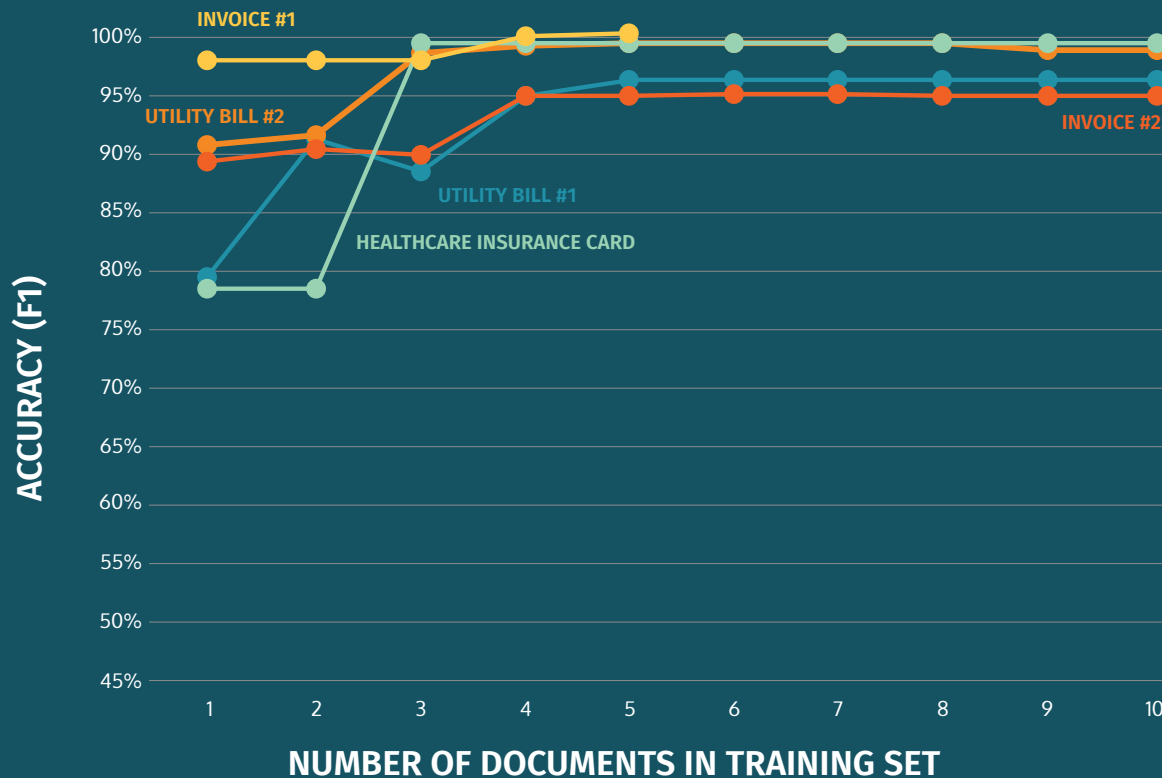
And here's another way to see GLYNT's improvement. What is Peak Accuracy for GLYNT? How many documents does it take to get to Peak Accuracy? The answers are in.

As the bar chart shows, GLYNT's F1 inched up to 99% in the past year. Previously strong results inched up.

And GLYNT is learning faster than ever. Provide five sample documents, and our AI will get to Peak Accuracy. Not every document set is the same, and these are our average results, but they are indeed impressive.

In the past year, GLYNT has reduced the number of documents needed to get to Peak Accuracy by 50%. This saves our users valuable time. And we're on a path to reduce this even further.





Digging deeper, here are some results by document type. As you know, a health insurance card is quite different in layout and data than a utility bill. Yet GLYNT has the same great results after training on three documents.

And look at how quickly GLYNT learns for any of the document types. To our users, this means that GLYNT can be quickly configured to the documents at hand, and that if the documents change, GLYNT will quickly learn that too.

Because GLYNT trains so quickly, it is a great solution for a huge range of documents. You're not limited by document type.

GLYNT is putting rapid learning under the control of the Business User. Improve our AI in just a few minutes, regardless of document type. And in every department.

Transforming Documents-to-Data: What is Needed for a Great ROI?

GLYNT's accuracy results naturally lead to the other drivers of a great ROI. These include customized training, fast time to data, and the ability to retrain quickly. And the user can select the fields they want, all without coding.

As you can see, accuracy alone is not enough. A great solution delivers that and more. The following table details how a full suite of features affects the ROI of a documents-to-data solution.

But accuracy is key, as it is the single most important factor in reducing expense. When the data result is correct and validated, it does not need a human review, and there is a reduced role for manual data review and clean up.

So the GLYNT team is focused on accuracy, and also how quickly our users can get to highly accurate data. The good news from our annual update is that GLYNT just keeps getting better.



Transforming Documents-to-Data: The Features that Deliver a Great ROI

FEATURE	KEY QUESTION	IMPACT ON ROI	GLYNT	OTHERS
Coverage Across Document Types	Can the system handle a variety of document types? Or is it configured for just one type, eg invoices?	Businesses have an enormous range of document types, and every department is different. Get coverage for all the document types in your department.	All Types	Limited to Document Types that it has been trained on
Level of Accuracy	How accurate is the system? How much review time will it take to get to 100% accuracy?	Higher accuracy means less review time	99%	60 - 85%
# Training Docs by Doc Type	How many sample documents does the system need to get to first results? How many more are needed to get to peak accuracy?	Customers hate hunting down documents for training. It is often very time consuming for them.	5	1000s
Level of Cruft	Does the system provide lots of unwanted data, eg cruft? Or just the items you want?	Cruft are the unwanted fields. Expensive to filter thorough them. And they keep changing.	None	20 to 1 ratio
Retraining Option?	What if the document changes? How does the system update?	An estimated 25% of documents in an incoming flow will change each year. Without correction, performance degrades significantly	Retrain and adapt in minutes	Can't update results
Staffing Requirements	Expensive data scientists needed? Lots of dev resources needed?	Enable the Business User to lower the total cost of ownership.	Business User	Developer
Time to First Data	How long does it take from "hello" to first results?	Data wrangling is 70 – 80% of the time spent on most AI projects.	20min	2 weeks

Study Details

Sample Documents

Documents from five sources were used in this study. For each document set, 10 documents were made available to GLYNT's training feature. The remaining documents were kept for the holdout data set; the number of documents in this set ranged from 12 to 147. Five to seven fields were extracted per document set. The five document sets are:

- Healthcare insurance cards
- Business invoices from two vendors
- Utility bills from Pacific Gas & Electric and San Diego Gas & Electric

With strict privacy laws in healthcare (known as HIPPA), synthetic data sets were developed for the insurance cards and clinic visit summaries. An image of each was taken from the internet, and then the data elements were randomly varied across 50 synthetic document images. It is expected that these document sets will have fewer OCR errors, but no particular advantage in the extraction of specific data fields through GLYNT's machine learning system.

The data items highlighted in blue are the fields to be extracted from this document. While the healthcare insurance card is a very compact document, the utility bills are 3 - 5 pages each and data items from several pages per bill were used in this study.

The Method



The following method was used for each document set:

- Create a database from the 50 documents of the data values for the desired fields, e.g. Ground Truth
- Prepare an inventory of fields by document
- Separate the documents into a Training Set (10 documents) and a Holdout Set (40 documents). The Training Set is curated so that all of the desired fields appear in the group of documents
- Train 10 machine learning models for each document set. The first model is based on a Training Set of one document. The second model is based on a Training Set of two documents, and so on
- Use each model to extract data from the Holdout Set, producing extraction results 1 through 10
- For each set of extraction results, report the Precision, Recall and F1 scores

Study Details, cont.

The Metrics

To measure performance, a Ground Truth data set was developed for each document type. Ground Truth is the 100% correct data for the desired fields, including whether a desired field was printed on the document. There is some natural variation on which data fields are printed on otherwise remarkably similar documents.

With Ground Truth in hand, each set of extracted data was compared to the Ground Truth and the following states per data item were noted:

Correct = the data item was provided by the ML and is 100% accurate (as measured at the character level)

Incorrect = the data item was provided by the ML and contains one or more character errors

Missing = the data item should have been provided

Not Printed = the data item was not available and should not have been returned

Incorrect errors largely arise through OCR errors. For example, the OCR engine confused an 0 with an O. As OCR engines improve, the amount of incorrect data has fallen significantly.

To summarize the accuracy results, the standard F1 score is calculated as follows:

Precision = # Data Items Correct / Total # Data Items

Recall = (Total # Data Items - # Data Items Missing) / Total # Data Items

F1: $[2 \times \text{Precision} \times \text{Recall}] / [\text{Precision} + \text{Recall}]$

Study Details, cont.

Results by Document Type

The following table shows the Precision, Recall and F1 results for four documents. Overall the F1 increases quickly as the first documents are added to the training set. After about five documents the Peak F1 is reached and the results do not improve. For Utility Bill #1, the recall drops with the addition of the last two documents in the test procedure. This happened because the two documents had a slight variation in the presentation of one field and the GLYNT system needs a bit more training on that presentation to return to the Peak Recall and Peak F1.

While no AI system produces perfect results, the details by document demonstrate a remarkable level of performance from training on just a handful of documents. 'Few Shot' learning is a different approach to ML and as expected, it delivers high accuracy across document types.

# OF DOCUMENTS IN TRAINING SET	UTILITY BILL #1			BUSINESS INVOICE #1			UTILITY BILL #2			HEALTHCARE INSURANCE CARD		
	PRECISION	RECALL	F1	PRECISION	RECALL	F1	PRECISION	RECALL	F1	PRECISION	RECALL	F1
1	100	86	92	100	93	97	81	77	79	93	66	78
2	100	88	94	100	93	97	98	86	92	93	55	78
3	99	98	99	100	93	97	98	80	88	100	98	99
4	100	99	100	100	100	100	98	92	95	100	98	99
5	100	99	99	100	100	100	99	94	96	100	98	99
6	100	99	100	100	100	100	99	94	96	100	98	99
7	100	99	100	100	100	100	99	94	96	100	98	99
8	100	99	100	100	100	100	99	94	96	100	98	99
9	100	98	99	100	100	100	99	94	96	100	98	99
10	100	98	99	100	100	100	99	94	96	100	98	99


Study Details, cont.

Comparison with Standard AI

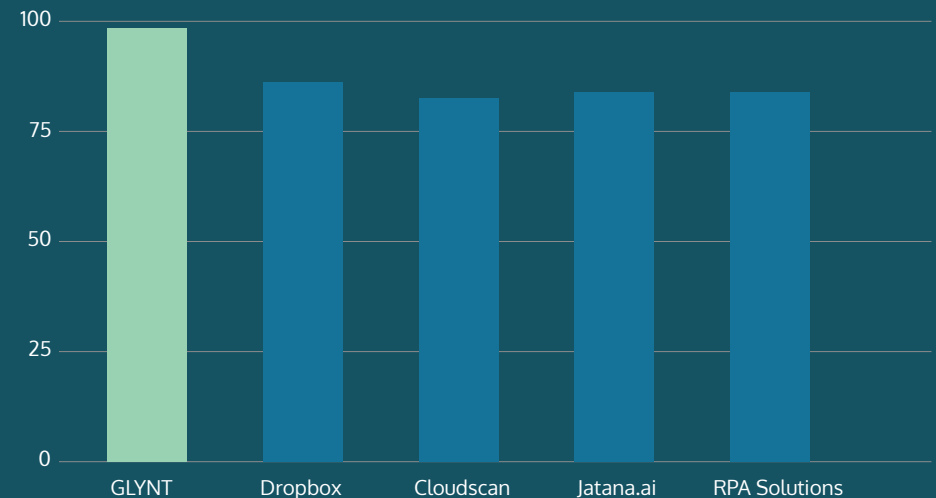
The GLYNT results are spectacular, and to provide context we compare GLYNT's performance with other systems that extract data from documents for document classification and data capture.

- Dropbox: 87% accuracy after training on more than 300,000 invoices
- Cloudscan: 84% accuracy after training on more than 326,000 invoices
- Jatana.ai: 85% accuracy after training on more than 14,000 documents
- RPA solutions: 65 – 85% accuracy after training on 1000+ documents
- GLYNT: 99% accuracy after training on 5 documents

NUMBER OF DOCUMENTS REQUIRED FOR FIRST DATA



REPORTED ACCURACY*



*Please note that vendors report various accuracy metrics and the results shown here may not be comparable


Technology

Here's a peek at GLYNT's advanced machine learning, and why we get such great results.

GLYNT uses an advanced technology known as 'Few Shot' machine learning (Machine Learning is a type of AI).

STANDARD ML

Automate high-volume standardized documents



50-1500

Training documents required

75%

Accuracy

'FEW SHOT' ML

Automate high-volume standardized documents



5

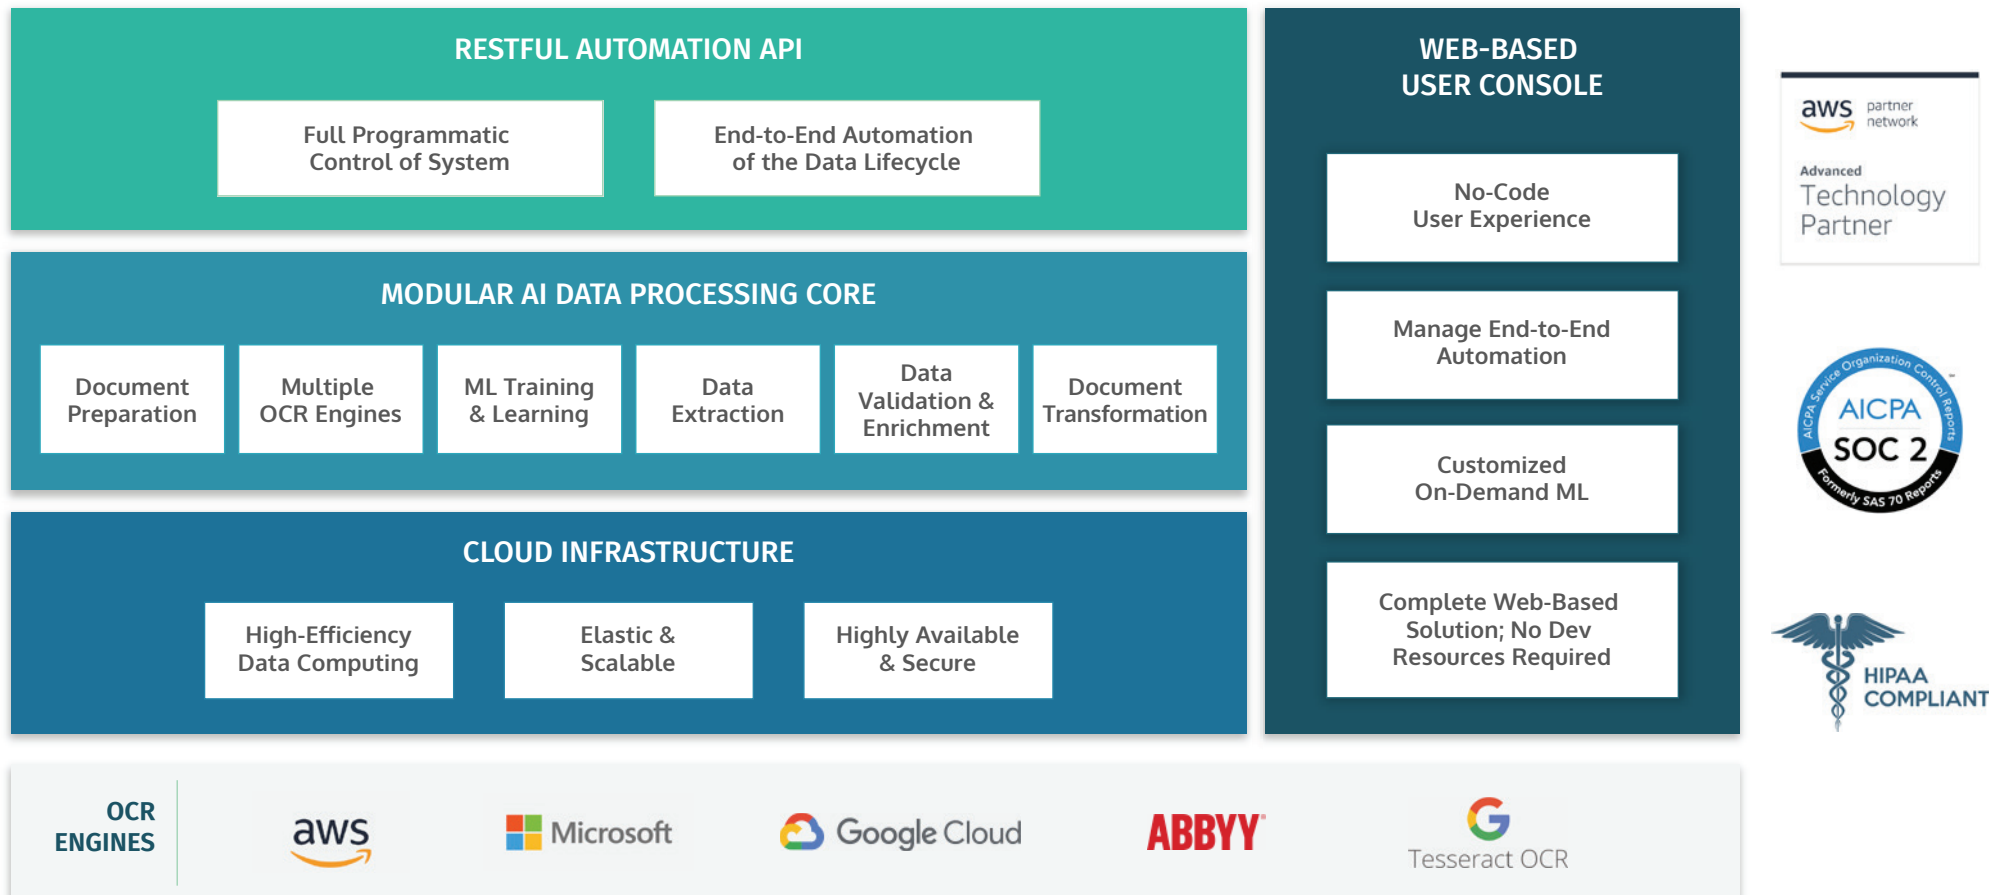
Training documents required

95%+

Accuracy

Technology, cont.

GLYNT self-contained and elastic tech stack, which makes deployments easy – just a REST API. With multiple OCR Engines at hand, GLYNT can give you the best-in-class results.



Every industry has treasure troves of data trapped in PDFs, scans and faxes. Liberating this data can change industry data flows by opening up access, speeding data processes and improving insights and understanding. GLYNT with 'Few Shot' machine learning makes the extraction of loosely structured data faster, better and cheaper than ever before.

We're constantly surprised at the persistence of manual data entry. In customer conversations we see a reliance on human teams to conquer document variety because legacy technologies have failed to automate complex document flows. This annual report from GLYNT reaffirms that our solution performs at a very high level and provides the accessible automation needed by Business Users everywhere.

READY TO TRY GLYNT?

If you would like to perform the Knee of the Curve analysis on your documents, get in touch with us. We have the tools ready to go!

GET IN TOUCH



530 Showers Dr, #7416, Mountain View, CA 94040

