



# Automating Paperwork

A practical overview for enterprise

Edward Benson

## Executive Summary

Read the full version at [www.automatingpaperwork.com](http://www.automatingpaperwork.com)

*Automating Paperwork is a practical, no-hype technical guide for business leaders, product managers, and operations teams who are pursuing a document automation initiative at their company. Its goal is to provide an end-to-end tour of the technical decisions and tradeoffs involved so that you can prepare for success, know what to expect, and ask the right questions of engineers and vendors along the way.*

## NLU CHANGES THE WAY WE DO BUSINESS

Natural Language Understanding (NLU) is one of the fastest moving areas of artificial intelligence today. NLU gives computers the ability to read documents, analyze emails, extract form fields, and make nuanced decisions about the information within to a degree that approaches human capability. While many Fortune 500 companies still rely exclusively on humans to handle business documents, those that integrate NLU fundamentally transform the way they do business. From investment banks, to law firms, to pharmaceutical companies, AI is helping teams find new ways to work with and extract value from the documents that power their business.

Beyond the mere efficiency of automation, NLU reduces the scarcity of knowledge as an operating resource. Information can be retrieved more quickly. Correlations can be drawn more broadly. Data can be reviewed more thoroughly. Information work considered high-touch and labor-intensive today will become more broadly available, changing business models and enabling new products and services.

There are many applications of NLU, from chat bots to sentiment analysis to language translation. But for large businesses who deal in information, document extraction is a high-value, low-risk way to start an NLU initiative. Document extraction is the process of reading documents intended for humans and locating information within them that provides the inputs to business processes. Account

onboarding, contract analysis, audit and diligence, customer support, and business intelligence are all areas of work that can be transformed by the ability to ask your filing cabinets very specific questions and have them answer you back.

## VETTING DOCUMENT EXTRACTION PROJECTS

The most important document extraction work happens before engineers are involved. It's the work of vetting a business process to learn if it's a good target for automation.

Automating Paperwork contains guides, worksheets, and checklists for a four-step vetting process that will develop a clear picture of whether a document extraction project is likely to succeed. These steps are:

- Defining a business process in terms of the inputs it requires
- Gathering and examining the documents that contain those inputs
- Simulating the information extraction from those documents
- Simulating the business process using only extracted data

Performing these steps will help you identify and resolve problems at the planning stage before they become expensive to correct at the implementation stage. And the resulting brief is an invaluable document to hand off to your technical team.

## SHOULD YOU BUILD OR BUY? WHO DOES WHAT?

Like most software categories, document automation inevitably involves using third-party software. The question isn't whether to build custom solutions or buy off the shelf, it is:

- **What style of software to use?** Industry-standard problems may be solvable with APIs and off-the-shelf software, whereas more bespoke problems may require an NLU development platform.
- **What modeling approach to prefer?** The more complex your data and extraction needs, the more likely deep learning may become necessary.
- **Who maintains and updates the system?** Rule-based systems generally require engineers to craft domain knowledge rules. Learning-based systems generally rely on domain experts labeling data.

## YOU MAY NEED A WHOLE PIPELINE

At scale, document extraction is more than just pulling information out of documents. You also have to convert document formats, identify the right part of the document from which to extract, route information to the right models, swap in new model versions, validate model results, and establish feedback loops with QA and human-in-loop systems. Think of this larger pipeline as the mailroom underneath the Manhattan skyscraper, making sure every parcel ends up in the right office. It might not be the piece of your business that makes the headlines, but without it everything else stops.

## WHAT DO I NEED TO KNOW ABOUT MODELS?

There is a lot of hype and fast movement in the world of AI models — the brains of a document automation pipeline. Pragmatically, the best model for you is the one that works. But it's also true that there are meaningful differences between types of models that may be important for you, both now and with respect to your future growth.

Rule systems have a long track record of producing explainable answers with minimal processing requirements. But as the complexity of the extraction task rises, rule systems often hit a performance and maintenance ceiling: there are too many possible rules for any team of humans to capture them all.

Deep learning systems have greater processing requirements and less intrinsic explainability, but they are capable of solving tasks of vastly more nuance and complexity. As the field of deep learning matures, practitioners are increasingly developing "one-size-fits-all" models that require less bespoke effort from engineers as well: domain experts need only provide a few examples of the correct inputs and outputs, and the system fine-tunes a pertained model to figure out the rest.

Unlike some areas of software, NLU is still a space in which hardware capabilities have a tremendous influence on software capabilities. The most advanced models push computer hardware to its limits, which means you can expect to encounter tradeoffs between quality, speed, convenience, and cost.

## WHAT CAPABILITIES ARE ON THE HORIZON?

Deep Learning has completely changed the way NLU practitioners design models, and it's reasonable to expect that advancements in Deep Learning are largely what will drive advancements in NLU in the coming decade. The following is a list of capabilities on the frontiers of NLU work today that are likely to become commonplace in industry over the next few years:

- **"One size fits all" neural architectures.** Standardized neural architectures that adapt to a broad range of usage scenarios will result in great engineering and IT efficiencies: less bespoke work, less specialized labor, and greater resource sharing between projects.
- **Lower training data requirements.** Pre-training models that can be fine-tuned for each customer will revolutionize the way in which bespoke models are developed and deployed. Customers will only need to gather enough data for last-mile adjustments rather than amassing the mountain of data traditionally required.
- **Multilingual by default.** Multilingual models are achieving accuracy levels that make them viable as replacements for monolingual models. This will allow global companies to have a unified technical approach across regions.
- **Multi-modal extraction.** Tomorrow's NLU models will analyze not just text, but also layout and graphics. This will provide improvements across the board, especially on tabular data, complex forms, and documents with graphical elements.
- **Domain modeling.** New systems will combine field-by-field information extraction with higher-level domain models that group fields together into objects and catch semantic errors that might otherwise have been invisible at the field-level.

## CONCLUSION

The Fortune 50 is hard at work incorporating NLU into their core operations, and the Fortune 500 is just getting started. Now is the time for business and operations leaders to begin crafting and implementing a strategy to leverage artificial intelligence across their company's information work.

Companies that do this now will benefit from a riches of momentum over the coming decade as NLU capabilities mature, with computers gaining greater ability to help humans read and make decisions about the paperwork that drives their business. This frees humans to spend time on more valuable and strategic tasks and opens new product possibilities for the enterprise at large.

—

*Edward Benson runs an AI prototyping studio called Heavy Foundry. He has a PhD from MIT, where he focused on Natural Language Processing and Human-Computer Interaction. He's spent his entrepreneurial career in Silicon Valley building the systems that process and understand documents at some of the world's largest companies.*

*Automating Paperwork is available to read at <http://www.automatingpaperwork.com>*