



Data Science White Paper

Automating Product Packaging Inspection with Deep Learning

White paper examining how firms should be automating product packaging inspection with deep learning.



Machine Learning

Mosaic data scientists collaborate with customers, digging deep into the data to inform design and deployment of custom ML tools that make a difference.



Artificial Intelligence

Mosaic integrates powerful AI tools into clients' existing technology stack to solve complex business challenges



Business Analytics

Mosaic helps corporations of all shapes and sizes take advantage of their data, transforming their decision-making processes.

THE PROMISE OF AI-DRIVEN AUTOMATION FOR QUALITY MANAGEMENT IN PRODUCT PACKAGING

Product packaging lines require their own custom inspection systems to perfect quality, minimize false rejects, improve throughput, and minimize the risk of recall. As larger players consolidate the industry, their product lines become increasingly more complex, leading to a need for automated inspection technology.

As with all applications of artificial intelligence (AI), there are winners & losers. Early adopters will realize hard cost benefits in significant operational cost reduction & fewer recalls. Soft benefits will materialize in the form of better customer experiences through better adherence to promised delivery dates & outpacing retailer shipment expectations, gaining more prominent shelf placement.

Deep learning, specifically computer vision and natural language processing, can be designed to identify defects during the product packaging process. These deep learning models can verify that a label on a package is present, correct, straight, and readable.



Artificial intelligence is the engine that powers an intelligent robotics packaging line; deep learning techniques like computer vision & NLP, are what powers AI

Packaging bottles, cans, cases, and boxes; present in industries like food and beverage, consumer products, and logistics, cannot always be accurately inspected by humans in an efficient and accurate manner. For applications that present variable, unpredictable defects on variable surfaces such as those that are highly patterned or suffer from specular glare, manufacturers have typically relied on the flexibility and judgment-based decisions of human inspectors. Yet human inspectors have some very large tradeoffs for the modern consumer packaged goods industry: they are not scalable.

For applications that resist automation yet demand high quality and throughput, deep learning technology is an effective new tool at the disposal of application engineers in the packaging industry. Deep learning technology can handle all different types of packaging surfaces, including paper, glass, plastics, and ceramics, as well as their labels. Be it a specific defect on a printed label or the cutting zone for a piece of packaging, deep learning algorithms can identify all these regions of interest simply



by learning the varying appearance of the targeted zone. Deep learning can then locate and count complex objects or features, detect anomalies, and classify said objects or even entire scenes. And it can recognize and verify alphanumeric characters using an optical character recognition (OCR) library.

WHY DON'T WE HEAR ABOUT MORE PACKAGING SUCCESS STORIES?

Sounds like every firm that runs a packaging line should utilize this technology, yet many fail to deploy correctly. Part of the reason is the sheer amount of data science knowledge it takes to train, deploy, and monitor a suite of deep learning models running in production.

There is no substitute for partnering with an experienced data science firm to learn the nuances of your business processes and data. An off-the-shelf product simply cannot provide the same level of customization and accuracy provided by training specific algorithms to solve different parts of the packaging line. Your data is going to be unique to you, and if you truly want to embrace the cost savings offered by automation, you need to treat the project as development and not a software purchase.

Computer vision, natural language processing, and other deep learning techniques need to be monitored and tuned over time to keep performance high. Failing to understand the intricacies needed at every step of this development & deployment leaves a business at a high risk of a failed AI endeavor.

QUALITY USE CASE ROAD MAP

The rise of machine learning & artificial intelligence, specifically computer vision & natural language processing techniques, offers any company that packages products a large opportunity for automation efficiencies and improved accuracy. Significant expertise in designing & deploying these advanced algorithms is needed to pick up the incredibly nuanced details in the data. In the following white paper, Mosaic will detail an AI-driven approach to solving this problem.

We will also call out specific process points we think are ripe for improvement and discuss the types of models best suited to solve them.

- Label Accuracy & Verification
- Defect Detection of Product and/or Packaging
- Lot Placement & Distribution Management

Each of these 'use cases' should not be completed in a vacuum. Each step in the road map provides critical quick wins while building toward a full-scale solution, allowing the production company to control for costs & analytic quality.

Label Accuracy & Verification

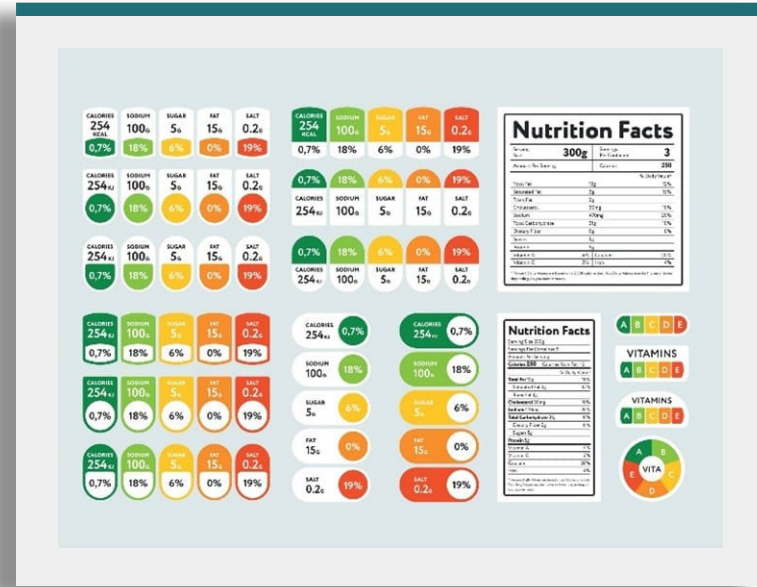
Any production company needs to provide accurate product packaging information in their catalog. Manually converting this type of information into usable data is incredibly time-consuming and error-prone. Think about a quality team tasked with updating product specs for millions of products. And think of the reverse, quality checking the information on the packaging against a



master data set to make sure the packaging is accurate instead of using the package as the source of truth.

Not only is this team under pressure to be extremely accurate, but they need to implement a process that is able to keep up with distributor demands. These firms face harsh financial penalties and even litigation risk when product specifications such as quantities, packaging dimensions, ingredients, and allergen or toxicity warnings are incorrect. This problem tends to grow in scale as you think about a company providing products in different countries, needing to integrate multiple languages and images.

Manually converting product specification information into usable data takes a lot of time. There is significant human involvement, as quality teams must find product information on packaging and then enter it into a dedicated data field in a schema used to record this information. Just look at how much information needs to be processed in the following image.




Deep learning can provide the key image processing capabilities to help automate this process. Deep neural networks can be trained to identify the regions of interest in the image, containing tables, paragraphs, images, etc. Supervised deep learning models such as [convolutional neural networks \(CNNs\)](#)¹ can provide state of the art performance. However, this approach requires a large training set – labeled example images of the information the neural network will be trained to extract – and significant compute resources such as graphical processing units (GPUs) for training and tuning the models.

Alternatively, when insufficient training data is available to train a supervised deep learning model, we can use unsupervised feature extraction methods to detect basic structural elements in the image, which can be leveraged to identify larger objects. For example, a Hough line transform can be used to identify pixels forming a line, and by identifying sets of parallel and perpendicular lines, we can identify tables or boxes containing structured information.



Once key objects and regions of interest are identified, [optical character recognition \(OCR\)](#)² can be used to digitize text in each region. OCR can automate the process of converting images of texts into digital documents, enabling automated parsing of the text. OCR is great to extract text from the raw packaging images, but additional text processing algorithms are needed to find, extract, and synthesize information of interest by matching keywords and phrases in the tables and raw text. Post-processing is then needed to clean up and structure results for acceptance criteria.



Defect Detection

Computer vision modeling is incredibly adept at learning what an acceptable product/package should look like as it moves down the line. A deep learning model should be trained on thousands of example images of accepted products/packages and thousands of example images of rejected products/packages. The true art of any AI system comes in the selection and tuning of the algorithm or set of algorithms that best fit the data, watching out for errors like overfitting and underfitting. It is imperative that the data scientist, automation engineers, and packaging quality team work together to train a computer

vision model early in the process to learn about what makes a product or package acceptable. If there are nuances in the data and decisions being made, they need to be identified early or else the model will not pick up on these details.

A properly trained model should easily detect wrinkles, rips, tears, warpage, bubbles, color, and printing errors. High-contrast imaging and surface extraction technology can capture defects even when they occur on less than ideal lighting and surface conditions. Once the model can accurately identify defects, there needs to be a scalable communication mechanism in place that alerts of the potential defect. This requires different analytical knowledge than algorithm development. Understanding the practices behind MLOps and agile software methodologies is imperative to developing an effective alerting system.



Lot Placement & Distribution Management

A completely robotic warehouse is a revolutionary concept today and a game-changer in several sectors. There is close integration between deep learning and making sure the processes are automated correctly. It seems that several of the robotic process automation (RPA) bot vendors tend to focus their efforts acutely on the manufacturing sector, especially the benefits it can add to packing. One of the critical first steps in this process revolves



around accurately identifying finished shipments and moving them to the correct distribution zone.

Computer vision and NLP can accurately identify lots when they come off the line and automate the next step in the process. Rather than having a human spend hours poring over each produced run, distribution and warehousing insights can be generated & communicated using many of the same techniques we have outlined above. A shipment might be at risk of showing up late to a key customer; deep learning can automatically flag a lot that needs to be prioritized based on current operating conditions. The more these models can learn your business, the better decision support they can provide.

For modern AI warehousing systems, packaging is a key player. AI tools can collect data from all data points and use it for recommending packaging updates. To perform above and beyond quotas, businesses need to ensure that the productivity of their workforce is a vital aspect of their business strategy, and AI can just do that, helping to identify events in real-time and recommend

optimal actions, from scheduling to safety. Companies that use AI to their advantage to build smarter warehouse processes will be able to free up cash flow that was previously spent on excess inventory expenses & costly penalty fees to spend on more productive business growth opportunities.

KEY BENEFITS OF AN AUTOMATED AI SOLUTION

Automating the work-intensive task of quality checking and rectifying information on product packaging can save huge amounts of repetitive work and increase process accuracy. This helps CPG companies and other manufacturers to maintain high quality standards under the aggressive manufacturing and delivery schedules necessary to keep distributors and retailers happy. The modeling process can be tuned to support a human-in-the-loop structure to identify packages at risk of poor placement.



HOW TO TAKE ADVANTAGE OF THIS AI | CALL MOSAIC!

Buying an out-of-the-box solution is simply not feasible due to the diversity of most firms' product catalogs and diverse packaging processes. An experienced data scientist can build an initial proof of concept or proof of value with the goal of scaling these models. Promising a plug-and-play solution does not account for the diversity of data feeding into the algorithm. Failure to understand the mechanics behind these deep learning algorithms leads to a failed project and a sharp drop in AI momentum.

Additionally, quickly understanding the art of the possible and what the exact requirements are is critical to a cost-effective implementation. Mosaic's Data Scientists understand the intricacies of different deep learning approaches and can help tailor the solution to what you need, helping to avoid thousands of dollars of ongoing commercial licensing fees and stalled modeling efforts.

Hiring a data scientist sounds expensive, but Mosaic frequently tackles these types of problems in an iterative engagement that starts with a proof of concept. Mosaic feels confident in being able to complete a Phase 1 pilot model delivered in 6-10 weeks. Future phases are typically gated to ensure quality control and clear ROI for a customer.

Before moving to future iterations of model development, Mosaic collaborates closely with customers to make sure the model insights make sense to the business. What good are model insights and recommendations if no one uses them? Mosaic develops all solutions with an eye towards usability, explainability, and continuous improvement.

Endnotes

1. <https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53>
2. <https://medium.com/states-title/using-nlp-bert-to-improve-ocr-accuracy-385c98ae174c>

