

# AI at Austin Data Labs

Solving Real Business Problems

SUSHIL VERMA, PHD



## Summary

Businesses that prosper optimize their operations with data science and AI to stay ahead of their competition. AI projects that succeed are led by teams that include industry experts that understand your business as well as they understand AI. In this paper, we'll demonstrate how we've used data science and AI to help a national commodity meat company, a global fast food chain, a mobile phone manufacturer, a healthcare company, and a national commodity dairy operation.

Management science has evolved to address increasingly complex business needs over the last few decades. New research is added every year to solve new challenges and create improvements for persistent existing challenges. With AI, the primary processes for success are digitization, quantification, and mathematical modeling and optimization of your complex business operations.

AI and data science are ever-evolving fields. New mathematics, statistics, and computation techniques must be implemented within your business on a regular basis to continue to see improvement in your operations with this new technology. It's important to work with a company who knows your business as well as they know data science, AI, and software, and understands which module will work best with our platform to quickly solve your problems. It's not always AI!

Many companies are struggling to move their AI initiatives past the "proof of concept" phase because they aren't working with the right vendor. They are chasing shiny objects they may not need. At Austin Data Labs, we are very excited by what AI has to offer and use it well when it's appropriate. Our passion is to put our customer first, however.

We understand what business problems AI is well-suited for and for what business problems it isn't. We understand that AI builds on the existing science of operations management, supply-chain planning, and demand planning, without replacing any of the prior knowledge. Any claim otherwise is at best misleading and at worst bound to fail. That's why we use a module system for our software. We start working with you to solve your challenges right away using our data science platform, and our AI module is as easy to slot in as a puzzle piece if your project warrants it.

How do you know when AI is right for you? Let's take a look:

## What problems in operations management are best suited for AI to solve?

If at least two of these attributes meet your situation, AI is appropriate for your project:

Attributes	Examples
<p>No clear theory underlies the prediction</p>	<p><b><i>How long would a piece of machinery last?</i></b>            This question, if answered well, can help manufacturing businesses greatly streamline their factories. But no fully developed theory exists to answer that question well. The best models are statistical in nature and have too large of a variance. <a href="#">AI deals with extreme variances better than your ERP.</a></p>
<p>Large number of often correlated variables impact the outcomes instead of just a few</p>	<p><b><i>Which products would be of most interest to a prospective consumer?</i></b>            Answering that question requires having deep knowledge about the consumer, their demographic profile, past purchases. This knowledge can be quantified but only through hundreds and if not thousands of variables. <a href="#">AI handles large numbers of variables well.</a></p>
<p>Variable values are either known or at least partially controllable in both historical and forward contexts</p>	<p><b><i>What does the demand and price for an agriculture (ag) commodity look like looking forward?</i></b>            Typical ag commodities' demand and supply is affected by many variables, both natural and man-made. But many of these variables such as rainfall, changing diet patterns, and the price and demand of other commodities that go into making this commodity are not fully known. However, partial knowledge is indeed available and can improve the accuracy of forward forecasts. <a href="#">AI is useful when dealing with variable knowledge sources and gaps in knowledge.</a></p>

## When is AI not suitable for solving your challenges in operations management?

Attributes	Examples
<p>A clear theory/model already exists that can be used to do the math of prediction or correlation</p>	<p><b><i>How much raw material do we need (and when) to meet the expected demand?</i></b></p> <p>There exists a well-developed theory of how to answer this question well in complex, real-world environments. A simple spreadsheet may suffice if the manufacturing process is a simple single-step. But a complex algebraic model may become necessary if the process has many steps, has many alternates with variable production and lead-times.</p> <p><i>This challenge is ideally suited for our platform but would not need the AI module.</i></p>
<p>Repeatability and explainability are of key importance</p>	<p><b><i>What would be the impact of a price change on demand?</i></b></p> <p>In most businesses, the planners would want to use a model that's easily explained and produces the same answers every time. Economists have studied this problem for a long time and have found certain simple models to be quite effective. Using AI for this purpose is possible but only at the expense of adding complexity while reducing transparency. <i>This is another challenge ideally suited for our SaaS platform and modules.</i></p>
<p>Problem is not of prediction but of making best trade-off decisions</p>	<p><b><i>What product-mix maximizes profits?</i></b></p> <p>In such situations, the predicting base demand may not be the key problem. Instead, what the business is trying to do is to guide the sales and operation to move towards a more profitable set of products. <i>Once again, this problem can be solved with our platform with no need for our AI module.</i></p>

## **ADL Case Study: Store-level Demand and Revenue Forecasting**

### **Customer**

- One of the world's largest global fast-food-chain brands

### **Overall Goal**

- Predict store-level revenue on a daily basis as a function of pricing, promotions, events, and other levers of demand
- This will become the foundation of a price-optimization tool for store managers

### **Challenges**

- A store's revenue on any given day is driven by many variables
- Store's often run promotions on multiple products at the same time, causing significant shifting of demand between products
- Competitive pricing and promotions have significant impact on the demand as well
- Area demographics, store's location in relation to competition, weather, local events etc. also play a big role on the revenue

### **Key Questions**

- What's a store's expected revenue under the planned promotions by day?
- How is a store revenue impacted by competitive pricing, area demographics and other major drivers?
- What promotions should I run in which store on which days to help reach my revenue and margin goals?

### **Solution**

- An AI-base demand forecasting system that takes into account hundreds of variables into account
- A revenue simulation engine that allows planner to study the impact of change in driver values on the revenue

### **Results**

- Delivered model that generates forecasts from transaction level sale history incorporating 20 external factors
- Increased forecast accuracy by 5%

# ADL Case Study: Commodity Workbench

## Customer

- A national retailer that runs their entire supply chain: from farm to processing to stores

## Overall Goal

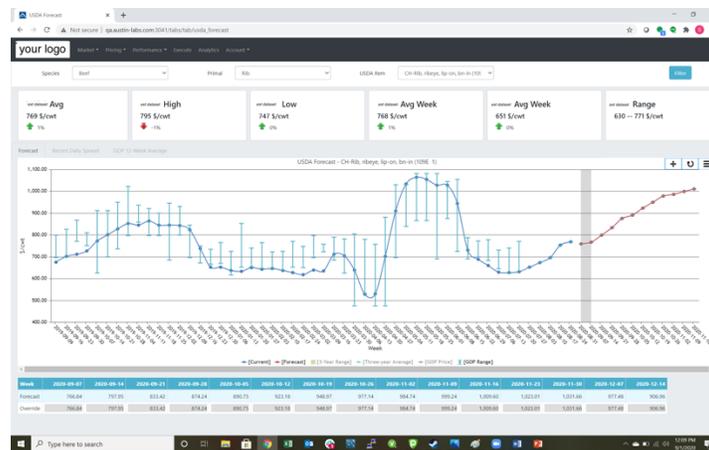
- Agricultural commodities are extremely volatile in price. A 5% price change in a week is not an uncommon occurrence
- Commodity processors struggle to price their products on a forward basis in presence of this volatility

## Key Questions

- How are commodity prices for one commodity related to other commodities?
- How do external events impact the commodity prices?
- What’s a good forecast for a specific commodity for next three months? What’s the degree of confidence in that forecast?

## Solution

- An AI-driven forecasting engine for agricultural commodity prices
- A workbench to help visualize all the related commodities in an integrated fashion, both historically and forward-looking



## Results

- 20% improvement in forecast accuracy over gut-feel decisions
- 1% improvement in revenue generated from forecast-driven prices

## ADL Case Study: Customer Retention

### Customer

- A fortune 500 global mobile handset manufacturer

### Background

- Unexplained return of phones in the warranty period and beyond
- Each returned phone eats away margin contribution of tens of phones
- Telcos share very little data about the consumers making it difficult to use marketing techniques such as surveys to pinpoint the root cause
- The manufacturer collected detailed usage data from the handsets  
Could the usage data unlock the key to understanding returns?

### Data

- A classical Big Data problem involving many terabytes of data. A Hadoop cluster was used for data storage and raw data processing.

### Solution

- A machine learning ensemble model to identify the most important predictors of returns. The model calculated a score for each consumer that represented the likelihood of that consumer returning the phone in near future
- The models sat between a Hadoop cluster and an end user interface and the results of the models fed a weekly root cause prioritization process aimed to continuously improve the product, support, and marketing communication

### Benefits

- A vastly more targeted process to address the primary cause of returns
- Lower returns, higher overall margins

### Results

- Processed data from 1 million cell phones, amounting to 4 terabytes each month
- Generated actionable insight into customer actions within 15 mins

## ADL Case Study: ICD-9 Coding

### Customer

- A leading healthcare natural language processing company

### Background

- Clinical notes contain a wealth of information about patients' medical history, present condition, undergone procedures, and prescribed medications
- The company's NLP engine converts the NL notes data into clear clinical concepts
- The goal put in front for us is to convert the output of the company's NLP engines into pinpoint inferences

### Opportunity

- The knowledge to be learned is highly distributed across many concepts
- The company has available millions of labeled and unlabeled clinical notes from newly mandated EHR systems
- The opportunity is to use artificial intelligence to help hospitals gain new levels of productivity and gain an unprecedented level of insights

### Data

- Terabytes of largely textual data, in a mix of semi structured and structured formats

### Solution

- A mixture of statistics, machine learning, and data mining techniques to identify the best combination of predictors of clinical inferences.
- The models being developed will sit between a no SQL data storage cluster and web-based user interface for hospital use

### Benefits

- Lower costs of deriving clinical inferences through artificial intelligence and automation
- Enable hospitals to cost effectively meet the new ACO quality guidelines

### Results

- 20% increase in productivity of human coders with the help of AI recommendations

## ADL Case Study: Document Classification

### Customer

- A multi-national oil and gas company

### Problem Background

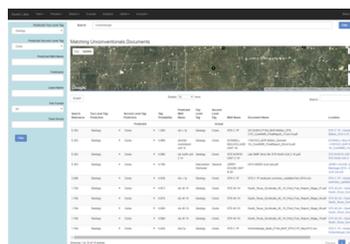
- Unstructured data trapped in millions of documents contained valuable information not available in any structured data systems
- Previous efforts to manually organize this information were found to be not scalable
- Structuring this data using modern machine-learning tools would open up a world of business applications (akin to how a search engine makes the world-wide-web data useful)

### Challenges

- Manually reorganizing the directories into a standard classification is prohibitively costly
- Inconsistent filing and directory design is the norm (engineers are free to organize their documents as they see fit)
- Wide variety of document formats and information presentation
- Duplication, misfiling, versioning etc. exacerbate the problem further

### Objective

- Test a machine-learning approach to partially tag and structure a sample of 400K documents
- Two aspects of structuring were tested on the key file formats (PDF, Excel, Word, Text)
- Identify well names from the information contained in the document
- Classify each document into a standard hierarchy of life-cycle and functional tags



The screenshot displays a software interface for document classification. On the left, there is a sidebar with navigation options. The main area is split into two panes: the top pane shows a satellite map of a geographical area, and the bottom pane displays a table of document classification results. The table has multiple columns, including file names, dates, and classification tags. The interface is clean and professional, typical of a data management or analytics tool.

### Results

- 90%+ accuracy in document classification with less than 10% of the training data set

## ADL Case Study: Failure Prediction

### Customer

- An international technology enterprise in the energy industry

### Overall Goal

- Investigate effectiveness of machine-learning models in enhancing reliability and efficiency of gas compressors
- Test the core problem in meeting the overall goal – How accurately can machine-learning models predict maintenance events in an real-world data set consisting of seven months' worth of data from six booster stations?

### Key Questions

- Is the data rich, clean, and sufficient enough at this stage?
- Where are the key gaps in the data?
- Is the sensor coverage ample to be able to provide accurate models?
- Which events are more difficult to predict?
- Which sensors are useful and which ones aren't?

### Solution

- Certain events can be predicted well-ahead of time with high degree of accuracy
- Accuracy of prediction varied significantly with the type of event
- Catastrophic failures are rare in the real-world. The focus shifts to predicting events that are leading indicators of needed maintenance

### Results

- 90% accuracy in anomaly detection
- 80% accuracy in event prediction

## **Why is industry expertise as important as AI knowledge when working with a vendor like Austin Data Labs?**

When you work with a vendor that doesn't understand your industry, your team ends up doing most of the heavy lifting for the project. It's taxing on your resources and creates resentment among your team members, which can cause them to be reluctant to adopt the new software as part of their routine. Vendors that don't understand your unique industry and your pain points create software for you that doesn't accurately and adequately address the challenges you need to solve, nor does it make your workload easier - how can it? That vendor has never shared your process and does not bring any domain expertise to the table.

You can't easily and quickly solve complex industry problems in industries like commodity meat, dairy, agriculture, specialty chemicals, fast-moving consumer goods, and food and beverage if you don't go into the project understanding how those industries work and where bottlenecks and breakdowns could be occurring in the processes. That's why Austin Data Labs has a team that includes both world-class data scientists and industry experts with decades of experience in your industry. We focus on customer experience and satisfaction as well as results, so our data science products and AI solutions are some of the best in the business.

# Don't Get Left Behind...

**CALL US TODAY.**

**Here's what you can expect:**

- Deep industry expertise
- Cutting edge AI and advanced data science
- A team focused on customer satisfaction

---

Email [info@austin-labs.com](mailto:info@austin-labs.com) today

