Towards a Sustainable Food Supply Chain Powered by Artificial Intelligence

Anonymous Author(s)
Affiliation
Address
email

Abstract

About 30-40% of food produced worldwide is wasted. This puts a severe strain on the environment and represents a $165B loss to the economy in the US alone. About 40% of all waste occurs at the retail level and downstream; this paper explores intelligent systems that help retail stores reduce food loss by forecasting demand and optimizing store decisions. Specifically, we describe a model-based reinforcement learning system for perishable inventory management; in an extensive simulation on real-world data from a US supermarket chain, the system provides reductions in retail waste of up to 80%.

1 Introduction

About 30-40% of food produced worldwide is wasted (Gunderson, 2012). This puts a severe strain on the environment and represents a $165B loss to the economy is the US alone.

The environmental impact of food waste is immense. Food production accounts for 92% of water use (Hoekstra et al., 2012 [5]) and 25% of green house gas emissions (Vermulen et al., 2012 [7]). According to a study by the Food and Agriculture Organization (FAO [1]), food waste generates greenhouse gas emissions comparable to those of a country the size of Russia.

In addition to its environmental impact, food waste also results in significant economic losses. According to the FAO study, the economic impact (aggregated across the world) of all the food lost to waste in the year 2007 represented total losses of about USD 750 billion (measured in 2009 prices). This amount is comparable to the gross domestic product of Turkey or Switzerland in 2011 [1].

Lastly, food waste represents a major societal challenge from a moral perspective. According to the United Nations, approximately one in every nine people in the world suffers from hunger, defined as not having sufficient access to food be healthy. Hunger is estimated to kill a greater number of people every day than diseases such as AIDS, malaria and tuberculosis. From a fairness perspective, society has a moral obligation to distribute fundamental resources such as food in a way that will satisfy the basic needs of the largest possible number of people.

1.1 The Causes of Food Waste

Food loss occurs at all stages of the supply chain, from the farm to the consumer. In industrialized countries, the retail and consumer stages are high contributors to waste; in developing countries, the most significant food losses occur at the farm, in part due to the the limited adoption of technology.

In industrialized countries, about 40% of food is lost at the retail and consumer stages of the supply chain. Consumer behavior and supply chain efficiency are important contributing factors to food waste. The 2018 edition of the Retail Food Waste Action Guide (published by a consortium of leading retail chains [2]) identifies the following five solutions as having the highest potential to reduce

Submitted to 32nd Conference on Neural Information Processing Systems (NIPS 2018). Do not distribute.

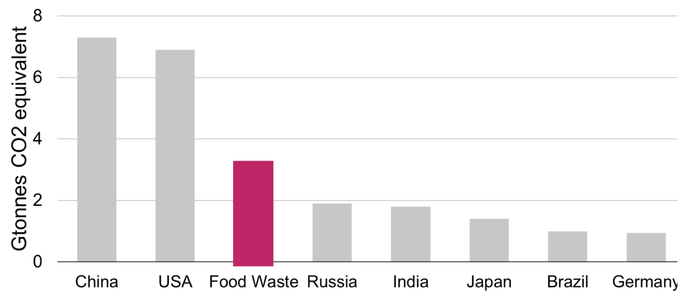

Figure 1: The environmental impact of food waste. Food production, especially the production of meat, greatly contributes to about 25% of green house gas emissions (Vermulen et al., 2012). In absolute terms, this amount exceeds the greenhouse gas emissions of Russia, India, or Japan. Source: Food and Agriculture Organization, 2013.

waste at the retail and consumer level: (1) enhanced demand forecasting, (2) dynamic pricing and markdowns, (3) dynamic routing, (4) cold chain management, (5) improved inventory management.

This paper explores intelligent systems that help reduce food loss. Specifically, we focus on applications of machine learning to demand forecasting as well as on model-based reinforcement learning methods for inventory management. The goal of this paper is to bring the food waste problem to the attention of the machine learning community, as well as to demonstrate how modern machine learning techniques have the opportunity to make a significant impact.

2 Intelligent Systems for Reducing Food Waste

This paper explores intelligent decision-making systems that assist store operators in managing perishable food items in the presence of uncertainty over future demand, weather, and other key deciding factors. We describe a system that provides recommendations on purchasing decisions. The system is comprised of two components. First, a prediction component generates demand forecasts; then, a model-based planning algorithm uses this learned demand to recommend purchasing decisions that optimize the store's overall utility function. The result of this system is a reduction in waste in retail stores, and an increase in item shelf-life and freshness.

2.1 Demand Forecasting

Food supply chain planning requires very accurate forecasting algorithms due of the extremely short shelf-life of perishable items. The accuracy of forecasts is much more critical compared to non-perishable goods, because over-ordering errors carry a much higher financial cost.

Multi-Task Learning. Retail food chains typically sell thousands of different items across hundreds of stores. This generates large amounts of time series data; fully leveraging this data requires multi-task learning across thousands of items across hundreds of stores. Specifically, predictions are highly impacted by rare events such as holidays or sales. These events are seen only a small number of times for each item. Achieving accurate predictions requires multi-task models that learn to forecast rare events across a large number of time series jointly.

Probabilistic Time-Series Forecasting. Accurate planning requires predicting not just a point forecast, but an entire distribution over model demand. We formulate probabilistic time series forecasting as multi-task quantile regression over multiple quantiles at once. In addition to enabling more accurate planing, probabilistic predictions are also a key component of interactive systems that can assess their confidence before making recommendations to a human operator.

2.2 Model-Based Planning

Given a set of demand forecasts, our system computes daily replenishment orders that minimize waste. We formalize this task as a Markov decision process (S, A, P, R). States $s \in S$ are sets of tuples $\{(q,l); l = 1,2,...,L\}$; each (q,l) indicates that the store carries q units of the item that expire in l days (L being the maximum shelf-life). Transition probabilities P are defined through the following process: on each day the store sells d units (a random quantity) which are removed from the inventory in s (items leave in a first-in first-out manner); the shelf-life of the remaining items is decreased (spoiled items are thrown away). Actions $a \in A$ correspond to orders: the store receives items with a shelf life of L before entering the next state s'. Finally, actions are chosen to optimize the reward R, which can account for both food waste and store profits. In our experiments, we set R to be the sum of waste and unmet demand due to stockouts.

In order to tractably solve this problem, we use online planning using Monte Carlo Tree Search (MCTS [3]). We perform a search over a finite time horizon using an epsilon-greedy algorithm to select new actions at each level of the search tree. New states are sampled from a transition model defined by the probabilistic forecasts learned by the prediction module of our system.

3 Experiments

We test our system on a public Kaggle dataset as well as on a real dataset obtained from a large US supermarket chain.

3.1 Kaggle Experiments

To demonstrate the effectiveness of our approach, with start with the publicly available grocery sales dataset from the Corporacion Favorita Kaggle contest.

Setup. We forecast the highest-selling item (#1503844) and use data from 2014-01-01 to 2016-05-31 in stores #1-4 for training and data from 2016-06-01 to 2016-12-31 for testing. We train a feedforward neural network with two layers of 128 hidden units with a dropout rate of 0.5 and parametric ReLU non-linearities. We feed the model autoregressive features from the past four days as well as binary indicators for the day of the week and the week of the year. We use dropout probabilities combined with the calibration method of Kuleshov et al. [6] to derive probabilistic forecasts.

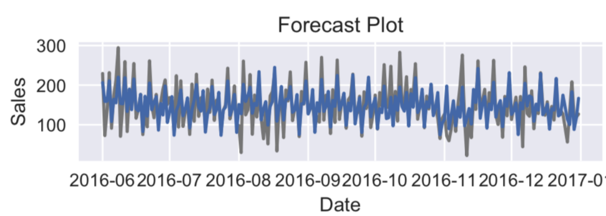

Figure 2: Historical sales of a perishable item (gray) from the Kaggle dataset, and our forecasts (blue)

Prediction Accuracy. We obtain mean absolute percent errors of 17.3-21.8% on the test set across the four stores. Our probabilistic forecasts are also calibrated: our 90% confidence interval correctly contains about 90% of the true outcomes.

Waste Reduction. We perform a simulation experiment in which we we use the Bayesian neural network from the previous section as our learned model of the environment (i.e., the state transition function). We then use Monte Carlo Tree Search to determine the action minimizing the sum of waste and out-of-stocks at each step. We evaluate the agent on the test set, using the historical sales as a proxy for unseen demand. Item prices and costs are set to 1.99 and 1.29 respectively; items can be ordered three days a week in packs of 12 and arrive on the next day; the shelf-life of new items is always five days.

Our system resulted in waste of 0.98-1.14% across the four stores, with only 0.61-0.89% of demand being unfulfilled due to under-ordering. These numbers suggest that our system has the potential to significantly impact food waste.

3.2 Real-World Experiments

We have been developing our technology jointly with a 100-store US grocery chain. Our pilot partner has released to us two years of historical data for the produce department in 27 stores, including past sales, shipments, prices, promotions, and other key data elements. We performed a historical simulation, and found that the our system dramatically improves the effectiveness of humans, increases store efficiency, and reduces food waste.

Setup. The first, historical, phase of testing involves training forecasting models on historical data up to 2017, and computing forecasts for 2017. We then run a simulation over 2017 data in which the store observes the true historical sales, and makes decisions based on the pre-computed predictions. Real store waste is estimated by looking at the inputs minus the outputs that are entering the store (this may slightly overestimate waste by including additional factors such as waste). Our system's waste is observed directly within the simulator, following the same methodology as for the Kaggle dataset.

We use a gradient boosting forecaster based on the popular XGBoost library [4]. In this initial experiment, we feed the model historical sales and shipments from the past four days, 7-, 14-, and 28-day rolling means of historical sales, binary indicators for the day of the week and the week of the year, sine and cosine features over the number of days elapsed in the year, features for popular US holidays, historical prices for the last four days, as well as indicators of price changes on the prediction day.

We pre-filtered our dataset to exclude items with potentially incorrect data (we filtered items whose shipment data exceeded sales data), and we filtered out data with sporadic sales, defined as having recorded sales on less that 70% of days. This process excluded 35% of input data by sales volume.

Waste Reduction and Store Efficiency. Our system incurred unit waste of 1.7% (aggregated across all stores and items), with only 1.1% of demand being unfulfilled due to under-ordering. These numbers are significantly lower than the 14% industry standard for food waste in supermarkets, and again suggest that our system has the potential to significantly impact food waste.

By examining the historical data from the supermarket chain, we can measure their historical levels of waste, and estimate the level of improvement offered by our machine learning system. In our experiments, we observed average reductions in food waste of 30%, with some items having 80% improvements.

Reducing food waste also increases store profits, which incentivizes store operators to adopt our system in practice. In our historical experiment, we were able to create on average $380K in additional profits per store per year. In addition, our inventory management system reduced stockouts by up to 80% and reduced inventory holding levels by up to 3-fold.

Long-Term Health Benefits. More generally, our system will increase the average freshness of items sold in the store, which will have tangible benefits on human health. In addition, food freshness will drive additional customers to stores, further incentivizing them to adopt waste reduction technology.

4 Conclusion

The goal of this work has been to bring the food waste problem to the attention of the machine learning community, as well as to demonstrate how modern machine learning techniques have the opportunity to make a significant impact. We have described an approach based on model-based reinforcement learning that has been shown in simulation to significantly reduce food waste in a real-world supermarket chain.

Our system has the potential to have a tangible impact on the environment (including water usage and green house gas emissions). In addition, a more efficient supply chain will reduce food prices and make food accessible to a larger number of people.

References

[1] Food wastage footprint. http://www.fao.org/docrep/018/i3347e/i3347e.pdf.

[2] Food waste retail action guide. http://www.refed.com/downloads/retailguideweb.pdf.

[3] Cameron Browne, Edward Powley, Daniel Whitehouse, Simon Lucas, Peter I. Cowling, Stephen Tavener, Diego Perez, Spyridon Samothrakis, Simon Colton, and et al. A survey of monte carlo tree search methods. *IEEE TRANSACTIONS ON COMPUTATIONAL INTELLIGENCE AND AI*, 2012.

[4] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 785–794, New York, NY, USA, 2016. ACM.

[5] Arjen Y Hoekstra, Mesfin M Mekonnen, Ashok K Chapagain, Ruth E Mathews, and Brian D Richter. Global monthly water scarcity: blue water footprints versus blue water availability. *PLoS One*, 7(2):e32688, 2012.

[6] Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning using calibrated regression. pages 2801–2809, 2018.

[7] Sonja J Vermeulen, Bruce M Campbell, and John SI Ingram. Climate change and food systems. *Annual Review of Environment and Resources*, 37, 2012.