Towards a Sustainable Food Supply Chain Powered by Artificial Intelligence

Volodymyr Kuleshov [1] Harlan Seymour [1] Danny Nemer [1] Sean Meador [1] Nathan Fenner [1] Matthew Schwartz [1]

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 US economy. This paper explores how artificial intelligence can be used to automate decisions across the food supply chain in order to reduce waste and increase the quality and affordability of food. We focus our attention on supermarkets — combined with downstream consumer waste, these contribute to 40% of total US food losses — and we describe an intelligent decision support system for supermarket operators that optimizes purchasing decisions and minimizes losses. The core of our system is a model-based reinforcement learning engine for perishable inventory management; in a real-world pilot with a US supermarket chain, our system reduced waste by up to 50%. We hope that this paper will bring the food waste problem to the attention of the broader machine learning research community.

1. Introduction

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

Specifically, food production accounts for 92% of water use (Hoekstra et al., 2012 (Hoekstra et al., 2012)) and 25% of green house gas emissions (Vermulen et al., 2012 (Vermeulen et al., 2012)). According to a study by the Food and Agriculture Organization (FAO (fao)), food waste generates greenhouse gas emissions comparable to those of Russia (Figure 1).

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 (fao).

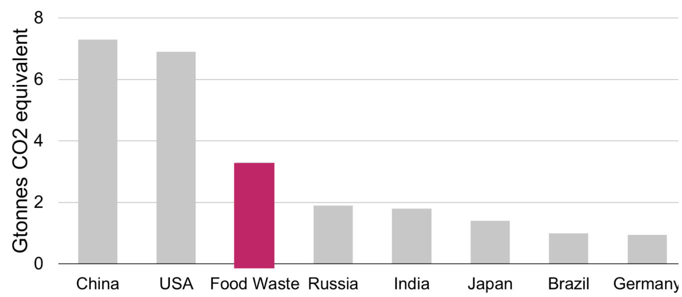

Figure 1. Food production, especially the production of meat, contributes to about 25% of green house gas emissions (Vermulen et al., 2012). The amount of green house gas emission associated with food waste exceeds emissions of Russia, India, or Japan.(fao)

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 AIDS, malaria and tuberculosis. From a fairness perspective, there is a moral imperative to distribute food in a way that satisfies the basic needs of large human populations.

1.1. Understanding the Sources 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, partly due to the the limited adoption of technology.

In the US, about 40% of food loss occurs at the retail and consumer stages of the supply chain. Supply chain ineffi-

[1]Afresh Technologies, San Francisco, CA. Correspondence to: Volodymyr Kuleshov <volodymyr@afreshtechnologies.com>, Matthew Schwartz <matt@afreshtechnologies.com>.

Proceedings of the 36^{th} International Conference on Machine Learning, Long Beach, California, PMLR 97, 2019.

ciency is a major contributing factor to these losses. Recently, the 2018 Retail Food Waste Action Guide (published by a consortium of leading supermarket chains (gui)) identified the following five solutions as having the highest potential to reduce 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.

2. Reducing Food Waste Using Intelligent Supply Chain Systems

In this paper, we argue that automating supply chain decisions in supermarkets (thus, implementing the five strategies outlined above) can significantly reduce food losses at the retail and consumer levels. Specifically, we describe a decision support system for supermarket operators that optimizes purchasing decisions and minimizes losses; the core of our system is a model-based reinforcement learning engine for perishable inventory management.

In this section, we outline a general approach to building such decision support systems using modern machine learning techniques. In the next section, we describe a simple implementation of these ideas and conduct experiments that demonstrate its ability to significantly reduce food waste.

More generally, we hope that this paper will bring the food waste problem to the attention of the machine learning community and demonstrate how modern machine learning techniques have the opportunity to make a significant impact.

2.1. A Decision Support System Powered by Model-Based Reinforcement Learning

We argue for an approach to perishable inventory management based on model-based reinforcement learning. We construct a model of the future (including demand, price elasticity, and other factors), and then train a decision-making agent that chooses store orders that maximize an objective over time (e.g. store profit or the sum of waste and out-of-stocks) given this model.

We describe below some of the real-world challenges involved in deploying such systems, and propose machine learning techniques to address these challenges.

Large-Scale Datasets. A typical mid-size grocery chain has hundreds of stores, each generating daily time series for thousands of items over the course of years. Effectively handling this data requires machine learning algorithms whose accuracy continues to improve with dataset size, such as modern methods based on deep learning These accuracy gains are especially important for perishable inventory management, where errors are particularly costly and can lead to waste and significant financial losses.

Multi-Task and Few-Shot Learning that Handles Rare Events. Although each product may have years of historical data, certain rare events (such as holidays) are only seen once a year (hence, rarely). Handling these rare events requires joint multi-task learning across thousands of items and across hundreds of stores. Multi-task techniques can be further improved with few-shot learning algorithms that quickly generalize to new products with little historical data.

Probabilistic and Bayesian Methods that Account for Uncertainty. Accurate planning requires predicting not just a point forecast, but an entire distribution over model demand. 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.

Planning Algorithms that Explicitly Minimize Waste. Given probabilistic demand forecasts, one needs to compute daily purchasing orders that balance minimizing waste and limiting out-of-stocks. Classical inventory management systems that target non-perishables do not adequately trade-off these two objectives. A more effective approach is to explicitly minimize an objective that accounts for both waste and out-of-stocks using modern planning approaches, such as Monte Carlo Tree Search (Browne et al., 2012) or Model Predictive Control (MPC).

Keeping Humans in the Loop. Finally, in order be useful, recommendations need to be surfaced to store operators along with confidence levels that correlate well with their performance.

3. Experiments

Next, we describe a simple implementation of the above ideas and demonstrate its effectiveness on historical data and also discuss the performance of another implementation in a real-time deployment.

3.1. Setup

We formalize perishable inventory management for one item using 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

	Calibrated	*Uncalibrated*	*Heuristic*
Shipped	332,150	319,692	338,011
Wasted	7,466	3,148	13,699
Stockouts	9,327	17,358	11,817
% Waste	2.2%	1.0%	4.1%
% Stockouts	2.8%	5.4%	3.5%
Reward	**-16,793**	-20,506	-25,516

Table 1. Performance of calibrated model planning on an inventory management task. Calibration significantly improves cumulative reward. Numbers are in units, averaged over ten trials.

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.

We use supervised learning algorithms to learn the model P from historical sales data. We perform planning using a simple MPC approach in which we sample 5,000 random trajectories over a 5-step horizon, and choose the first action of the trajectory with the highest expected reward under the model. We estimate the expected reward of each trajectory using 300 Monte Carlo samples from the model. We also compare the planning approach to a simple heuristic rule that always sets the inventory to the expected demand multiplied by a safety factor of 1.5.

3.2. Kaggle Experiments

To demonstrate the effectiveness of the simple system described above, we perform experiments on the publicly available grocery sales dataset from the Corporacion Favorita Kaggle contest. We experiment on the 100 highest-selling items and use data from 2014-01-01 to 2016-05-31 for training and data from 2016-06-01 to 2016-08-31 for testing.

We train a Bayesian DenseNet (Huang et al., 2017) to predict sales on each of the next five days based on features from the current day (sales serve as a proxy for demand). We use autoregressive features 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, and sine and cosine features over the number of days elapsed in the year. The Bayesian DenseNet has five layers of 128 hidden units with a dropout rate of 0.5 and parametric ReLU nonlinearities. We use variational dropout (Gal & Ghahramani, 2016) combined with the calibration method of Kuleshov et al. (Kuleshov et al., 2018) to compute probabilistic forecasts from the model.

Prediction Accuracy and Waste Reduction. We observe an average mean absolute percent error of 25.8% on the test set across all the items. Our probabilistic forecasts are also calibrated: our 90% confidence interval correctly contains about 90% of the true outcomes.

We evaluate the agent within the inventory management MDP; the demand is instantiated with the historical sales on test day (which the agent did not observe). We measure total cumulative waste and stockouts over the 100 items in the dataset, and we report them as a fraction of the total number of units shipped to the store.

Table 1 shows our results. The calibrated model incurs waste and out-of-stocks ratios of 2.2% and 2.8%, respectively, compared to 1.0% and 5.4% for the uncalibrated one. These values are skewed towards a smaller waste, while the objective function penalizes both equally. The heuristic has ratios of 4.1% and 3.5%.

3.3. Real-World Experiments

We have been co-developing our technology jointly with a 300-store regional US grocery chain. Our pilot partner has released to us five years of historical data for all products and stores, including past sales, shipments, prices, promotions, and other key data elements. We initially performed a similar historical simulation, and found that the our system increases store efficiency, and reduces food waste. Next, we deployed our system for four months in a real-time pilot within a store, making daily recommendations for a number of select key items.

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. Specifically, 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).

Waste Reduction. In simulation, our system incurred unit waste of 1.7% (aggregated across all stores and items), with only 1.1% of demand being unfulfilled due to underordering. 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. In the real-time pilot, we observed average reductions in food waste of 50%, with some items having 80% improvements.

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 RL that has the potential to have a tangible impact on the environment and to increase access to fresh food.

References

Food wastage footprint. http://www.fao.org/docrep/018/i3347e/i3347e.pdf. URL http://www.fao.org/docrep/018/i3347e/i3347e.pdf.

Food waste retail action guide. http://www.refed.com/downloads/retailguideweb.pdf. URL http://www.refed.com/downloads/RetailGuideWeb.pdf.

Browne, C., Powley, E., Whitehouse, D., Lucas, S., Cowling, P. I., Tavener, S., Perez, D., Samothrakis, S., Colton, S., and et al. A survey of monte carlo tree search methods. *IEEE TRANSACTIONS ON COMPUTATIONAL INTELLIGENCE AND AI*, 2012.

Gal, Y. and Ghahramani, Z. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in neural information processing systems*, pp. 1019–1027, 2016.

Hoekstra, A. Y., Mekonnen, M. M., Chapagain, A. K., Mathews, R. E., and Richter, B. D. Global monthly water scarcity: blue water footprints versus blue water availability. *PLoS One*, 7(2):e32688, 2012.

Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2261–2269. IEEE, 2017.

Kuleshov, V., Fenner, N., and Ermon, S. Accurate uncertainties for deep learning using calibrated regression. pp. 2801–2809, 2018.

Vermeulen, S. J., Campbell, B. M., and Ingram, J. S. Climate change and food systems. *Annual Review of Environment and Resources*, 37, 2012.