

---

# Retail: Open your own grocery store to reduce waste

---

Sami Jullien<sup>1</sup> Sebastian Schelter<sup>1,2</sup> Maarten de Rijke<sup>1,2</sup>  
<sup>1</sup>AIRLab, University of Amsterdam <sup>2</sup>Ahold Delhaize  
{s.jullien,s.schelter,derijke}@uva.nl

**Introduction.** Food waste is a major societal, environmental, and financial problem - it represents USD\$ 1 Trillion per year<sup>1</sup>, and 8% of the world's CO<sub>2</sub> emissions.<sup>2</sup> One of its main actors are grocery stores, most of which are being restocked through orders passed to warehouses. Order quantities are often the result of heuristics, combined with domain knowledge. This leaves little room for experimentation: indeed, increasing or decreasing order levels might lead to more waste or lower customer satisfaction. Moreover, with the rise of Reinforcement Learning (RL), a few papers have hinted that more efficient order quantities are possible. As RL is known to be quite sample inefficient, it is clear that field experiments are unlikely to happen - this is why simulations are needed to improve ordering policies [Ingallis, 1998]. A few frameworks have been developed to experiment on supply chains in simulated environments [Brockman et al., 2016, Sterman, 1984]. Yet, they only take into account sales levels and storing costs, but not waste. This is why we introduce *Retail*, a framework to simulate the interaction between a grocery store and a warehouse. Our goal is to help data scientists, researchers and supply chain practitioners experiment with ordering policies to reduce food waste, while optimizing for their usual metrics such as sales and availability. *Retail* makes it very easy to build multiple stores and confront them to different demand scenarios, allowing our users to conduct their experiments easily. For instance, it is straightforward to compute the out-of-stock probability for a given restocking pattern, allowing for a risk-based approach to inventory management.

## Novelty elements.

- We showcase *Retail*, a framework that *simulates the ordering of items from a store to a warehouse* and subsequent customer purchases. *Retail*'s main innovation is *waste tracking*.
- We provide a *synthetic product assortment generator*, based on real product data from a large European grocery chain. This generator allows us to ensure that ordering policies are robust and can accommodate for new items.
- *Retail* also introduces a *new control task for RL*, where the agent acts on item orders to balance waste, sales and availability level through a user-defined reward function, based on utility theory from economics.

**Technology demonstrated.** *Retail* simulates interactions between a store and a warehouse. Using this simulation, the user can test her restocking policies. A policy usually looks like “given the forecast and current stock, order  $N$  items”. First, the initialization phase creates a grocery store and its associated characteristics. Second, the simulation phase returns metrics used to assess the user's policy performance. Using the store as an environment and the restocking policy as an action allows to frame this as a Reinforcement Learning problem, where the agent has to restock the store correctly every day to receive a reward.

During the initialization phase, a copula model is called to generate synthetic product data [Sun et al., 2019]. This copula models the multivariate dependency between the univariate distributions of item price, cost shelf life, base demand, and dimensions. The items are represented as numerical vectors whose multivariate distribution is based on real item data from a large European retailer. A store is modeled as a set of items - it can be restocked, and items can be consumed. Items have a limited duration to be consumed (the shelf life), otherwise they are declared as waste and removed from the

---

<sup>1</sup>FAO : <http://www.fao.org/policy-support/policy-themes/food-loss-food-waste/en/>

<sup>2</sup>European Commission: [https://ec.europa.eu/food/safety/food\\_waste\\_en](https://ec.europa.eu/food/safety/food_waste_en)

store. Items are going in and out in a last-in-first-out manner - as customers tend to purchase items with a longer remaining shelf life. Our simulated store operates over several days (or time periods) in order to evaluate the long term performance of the restocking policy.

In the simulation phase, RetailL generates an item-specific demand, and adds a random seasonality element to the items' base demand in order to represent various seasonal demands (such as ice-cream, glühwein...). The store takes an ordering policy as input, and matches it against this simulated customer demand. This matching occurs a few times per period, to mimic stores being restocked several times per day to accommodate for the time it takes to transport an item from the warehouse to the store. At the end of each time period, items see their remaining shelf life reduced by one day. Perished items are then written off. This allows RetailL to aggregate sales, availability and waste levels for the day, and return it to the user. These quantities are aggregated through an utility function, summarizing the overall store performance, facilitating customization for the user and learning for Reinforcement Learning algorithms, that can use this utility as a reward function.

**Audience-interactive part.** We will run our demonstration live, offering the audience the opportunity to “open their own grocery store” - an example is available on <https://vimeo.com/452514852>:

1. The first step is choosing various details about the store, such as its size, number of customers or even restocking frequency. Attendees can create a grocery store that looks like the one they have close to home. They can also define what they see as a success, using either pre-defined customisable or their entire own metrics to aggregate availability, sales and waste. (Figure 1a)
2. Second, by defining the number of items ordered to the warehouse, users can see their orders matched to a generated customer demand, and investigate how well they perform over the course of a horizon that they themselves chose. (Figure 1b)

Our demonstrator allows the attendees to modify their stores on the fly and to test various combinations of parameters. Combined with a visualisation of the store's performance, attendees will get an intuition of which ordering policy performs best for various stores, and if they indeed reduce waste. This should highlight the difficulties of item ordering as a balancing task.

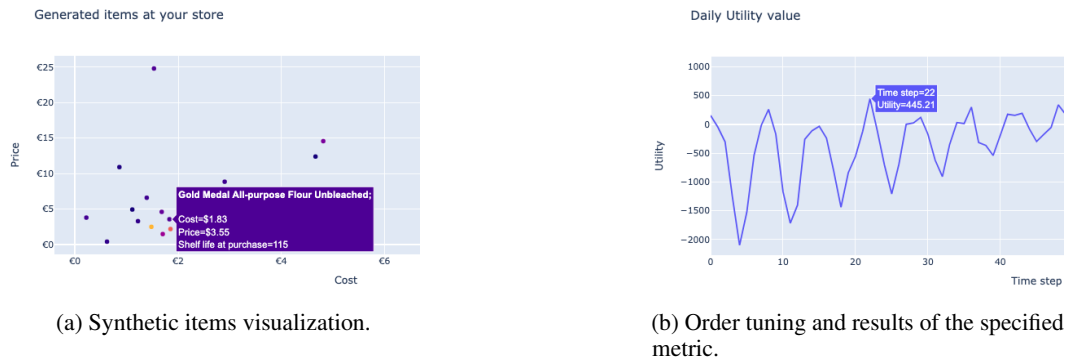


Figure 1: Interface of our demo: attendees can customize their store as well as the order formula, and visualize the results

**Broader Impact.** We expect our work to help in the field of waste reduction, possibly driving prices down for customers, and lowering strain on the production demanded from the fields. Overall, waste reduction in supermarkets can reduce global CO2 emissions by up to 2%.

## References

- R. G. Ingallis. The value of simulation in modeling supply chains. In *1998 Winter Simulation Conference.*, volume 2, 1998.
- G. Brockman et al. OpenAI Gym. *arXiv:1606.01540*, 2016.
- J.D. Sterman. Instructions for running the beer distribution game. In *System Dynamics Group Working Paper D-3679*. Sloan School of Management, MIT, 1984.
- Yi Sun et al. Learning vine copula models for synthetic data generation. *AAAI*, 33, 07 2019.