Carbon Neutral Supply Chain Management

Anna Ionescu, Max White, Sabrina Alvarez, Will Vejcik

Introduction

The problem of cost-effective supply chain management is one of the most studied and developed problems in the field of operations research and for good reason. It is widely applicable to many large corporations in maximizing their profits from the resources available to them. It can be applied to shipping problems, trucking routes, network traffic, and more. In a world with an ever-growing concern for carbon neutrality and environmental causes a novel problem poses itself. Many corporations anticipate US government regulations for reduced carbon emissions similar to those seen in the EU. How do companies tackle the supply chain issue of minimizing dollar costs while also minimizing the carbon cost they endure from the price of offsetting their emissions?

This paper aims to tackle the model of optimizing such supply chain routes while minimizing the costs associated with operation and carbon emissions. We consider breaking down the supply chain into smaller subsections from production centers to distribution centers to end clients, and at each step include another variable into the objective: carbon cost. This carbon cost was estimated using recent research from the University of Michigan in 2023 as well as secondary market research of European companies offering carbon removal services for a fee.

Assumptions

The proposed method to tackle this supply chain problem needs to come with a few assumptions. Company emissions are broken down into 3 scopes. Scope 1 emissions come from sources that the company directly owns and controls; Scope 2 emissions are a result of electricity, steam, heat, and cooling that a company purchases; Scope 3 emissions occur indirectly during a company's value chain. We will focus on Scope 1 emissions, more specifically, on the carbon emissions from company vehicles. We aim to target the truck routing problem of supplying each retailer within a large corporate supply chain network.

We first assume that retailer demand is constant across the supply chain, as the demand problem aspects are a uniquely different focus than our project. We also will assume that gas pricing and fuel efficiency are uniform across the supply chain. Simply put, the data gathering required for gas pricing and efficiency in a large-scale network is out of the scope of this project. We also assume that each distribution depot has the capability to serve any number of retail locations as well as trucks.

Note as well that we also consider tackling the problem of maximizing company profits while reducing carbon emissions by also considering which products the company is interested in buying and selling. When considering this part of the problem we again make several assumptions. Firstly, we assume that we are only dealing with a fixed number of possibilities in terms of which product we are interested in producing. These products we consider are limited to the following: beef, lamb, cheese, poultry, fish, bananas, nuts, chocolate, coffee, shrimp, palm oil, olive oil, eggs, rice, and milk. We also assumed fixed standard prices for each of these products in our calculations which were based on average prices for these products from the company researched during the month of December 2023. Finally, we also assume a fixed 30% markup in prices for the company between acquiring these products and selling them to consumers.

Approach

Since this problem is formulated under the guise of carbon neutrality and a reduction in carbon emissions, we need an approach and objective function that will take this consideration into account. Based on the assumptions previously stated, we assume that each truck will have uniform gas mileage across the routes. In our approach, we consider both the problem from a traveling salesman point of view and then expand upon this idea by considering using multiple traveling salesmen. From our traveling salesman point of view, we first consider a naive greedy algorithmic approach. Following this we develop a greedy algorithm that minimizes the distance of the total round trip over all possible starting nodes (retailers). Finally, we consider building on our Greedy Min-Path algorithm using 2-Opt for path assignments.

We next expand upon our traditional traveling salesman approach by considering using multiple traveling salesmen to solve the distribution problem. For this approach, we begin using two trucks for each distribution center which we wish to optimize such that the two trucks take distinct paths to minimize overlapping paths. We consider using a 2-opt algorithm where we minimize path distance in the x direction. We similarly consider using a 2-opt algorithm where we minimize path distance in the y direction. Finally we run the algorithm seeking to simultaneously minimize both directions.

Note for considering which products to produce in the first place, we considered first a typical integer programming approach which includes the solving of a knapsack problem and then increasing the complexity of our model by also allowing for an exponentially decreasing demand curve for the products sold.

Solving Method

We aim to break this problem down into subsections highlighting unique algorithmic techniques to solve each individual problem. Note the objective remains the same throughout the entire process.

Section 1: Distribution Routes

[Link to Google Colab]

Here we model the truck routes originating from distribution centers that travel to retail centers. This takes form in the traveling salesman problem, by which trucks must minimize their total distance traveled while ensuring each retailer is visited with stock. We initially began with uniformly distributed data points (5 distribution centers and 80 retailers). We will later apply this to a real-world example, but want to first illustrate our algorithms on a smaller data set. We start with a simple nearest-neighbor algorithm to allocate retailers to distribution centers by minimizing retailer-to-distributor distance.

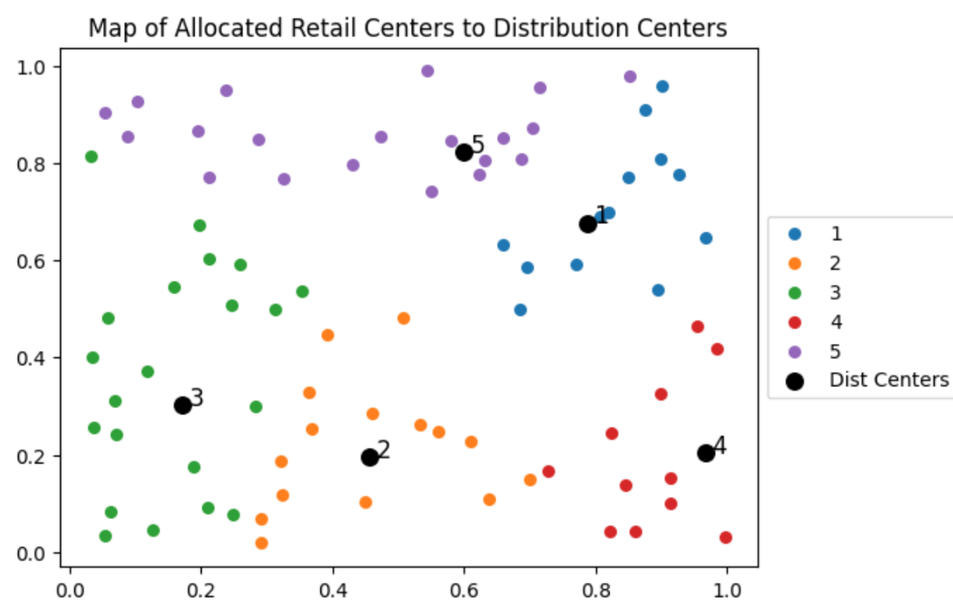

Once allocated, we utilize a greedy algorithm to determine a viable tour from each distribution center. We assume that trucks are leaving each distribution center and taking the same route daily in order to make deliveries. In doing so, we obtain this map of the routes:

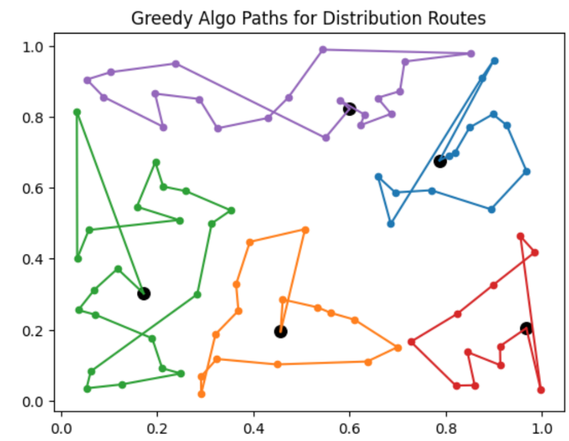

Center [0.787 0.677] has Distance 1.702
Center [0.457 0.194] has Distance 1.73
Center [0.172 0.302] has Distance 3.02
Center [0.968 0.204] has Distance 1.521
Center [0.6   0.822] has Distance 2.315

In this case, the greedy algorithm has many shortcomings and clearly does not take the shortest route. We begin by iterating over all the starting points used to begin the algorithm and take the path that provides the shortest route for each distribution center.

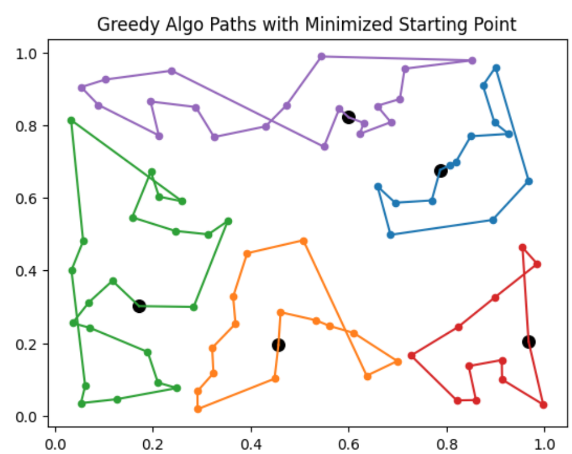

Center [0.787 0.677] has Distance 1.414
Center [0.457 0.194] has Distance 1.687
Center [0.172 0.302] has Distance 2.665
Center [0.968 0.204] has Distance 1.371
Center [0.6   0.822] has Distance 2.3

While this has improved the paths, we want to eliminate crossovers within routes to the best of our ability. Thus, we utilize the 2-opt algorithm (aka hill climbing) in order to further optimize routes and obtain the following. Note that while it is still not perfect, it has improved the distance traveled from most centers via 500 iterations of hill climbing per route.

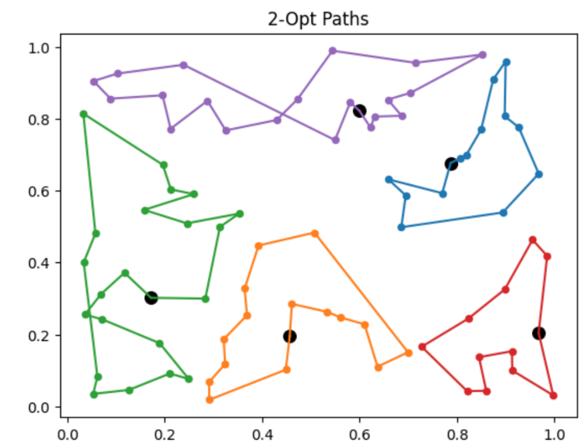	Center [0.787 0.677] has Distance 1.337 Center [0.457 0.194] has Distance 1.671 Center [0.172 0.302] has Distance 2.603 Center [0.968 0.204] has Distance 1.349 Center [0.6 0.822] has Distance 2.263

At this point, we could explore algorithms such as simulated annealing to perfect the path and ultimately get rid of any intersections. However, what is more important to consider is that distribution centers rarely send multiple trucks down the same exact route. Thus, we begin to look at methods of partitioning retailers within each center's neighborhood so that we can optimally conceive two separate truck routes per center. We solve the multiple traveling salesman problem by splitting retailers into two equally sized groups across an axis. The X-partition splits rightmost and leftmost retailers and runs the aforementioned methods on what now are two separate routes, ending the optimization with hill climbing on each sub-route. The Y-partition does the same on the opposite axis, splitting uppermost and lowermost retailers.

X-Partition Distance X-Partition: 10.9696	**Y-Partition** Distance Y-Partition: 10.3199
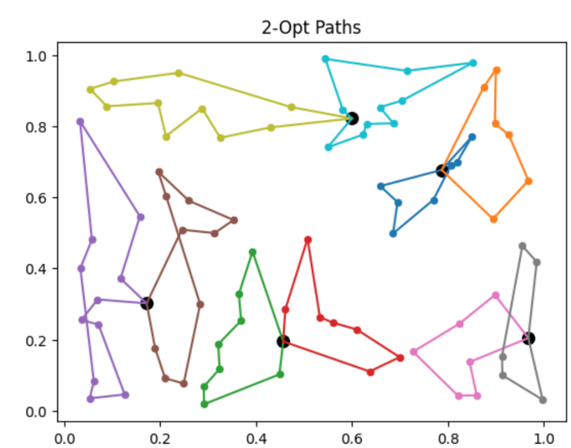	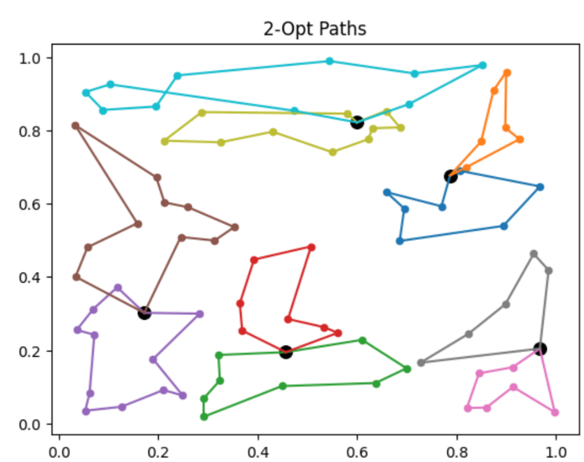

To obtain the “best of both worlds,” we look at each center and compare its X-Partition to its Y-partition, taking one in order to further minimize total distance.

Distance of Best Partitions: 9.86063

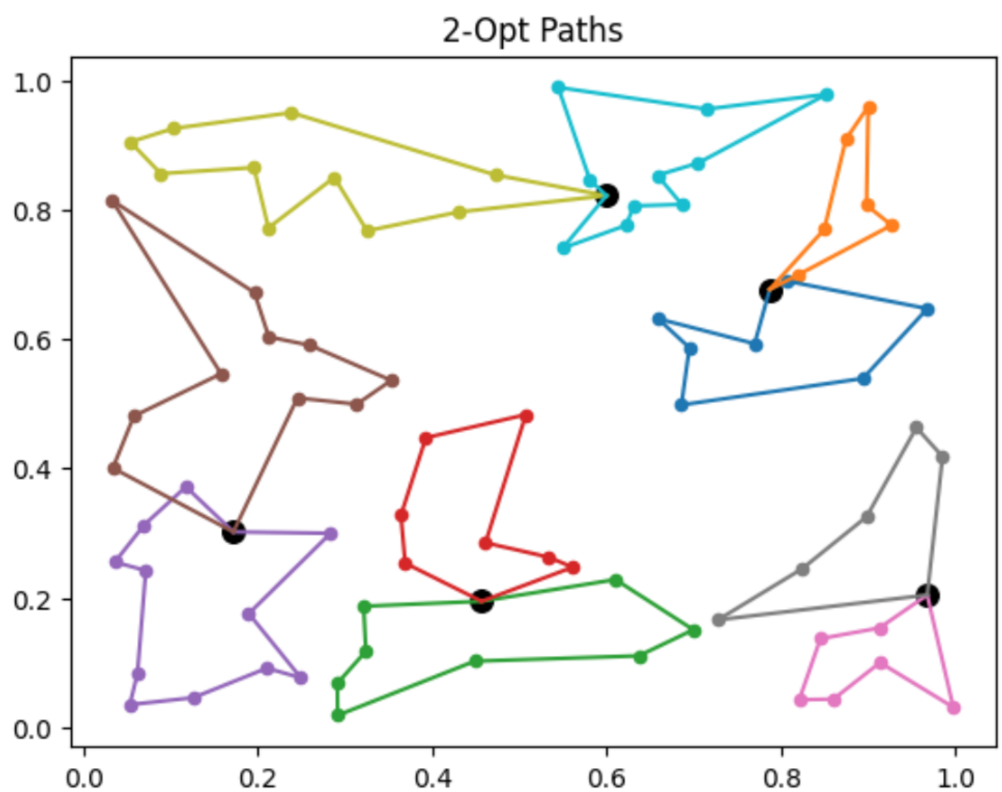

Note that only our uppermost center took the X-partition, while the rest took the Y-partition. This does not imply a trend as both the distribution and retailer centers are derived from a uniform distribution on [0,1]^2. We can create any number of axes to partition our retailer centers on, for instance, diagonals and quadrants. Additionally, we can further apply hill climbing between partitions, examining how non-equally sized partitions may affect overall distance. This may be beneficial in the context of our leftmost center. Ultimately, when comparing our retailer sets and taking the optimal groups, we are guaranteed to obtain a total distance at most that of the minimum between our partition options.

We can put this algorithmic proof of concept into practice on a real-world large-scale corporation: Walmart. Walmart owns and operates 42 distribution centers across the continental United States, each of which will service around 100-150 of the over 4700 Walmart retail stores. Considering the massive amount of data gathered from Walmart it makes more sense to provide a visual solution rather than the raw data.

Distance X-Partition: 85,745.4522 Miles

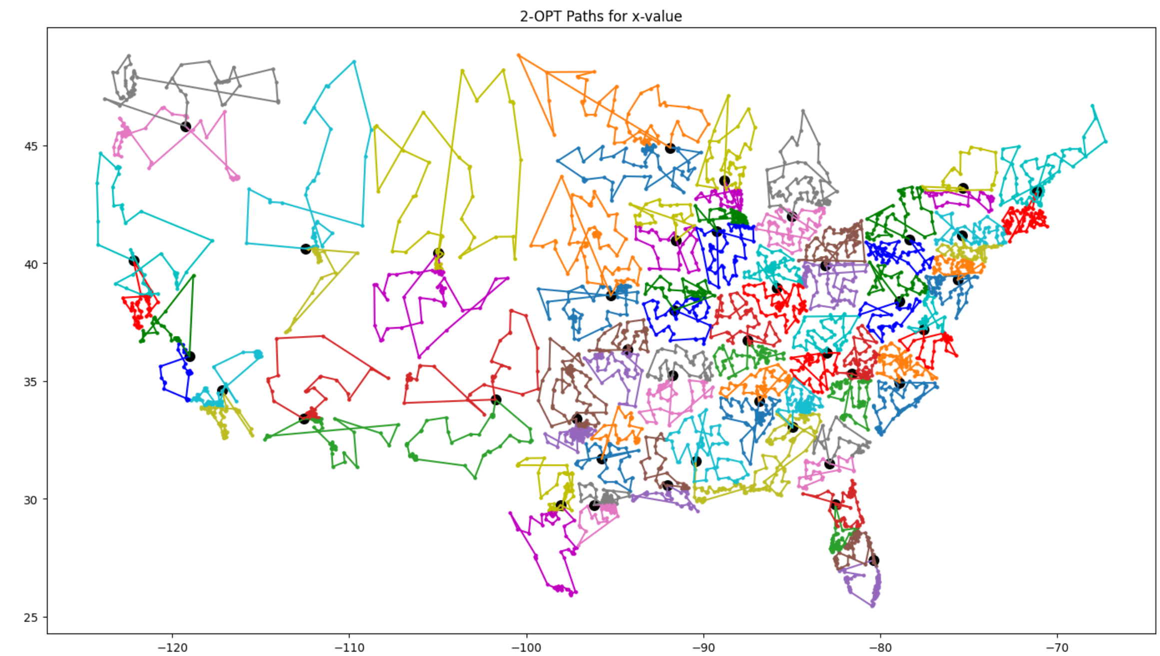

Results of the 2-OPT multiple traveling salesman problem minimizing for Latitude

Distance Y-Partition: 86,791.5946 Miles

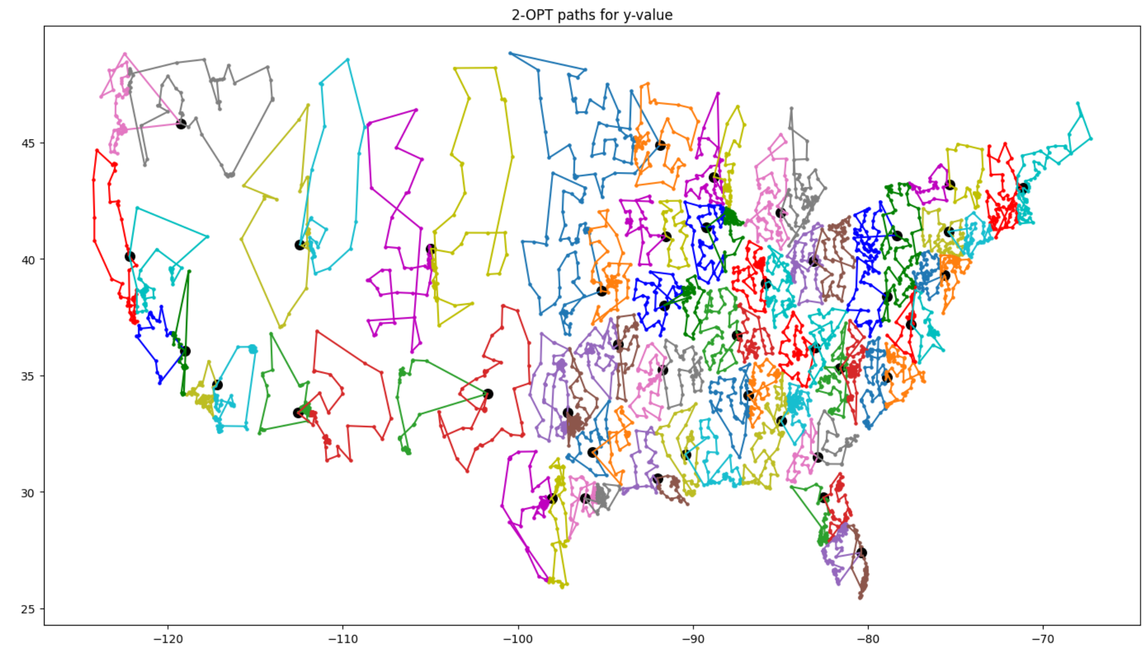

Results of the 2-OPT multiple traveling salesman problem minimizing for Longitude

Distance Best Partitions: 83,902.5165 Miles

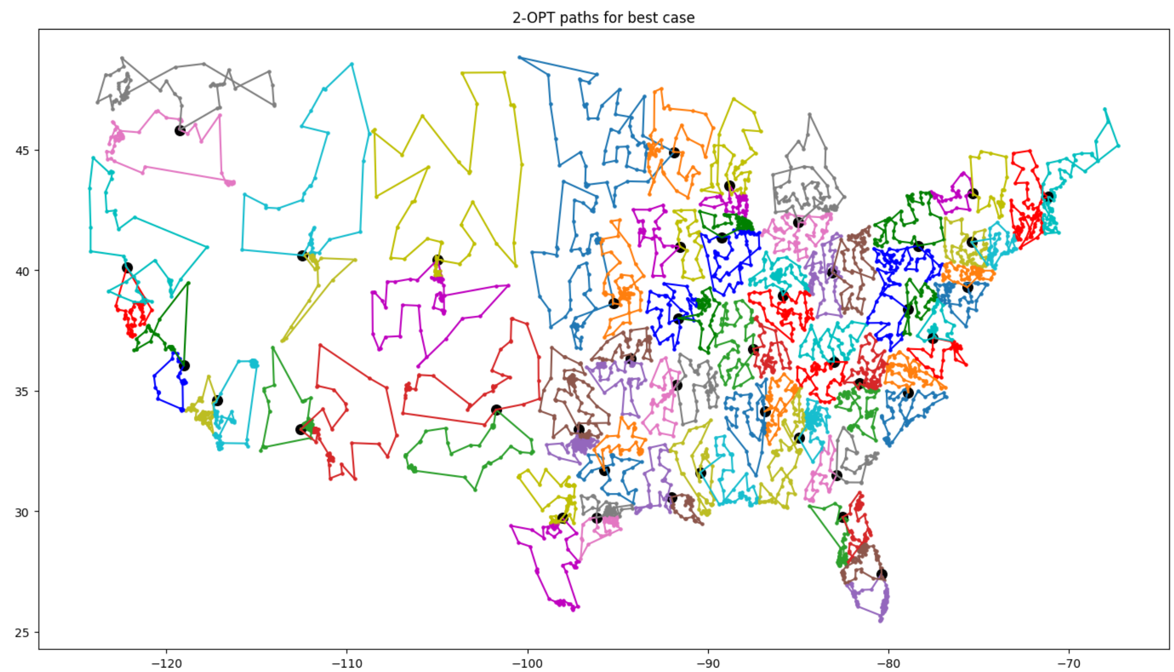

Results of the 2-OPT multiple traveling salesman problem minimizing Longitude and Latitude

Section 2: Product Selection

Note we choose to formulate the problem of problem selection as a linear integer program which is analogous to the traditional integer programming knapsack problem. The weight of each product is its weight in pounds while the value of each product is the total profit made by selling each individual product while offsetting this profit by the carbon emissions produced by the production of this product. Specifically, we look at the use of carbon credits to account for the carbon emissions of each product in terms of U.S. Dollars. We used a simple Python program to first compute which products based on this setup should be included in a standard 100 lbs. load. Our results showed that 45 palm oil and 11 coffee produced the maximum value while accounting for carbon emissions given the constraint on total product weight.

However, given that this initial model makes several simplistic assumptions, we then introduced a more realistic model that attempts to limit the maximum amount selected of any individual product. This would take into account the need for diversification of products chosen given a diminishing marginal utility for each product chosen from the same category. In our model, we simply used our Python program to account for this new assumption by running the knapsack solver several times. On each iteration, the algorithm evaluates the total value of each product category and multiplies it by an exponential constant if the number of products chosen exceeds n = 40% of the total weight. Specifically, we used the constant A = e^-0.1. We then recomputed the value of an individual product from the new aggregate value of the products. This value for the product was then introduced as its initial value in the next iteration of our knapsack problem. We then ran the algorithm until stabilization or at least until the results of each iteration differed by less than epsilon = 10^-1. Our results showed that the optimal load in our second algorithm was 20 palm oil and 11 coffee and 12 fish and 20 chocolate. Although our net value produced shrank from 313.50$ in our first algorithm to 244.30$ in our second algorithm, these results were based upon more realistic and less simplistic assumptions.

Conclusion and Further Research

To continue our project, we could expand the distribution methods of our supply chain by adding cargo aircraft and ship routes. Additionally, we could account for aspects that we assumed were uniform in our project, including differences in gas pricing and fuel efficiency on carbon emissions throughout our supply chain. For our distribution routes, we could improve our findings by adding comparisons and crossovers to optimize the nearest-neighbor algorithm.

Additionally, our Section 2 results can continue to be expanded upon in further research. For one, the results still rely on a fixed sample of products with fixed input prices. Further research could instead use variable prices and a variable number of products. Perhaps more practically, further research should use a probabilistic model to create a distribution for each product and its demand based on historical data on that product's demand. Additionally, we use several parameters which seemed to have been selected arbitrarily and could further be refined.

However, our research ultimately demonstrated several key results. Our research demonstrates that by using a 2-OPT multiple traveling salesman problem minimizing Longitude and Latitude we can solve for the best results in terms of efficiently picking distribution routes from distribution centers to retailers within the United States. Furthermore, by choosing to produce more products such as coffee, fish, and chocolate when compared to other products we can still maximize our profit while minimizing our carbon emissions. Both of these approaches combined allow the company to ensure a maximization of profit in its overall supply chain process while still minimizing carbon emissions.