ISYE6501 Homework 9

Dylan Peters

July 16, 2017

Retail shelf space optimization problem

The retailer wants to optmize shelf in a way that maximizes sales. They have the following hypotheses: First, that the more shelf space they give to a type of product, the more of it they'll sell. Second, the more of a product type they sell,the more they'll sell complementary products. And third, if two complementary product types are next to each other on the shelves, then the complementary effects will be even greater.

Some assumptions are: we are looking at product types/categories and not individual product sales; we have sales data at a receipt level; there are multiple stores that we can reconfigure simultaneously.

The data

We need to start with a baseline set of sales data. We can take existing sales data as a base in order to compare any changes against. The null hypothesis will be that our changes in shelf space and proximity have no effect. We also need to track the amount of shelf space each product type (in square feet) and the proximity of product types to each other. Proximity could be measured in several ways. First, actual distance from each other. Second, whether they are on the same aisle. Third, a flag to indicate they are next to each other.

Does shelf space affect sales?

This question can be addressed with a design of experiments approach. We can take the existing store layout and tweak it store by store. We can divide the total number of stores randomly into 4 sets. One set would be our control group. This would help control for effects such as sales or other promotions. We can use a fractional factorial design to determine how the other stores would adjust the shelf space for different product types, one configuration per group. After a given time period, say one month, we can analyze the sales. If we run a linear regression on sales (y) as a factor of shelf space (x) then we can measure the effect, if any. It may be that the retailer already has all this data organized properly, in which case we can analyze historical data in a similar fashion.

Do sales of one product type affect sales of a complementary type?

In order to solve this question, we have to identity complementary product types, which is not as easy as it sounds. One approach would be to create a graph of the product types. Each node is one product type with a value equal to its total sales for a given period. Each vertex is a link defining how often the two product types are bought at the same time. Any vertex over a certain threshold could be considered part of a complmentary pair. This threshold could be adjusted, or the strength kept as a variable to feed into later analysis. We could also use Louvaine analysis to find groups of product types that are often bought together.

Determining the strength of complementary sales is basically a linear regression analysis. For every product sale, we can examine whether a compementary product was sold, and how often. We could examine weekly totals and see how closely items track each other. For example, a 10% increase in TV sales relates to a 12% increase in DVD player sales.

How does product proximity affect sales for complementary products?

Once we have identified product types that are sold together and their closeness, we can run another experiment do analyze product placement and proximity. The variables are the product types and their relationships, and we can come up with variations based on the different combinations. For example, if product A is complementary with product B and product B is complementary with product C, we can place product B next to product A in one set of stores and next to product C in another set. One assumption is that there are some product types that are considered "primary", probably based on total sales, around which other types should be considered "complementary".

After a set period, for example one month, we can analyze the data and do a regression to determine the strength of sales based on proximity. For example, if sales of Product A gave Product B a 10% boost in a distant configuration, but 20% in a closer configuration, then we'd say the proximity had a 10% effect on sales.

Putting it all together

The bottom line is finding the optimal configuration of product types on shelves. For this we can use an optimization model with the variables from the above steps.

Variables:

Space_i: the amount of shelf space allocated to product i

Price_i: the price of product i

$\text{Proximity}_{i,j}$: the ideal proximity between complementary products i and j (this could be a vector representing a range of strengths)

SalesMatrix_i: the observed effect of Space_i on Sales_i

Constraints:

$\text{SUM}(\text{Space}_i) ==$ <the total shelf space available>

$\text{ActualProximity}_{i,j} == \text{ActualProximity}_{j,i}$ (intuitive, but not obvious!)

$\text{SUM}(\text{IsNextTo}_{i,j,k}) <= 2$ for all i/j/k, in other words one product type can only be next to two others in a linear space

$\text{Space}_i >=$ some minimum space for each product type

Optimize:

$\text{SUM}(\text{Sales}_i)$

I specifically don't recommend a Multi-Armed Bandit approach, mainly because the cost of transitioning a store configuration is pretty high, and usually done across the corporation. However, it might be possible to accomodate this incrementally if we iterate monthly and incorporate new information when we try a new store configuration on a new set of stores. I.e., last months control store is a new experimental store, so include any lessons learned about product placement. In any case, the model will need to be monitored and revisited on a regular basis.