Stochastic Modeling of Retail Stores for Workforce Management

Nicolas Chapados
Université de Montréal
chapados@iro.umontreal.ca

Marc Joliveau, Louis-Martin Rousseau
École Polytechnique de Montréal and CIRRELT
{Marc.Joliveau,Louis-Martin.Rousseau}@cirrelt.ca

Abstract

We address the problem of retail store sales personnel scheduling by casting it in terms of an expected operating income maximization. In this framework, salespeople are no longer only responsible for operating costs, but also contribute to operating revenue. We model the marginal impact of an additional staff by making use of historical sales and payroll data, conditioned on a store-, date- and time-dependent traffic forecast. The expected revenue and its uncertainty are can then fed into a mathematical program which builds an operational schedule maximizing the expected operating income. A case study with a medium-sized retailer suggests that revenue increases of 7% and operating income increases of 3% are possible with the approach.

Workforce Management

- Satisfy the **demand curve**: must have right nb of employees at the right time.
- Retail context: quality of service driven by nb of salespeople.

- Sales staff **contribute to revenues**, not only to expenses.
- **Goal**: schedule employees to maximize expected profit.

Stochastic Models in Retail

- **Goals**: forecast traffic and sales; optimize salespeople schedules to maximize expected profit
- Sales decomposition:

$$\mathbb{E}[S_t \mid E_t, \tilde{T}_t, \mathbf{X}_t] = \mathbb{E}[V_t P_t \mid E_t, \tilde{T}_t, \mathbf{X}_t],$$

- Implementation by forecasting models

$$\mathbb{E}[S_t \mid \mathcal{I}_t] = \mathbb{E}[V_t \mathbb{E}[P_t \mid V_t, \mathcal{I}_t] \mid \mathcal{I}_t]$$
$$= \sum_{v_t} P(V_t = v_t \mid \mathcal{I}_t)\, v_t\, \mathbb{E}[P_t \mid V_t, \mathcal{I}_t],$$

- Need three basic modeling building blocks:
  1. Traffic forecasting
  2. Volume distribution (number of items sold)
  3. Average price per item
- This decomposition empirically performs much better than direct forecasting of intraday sales.

Typical Data Sources

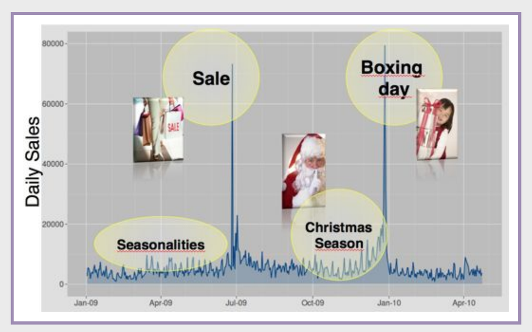

- Transactions: # of items, price per item
- People Counters: traffic
- Past Schedules: how many salespeople on the floor at each period
- Other sources: special events, weather, macro

Traffic Forecasting: Model Hierarchy

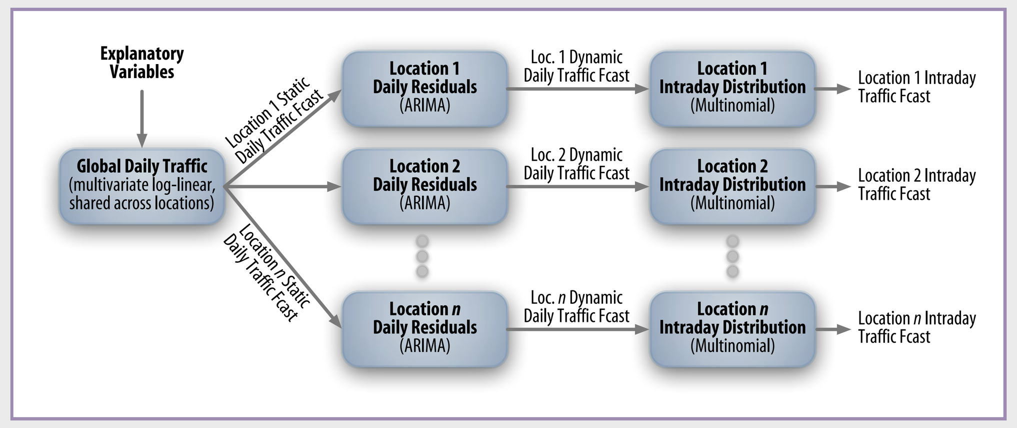

Traffic Forecasting: Daily

- **Global Daily Traffic**: log-linear model of total daily traffic $T_{t,j}$ at store j; shared across all stores with a store-specific intercept ($\beta_{0,j}$):

$$\log \tilde{T}_{t,j} = \beta_{0,j} + \boldsymbol{\beta}' \mathbf{x}_{t,j} + \epsilon_{t,j}$$

- Predictive distribution is **lognormal**, with $\mathbb{E}[\tilde{T}_{t,j} \mid \mathbf{x}_{t,j}] = \exp\left(\mu + \frac{\sigma^2}{2}\right)$.
- **Store-Specific Residuals**: store-specific univariate $\mathrm{ARMA}(p,q)$:

$$\epsilon_{t,j} + \sum_{k=1}^{p} \gamma_{j,k} \epsilon_{t-k,j} = \sum_{i=1}^{q} \alpha_{j,i} \nu_{t-i,j},$$

- ARMA model order determined by AIC tests.
- Significant residual fat tails; variance-stabilizing transformations (used to normalize call center traffic) not much helpful.
- Regressors: seasonalities, special events, weather

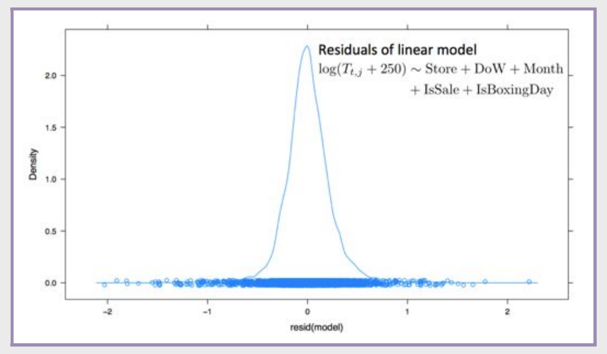

Traffic Forecasting: Intraday

Idea: **spread** the total daily traffic into intraday periods (15- to 60-minute); conditional multinomial model.

Let $y_{t,\tau,j} = \beta_\tau' \mathbf{z}_{t,j}$; the intraday probability attributed to interval τ, $P(\tau \mid \mathbf{z}_{t,j})$, is given by

$$P(\tau \mid \mathbf{z}_{t,j}) = \frac{\exp y_{t,\tau,j}}{\sum_{\tau'} \exp y_{t,\tau',j}}.$$

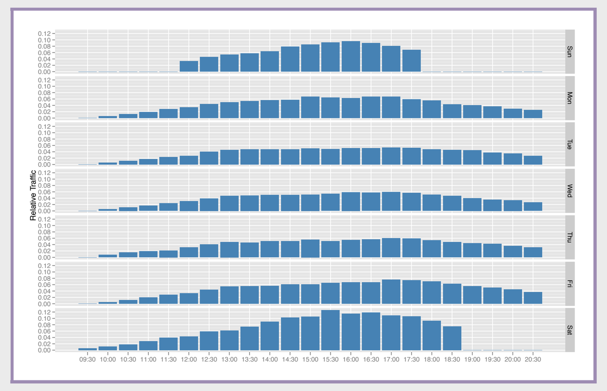

Traffic Modeling Results

- Studied chain of 15 upscale clothing and apparel stores in Canada
- Train: 2009-01-01 – 2010-02-28   Test: 2010-03-01 – 2010-08-31

Statistic	Daily Traffic		Intraday Traffic (30-minute intervals)	
	Train	Test	Train	Test
RMSE (persons)	193.12	235.20	15.24	16.53
MAE (persons)	91.33	121.46	8.16	9.09
MAPE (%)	19.94	25.61	41.17	42.81
Quantile 0.025		0.020		0.096
Quantile 0.1		0.067		0.188
Quantile 0.9		0.946		0.800
Quantile 0.975		0.978		0.878

Item Volume Forecasting

- For small- and medium-sized retail stores, the number of items sold during an intraday interval is a small integer (e.g. below 30).
- Need the **full distribution** of sales volume (not just the expectation).
- Standard parametric forms, such as the conditional Poisson distribution, provide a bad fit to the realized distribution.
- Use the statistical framework of **ordinal regression**.
- Define a latent real variable Z, discretized according to ordered cutoff points $-\infty = \zeta_0 < \zeta_1 < \cdots < \zeta_K = \infty$. We observe $V = k$ if and only if $\zeta_{k-1} < Z \leq \zeta_k, k = 1, \ldots, K$. The proportional odds model assumes that the cumulative distribution of V on the logistic scale is modeled by a linear combination of input variables $\mathbf{x}$, i.e.

$$\mathrm{logit} P(V \leq k \,|\, \mathbf{x}) = \mathrm{logit} P(Z \leq \zeta_k \,|\, \mathbf{x}) = \zeta_k - \boldsymbol{\theta}'\mathbf{x},$$

where $\boldsymbol{\theta}$ are regression coefficients and $\mathrm{logit}(p) \equiv \log \frac{p}{1-p}$.

- Produce flexible and sensible estimates of the conditional distribution of the volume of items sold, as a function of important determinants (traffic, seasonalities, number of salespeople at work).

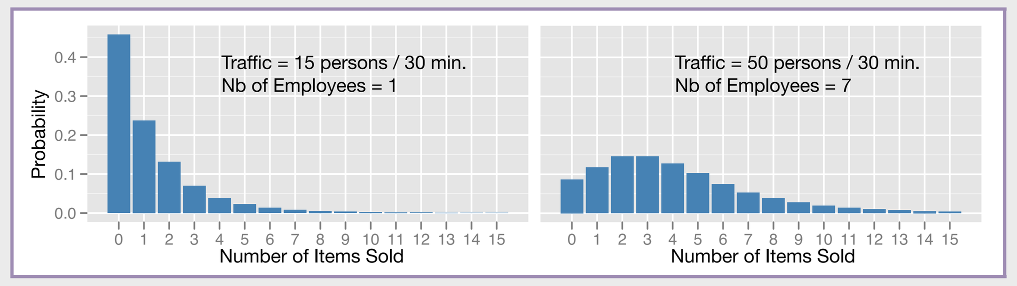

Average Item Price Model

- Average item price is handled by a simple log-linear model.
- Account for store-level differences and monthly seasonalities.
- Account as well for **number of items sold**:
  - ▶ Larger number of items sold reduce average item price.

Profit Curve

Combine 3 models to yield a **conditional sales forecast**, for each store and each period within the day:

1. For store j, day t, intraday interval τ, forecast **expected traffic** $\mathbb{E}[\tilde{T}_{t,\tau,j} \mid \mathbf{x}_{t,j}, \mathbf{z}_{t,j}]$.
2. Given expected traffic, forecast the distribution of **item volume**.
3. For each possible volume, forecast the expected **item price**.

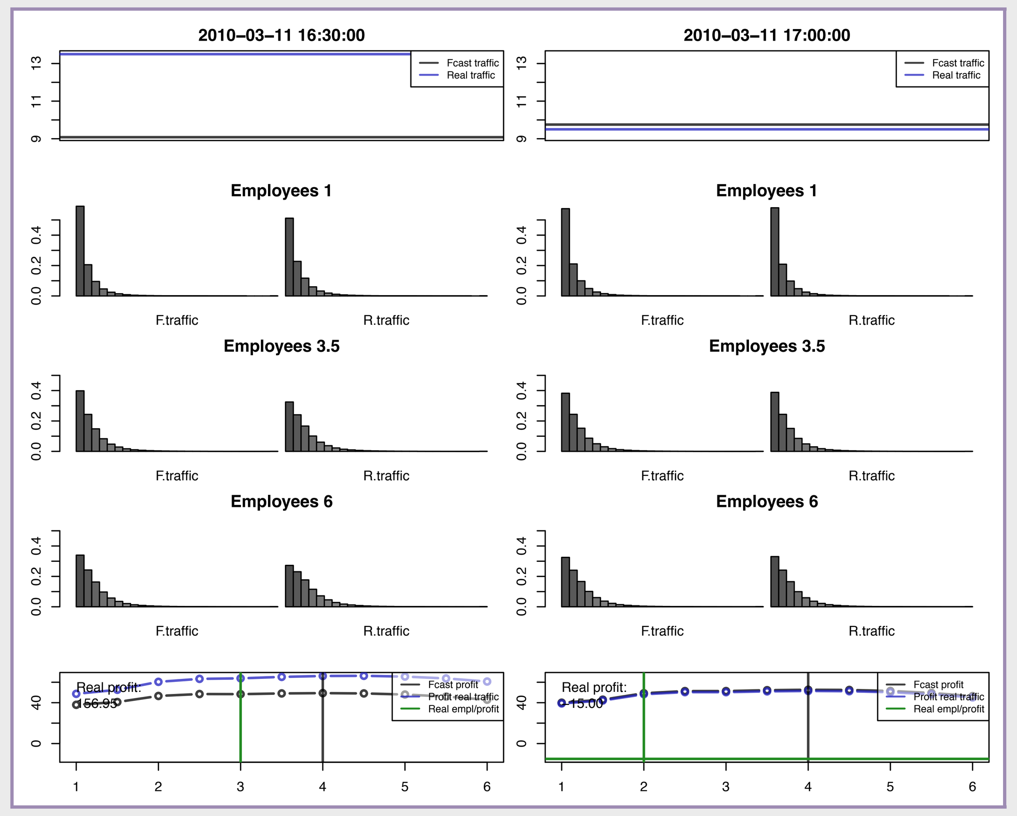

Direct Sales Model

- Comparison of the proposed product model to a **direct model**: using the same topology as the daily/intraday traffic model combination, directly estimate the intraday sales (for each 30-minute interval).
- Input variables in direct model:
  - ▶ Store dummy
  - ▶ Seasonals (month, DoW)
  - ▶ Traffic (w/ splines)
  - ▶ Nb of employees (w/ splines)

Statistic	Proposed Product Model		Direct Model	
	Train	Test	Train	Test
RMSE ($)	291.94	308.76	434.48	398.50
MAE ($)	186.42	203.57	261.16	262.77
MAPE (%)	67.80	68.16	105.36	98.87

Employee Scheduling

Projected impact on revenue (% increase) by day-of-week and time of the day, across all stores participating in the study. (Theoretical maximum average revenue increase is 7%, resulting in an operating income increase of 3%.)

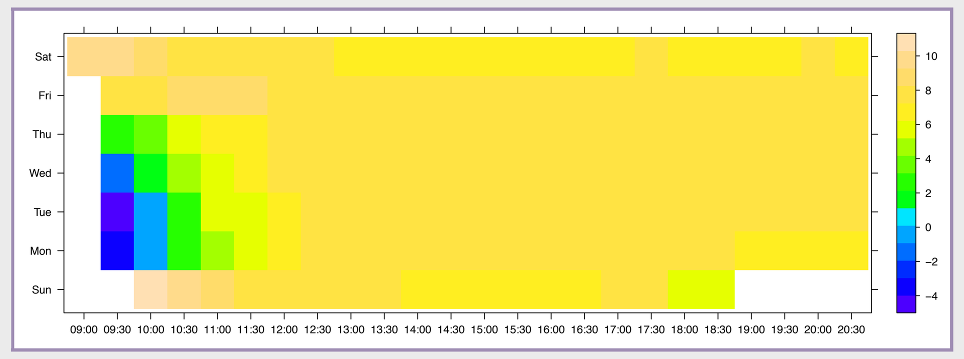

Staffing differences (in # of employees) for one store across 3 weeks, after implementing full employee schedules. Average staffing increases by $1/4$ employee per period.

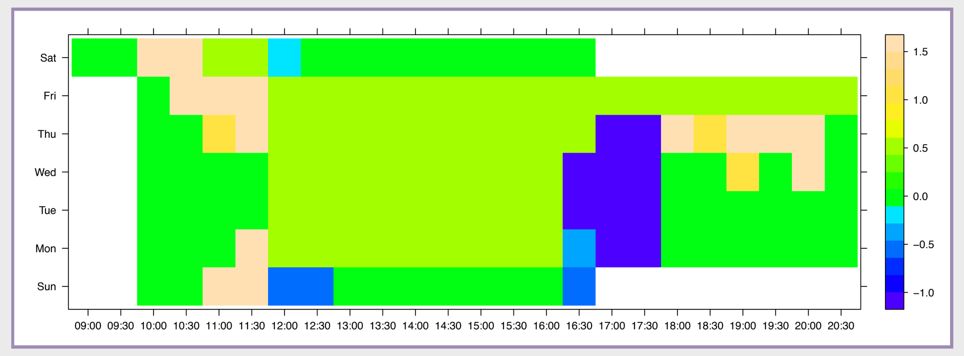

Conclusion

We presented a new methodology to forecast the expected revenue function associated to the presence of a given number of employees in a retail store. This function can be incorporated into a mathematical programming framework to build schedules that maximize the expected store operating profit, and not only minimize payroll expenses. Significant real-world savings are being demonstrated.

- Aldor-Noiman, S., Feigin, P.D., Mandelbaum, A.: Workload forecasting for a call center. Annals of Applied Statistics 3(4), 1403–1447 (2009)
- Avramidis, A.N., Deslauriers, A., L'Ecuyer, P.: Modeling daily arrivals to a telephone call center. Management Science 50(7), 896–908 (July 2004)
- Chapados, N., Joliveau, M., Rousseau, L.M.: Retail Store Workforce Scheduling by Expected Operating Income Maximization. In: Lecture Notes in Computer Science vol. 6697, 53–58. Springer (2011).
- McCullagh, P.: Regression models for ordinal data (with discussion). Journal of the Royal Statistical Society B 42(2), 109–142 (1980)
- Quimper, C.G., Rousseau, L.M.: A large neighbourhood search approach to the multi-activity shift scheduling problem. Journal of Heuristics 16, 373–392 (2010)
- Taylor, J.W.: A comparison of univariate time series methods for forecasting intraday arrivals at a call center. Management Science 54(2), 253–265 (February 2008)