Retail Store Workforce Scheduling by Expected Operating Income Maximization

Nicolas Chapados[1], Marc Joliveau[2], and Louis-Martin Rousseau[2]

[1] Université de Montréal
Montréal, Canada
[2] École Polytechnique de Montréal
Montréal, Canada

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 then fed into a constraint 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.

Keywords: Shift Scheduling, Statistical Forecasting, Retail, Constraint Programming

1 Introduction

The objective of work schedule optimization is to allocate the right number of employees at the right time and the right place, to respond to an expected demand, while satisfying organizational constraints and—as much as possible—employee preferences [6]. The ability to forecast future demand constitutes one of the most important ingredients in the quality of the eventual schedule: an ideal schedule is capable of perfectly matching changes in employee demand with commensurate changes in employee supply. In the area of call center management, a coherent methodology has started to crystallize, combining statistical models for call arrival, duration and drop rate forecasting [2, 3, 12, 11, 13, 1], together with a *dimensioning* aspect that determines the optimal number of employees to schedule in each daily segment to satisfy *quality of service* criteria [7].

The retail sector is another instance of prime economic significance where a combination of statistical modeling and mathematical programming techniques could yield substantial operational efficiencies. As with call centers, there is often considerable uncertainty in the traffic arrival process. However, and in contrast to call centers, the sales staff plays a role in the *generation of operating revenues*, and not only accrues to operating expenses. This means that we can view the introduction of an additional employee in terms of its marginal impact on operating income: the optimal number of employees at any point in time is that which balances revenue increase with additional payroll costs. This is the number that the scheduling process would aim to reach.

For the retail industry, these considerations have received only very little attention from the academic literature [5]. In particular, to the authors' knowledge, sales staff dimensioning based on formal business metrics has not been addressed. This paper introduces a statistical modeling and mathematical programming methodology that intends to bridge this gap. We start by introducing the overall approach, followed by an account of the individual statistical models that constitute the *forecasting* aspect of the methodology. Next, we provide results of a case study with a medium-sized clothing and apparel retailer showing the potential of the approach. We close with a preliminary overview of the constraint programming aspect of the methodology that would introduce the results of forecasting uncertainty into the creation of the final employee schedule, with the goal of increasing schedule robustness.

2 Forecasting in Retail Workforce Management

It is traditional to analyze retail sales in terms of decompositions along more fundamental quantities. A simple one—amenable to robust statistical modeling—is to write the expected sales during period t, S_t, in terms of the number items sold during the period and their average price:

$$\mathbb{E}[S_t \mid E_t, \tilde{T}_t, \mathbf{X}_t] = \mathbb{E}[V_t P_t \mid E_t, \tilde{T}_t, \mathbf{X}_t], \tag{1}$$

where V_t is the *sales volume*, i.e. the number of items sold during period t, and P_t is the average price of an individual item. We condition the expectation on three quantities: E_t is the number of *selling employees*,[3] which we posit has a causal effect on sales—this is the key element linking the forecasting model with workforce optimization. The quantity $\tilde{T}_t$ represents the store traffic during period t (the number of shoppers who can become potential buyers); it can either be the actual traffic, or a *traffic proxy*, somehow related to the traffic but that might be less precisely measured, such as the tick count from unidirectional people counters. Finally, $\mathbf{X}_t$ is a vector of other explanatory variables, such calendar regressors or indicators for special events. For brevity, we denote $\{E_t, \tilde{T}_t, \mathbf{X}_t\}$ by $\mathcal{I}_t$. Expression (1) can be rewritten as

$$\begin{aligned}\mathbb{E}[S_t \mid \mathcal{I}_t] &= \mathbb{E}[V_t \mathbb{E}[P_t \mid V_t, \mathcal{I}_t] \mid \mathcal{I}_t] \\ &= \sum_{v_t} P(V_t = v_t \mid \mathcal{I}_t)\, v_t\, \mathbb{E}[P_t \mid V_t, \mathcal{I}_t],\end{aligned} \tag{2}$$

where the second equation is written assuming that the number of items sold is a discrete quantity. This allows two separate models to be brought to bear: (i) a direct model of the conditional expectation of the item price given the number of items sold and other variables $\mathcal{I}_t$,[4] and (ii) a model of the conditional distribution of the number of items sold given only the variables $\mathcal{I}_t$. These models are detailed in section 3.

[3] Defined, in a retail store context, as employees assigned to sales and direct customer assistance on the floor; it would exclude, for instance, cashiers that cannot have a direct impact in turning shoppers into buyers, or employees tasked with replenishing shelves if they do not actively seek to help out shoppers.

[4] Empirically, it is often noted that the average item price goes down with the number of items sold during a given period; this may be attributable to the propensity of shoppers to buy more

From this model, a *sales curve* can be established by varying the number of employees across a reasonable range to determine how the store sales are impacted. A further step, assuming homogeneity in staff cost and abilities, is to compute the functional dependency between the operating income (profit) during period t and the number of employees:

$$\mathrm{Op.\,Income}_t(E_t) = \mathbb{E}[S_t \mid E_t, \tilde{T}_t, \mathbf{X}_t] \;-\; E_t W_t, \tag{3}$$

This expression ignores other costs that are not directly affected by the number of employees, and assumes that the gross margin on items sold and an eventual commission rate have been absorbed into S_t. For simplicity, the dependency on other conditioning variables has been suppressed from the notation.

3 Statistical Forecasting Models

The above framework for expected operating income maximization can cast into three separate statistical forecasting problems: (i) traffic estimation, (ii) average price evaluation, (iii) sales volume forecasting. For the first part, we can rely on the extensive methodological toolset proposed for call center traffic forecasting (e.g. [2, 3]), whereas we found that for the second, a linear regression model proves sufficient. We cover in more depth the third model.

At a typical small- to medium-sized retail store, the number of items sold during a daily interval (e.g. 30 minutes) will be a small integer. Due to the decomposition of eq. (2), one needs to estimate the entire distribution of sales volume, not only the expected value. Empirically, we find that a standard parametric forms, such as a conditional Poisson distribution, generally provide a bad fit to the realized distribution. We propose to use a more flexible framework, that of *ordinal regression* to estimate the volume distribution.

Ordinal regression models [8, 9] attempt to fit measurements that are observed on a categorical *ordinal* scale. Let $Z \in \mathbb{R}$ be an unobserved variable and $V \in \{1, \ldots, K\}$ be defined by discretizing Z 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 $\mathrm{logit}(p) \equiv \log \frac{p}{1-p}$. Model parameters $\{\zeta_i, \boldsymbol{\theta}\}$ can be estimated by maximum likelihood.

In our application, the variable V represents the sales volume during a sub-daily interval (e.g. a 30-minute period). Incorporating the number of employees and the store

items during sales, and the fact that for some store types, for instance clothing and apparel, additional items beyond the first few are often lower-priced accessories that complement main purchases.

traffic among the inputs (along with calendar regressors for representing monthly, day-of-week and intraday seasonalities), we obtain a flexible distributional forecast of the sales volume. Figure 1 illustrates two different operation points for a store that was part of our case study (see section 4).

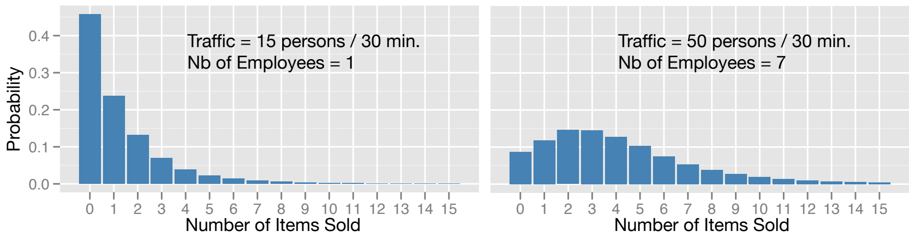

Fig. 1. Distributional sales volume forecast at two operation points. **Left:** With only a single employee in the store, and relatively low store traffic, the most likely number of items sold during a 30-minute period is zero (with 45% probability). **Right:** With seven employees and higher traffic, the distribution shifts rightward, exhibiting a mode between two and three items sold.

4 Case Study

We carried out a case study of the proposed methodology on data provided by a medium-sized chain of upscale clothing and apparel retail stores. Store locations are present in every major Canadian cities and obey normal retail opening hours and sales cycle, including increased holiday activity and the presence of occasional promotional events. A total of 16 months of historical data was used, covering the 2009–early-2010 period. Confidentiality agreements preclude from giving additional details.

We applied the forecasting methodology of section 3 to drive an operating income maximization based on eq. (3), for all store locations and time periods covered in the data. Using realized traffic, an improved staffing process would have resulted in an average sales increase of 7% across all stores, resulting in an operating income increase of 3%. Surprisingly, the model also indicates that stores were, on average, slightly *understaffed*, and that those top- and bottom-line increases would have resulted from increased staffing. Figure 2 shows how these revenue increases would have been distributed across the week, on average.

5 Constraint Programming for Schedule Construction

We are now at the stage of building feasible schedules that maximize a store net income, as opposed to the traditional minimization its labor expenses. Given a set of employees E, a set of time periods T and a set of *activities* $A = \{w, r, m, b\}$ (capturing the occupation of an employee, who is either *working*, *resting* at home, taking a *meal* or on *break*), we propose the following Constraint Programming model:

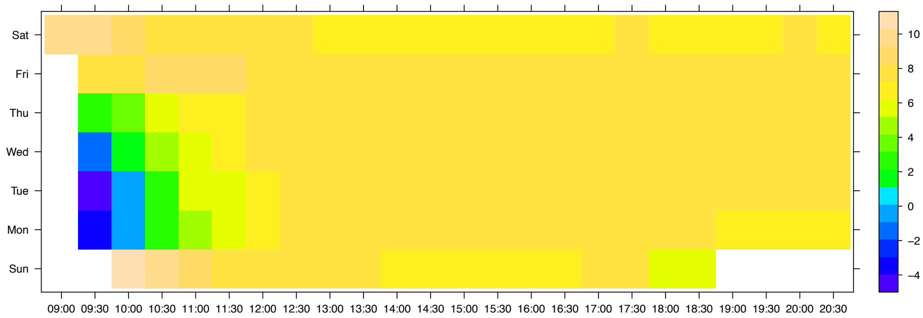

Fig. 2. Average revenue increase by day-of-week and time-of-day across all stores covered by the case study.

$$\max \sum_{t \in T} R^t_{y^w_t} - \sum_{a \in A} C^a_t y^a_t \tag{4}$$

$$s.t. \quad \texttt{Regular}(\Pi, x^e_{(\texttt{all}\, t \in T)}) \quad \forall e \in E \tag{5}$$

$$\texttt{GCC}(x_t^{(\texttt{all}\, e \in E)}, A, y_t^{(\texttt{all}\, a \in A)}) \quad \forall t \in T \tag{6}$$

$$y^a_t \in \{0, \ldots, |E|\} \quad \forall a \in A, t \in T \tag{7}$$

$$x^e_t \in A \quad \forall e \in E, t \in T \tag{8}$$

The model is based on two sets of variables: x^e_t specifies the *activity* (among those in A) performed by employee e at time t, and y^a_t represents the *number* of employees performing each activity a at time t. The objective function (4) maximizes the revenue forecasted at each time period given the chosen number of working employees (those forecasted values are stored in table R) minus the salary expenses associated to each type of activity (provided in table C). The values in table R are obtained from the statistical forecasting model (*cf.* eq. (2)) for *working* employees; we can assume that the revenue associated with other activities is zero. Constraints (5) enforce that each employee's schedules follows a certain number of regulations, which are captured in Π (as proposed in [4, 10]), constraints (6) link the x and y variables, while constraints (7)–(8) define the domain of all variables. We plan to implement this straighforward model in the coming weeks in order to present detailed results during the conference.

6 Conclusion

We have proposed a new methodology to forecast to expected revenue function associated to the presence of a given number of employees in a retail store. The use of a Constraint Program, as opposed to an Integer Program for instance, allows to easily incorporate this non-linear function into a shift scheduling model. The overall approach will thus allow to build schedules where employees are considered as a source a revenue rather than a sole expense.

References

1. Aldor-Noiman, S., Feigin, P.D., Mandelbaum, A.: Workload forecasting for a call center: Methodology and a case study. Annals of Applied Statistics 3(4), 1403–1447 (2009)
2. Avramidis, A.N., Deslauriers, A., L'Ecuyer, P.: Modeling daily arrivals to a telephone call center. Management Science 50(7), 896–908 (July 2004)
3. Brown, L., Gans, N., Mandelbaum, A., Sakov, A., Shen, H., Zeltyn, S., Zhao, L.: Statistical analysis of a telephone call center: A queueing-science perspective. Journal of the American Statistical Association 100(469), 36–50 (March 2005)
4. Demassey, S., Pesant, G., Rousseau, L.M.: Constraint programming based column generation for employee timetabling. In: Proceedings of the Seconc International Conference on Intergretion of AI and OR techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR 2005). vol. 3524, pp. 140–154 (2005)
5. Ernst, A., Jiang, H., Krishnamoorthy, M., Owens, B., Sier, D.: An annotated bibliography of personnel scheduling and rostering. Annals of Operations Research 127, 21–144 (2004)
6. Ernst, A., Jiang, H., Krishnamoorthy, M., Sier, D.: Staff scheduling and rostering: A review of applications, methods and models. European Journal of Operational Research 153, 3–27 (2004)
7. Gans, N., Koole, G., Mandelbaum, A.: Telephone call centers: Tutorial, review, and research prospects. Manufacturing and Service Operations Management 5(2), 79–141 (2003)
8. McCullagh, P.: Regression models for ordinal data (with discussion). Journal of the Royal Statistical Society B 42(2), 109–142 (1980)
9. McCullagh, P., Nelder, J.A.: Generalized Linear Models. Chapman & Hall, London, second edn. (1989)
10. 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)
11. Shen, H., Huang, J.Z.: Forecasting time series of inhomogeneous poisson processes with application to call center workforce management. Annals of Applied Statistics 2(2), 601–623 (2008)
12. Shen, H., Huang, J.Z.: Interday forecasting and intraday updating of call center arrivals. Manufacturing and Service Operations Management 10, 391–410 (2008)
13. 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)