# AI Technology Embedded in PeopleScheduler

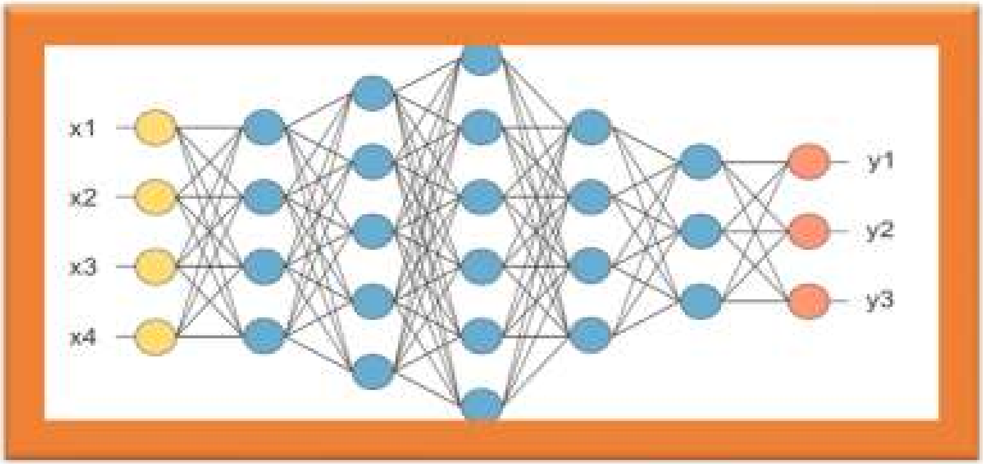


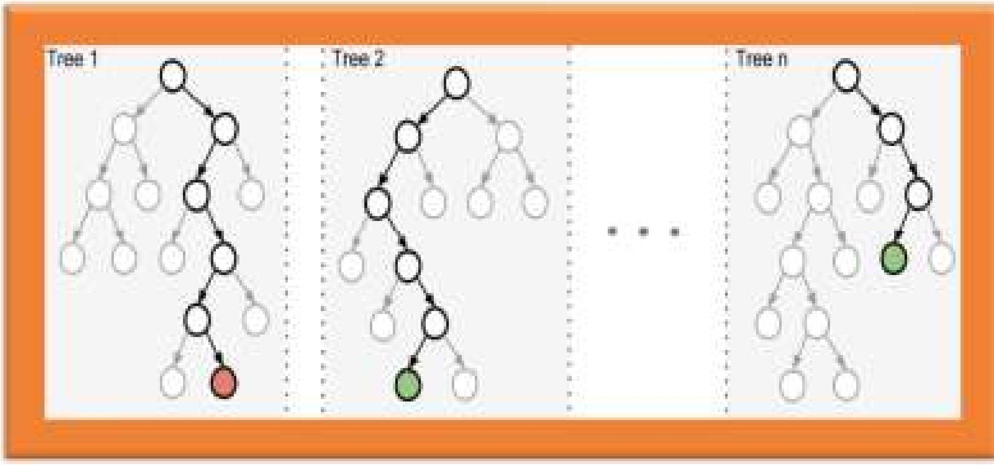


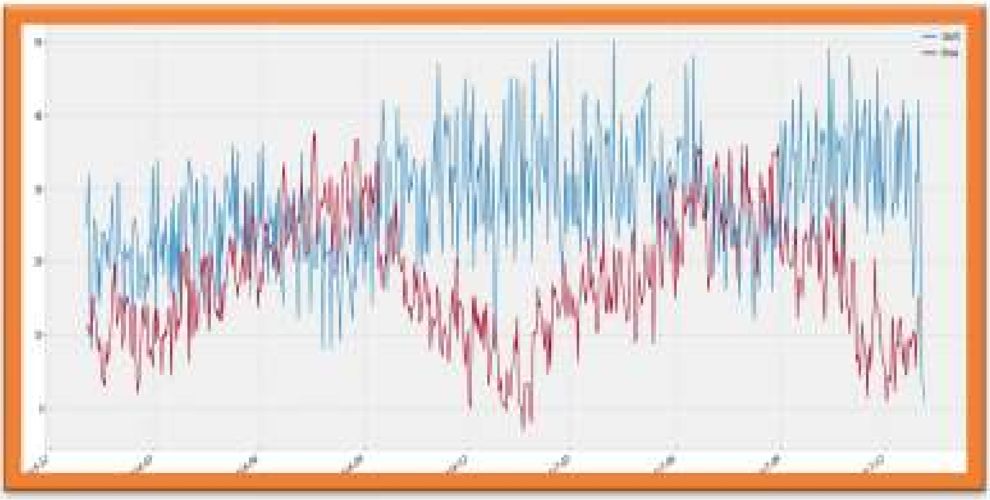


**By: Knowledge Touch technical team and advisors**



**Date: 15 November 2018**

# TABLE OF CONTENTS

# Introduction

## 1 PEOPLESCHEDULER = AI BEHIND THE SCENE

PeopleScheduler combines state-of-the-art AI technologies to provide some of the most advanced capabilities in the WorkForce Optimization software market. This was done to ensure that the most powerful technology delivers the best possible planning results for our customers as the same time making it the main diffentiator from our competitors in this field. This represented a substantial investment over time but has been essential to achieve:

- High ROIs for our customers
- A clear gap with our competition
- A moat against any software company getting into this market

These advanced algorithms used inside PeopleScheduler belong to the field of Artificial Intelligence with 2 core algorithms implementing:

- *Machine Learning* for PeoleScheduler's forecasting engine
- *Prescriptive Analytics* for PeopleScheduler's mathematical planning engine

The two AI engines described in this document fit within the Advanced Analytics vision described below by Gartner. The first engine: our Forecast ML Engine belongs to the Predictive Analytics step. Whilst our optimization engine belongs to the Prescriptive Analytics step.

The first echelon (ie Descriptive Analytics) is simpler and is also addressed by PeopleScheduler in the form of BI Reports: for instance, the gap analysis report comparing *Planned* versus *Actual*.

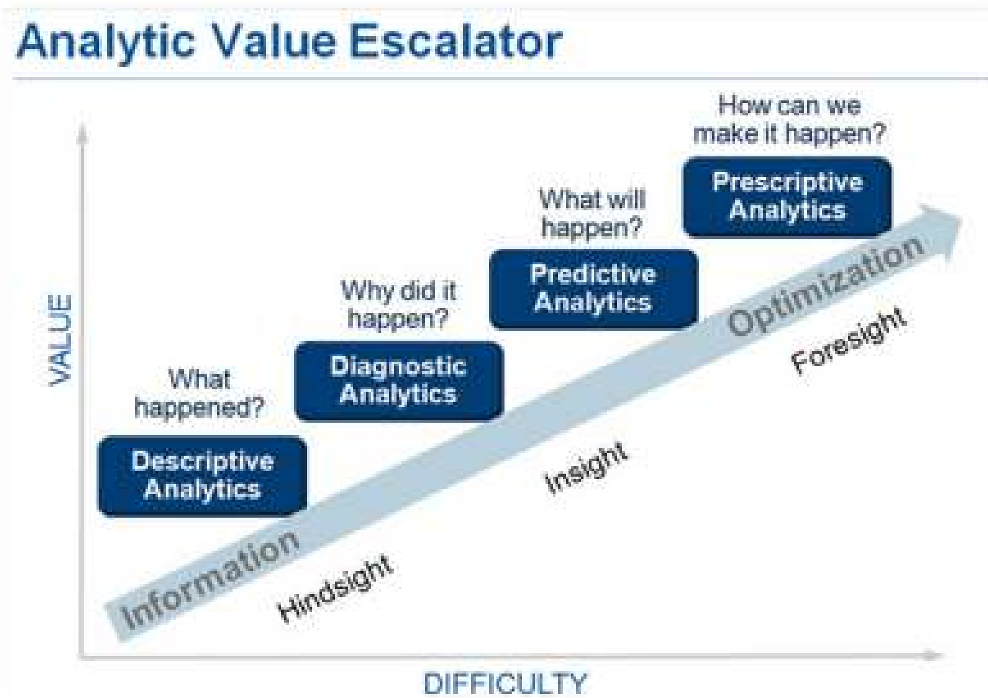


**GARTNER'S ANALYTIC VALUE ESCALATOR**

# 2 OUR FORECASTING ENGINE

## 2.1 PRESENTATION

The forecasting engine is the first engine of the process flow.

1. The *Forecasting Engine* predicts the sales drivers (number of customers, number of items sold, sales revenue, etc) for the next period of time (eg 1 week in advance)

2. The forecast is then converted into workforce requirements (number of staff required) over time.

3. This workforce requirement is the main input for the *Planning Engine* which will optimize the personnel it assigns cover as closely as possible this demand.

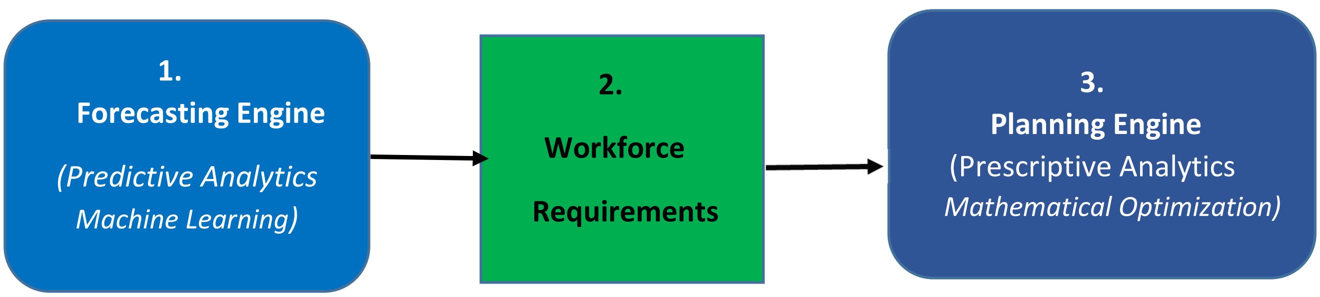


## 2.2 DESCRIPTION

The forecasting engine uses a fully-fledged Machine Learning approach containing 2 major steps:

- The training phase (aka as Training pipe-line)
- The scoring pipe-line

### 2.2.1 The training phase

This is the phase where the algorithm “learns” automatically from the data it is provided with. It consists of the following steps:

- Data Preparation step (Cleaning, Feature Creation/Engineering, etc)
- Training/Validation Step


*The training Phase-> runs once a week*

### 2.2.2 The scoring phase

This is the actual generation of the predictions for the coming week(s) using the algorithm (ie forecast engine) that has already been "trained".

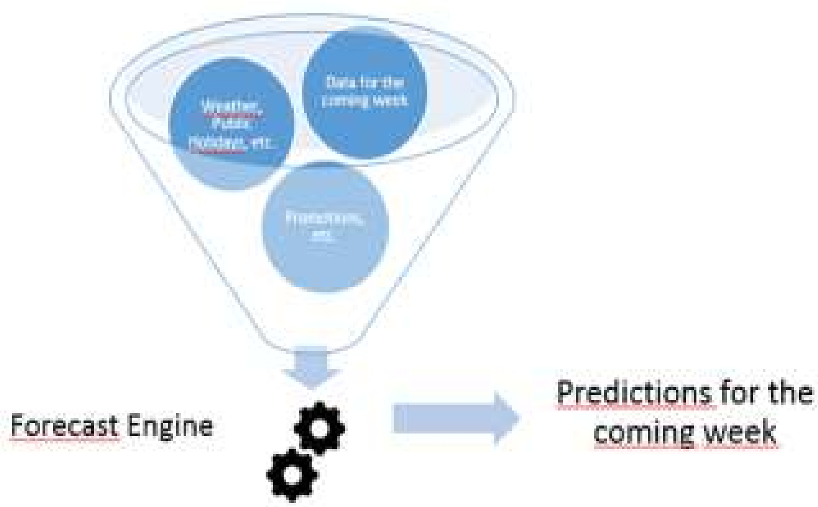


*Prediction Phase-> executed once a week (possibly once a day)*

The Forecasting engine has been designed specifically:

- to predict time-series situations (meaning that the various dimensions to be predicted are all *time based)*: the forecast and historical data are time-stamped and will have typically a precision of 15mins, 30mins or 1 hour depending on the situation/customer.

- for short term forecasting, it excels at predicting for a horizon starting from a few days ahead to a few weeks. It is not meant to forecast sales drivers for the next 12 months.

The 2 characteristics above are what is needed/required for the retail and F&B industry.

The software uses a mix of Deep Learning *Neural Networks* techniques together with *Ensemble Methods* which are currently the most performing machine learning techniques for these types of problems.

## 2.3 Quality of the results

The results obtained with this technology are clearly superior to the traditional *univariate* statistical forecast used. Most of our prospects use various forms of *univariate* statistical forecast from simple *Moving Average Smoothing*, *Exponential Smoothing* or more sophisticated *ARIMA/SARIMA* models.

Our machine learning platform embeds full scale *Machine Learning* models combined with insights that the *univariate* models can bring. Combining smartly the 'new' ML techniques with 'old' techniques (univariate analysis) has provided the best results for our customers use cases.

When we compare the results of our forecasting engine with a typical forecast generated by customers, we improve the forecast quality in a significant way on all drivers/dimensions that we predict and for all stores/outlets.

## 2.4 Performance of the forecasting engine

In general, the ***training*** phase of a Machine Learning engine can consume a lot of CPU power and time. This would have been an issue since the Training phase needs to be launched on a regular basis for all stores (some of our customers have thousands of stores).

We have also tuned the engine so that the training phase is not too heavy, avoiding draining a lot of CPU power, etc. It typically runs in a few minutes.

Once the *learning/training* phase is complete (the engine is ready), it can be used for generating the forecast: this is called the *scoring* phase. The scoring (i.e. the generation of the forecast) just takes a few seconds for each store.

## 2.5 More Insights

Our forecasting engine implements a '*supervised'* machine learning model. It has been designed specifically for the retail and food industry. It uses endogenous data (i.e. data provided from the customer such as sales historical files, stores geolocation, promotions, etc.) as well as exogeneous data (ie weather, public holidays, school holidays, information from the competition, etc.). From all these raw data sources, a modelling stage is required to transform these data into variables/dimensions that enable a high performance of the Machine Learning engine.

## 2.6 Benefits of Machine Learning

Beside the quality of the forecasts generated by our forecasting engine, it will not require new development should new data sources become available. It is very common that our customers may not have all the data sources available at the start (for instance the promotions information is missing), this is not an issue: the engine still gets deployed. When additional information/data becomes available, it will just require some minimal testing (no recoding) to re-deploy the engine. This is one of the strongest advantages of using Machine Learning: it can adapt to new data information without additional coding effort.

# 3 Our Optimization Engine

PeopleScheduler uses a powerful optimization engine that generates optimized workforce schedules. It uses a highly tuned Mathematical Programming model (Integer Programming). The model belongs to the field of Prescriptive Analytics (as opposed to Predictive Analytics) as described in the Introduction of this white paper.

At a high level, the model consists of 3 main components:

1. Decision variables
2. Constraints
3. Objectives

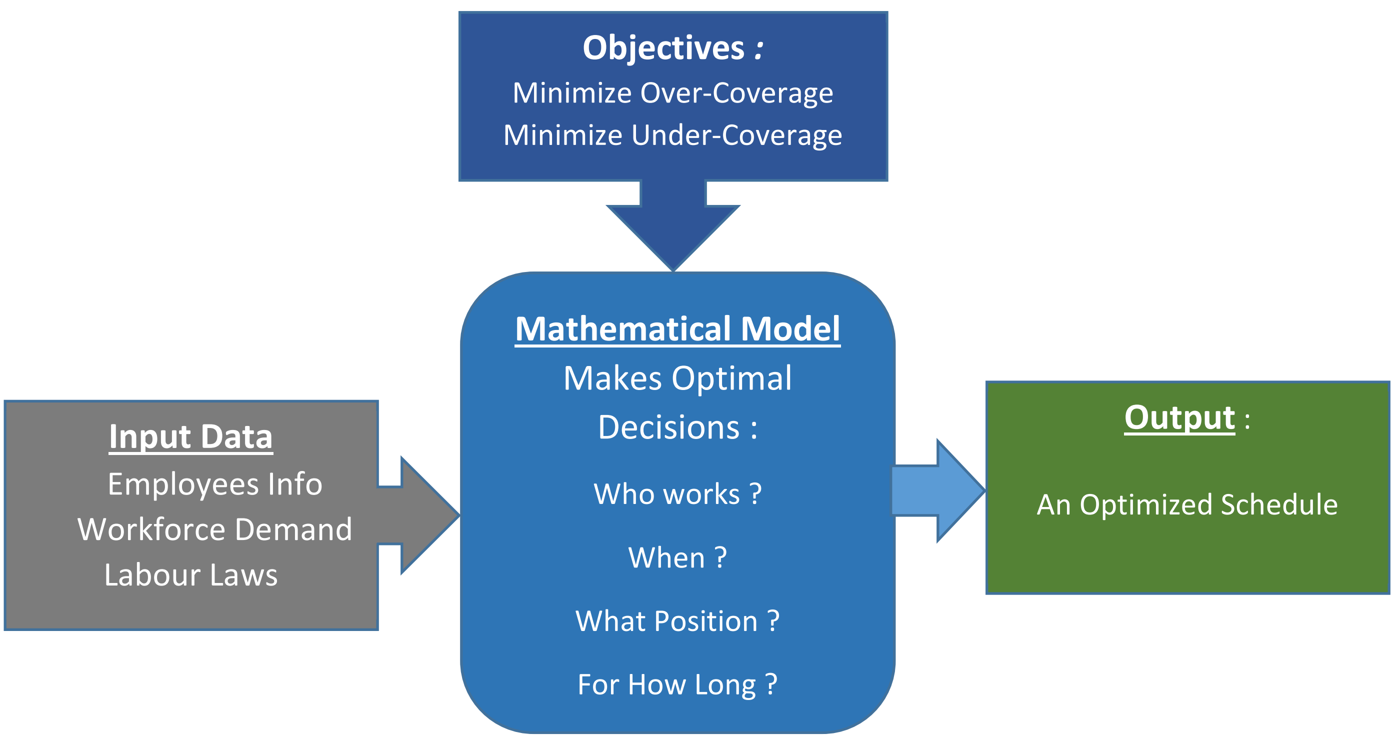


## 3.1 Decision Variables:

The decision variables are what the model must decide upon.

Here the Decision Variables are of multiple types:

- Which staff works today?
- When will shift start?
- What position to cover during the shift?

- How Long will the shift be?
- When should the break occur?
- Etc.

The more types of decisions there are, the larger the problem becomes in terms of complexity and difficulty to solve.

## 3.2 CONSTRAINTS

The constraints are quite numerous and refer to:

- All legal constraints and HR constraints inherent to manpower planning (max hours per day, per week, minimum rest time between shift, etc.)
- Contractual constraints (full timers, part timers, etc.)
- Staff Availabilities
- Staff Preferences
- Position constraints: eg. an employee must be trained, or cannot be assigned customer facing positions after an assignment in the kitchen
- Specifics: e.g. breaks should be staggered

## 3.3 OBJECTIVES

The main objectives are about finding the best coverage to meet the demand. In other words, to avoid:

- Over-coverage: too many staff compared to the demand at a given time
- Under-coverage: not enough staff compared to the demand

It is important to note that under-coverage usually has a higher priority than over-coverage: not meeting customer demand is often more important than over-staffing.

Other secondary objectives come into play such as fairness amongst staff, assigning better rated employees during peak periods, etc.

## 3.4 A POWERFUL ENGINE

We believe that the mathematical model behind KT's optimization engine is one of the most powerful that exists in the market place today (in terms of workforce scheduling software). The term "*Powerful*" refers to the fact that PeopleScheduler can find much better solutions within a much faster response time than our competitors.

### 3.4.1 Why is it complex?

It is important to realize that solving these problems efficiently is extremely complex due to the huge combinational search space.

Just consider a dead simple problem of 2 employees, both having 2 positions/skills (the same 2 skills) having to cover the demand for Monday (a single day) from 9:00 to 4pm with a precision of 15 mins. For instance, the 1st employee can start anytime between 9:00 and 3pm, then for each 15 mins time step, one needs to decide if she/he works on position 1 or position 2, etc. The overall search space already contains billions of possibilities and this is just a tiny problem…

### 3.4.2 A Historical Perspective

There are several reasons that explain the power behind our optimization engine.

**First,** the technical team that has been involved in the design of the solver has more than 20 years of experience in the design of planning/optimization engines.

**Second**, with one of our key and long-term customer McDonald's, we had the opportunity right from the start to design/test our engine on one of the hardest problems there is to solve in terms of manpower planning. Manpower Planning at McDonald's is more complex than all other planning processes in the retail and F&B industry simply because they haves a high level of multi-skilling and allows for so much flexibility compared with other stores.

- The fact that an employee can change position/skill 2 or 3 times during his shift is beyond the reach of many solvers.
- The sheer flexibility in terms of flexi start time, flexi shift duration times, etc creates many more possibilities than other planning processes in other environments (retailers for instance).

**Third**, the technical team made a breakthrough with our own novel decomposition method which represents a leap frog in terms of performance. This algorithm has been kept as KT's IP.

The superiority of PeopleScheduler engine has been demonstrated many times in the context of independent benchmarks.

# 4 CONCLUSION

This strategy of using powerful AI technologies right from the design of PeopleScheduler combined with constant improvement over the years has enabled KnowledgeTouch to maintain its technological lead.

Another benefit of this technological edge has been our capability to scale the deployment of PeopleScheduler over a very large number of stores and a great variety of store size.

We intend to continue this approach over the years. We are constantly on the look out for new AI technologies that may bring additional benefits to our customers.

For this purpose, we have also a panel of senior technical advisors that provides KnowledgeTouch valuable technology guidance.