

The Technician Routing Problem with Experience-based Service Times

Xi Chen, Barrett W. Thomas*

Department of Management Sciences

Tippie College of Business

University of Iowa

Iowa City, Iowa, USA 52242

319-335-0938 (P), 319-335-0297(F)

xi-chen-3@uiowa.edu, barrett-thomas@uiowa.edu

Mike Hewitt

Information Systems and Operations Management

Quinlan School of Business

Loyola University Chicago

Chicago, Illinois, USA 60611

312-915-7394 (P)

mhewitt3@luc.edu

*Corresponding author.

Abstract

While home services is a fast growing industry, little attention has been given to the management of its workforce. In particular, the productivity of home-service technicians depends not only on efficiently routing from customer-to-customer, but also the management of their skillsets. This paper introduces a model of technician routing that explicitly models individualized, experience-based learning. The results demonstrate that explicit modeling and the resulting ability to capture changes in productivity over time due to learning lead to significantly better and different solutions than those found when learning and workforce heterogeneity is ignored. We show that these differences result from the levels of specialization that occur in the workforce.

Keywords: Workforce Planning, Routing, Learning

Highlights

- Model multi-period technician routing problem with experience-based learning.
- Explicitly modeling learning and workforce heterogeneity impacts solutions.
- Taking advantage of the capacity of fast learners is important.
- Technicians specialize more when they have less experience.

1. Introduction

Home services is one of the fastest growing industries in the US. For example, revenue from heating, ventilation and air conditioning service is expected to rise at an average annual rate of 5.9% between 2012 and 2017, reaching \$2.5 billion by 2017 (Panteva 2012). To maintain growth, a key challenge for home-service companies is managing their expensive and limited labor resources. In particular, the time an employee needs to provide high quality service often depends on his/her experience. Importantly, experience increases over time, thus gradually decreasing the time required to provide service. By accounting for employee experience and the accompanying learning, managers can take advantage of capacity increases that result from experience, improving efficiency and enabling further growth.

While home service workers have many job titles, we will generically refer to them as technicians throughout this paper. As such, it is easy to see that the problem discussed in this paper is a variant of the technician routing and scheduling problem (TRSP), a problem first introduced by Dutot et al. (2006). In the TRSP, a set of technicians serves a set of customer requests. The key difference between the TRSP and traditional routing problems is that, in the TRSP, customers are associated with certain tasks and different tasks can have different service times associated with them. In our version of the problem, in addition to the task, technicians can have different service times depending on their experience with the skill required for the task. These differences in service times are a reflection of each technician’s experience.

In this paper, we consider the TRSP over multiple periods or days and account for the fact that productivity increases (or service time decreases) as technicians gain experience. These increases in productivity are often referred to as “learning.” We assume that the time that it takes a technician to complete a task depends on the technician’s experience in the skill associated with the task and how quickly the technician learns. How quickly a technician learns is known as the technician’s learning rate. We assume that we have a set of heterogeneous technicians whose learning rates and initial experience are known. For

this problem, the experience of a technician with a skill depends on the number of times the technician has performed the task.

We assume that daily demand is not revealed until the day of service. Each day, the technicians serve the day’s known demand, starting and ending each day at the depot. In this work, we seek to minimize the sum of each day’s makespan over a finite horizon, accounting for both travel and service times on each individual day. The objective accounts for the desire to increase the capacity available to grow the business. We call our problem variant the technician routing problem with experience-based service times (TRSP-EST).

To solve the problem, we implement a rolling-horizon procedure, creating routes for each day’s known demand without regard for future demand. In the rolling-horizon framework, the objective becomes simply the minimization of the makespan for a given day. To solve the daily routing problem, we use a variant of the record-to-record travel algorithm (RTR), a heuristic first introduced by Li et al. (2005). At the end of each day, we update each technician’s accumulated experience.

As the first to explicitly model the impact of experience-based learning on technician productivity, this paper makes several contributions to the literature. First, we introduce to the literature a Markov decision process model of the problem and introduce a myopic solution approach. In addition, this paper presents several important insights. These are:

1. Explicitly modeling workforce heterogeneity and learning offers better solutions in comparison to assuming homogeneous learning curves and/or static productivity.
2. Importantly, modeling workforce heterogeneity and learning captures that fast learners have more capacity that can be used to improve solution quality.
3. Regardless of the learning rate, inexperienced technicians specialize more than more experienced technicians.

Further, we show that, in the presence of workforce heterogeneity and human learning, technician routing solutions trade-off routing and scheduling. We introduce “rules of thumb”

that demonstrate which aspect is more important based on the individual characteristics of a technician.

The remainder of this paper is organized as follows. Section 2 reviews the literature on problems related to the TRSP-EST. Section 3 presents a model for the problem. Section 4 describes the solution approach. Section 5 introduces the datasets used in this paper, and Section 6 presents our computational results. Finally, Section 7 concludes this work and suggests areas of future research.

2. Literature Review

Two major fields of literature are related to the problem studied in this paper: the technician routing and scheduling problem (TRSP) and learning.

2.1 Technician Routing and Scheduling

The existing literature contains a variety of technician routing and scheduling problems. The TRSP was first introduced by Dutot et al. (2006) based on a real problem in the telecommunications industry. In problem as introduced in Dutot et al. (2006), technicians are grouped into teams, and tasks are assigned to teams so that skill requirements and the skill level can be matched. However, neither learning nor the extended horizon over which learning occurs is considered. In 2007, the French Operations Research Society introduced a challenge (<http://challenge.roadef.org/2007/en/>) based on Dutot’s work and offered a real-world data set for technician scheduling. The challenge resulted in a stream of papers. The papers are largely algorithmic, and none of the papers resulting from the challenge consider routing. Hurkens (2009) uses mixed integer programming to construct a day schedule and demonstrates the effectiveness of the linear programming techniques in solving scheduling problems. Firat and Hurkens (2011) propose a solution methodology that uses a flexible match model for a special multi-skill workforce scheduling problem, in which a set of combined technicians

stays together for the duration of a work day. Cordeau et al. (2010) propose a construction heuristic and an adaptive large neighborhood search heuristic for the technician and task scheduling problem arising in a large telecommunications company. The objective is to minimize a weighted combination of makespans of each priority class. Hashimoto et al. (2011) present a variant of the Greedy Randomized Adaptive Search Procedure (GRASP) for solving the technician and interventions scheduling problem for telecommunications. The authors also introduce a lower bounding procedure for the problem.

Other literature has considered both the travel and service time aspects of the problem, but again, does not consider learning and a multi-period horizon. Kovacs et al. (2012) define the service technician routing and scheduling problem with and without team building. The objective is to minimize the sum of total routing and outsourcing costs. Tsang and Voudouris (1997) and Pillac et al. (2012a) propose heuristics for related problems. Alsheddy and Tsang (2011) consider a bi-objective optimization problem in which both the technician routing costs and the employees' interests are considered.

Additional papers incorporate dynamic and stochastic service requests. Similar to the work in this paper, Bostel et al. (2008) consider a multi-period planning and routing problem of technicians in the field. However, Bostel et al. (2008) do not consider learning that takes place over time. Also similar to this work, the problem is solved without incorporating information about future information. Other work considers single-day problems. Inspired by British Telecommunications plc, Lesaint et al. (2000) describe a dynamic scheduler based on a combination of heuristic search and constraint-based reasoning for dynamic workforce scheduling problem. Weintraub et al. (2012) address the routing and scheduling of service technicians for energy providers in Chile. Customers service requests are considered to be stochastic and priorities of different tasks are taken into consideration. The objective is to minimize the response time to these requests. Pillac et al. (2012b) study the Dynamic TRSP in which new requests appear over time by proposing a fast reoptimization approach based on a parallel Adaptive Large Neighborhood Search and a Multiple Plan Approach.

Home healthcare scheduling and routing is a special case of the technician routing problem. A key feature of the home healthcare literature is the need to respect patient preferences for particular healthcare workers. These preferences usually result from a patient’s prior experience with a particular worker. In a sense, these preferences capture the fact that the home healthcare problems are multi-period, even if they are not explicitly modeled as such. Examples include Bertels and Fahle (2006), Bard et al. (2014), Rasmussen et al. (2012).

The authors are aware of only limited work that incorporates learning in a routing context. Zhong et al. (2007) explicitly models driver learning, but in the context of familiarity with a particular geographic area and the customers found in that area. Unlike this work, the heterogeneity of tasks at the individual customers is ignored. While learning is not explicitly modeled, work on consistency in multi-day vehicle routing often cites advantages of repeat visits to the same region or same customers. For example, Smilowitz et al. (2012) suggests that repeated visits may allow a delivery driver to “more efficiently serve her customer base.”

2.2 Models of Learning

The impact of experience on service or production times is often called “learning” in the literature. There exists an extensive body of literature that develops mathematical representations for the improvement in service and production times as experience increases. These representations are often called learning curves. Detailed discussions of various learning curves and their applications are available in Dar-El (2000), Jaber and Sikström (2004), Jaber (2006), and Anzanello and Fogliatto (2011). A review of learning curves in optimization models other than routing can be found in Hewitt et al. (2014). To the best of our knowledge, the effect of learning has never been investigated in the context of the technician routing.

Wright (1936) first quantified learning curves with the observation that the cost of assembling airplanes decreases as the number of airplanes manufactured increases. To formally describe this phenomenon, Wright introduced the now well known Power Model (also called

the Log-linear Model). Due to its simplicity and flexibility, the Power Model has been the most widely used model (Dar-El 2000). The Power Model has been applied to various areas including the service sector and job rotation in manufacturing systems. In this research, we use the modified Power Model or De Jong Model introduced by De Jong (1957). In contrast to the original Power Model, the De Jong model accounts for the empirically observed phenomenon that there are limits to the degree to which production or service times can be reduced, and notably that times do not go to zero regardless of the amount of experience that has been accumulated. We note that, while we employ the De Jong model, most learning curves have similar shapes and would support conclusions similar to those discussed in Section 6.

3. Problem Description, Model, and Solution Approach

In this section, we first present a formal description of the TRSP discussed in this paper. We then present a model of the problem and describe our solution approach.

3.1 Problem Description

In the TRSP, we assume a horizon of T days, and let the set $\mathcal{T} = \{1, 2, \dots, T\}$ index the days of the planning horizon. We let $\mathcal{K} = \{1, 2, \dots, K\}$ be a set of technicians and assume \mathcal{K} is invariant in t . Let $\mathcal{R} = \{1, 2, \dots, R\}$ be the set of all possible skill types.

Associated with each customer is a task requiring skill r in \mathcal{R} , and associated with each technician k in \mathcal{K} is a set of parameters related to the technician's ability to learn that skill. Let d_{r0} be the service time for skill r for any technician with a minimum level of training. Let D_r^k be the steady state service time for technician k performing task requiring skill r . The parameter L_r^k is the individual learning factor for technician k on skill r . These parameters are typically estimated from empirical data.

For each day t in the planning horizon, we capture the experience of technician k at

the start of day in the R -dimensional vector $P_t^k = (p_{1t}^k, p_{2t}^k, \dots, p_{Rt}^k)$, where the r^{th} entry p_{rt}^k indicates the technician k 's experience with the r^{th} skill at the start of day t . Let P_0^k represent technician k 's experience at the beginning of the horizon. Each skill r is associated with a service time d_{rt}^k for day t and technician k , which is negatively related to the technician's skill experience. Assuming the De Jong Model of learning, we compute d_{rt}^k as

$$d_{rt}^k = D_r^k + d_{r0}(p_{rt}^k)^{-L_r^k}.$$

As suggested previously, this learning function accounts for both learning and the fact that service or production times do not go to zero. Thus, the service time is divided into two parts. The first term represents the “incompressible” part of the task. As noted by De Jong (1957), the incompressible time is not related to the technician's experience level. The second term describes the learning process.

On a given day t , we seek to serve a set of customers $\mathcal{C}_t = \{1, 2, \dots, C_t\}$. We assume that this set of customers becomes known only at the beginning of day t . Each customer requires the completion of a single task requiring skill r in \mathcal{R} . Every technician departs from and returns to the depot, denoted by 0 at the beginning of the day and by $C_t + 1$ at the end of the day. An arc (i, j) is associated with each pair of elements in $\mathcal{C}_t \cup \{0, C_t + 1\}$, and each arc (i, j) is associated with a travel time τ_{ij} . We assume that no driver incurs service time at the depot. Our problem objective is to minimize the expected sum of latest task completions time for each day.

3.2 Markov Decision Process Model

The TRSP described above is a sequential and stochastic decision making problem. A natural modeling framework for such a problem is a Markov decision process (MDP). In this section, we present a formal MDP model for the problem.

- States

In this problem, decisions are made at the beginning of each day during the planning horizon. Let $t = 1, 2, \dots, T$ be the decision epochs, where day T is the last day in the problem horizon. The state of the system at the beginning of day t captures all the information that is needed to make a routing and scheduling decision. For this problem, the state needs to capture the experience of the technicians at the start of day t as well as the service requests on day t . Let Q_t be the matrix of the technicians' experience on day t with elements q_{rt}^k for all $k \in \mathcal{K}$ and $r \in \mathcal{R}$. To capture the customer requests, we consider a vector of tuples. We represent each tuple in the vector as $(lat_{ct}, long_{ct}, r_{ct})$, where lat_{ct} is the latitude of the c^{th} customer requesting service on day t , $long_{ct}$ the longitude of the c^{th} customer requesting service on day t , and r_{ct} is the skill required to perform the task at the c^{th} customer requesting service on day t . We denote the vector of tuples representing the requests on day t as W_t . Then, the state of the system on day t is given by s_t be the state of the system on day t , $s_t = (Q_t, W_t)$.

- Actions

Given state s_t , an action is a set of routes that serves the day t requests. For convenience we convert the vector of tuples W_t into the set of customers \mathcal{C}_t . We formally represent the action for day t as $a_t(s_t) = \{x_{ijt}^k : i \in \mathcal{C}_t \cup \{0\}, j \in \mathcal{C}_t \cup \{C_t + 1\}, k \in \mathcal{K}\}$ where binary variables x_{ijt}^k equal to 1 if arc (i, j) is traversed by technician k in day t . A

feasible action satisfies the following constraints:

$$\sum_{i \in \mathcal{C}_t \cup \{0\}} \sum_{k \in \mathcal{K}} x_{ijt}^k = 1 \quad \forall j \in \mathcal{C}_t, \quad (1)$$

$$\sum_{j \in \mathcal{C}_t \cup \{C_t+1\}} x_{0jt}^k = 1 \quad \forall k \in \mathcal{K}, \quad (2)$$

$$\sum_{i \in \mathcal{C}_t \cup \{0\}} x_{i(C_t+1)t}^k = 1 \quad \forall k \in \mathcal{K}, \quad (3)$$

$$\sum_{j \in \mathcal{C}_t \cup \{0\}} x_{jit}^k - \sum_{j \in \mathcal{C}_t \cup \{C_t+1\}} x_{ijt}^k = 0 \quad \forall i \in \mathcal{C}_t, \forall k \in \mathcal{K}, \quad (4)$$

$$d_{rt}^k = D_r^k + d_{r0}(q_{rt}^k)^{-L_r^k} \quad \forall r \in \mathcal{R}, \forall t \in \mathcal{T}, \forall k \in \mathcal{K}, \quad (5)$$

$$B_j \geq \sum_{k \in \mathcal{K}} (B_i + \sum_{r \in \mathcal{R}} z_{ir} d_{rt}^k + \tau_{ij}) x_{ijt}^k \quad \forall i \in \mathcal{C}_t \cup \{0\}, \forall j \in \mathcal{C}_t, \quad (6)$$

$$\begin{aligned} x_{ijt}^k &\in \{0, 1\} & \forall i \in \mathcal{C}_t \cup \{0\}, \forall j \in \mathcal{C}_t \cup \{C_t+1\}, \\ & & \forall k \in \mathcal{K}, \end{aligned} \quad (7)$$

$$B_i \geq 0 \quad \forall i \in \mathcal{C} \cup \{0, C_t+1\}. \quad (8)$$

Constraints (1) ensure that a customer is assigned to exactly one technician and that it is assigned only once. Constraints (2) and (3) guarantee that every technician starts and ends a day's working at the depot. Flow constraints (4) require the technician to enter and leave a customer if the customer has been assigned to that technician. Constraints (5) state the negative relationship between service time and experience level. Constraints (6) state precedence relationship between two consecutive customers for every technician. Constraints (7) through (8) ensure integrality and non-negativity. Let $\mathcal{A}_t(s_t)$ be the set of actions available on day t , the set of actions is the set of sets of routes that serve the day t requests.

- Transition Function

We consider the transition in two parts. The first is a deterministic transition governed by the action selected on day t . We call this new state the post-decision state. Given that the state is currently s_t and that action a_t in $\mathcal{A}_t(s_t)$ is selected, a deterministic transition is made to post-decision state $s_t^a = Q_t^a$ by updating technicians' experience

as follows:

$$q_{r(t+1)}^k(s_t, a_t) = q_{rt}^k + \sum_{i \in \mathcal{C}_t} \sum_{j \in \mathcal{C}_t \cup \{C_{t+1}\}} x_{ijt}^k z_{ir} \quad \forall k \in \mathcal{K}, \forall r \in \mathcal{R}, \quad (9)$$

where $z_{ir} = 1$ if customer i requires task r and set $z_{ir} = 0$ otherwise.

On day $t + 1$, a transition is made from post-decision state s_t^a to pre-decision state $s_{t+1} = (Q_{t+1}, W_{t+1})$ by observing new service requests W_{t+1} arriving at the beginning of day t_1 . In this transition, technicians' experience remain unchanged from the post-decision state, and thus $Q_{t+1} = Q_t^a$.

- Contribution function

At decision epoch t , given state s_t and action $a_t(s_t)$, a transition from pre-decision state s_k to post-decision state s_k^a results in a contribution

$$c(s_t, a_t) = e_{max}^t \quad \forall a_t \in A_t(s_t), \quad (10)$$

which is the time required to complete the last task on day t .

- Objective function

The problem objective is then $\min_{\pi \in \Pi} E \left[\sum_{t=1}^T c(s_t, a_t^\pi(s_t)) \right]$, where π is a policy that determine actions for all days t over the problem horizon T and Π is the set of all policies.

3.3 Myopic Solution Approach

To solve MDPs, authors often turn to the well known Bellman equation:

$$V(s_t) = \min_{a_t \in A_t(s_t)} \{c(s_t, a_t) + E[V(s_{t+1}) \mid s_t, a_t]\}. \quad (11)$$

Because of the size of the state space and the challenges associated with even computing the expectation in the Bellman equation, however, we propose a myopic solution approach. In the language of MDPs, we are seeking the optimal myopic policy (Powell 2011). A myopic policy is constructed by, at each decision epoch or in our case on each day, choosing an action that minimizes the current state costs while ignoring information about the future. This approach is equivalent to using the decision rule $\operatorname{argmin}_{a_t \in A_t(s_t)} \{c(s_t, a_t)\}$. Thus, the problem becomes one of solving a series of daily routing problems.

For our problem, this approach offers several important computational advantages. First, by not considering the cost-to-go in the Bellman equation, we solve a daily routing problem that is a deterministic problem. Second, in the fashion of rolling horizon or rollout methods (for discussion, see (Goodson et al. 2015)), rather than solving for every state as is necessary in traditional backward dynamic programming, we can step forward in time and solve for only the observed demand realizations. A sketch of our solution approach can be found in Algorithm 1. In the next section, we describe how we solve the daily routing problem and update technician productivity based on accumulated experience.

Algorithm 1 Myopic Solution Approach

- 1: **for** day $t = 1$ TO T **do**
 - 2: Observe demand for day t
 - 3: Solve daily routing problem, minimizing the completion time of last task
 - 4: Update technician productivity based on experience gained on day t
 - 5: **end for**
-

4. The Daily Routing Problem

In this section, we present a model and solution approach for the daily routing problem. Throughout, to make explicit the dependence of service times on accumulated experience, where appropriate, we subscript our parameters and variables with the day t .

In addition to the notation introduced in the previous section, we let $r(i) \in \mathcal{R}$ indicate the skill needed to perform the task required by customer $i \in \mathcal{C}_t$. For customers i and j

in \mathcal{C}_t , our model uses binary variables x_{ijt}^k equal to 1 if arc (i, j) is traversed by technician k . Continuous variable $B_i \geq 0$ is the start time of service at customer i , and e_t^{max} is the completion time of the last task on day t . With this notation, the daily technician routing problem with experience-based service times can be modeled as follows:

$$\begin{aligned}
& \min && e_t^{max} \\
& \text{s.t.} && \sum_{i \in \mathcal{C}_t \cup \{0\}} \sum_{k \in \mathcal{K}} x_{ijt}^k = 1 && \forall j \in \mathcal{C}_t, && (12) \\
& && \sum_{j \in \mathcal{C}_t \cup \{C_t+1\}} x_{0jt}^k = 1 && \forall k \in \mathcal{K}, && (13) \\
& && \sum_{i \in \mathcal{C}_t \cup \{0\}} x_{i(C_t+1)t}^k = 1 && \forall k \in \mathcal{K}, && (14) \\
& && \sum_{j \in \mathcal{C}_t \cup \{0\}} x_{jit}^k - \sum_{j \in \mathcal{C}_t \cup \{C_t+1\}} x_{ijt}^k = 0 && \forall i \in \mathcal{C}_t, \forall k \in \mathcal{K}, && (15) \\
& && B_j \geq \sum_{k \in \mathcal{K}} (B_i + d_{r(i)t}^k + \tau_{ij}) x_{ijt}^k && \forall i \in \mathcal{C}_t \cup \{0\}, \forall j \in \mathcal{C}_t, && (16) \\
& && e_t^{max} \geq B_i + \sum_{j \in \mathcal{C}_t \cup \{C_t+1\}} \sum_{k \in \mathcal{K}} \sum_{r \in \mathcal{R}} x_{ijt}^k (d_{r(i)t}^k + \tau_{ij}) && \forall i \in \mathcal{C}_t, && (17) \\
& && x_{ijt}^k \in \{0, 1\} && \forall i \in \mathcal{C}_t \cup \{0\}, \forall j \in \mathcal{C}_t \cup \{C_t+1\}, && \\
& && && \forall k \in \mathcal{K}, && (18) \\
& && B_i \geq 0 && \forall i \in \mathcal{C} \cup \{0, C_t+1\}. && (19)
\end{aligned}$$

The objective of the model is to minimize the completion time of the last task. Constraints (12) ensure that a customer is assigned to exactly one technician and that it is assigned only once. Constraints (13) and (14) guarantee that every technician starts and ends a day's working at the depot. Flow constraints (15) require the technician to enter and leave a customer if the customer has been assigned to that technician. Constraints (16) state precedence relationship between two consecutive customers for every technician. Constraints (17) impose that e_t^{max} is no less than the completion time of any task. Constraints (18) through (19) ensure integrality and non-negativity.

To solve the daily routing problem, we modify the record-to-record travel (RTR) algo-

rithm presented in Li et al. (2005). We chose the RTR algorithm for our problem for two reasons. First, the algorithm is known to return high-quality solutions for vehicle routing problems and does so in relatively short run times. Second, code for the algorithm is open source and available in the COIN-OR repository (<https://projects.coin-or.org/VRPH>). This open-source code shortens development times.

In the algorithm, the record represents the best solution found so far. The RTR is a multi-phase local search algorithm. The algorithm alternates between two phases. The first phase is a diversification phase in which an uphill criterion is applied. That is, non-improving neighboring solutions are accepted if they are within a particular threshold of the record. The diversification phase ends after a fixed number of iterations. In the second phase, the improvement phase, only improving moves are accepted. The improvement phase continues until a local minimum is found.

To overcome convergence to local minima, a perturbation step follows the improvement phase (see Li et al. (2005)). The sequence of the diversification, improvement, and perturbation phases is run for a minimum number of iterations before the algorithm terminates. The algorithm terminates when no solution can be found that, after a fixed number of attempts, improves the best known solution by more than ε . In our implementation, the threshold is $(1 + 0.01)$ times the value of the record or best known solution, the number of iterations in the diversification phase is 30, ε is 0.00001, and the number of attempts is 5. These values are suggested by Li et al. (2005). A detailed algorithmic description can also be found in Li et al. (2005).

Because the daily routing problem must be solved multiple times over the horizon \mathcal{T} , we update each technician's experience and service times based on the assignments of the previous day. For each $t \in \mathcal{T}, t \neq T$, for each $k \in \mathcal{K}$, we have:

$$p_{rt}^k = p_{r(t-1)}^k + \sum_{i \in \mathcal{C}_t: r(i)=r} \sum_{j \in \mathcal{C}_t \cup \{C_{t+1}\}} x_{ij(t-1)}^k \forall k \in \mathcal{K}, \forall r \in \mathcal{R} \quad (20)$$

and

$$d_{rt}^k = D_r^k + d_{r0}(p_{rt}^k)^{-L_r^k} \quad \forall r \in \mathcal{R}, \forall k \in \mathcal{K}. \quad (21)$$

Equation (20) updates a technician’s experience for a particular skill. Equation (21) computes the service times based on previous experience.

Our implementation of the RTR is a modification of existing RTR code, called VRPH, available in the COIN-OR repository (<https://projects.coin-or.org/VRPH>) and described in Groër et al. (2010). The code was originally developed to solve the capacitated vehicle routing problem. To solve the daily routing problem, we modify the code to account for service times. We can solve for each day in our horizon by updating the service times using the solution from the previous day (see equations (20) and (21)) and calling RTR with data for each new day.

5. Instances

This section describes the instances that were used to examine the impact of including learning in technician routing models. There are three sets of attributes that define an instance of the TRSP-EST: (1) the number and geographic diversity of the customers requiring service, (2) the number of different skills required to perform tasks requested by customers, and, for each skill, the length of time a novice technician needs to perform a task requiring that skill, and, (3) the number of technicians and each technician’s individual traits.

The TRSP-EST is a multi-period/day problem, with the number of customers and their geographic distribution not known until the beginning of each day. For our experiments, we limit our analysis to a 29 day period and generate an instance by creating a series of 29 daily customers sets, where the set for day t is \mathcal{C}_t . The number of customers and their geographical distributions in these customer sets is based on instances taken from the Symmetric CVRP instances found in VRPLIB: A Vehicle Routing Problem LIBrary. This library of in-

stances is provided by the Operations Research Group at the University of Bologna, Italy and can be found at: http://www.or.deis.unibo.it/research_pages/ORinstances/VRPLIB/VRPLIB.html. Specifically, the 29 instances “E072-04f” to “E151D14r” are used, with the number of customers in these instances ranging from 71 to 150. To generate a 29-day instance, we randomly order these 29 daily instances. We generate 10 such instances by considering 10 different orderings of the 29 daily instances. In all instances, travel times correspond to Euclidean distances.

Regarding the second set of attributes, in each instance, each customer $i \in \mathcal{C}_t$ on day t requires a task to be performed and performing that task requires a skill $r \in \mathcal{R}$. We refer to the customer diversity of an instance as the number of different skills required to serve all the customers, or $|\mathcal{R}|$. In our experiments, each instance can have one of four diversity levels: 5, 10, 25, and 50. For each of the 29 instances “E072-04f” to “E151D14r” and each diversity level $|\mathcal{R}|$, we randomly assign each customer a skill from the set $\{1, \dots, |\mathcal{R}|\}$.

As discussed in Section 3, we model the relationship between experience and productivity using the function $d_{rt}^k = D_r^k + d_{r0}(p_{rt}^k)^{-L_r^k}$. Thus, each instance requires values for the parameters d_{r0} , $\forall r \in \mathcal{R}$, which represents the time it would take a novice technician to perform a task requiring skill r . We assign the parameter d_{r0} one of five initial service time values: 100, 200, 300, 400, and 500 (time units). These values were chosen to represent values that were on the order of travel times (100 units) versus those that were significantly greater (500 units).

These values are assigned to skills in a round-robin manner. For example, with a customer diversity of ten, skills one and six ($r = 1, 6$) will be assigned an initial service time of 100 ($d_{10} = d_{60} = 100$), skills two and seven ($r = 2, 7$) will be assigned an initial service time of 200 ($d_{20} = d_{70} = 200$) and so forth. As noted earlier, the skill associated with any particular customer \mathcal{R} is assigned randomly among the customers in the geography.

Regarding the third set of attributes, in each instance we consider a workforce that consists of 18 technicians. Again recalling the equation $d_{rt}^k = D_r^k + d_{r0}(p_{rt}^k)^{-L_r^k}$, we see

that an instance requires values for the parameters D_r^k, p_{r0}^k , and L_r^k , $\forall r \in \mathcal{R}, k \in \mathcal{K}$. In all experiments, we set $D_r^k = 5 \forall r \in \mathcal{R}, k \in \mathcal{K}$. As noted previously, this value is unrelated to the technician’s experience and ability to learn and reflects that fact that task times do not tend to zero. As such, while the value is chosen arbitrarily, a different parameter value would change the values of our solutions, but not the trends discussed subsequently.

We assume that an individual technician k has the same learning rate on all skills, or, $L_r^k = L_{r'}^k, \forall r, r' \in \mathcal{R}$. This choice reflects the literature on human learning that suggests that individuals tend to learn different skills at the same rate (Ree and Earles 1991). As a result, we refer to L_r^k as L^k , and the parameter values used are given in Table 1. According to Dar-El (2000), the learning rates 0.515, 0.321, and 0.152 are associated with “fast,” “medium,” and “slow” learners, respectively. To contextualize these parameter values, Table 1 also notes the reduction in service times resulting from doubling a technician’s experience level on a skill.

	Label		
	<i>Fast</i>	<i>Medium</i>	<i>Slow</i>
L^k	.515	.321	.152
Reduction in d_{rt}^k when experience doubled	30%	20%	10%

Table 1: Learning rates used in experiments

Finally, similar to parameter L_r^k , we have constructed our experiments so that technician k has the same initial experience, p_{r0}^k in all skills. That is, $p_{r0}^k = p_{r'0}^k, \forall r, r' \in \mathcal{R}$. Thus, we can refer to this parameter as p_0^k . We consider three values for this parameter, $p_0^k = 1$, which we call *Low*, $p_0^k = 25$, which we call *Medium*, and $p_0^k = 50$, which we call *High*. We illustrate in Figure 1 each of the three learning curves, as well as where each initial experience level puts a technician on his/her curve.

With three possible learning rates and initial experience levels, there are nine possible combinations of these traits, and the workforce of 18 technicians that we use in our experiments consists of two technicians for each combination. We present the workforce in detail

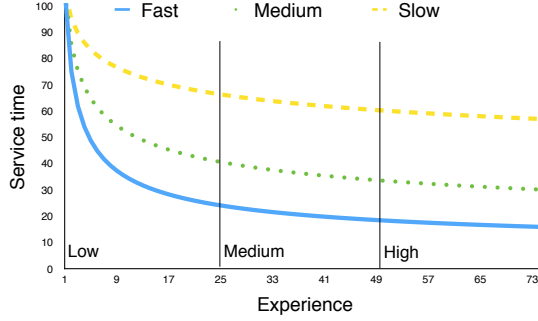


Figure 1: Learning curves

Technician		Learning rate		
		<i>Fast</i>	<i>Medium</i>	<i>Slow</i>
Initial experience	<i>Low</i>	1,10	4,13	7,16
	<i>Medium</i>	2,11	5,14	8,17
	<i>High</i>	3,12	6,15	9,18

Table 2: Workforce

in Table 2. Throughout the rest of the paper, “X-Y” refers to technicians with learning rate “X” and initial experience level “Y.” For example, “Slow-Low” refers to technicians 7,16 (see Table 2), who are slow learners ($L^k = .152$) and have low initial experience ($p_0^k = 1$).

In summary, we consider 10 different orderings of the 29 daily customer sets, and for each ordering, we consider four diversity levels. For each ordering and diversity level, we consider one random assignment of skills to customers. For all experiments, we consider the same workforce. As such, our experimental setup consists of 40 different instances.

6. Computational Analysis

In this section, we perform a computational study to analyze whether, and at what fidelity, learning should be considered when making daily planning decisions. Specifically, we begin our analysis by running a series of experiments to quantify the magnitude of the benefits associated with recognizing both that learning occurs and that each technician’s individual traits impact how much they learn when making daily assignments of technicians to tasks.

We then seek to understand which is the source of these benefits: recognizing that learning occurs or that each technician has individual traits that impact how they learn.

We then turn our analysis to specialization. Specifically, we study whether when making daily assignments a planner should focus on giving a technician a high level of experience in a limited set of skills (making them “experts”) or some experience in many skills (making them “jacks-of-all-trades”). We also study whether a technician’s traits (learning rate, initial experience levels) should determine whether they become “experts” or “jacks-of-all-trades.” Finally, we study whether the number of different skills on which technicians can gain experience impacts whether technicians should specialize and/or which technicians should specialize.

6.1 Value of Heterogeneity and Learning

In this section, we first consider the value of our model versus models that do not incorporate workforce heterogeneity and individual learning. We then demonstrate the value of our model versus models that consider workforce heterogeneity but not individual learning

6.1.1 Comparison versus Not Modeling Heterogeneity and Learning

Our analysis begins by studying whether a service organization should track their technicians’ experience levels on different skills and then recognize both experience level and each technician’s learning rate when performing daily assignments of technicians to tasks. To explore this question, we first run Algorithm 1 on each of the 40 instances described above and record the makespan e_t^{max} for each day. As learning is explicitly recognized when Algorithm 1 is executed, we label these values e_t^{max-L} .

Next, we compare the values e_t^{max-L} with the daily makespan values seen when both learning and a technician’s individual traits (learning rate, experience levels) are ignored when making daily planning decisions. To do so, we run a variant of Algorithm 1 that ignores individual learning rates and individually accumulated experience when determining daily

assignments. Specifically, recalling that technician k 's time to complete a task requiring skill r in period t is modeled by the equation $d_{rt}^k = D_r^k + d_{r0}(p_{rt}^k)^{-L_r^k}$, in this variant of Algorithm 1, daily routes are created with the assumption that $d_{rt}^k = \bar{d} \quad \forall r \in \mathcal{R}, t = 1, \dots, \mathcal{T}, k \in \mathcal{K}$. To calculate the value \bar{d} , we first assume that all 18 technicians described in Table 2 have *Medium* initial experience ($p_0^k = \bar{p}_0 = 25, \forall k \in \mathcal{K}$) and learn at the *Medium* rate ($L^k = \bar{L} = 0.321, \forall k \in \mathcal{K}$). Next, to calculate \bar{d} , we model that, while an organization may not track each technician's actual accumulated experience level, a planner may still use a forecast of accumulated experience when making daily decisions. We presume that this forecast of accumulated experience in each skill depends on the number of different skills required by the customers, or $|\mathcal{R}|$, and report in Table 3 the values $\bar{p}_{|\mathcal{R}|}$. As such, we have that technician k 's time to complete a task requiring skill r in period t is dictated by the equation $d_{rt}^k = D_r^k + d_{r0}(\bar{p}_{|\mathcal{R}|})^{-\bar{L}}$.

	$ \mathcal{R} $			
	5	10	25	50
$\bar{p}_{ \mathcal{R} }$	43	33	28	27

Table 3: Fixed experience levels, \bar{p} , used in experiments

While this variant of Algorithm 1 determines daily routes without acknowledging that technicians accumulate experience or differ in their individual traits, to make a fair comparison to e_t^{max-L} , we calculate daily makespans for those routes using the appropriately updated values for p_{rt}^k, d_{rt}^k as specified by Equations (20) and (21). That is, when calculating the daily makespan for the solutions returned by the variant of Algorithm 1, we recognize that experience accumulates and that the time a technician needs to perform a task requiring a specific skill decreases at a rate that depends on his/her individual traits. Because in these experiments we do not consider learning and essentially assume the workforce is homogeneous when making daily planning decisions, we refer to these daily makespan values as $e_t^{max-NL-H}$.

As an example, consider an instance of the TRSP-EST for which $|\mathcal{R}| = 5$. Per Table 3,

on all days, all technicians are assumed to have the experience level 43 on all skills. Suppose that the variant of Algorithm 1 described above prescribed that Technician 3 performs three tasks that required skill 2 on day 0 and is currently determining the task assignments for day 1. When executing the RTR code to solve the routing problem for day 1, d_{23}^1 , or the time Technician 3 needs to perform a task requiring skill 2, is set to the value $D_2^3 + d_{20}(43)^{-.321}$ (i.e. the experience accumulated on day 0 is ignored and the *Medium* learning rate is assumed). However, after the routes have been created, the makespan value, $e_1^{max-NL-H}$ is calculated with the value d_{23}^1 set to the value $D_2^3 + d_{20}(50 + 3)^{-.515}$ (recall that Technician 3 is a fast learner with high initial experience).

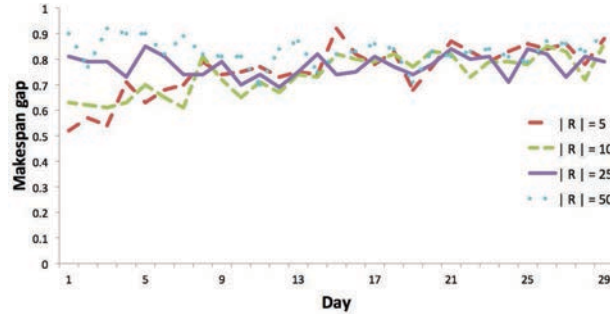


Figure 2: Impact of planning without recognizing individual traits or learning

Figure 2 presents the daily relative gap in the makespans (calculated as $\frac{e_t^{max-NL-H} - e_t^{max-L}}{e_t^{max-NL-H}}$), averaged over instances with the same customer diversity level. We see in this figure that assuming a homogeneous workforce that does not learn leads to significantly worse routes across all diversity levels and that these routes tend to get worse over time.

To try to understand the source of these large gaps, Figures 3(a)-3(d) present the degree to which the daily assignment of tasks to technicians is imbalanced. Specifically, for the case in which the workforce is assumed to be homogeneous and learning is not recognized, the line “NL-H-Low” represents the number of tasks assigned to the technician that is assigned the smallest number of tasks, averaged over the 10 instances associated with the respective diversity level. The line associated with “L-Low” is the analogous result for the workforce when recognizing learning and technician heterogeneity. For the case in which the workforce

is assumed to be homogeneous and learning is not recognized, the line “NL-H-High” represents the number of tasks assigned to the technician that is assigned the largest number of tasks, averaged over the 10 instances associated with the respective diversity level. The line “L-High” is the analogous result for the workforce when recognizing learning and technician heterogeneity. The results show that, across all diversity levels, there is a larger imbalance in assigned work when learning is recognized.

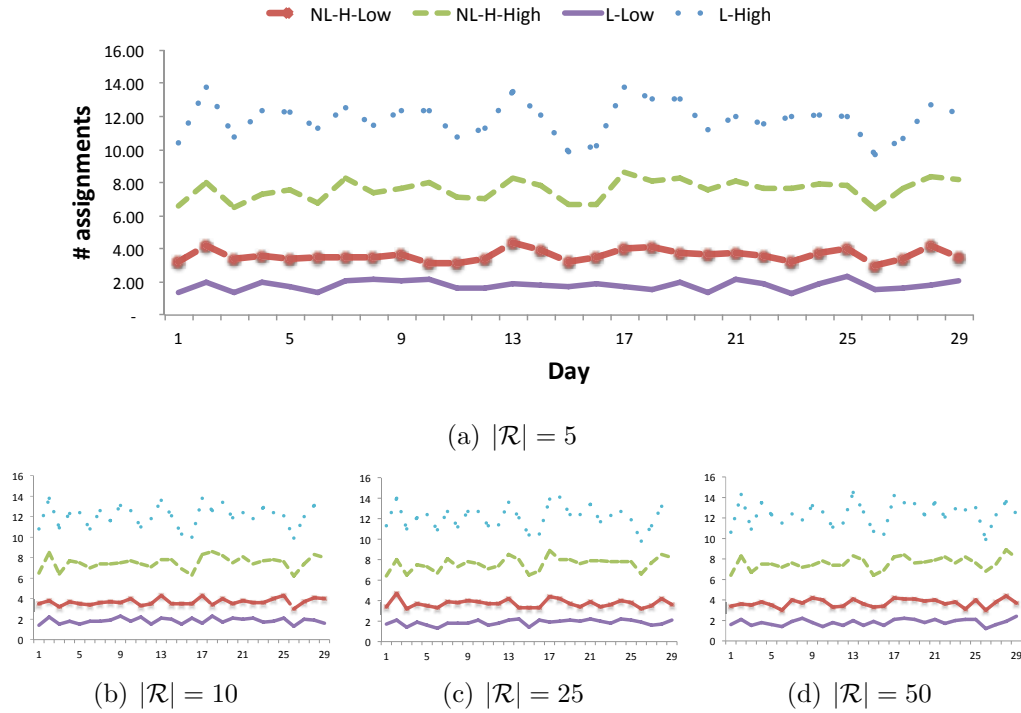


Figure 3: Workload imbalance by diversity

This imbalance results from assignments in the case of recognizing learning and heterogeneity that take advantage of the productivity associated with fast learners. Figure 4 shows, by technician type (combination of learning rate and initial experience), the total number of tasks performed. The results show that the solutions considering both learning and workforce heterogeneity assign more tasks to technicians with the *Fast* learning rate and fewer to those that have a *Slow* learning rate. Alternatively, the poorer quality solutions resulting from not considering heterogeneity distribute tasks evenly across the workforce. In

summary, when not recognizing learning or individual traits, one ignores the extra capacity associated with *Fast* learners.

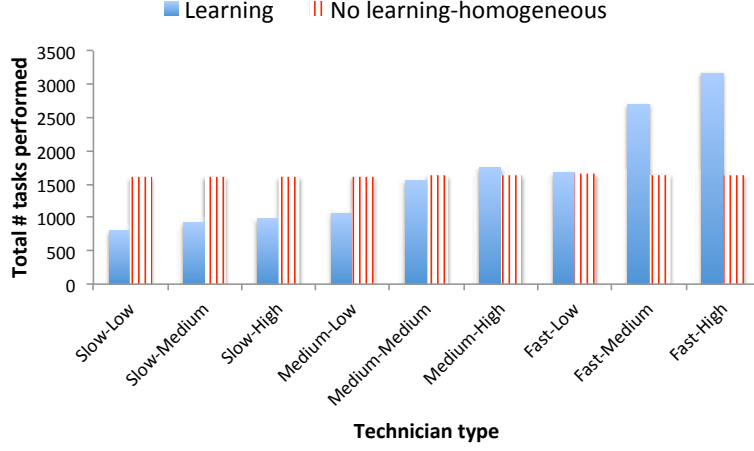


Figure 4: Total number of tasks performed by technician type

While our results are from a 29-day planning horizon, we believe the problems associated with ignoring learning and assuming a homogeneous workforce will persist over longer horizons. As evidence, Figures 5(a)-5(d) present the distribution of technicians' experience levels at the end of the planning horizon, averaging over all 10 instances with a task diversity of 5 and by technicians with the same learning rate and initial experience. Specifically, to calculate the "Skill 1" point on the "Slow-Low" curve, we first calculate the quantities $f_7^1 = \frac{p_{1,7}^{29}}{\sum_{r=1}^5 p_{r,7}^{29}}$ and $f_{16}^1 = \frac{p_{1,16}^{29}}{\sum_{r=1}^5 p_{r,16}^{29}}$, with the first representing the fraction of technician 7's assignments over the 29-day period that were to a task that required Skill 1 and the second is the same only for technician 16. Then, the "Skill 1" point for a "Slow-Low" technician is calculated as $f_{SL}^1 = (f_7^1 + f_{16}^1)/2$. Other points in the figures are calculated similarly.

The figures show that the distribution of experience levels for technicians are different when learning and individual traits are recognized from when they are not. For example, when learning and individual traits are recognized, the experience levels for technicians with *Low* initial experience (Figure 5(b)) are not evenly distributed across all skills as they are when they are ignored. As these experience levels impact the productivity of technicians

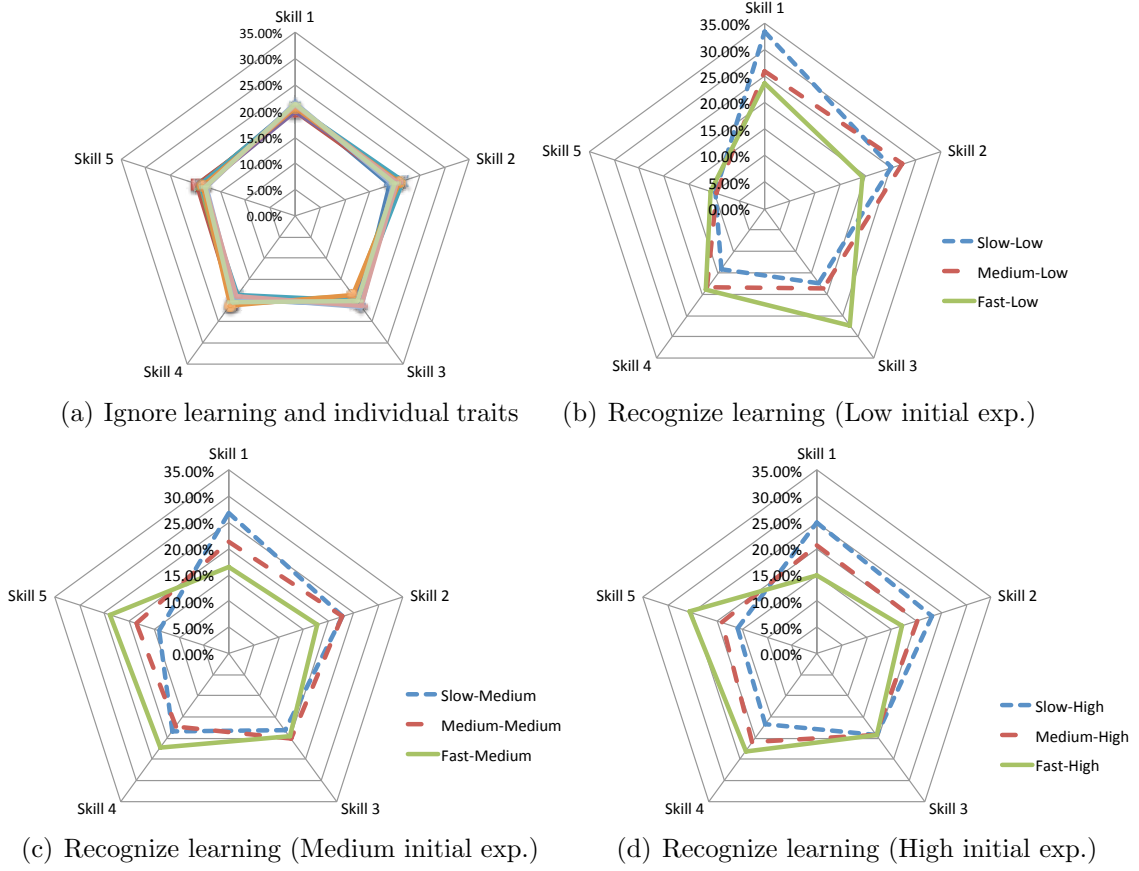


Figure 5: Distribution of experience levels at end of 29 days ($|\mathcal{R}| = 5$)

in subsequent days, we conclude that ignoring learning and individual traits will leave a workforce poorly prepared for the future. We note that similar graphs for other diversity levels ($|\mathcal{R}| = 10, 25, 50$) exhibit a similar pattern. Recognizing learning and individual traits leads to a workforce with experience unevenly distributed across all skills whereas not recognizing these factors does not. We explore this issue further in the next subsection.

6.1.2 Comparison versus Modeling of Heterogeneity but not Learning

In the previous analysis, recognizing learning and each technician’s individual traits leads to much better daily decisions and a different workforce at the end of the 29-day planning horizon, in terms of skill proficiency, than not doing so. We next discuss experiments that seek to understand whether recognizing workforce heterogeneity, but still not recognizing learning, will close the gaps seen in Figure 2.

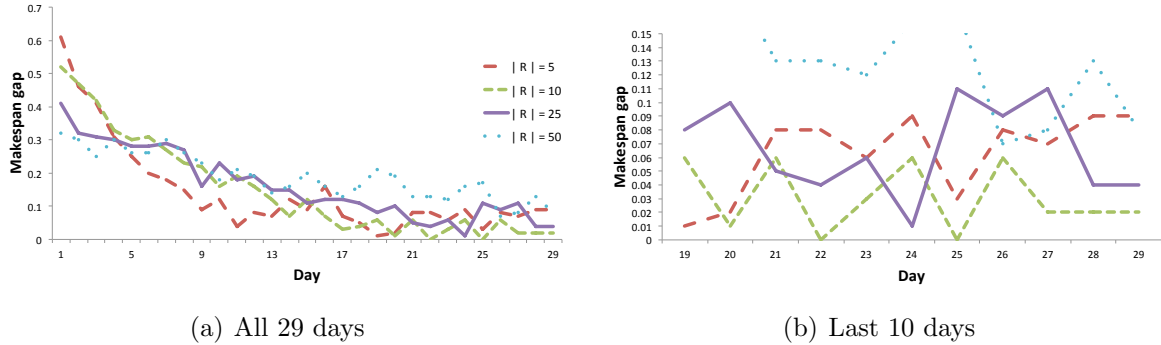


Figure 6: Impact of planning without recognizing learning but recognizing individual traits

Specifically, we again execute a variant of Algorithm 1 where daily routes are created with the assumption that each technician has a static experience level on each skill. In contrast to the previous experiments in which experience was fixed based on the assumption that everyone began with a *Medium* experience level, in these experiments, the fixed experience level for a technician is based on the appropriate value p_0^k according to Table 2. Similarly, in the previous experiments, the time a technician needed to perform a task \bar{d} was calculated based on the assumption that the technician learned at the *Medium* learning rate. In these

experiments, the value \bar{d}_k for technician k is calculated based on their individual value L^k from Table 1 and according to Table 2. As in the previous experiment, we also assume that each technician reaches a level of forecasted experience. Finally, analogous to the previous experiment, while we generate solutions that ignore learning, we evaluate the solutions returned by the algorithm by including learning and accumulated experience. We denote these solutions values by e_t^{max-NL} for each day t in the horizon.

We now perform an analysis similar to what was done for the previous experiments. Figures 6(a) and 6(b) present results calculated in a manner similar to those presented in Figure 2. We see that, while recognizing individual traits reduces the gaps seen in Figure 2, it does not close them. The gaps still range from 2% to 10% at the end of the 29-day planning horizon.

Evidence for why the gaps close can be seen in Figures 7(a) and 7(b), which display results similar to those seen in Figures 3(a) and 4. The figures show that recognizing workforce heterogeneity, yet still ignoring learning, leads to a much better appreciation of the capacity associated with *Fast* learners and thus a similar imbalance of work as when learning is recognized. While we only show figures for $|\mathcal{R}| = 5$, results for other diversity levels exhibit a similar pattern.

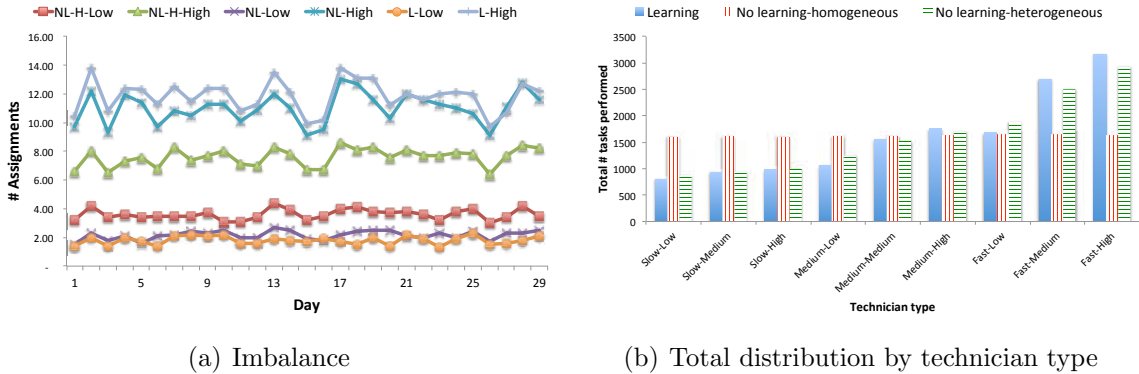


Figure 7: Workload analysis when recognizing individual traits but not learning ($|\mathcal{R}| = 5$)

The reason that the gaps do not close completely can be seen by analyzing two differences in the distributions of experience levels at the end of the 29-day period: (1) the differences

between when learning is recognized and when both learning and individual traits are ignored, and, (2) the differences between when learning is recognized and when learning is ignored but individual traits are not. We calculate the quantities f_{XY}^r (e.g. f_{SL}^1 as discussed previously) for each of the three runs of Algorithm 1: (1) when learning is recognized, which we label f_{XY}^{r-L} , (2) when learning and individual traits are ignored, which we label f_{XY}^{r-NL-H} , and (3) when learning is ignored but individual traits are not, which we label f_{XY}^{r-NL} . Figure 8(a) presents the quantities $|f_{XY}^{r-L} - f_{XY}^{r-NL-H}|$, $\forall r \in \mathcal{R}$ and all technician types, “X-Y” (e.g. “Slow-Low”). Figure 8(b) presents the quantities $|f_{XY}^{r-L} - f_{XY}^{r-NL}|$, $\forall r \in \mathcal{R}$ and all technician types. Similar to what was seen in Figures 5(a) to 5(d), we see that ignoring

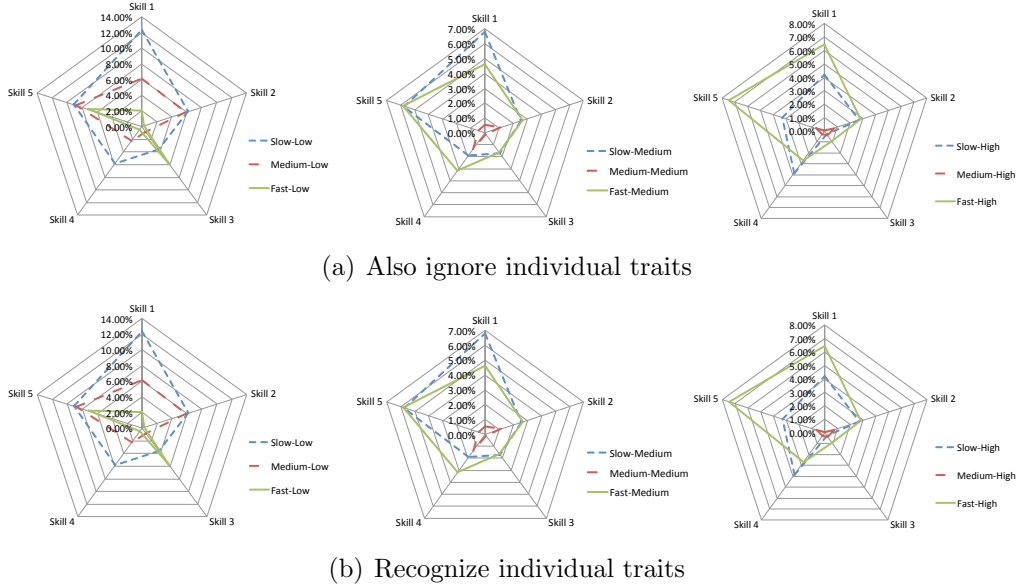


Figure 8: Differences in distribution of experience levels at end of 29 days when recognizing learning and not ($|\mathcal{R}| = 5$)

learning and individual traits leaves a workforce with a very different set of skill levels than recognizing learning does. Figure 8(b) suggests that recognizing individual traits reduces the differences in skill levels but does not eliminate them. Again, while we only display results for $|\mathcal{R}| = 5$, the pattern is similar for other diversity levels. In summary, while recognizing individual traits is better than not doing so, it is best to recognize learning as it has benefits in the short term and the impact of doing so persists through differences in the distribution

of final experience levels.

6.2 Workforce Specialization

Having established that an organization should consider both an up-to-date accounting of their technician’s experience levels at different skills as well as their individual learning rates, we next turn our attention to what insights can be gained into how experience levels and learning rates should be mapped to daily scheduling and routing decisions. Our analysis primarily focuses on the concept of *specialization* and looks at whether a technician should become an expert in a small set of skills and whether a technician’s traits impact this decision.

To measure specialization, we use the Coefficient of Variation ($CV = \frac{\text{StandardDeviation}}{\text{Mean}}$) of the number of times a technician performs tasks requiring each skill. For example, suppose that at the end of 29 days, a technician has performed tasks that require Skill 1 five times, tasks that require Skill 2 one time, tasks that require Skill 3 two times, tasks that require Skill 4 three times, and tasks that require Skill 5 four times. Then, the expected number of times that technician performs a task requires a given skill is $\mu = (5 + 1 + 2 + 3 + 4)/5 = 3$. Similarly, the standard deviation of the number of times a task requiring a skill can be calculated as $\sigma = \sqrt{\frac{(5-3)^2 + (1-3)^2 + (2-3)^2 + (3-3)^2 + (4-3)^2}{5}} = 1.414$, leaving a $CV = \frac{1.414}{3} = 0.471$. In the analysis that follows, we calculate CV values based on a technician’s experience at the end of the 29-day planning horizon. Specifically, $CV_k = \frac{\sigma_k}{\mu_k}$ where $\mu_k = \frac{\sum_{r \in \mathcal{R}} p_{r29}^k}{|\mathcal{R}|}$ and $\sigma_k = \sqrt{\frac{\sum_{r \in \mathcal{R}} (p_{r29}^k - \mu_k)^2}{|\mathcal{R}|}}$. When referring to a CV value for a technician type, such as a “Slow-Low” technician, we refer to the average of the CVs of the technicians that are of that type; e.g. $CV_{Slow-Low} = (CV_7 + CV_{16})/2$. Higher CV values represent greater specialization as they indicate a wider disparity between the tasks done the most times and those done the least.

Figures 9(a)-9(d) present CV values averaged over all 10 instances of each customer diversity and by technician type. The figures show that, regardless of diversity level, technicians with Low initial experience specialize the most, and their degree of specialization increases as

diversity increases (note the scale is different for each figure). On the other hand, technicians with High initial experience specialize the least.

To understand the result, consider that workers with low initial experience are relatively inefficient in all tasks, and more importantly, see a very large marginal benefit from additional experience on a task. As a result, the workers with Low initial experience see significant gains in productivity on the tasks to which they are assigned in the first few days of the horizon. These gains make Low initial experience workers relatively more attractive on those tasks relative to workers with High initial experience. Thus, the best solutions see Low initial experience workers frequently assigned to the tasks on which they gain the most experience in the first few days.

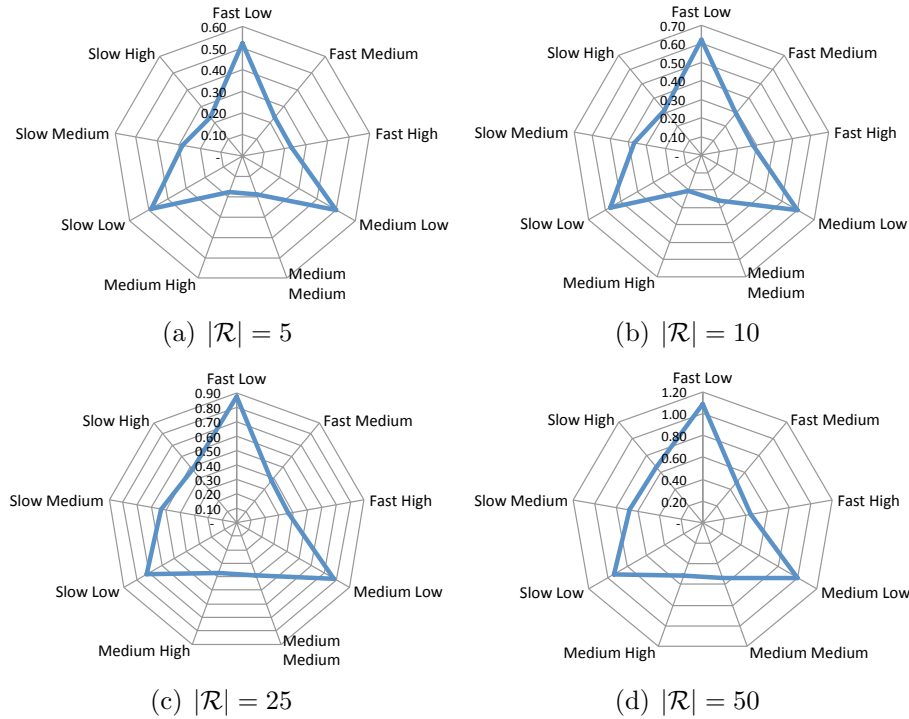


Figure 9: Specialization by customer diversity and technician type

We next look at specialization from a different perspective. If daily routes were evaluated based only on travel times, one expects specialization to only occur due to the geographic distribution of customers (e.g. a cluster of customers requesting tasks that require the same

skill). Conversely, if daily routes were evaluated based only on service times, one would expect a high degree of specialization. We next seek to understand whether, when making daily task assignment decisions, some technician types can be treated as individuals for whom one should focus on routing (minimizing travel time) or scheduling (minimize service time).

To do so, we run two different variants of Algorithm 1 for each instance; each variant differs from Algorithm 1 only in how e_t^{max} is calculated. The first variant only counts travel times when minimizing e_t^{max} . We treat the CV values (which we label CV_{X-Y}^{TT}) in solutions produced with this variant as lower bounds on the amount of specialization one should see in a solution produced by Algorithm 1. Similarly, the second variant only counts service times when minimizing e_t^{max} , and we interpret the CV values (which we label CV_{X-Y}^{ST}) in solutions produced with this variant as upper bounds. We present an example of these three sets of CV values in Figure 10. With these values, we hypothesize that a planner should focus on routing technician type “X-Y” when

$$CV_{X-Y} - CV_{X-Y}^{TT} < CV_{X-Y}^{ST} - CV_{X-Y}, \quad (22)$$

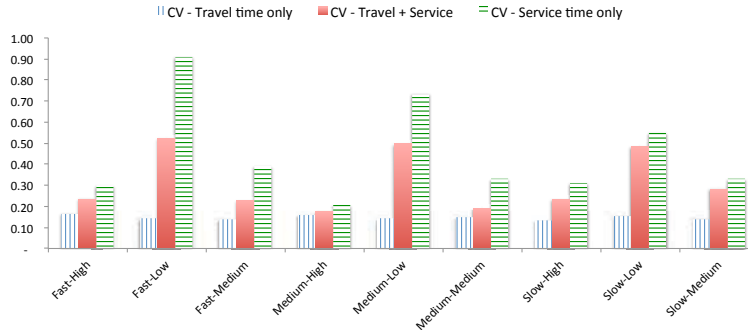


Figure 10: Determining assignment rules ($|\mathcal{R}| = 5$)

or, the CV value produced when running Algorithm 1 is closer to the one produced when only counting travel times than the one produced when only counting service times. We

then say that a planner should focus on scheduling technician type “X-Y” when Inequality (22) does not hold. For example, referring to Figure 10, we conclude that a planner should focus on routing “Medium-Medium” technicians and scheduling “Slow-Low” technicians. Essentially, specializing matters more for “Slow-Low” technicians than “Medium-Medium” technicians.

We then calculate average CV, CV_{X-Y}^{TT} , and CV_{X-Y}^{ST} values over all 10 instances for each diversity level and apply this line of reasoning to derive heuristics or “rules of thumb” for what a planner should focus on when making daily task assignments. Tables 11(a)-11(d) present these rules for all technician types and by customer diversity level. We see that the rules are fairly static across customer diversity levels, with the only differences being the degree to which a planner should focus on scheduling individuals with “Low” initial experience. For all cases of a “High” or “Medium” initial experience and either “Fast” or “Medium” learning rate, routing is the dominant factor in determining the assignment. Further, when initial experience is “Low” and the learning rate is “Slow,” the scheduling dominates. Further, for diversity levels of 25 and 50, the assignments of the technicians with “Low” initial experience are dominated by the scheduling concerns. Differentiation comes at the lower diversity levels and the case of “Low” initial experience. For the case of a diversity of 10, technicians with “Medium” initial experience and “Slow” or “Medium” learning rates are dominated by routing. For the diversity of five, technicians with “Medium” initial experience, and “Slow” or “Medium” learning rates, the assignments are dominated by routing and scheduling, respectively. We believe that the difference in the “Low” initial experience and “Medium” learning rate is the result of sampling error and that there is no significant difference to focusing assignments on the routing or scheduling of “Slow” or “Medium” learners with “Low” initial experience.

Assignment rule		Learning rate		
		<i>Fast</i>	<i>Medium</i>	<i>Slow</i>
Init.	<i>Low</i>	Route	Schedule	Schedule
exp.	<i>Medium</i>	Route	Route	Schedule
	<i>High</i>	Route	Route	Schedule

(a) $|\mathcal{R}| = 5$

Assignment rule		Learning rate		
		<i>Fast</i>	<i>Medium</i>	<i>Slow</i>
Init.	<i>Low</i>	Route	Route	Schedule
exp.	<i>Medium</i>	Route	Route	Schedule
	<i>High</i>	Route	Route	Schedule

(b) $|\mathcal{R}| = 10$

Assignment rule		Learning rate		
		<i>Fast</i>	<i>Medium</i>	<i>Slow</i>
Init.	<i>Low</i>	Schedule	Schedule	Schedule
exp.	<i>Medium</i>	Route	Route	Schedule
	<i>High</i>	Route	Route	Schedule

(c) $|\mathcal{R}| = 25$

Assignment rule		Learning rate		
		<i>Fast</i>	<i>Medium</i>	<i>Slow</i>
Init.	<i>Low</i>	Schedule	Schedule	Schedule
exp.	<i>Medium</i>	Route	Route	Schedule
	<i>High</i>	Route	Route	Schedule

(d) $|\mathcal{R}| = 50$

Figure 11: Assignment rules by customer diversity and technician type

7. Conclusions

In this paper, we present a model and approach for accounting for on-the-job learning in the daily routing of technicians. Our objective minimizes the completion time of the last task. We solve the daily routing problems using the record-to-record travel (RTR) heuristic. Our results offer the following key insights:

1. Explicitly modeling both learning and technician heterogeneity leads to better and different solutions in comparison to assuming homogeneous learning curves and/or static productivity.
2. Relatedly, explicitly modeling both learning and technician heterogeneity leads to different distributions of skills in comparison to assuming homogeneous learning curves and/or static productivity.
3. Inexperienced technicians specialize the most and experienced technicians the least.

The differences in the solutions to the model that incorporates workforce heterogeneity and human learning versus those that do not demonstrates the importance of incorporating workforce characteristics into the model when applicable. While the importance of such characteristics have been discussed in the manufacturing literature, this paper is the first to show it for routing problems.

Ultimately, our results reveal the need to balance routing and scheduling when making assignments for a heterogeneous workforce that learns. Our experiments and “rules of thumb” identify that the scheduling aspect is particularly dominant for workers who are “Slow” learners or who are of “Low” initial experience. Routing dominates the assignments for “Fast” or “Medium” learners coupled with either “High” or “Medium” initial experience. This rule-based approach can help managers develop effective daily plans with less computational efforts than what is required by our optimization.

Because the on-the-job learning is a new addition to the technician routing and scheduling literature, there are many directions for future research. One direction is to incorporate information about the future into the solution approach. Currently, we decompose the problem into a series of myopic, daily routing problems. By incorporating some estimate of the future, we can direct the effects of learning. In particular, by including an estimate of the future, the solution for today can trade off near-term task completion times for experience enhancement that will yield benefits in later days. This global consideration over the planning horizon may better demonstrate the difference between fast and slow learning technicians in improving the routing and scheduling decisions. At the same time, this work considers a min-max objective. In some circumstances, it may be more appropriate to minimize the sum of service and routing times or of cost. Similarly, because the min-max objective can lead to workload imbalances, the addition of workload balancing constraints may be desirable. Finally, on any given day, a technician may work slower or faster than suggested by his/her experience and the learning curve. In such cases, it might make sense to dynamically adjust assignments throughout the day. Another avenue of future work would be to consider such dynamic routing.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. CMMI-1266010

References

- Alsheddy, A., Tsang, E. P. K., 2011. Empowerment scheduling for a field workforce. *Journal of Scheduling*, 14 (6), 639–654.
- Anzanello, M. J., Fogliatto, F. S., 2011. Learning curve models and applications: Literature review and research directions. *International Journal of Industrial Ergonomics*, 41 (5), 573–583.
- Bard, J. F., Shao, Y., Qi, X., Jarrah, A. I., 2014. The traveling therapist scheduling problem. *IIE Transactions*, 46 (7), 683–706.
- Bertels, S., Fahle, T., 2006. A hybrid setup for a hybrid scenario: combining heuristics for the home health care problem. *Computers & Operations Research*, 33 (10), 2866–2890.
- Bostel, N., Dejax, P., Guez, P., Tricoire, F., 2008. Multiperiod Planning and Routing on a Rolling Horizon for Field Force Optimization Logistics. In B. Golden, S. Raghavan, E. Wasil (eds.), *The Vehicle Routing Problem: Latest Advances and New Challenges*, Springer US, Boston, MA, volume 43 of *Operations Research/Computer Science Interfaces*, 503–525.
- Cordeau, J.-F., Laporte, G., Pasin, F., Ropke, S., 2010. Scheduling technicians and tasks in a telecommunications company. *Journal of Scheduling*, 13 (4), 393–409.
- Dar-El, E. M., 2000. *Human Learning: From Learning Curves to Learning Organizations*, volume 29 of *International Series in Operations Research & Management Science*. Kluwer Academic Publishers, Boston.
- De Jong, J., 1957. The effects of increasing skill on cycle time and its consequences for time

- standards. *Ergonomics*, 1 (1), 51–60.
- Dutot, P.-F., Laugier, A., Bustos, A.-M., 2006. Technicians and Interventions Scheduling for Telecommunications. Technical report, France Telecom Research and Development.
- Firat, M., Hurkens, C. A. J., 2011. An improved MIP-based approach for a multi-skill workforce scheduling problem. *Journal of Scheduling*.
- Goodson, J. C., Thomas, B. W., Ohlmann, J. W., 2015. A Rollout Algorithm Framework for Heuristic Solutions to Finite-Horizon Stochastic Dynamic Programs. Available from <http://slu.edu/~goodson/papers/GoodsonRolloutFramework.pdf>.
- Groër, C., Golden, B., Wasil, E., 2010. A library of local search heuristics for the vehicle routing problem. *Mathematical Programming Computation*, 2 (2), 79–101.
- Hashimoto, H., Boussier, S., Vasquez, M., Wilbaut, C., 2011. A GRASP-based approach for technicians and interventions scheduling for telecommunications. *Annals of Operations Research*, 183 (1), 143–161.
- Hewitt, M., Chacosky, A., Grasman, S., Thomas, B. W., 2014. Integer Programming Techniques for Solving Non-linear Workforce Planning Models with Learning, submitted for publication, available from <http://myweb.uiowa.edu/bthoa/iowa/Research.html>.
- Hurkens, C. A., 2009. Incorporating the strength of MIP modeling in schedule construction. *RAIRO - Operations Research*, 43 (04), 409–420.
- Jaber, M. Y., 2006. Learning and Forgetting Models and Their Applications. In A. B. Badiru (ed.), *Handbook of Industrial and Systems Engineering*, CRC Press, Boca Raton, FL, USA, Industrial Innovation Series, chapter 30, 30–1–30–27.
- Jaber, M. Y., Sikström, S., 2004. A numerical comparison of three potential learning and forgetting models. *International Journal of Production Economics*, 92 (3), 281–294.
- Kovacs, A., Parragh, S., Doerner, K., Hartl, R., 2012. Adaptive large neighborhood search for service technician routing and scheduling problems. *Journal of Scheduling*, 15, 579–600.
- Lesaint, D., Voudouris, C., Azarmi, N., 2000. Dynamic Workforce Scheduling for British

- Telecommunications plc. Interfaces, 30 (1), 45–56.
- Li, F., Golden, B., Wasil, E., 2005. Very large-scale vehicle routing: new test problems, algorithms, and results. Computers & Operations Research, 32 (5), 1165 – 1179.
- Panteva, N., 2012. Heating up: Demand for HVAC services will rise along with construction market recovery. IBISWorld Industry Report OD5583.
- Pillac, V., Guéret, C., Medaglia, A., 2012a. A parallel matheuristic for the technician routing and scheduling problem. Optimization Letters, 7 (7), 1525–1535.
- Pillac, V., Guéret, C., Medaglia, A., 2012b. On the dynamic technician routing and scheduling problem. Technical Report 12/5/AUTO, École des Mines de Nantes, Nantes, France.
- Powell, W. B., 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley Series in Probability and Statistics, Wiley, Hoboken, New Jersey, second edition.
- Rasmussen, M. S., Justesen, T., Dohn, A., Larsen, J., 2012. The home care crew scheduling problem: Preference-based visit clustering and temporal dependencies. European Journal Of Operational Research, 219 (3), 598–610.
- Ree, M. J., Earles, J. A., 1991. Predicting training success: Not much more than g. Personnel Psychology, 44 (2), 321–332.
- Smilowitz, K., Nowak, M., Jiang, T., 2012. Workforce Management in Periodic Delivery Operations. Transportation Science, 47 (2), 214–230.
- Tsang, E., Voudouris, C., 1997. Fast local search and guided local search and their application to British Telecom’s workforce scheduling problem. Operations Research Letters, 20 (3), 119–127.
- Weintraub, A., Aboud, J., Fernandez, C., Laporte, G., Ramirez, E., 2012. An emergency vehicle dispatching system for an electric utility in chile. Journal of Operations Research Society, 44 (10), 690–696.
- Wright, T. P., 1936. Factors affecting the cost of airplanes. Journal of Aeronautical Sciences,

4 (3), 122–128.

Zhong, H., Hall, R. W., Dessouky, M., 2007. Territory planning and vehicle dispatching with driver learning. *Transportation Science*, 41 (1), 74–89.