European Journal of Operational Research 242 (2015) 942–950

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Integer programming techniques for solving non-linear workforce planning models with learning

Mike Hewitt[a,*], Austin Chacosky[b], Scott E. Grasman[b], Barrett W. Thomas[c]

[a] Department of Information Systems and Operations Management, Loyola University Chicago, Chicago, IL 60611, USA
[b] Department of Industrial and Systems Engineering, Rochester Institute of Technology, Rochester, NY 14623, USA
[c] Department of Management Sciences, University of Iowa, Iowa City, IA 52242, USA

ARTICLE INFO

Article history:
Received 30 October 2013
Accepted 28 October 2014
Available online 15 November 2014

Keywords:
Production planning and scheduling
Human learning
Nonlinear programming

ABSTRACT

In humans, the relationship between experience and productivity, also known as learning (possibly also including forgetting), is non-linear. As a result, prescriptive planning models that seek to manage workforce development through task assignment are difficult to solve. To overcome this challenge we adapt a reformulation technique from non-convex optimization to model non-linear functions with a discrete domain with sets of binary and continuous variables and linear constraints. Further, whereas the original applications of this technique yielded approximations, we show that in our context the resulting mixed integer program is equivalent to the original non-linear problem. As a second contribution, we introduce a capacity scaling algorithm that exploits the structure of the reformulation model and reduces computation time. We demonstrate the effectiveness of the techniques on task assignment models wherein employee learning is a function of task repetition.

1. Introduction

Peter Senge famously wrote that the learning organization is "continually expanding its capacity to create its future" (Senge, 2006). Citing employee learning as an effective way to create capacity from existing resources, Levinthal and March (1993) argue that effective management of employee learning can lead to a sustainable competitive advantage. Some go further and claim that, in the modern business environment, the only sustainable competitive advantage will come from a company's ability to learn more effectively than its competitors (Kapp, 1999). Even if it is not the only source of competitive advantage, Moustaghfir (2009) identifies learning as a core component of organizational capabilities that are "immune to competitive duplication." In this paper, we address the issue of learning at the operational level by developing solution techniques for workforce planning models that incorporate individual on-the-job learning.

In particular, this paper focuses on a class of multi-period workforce planning models that recognize human learning and forgetting over a fixed planning horizon in environments that can be characterized as task assignment. We consider a workforce that produces a product, the production of which requires the completion of a sequence of tasks. Inventory is allowed to accumulate in between each task in the sequence. We assume that the workforce's current skill levels and thus their productivity are heterogeneous. As workers perform a particular task, they gain experience that correspondingly leads to a productivity improvement as they learn on the job. In addition, when workers are not doing a particular task, they forget some of what they have learned and their productivity on that particular task erodes. We assume that each worker has his or her own rate of learning and forgetting for each task. The objective is to maximize production over a fixed time horizon.

In humans, the relationship between experience and productivity, also known as learning (possibly also including forgetting), is non-linear. Thus, prescriptive models incorporating learning are difficult to solve. This paper helps overcome this challenge in two ways.

First, it adapts a reformulation technique for non-convex optimization problems that enables linearization of the learning curves. For general non-convex optimization, the reformulation yields an approximation of the original non-linear program (NLP). However, in our application, the non-linear functions have a specific structure that has not previously been explored. Specifically, they are functions of the number of times an individual has performed a task over a fixed number of periods, and thus have discrete and finite domains. In this case, solutions to the mixed integer program (MIP) resulting from the reformulation are optimal for the original non-linear program. As a result, we can solve MIPs instead of NLPs, and with the superior

* Corresponding author. Tel.: +1 585 789 8034.
E-mail addresses: mhewitt3@luc.edu (M. Hewitt), atc7417@mail.rit.edu (A. Chacosky), segeie@rit.edu (S. E. Grasman), barrett-thomas@uiowa.edu (B. W. Thomas).
URL: http://tippie.uiowa.edu/barrett-thomas (B. W. Thomas)

http://dx.doi.org/10.1016/j.ejor.2014.10.060
0377-2217/

capabilities of MIP solvers to those of NLP solvers, we are able to solve problems that are significantly larger than what has been previously possible.

The reformulation technique associates a binary variable with each element of the domain of the non-linear function we wish to linearize. In our application this translates to representing each potential skill level for each worker in each period of the planning horizon with a binary variable. As a result, as instance sizes grow, the resulting MIPs become large, increasing solve times. Thus, as a second contribution, we derive structural properties of the feasible region of the MIP that in turn motivate a capacity scaling algorithm wherein these variables are generated dynamically. With an extensive computational study, we see that this algorithm reduces solution time. The capacity scaling algorithm can be applied to any MIP resulting from the reformulation technique.

The remainder of this paper is organized as follows. In Section 2, we review the relevant literature. Section 3 presents our reformulation technique, and Section 4 introduces the capacity scaling heuristic that exploits the structure of the reformulation. Section 5 presents the non-linear and reformulated math programs for the workforce planning model on which we perform computational tests. The section also demonstrates how to implement the capacity scaling algorithm with the reformulated workforce planning model. Using instances of different sizes and characteristics, Section 6 studies the computational times associated with the reformulated math program and the scaling algorithm. Section 7 offers conclusions and opportunities for future work.

2. Literature review

Given its importance, a significant body of research has been devoted to developing quantitative models of individual learning and forgetting. These models are often called learning curves. Of particular interest for this paper is research devoted to on-the-job or experiential learning, which is sometimes also called autonomous learning. Dar-El (2000) provides a comprehensive review of both learning and forgetting models as well as parameter estimation for learning models developed before 2000. Jaber and Sikström (2004) and Anzanello and Fogliatto (2011) provide overviews that account for more recent work on learning and forgetting models. Taking advantage of improved data collection capabilities resulting from bar code readers and other similar devices, Nembhard and Uzumeri (2000a) conclude that a three-parameter hyperbolic function provides best fits observations of individual learning. Nembhard and Uzumeri (2000b) extend the three-parameter hyperbolic function to incorporate forgetting. Nembhard (2001) examines several models of forgetting and identifies those with robust performance. Shafer, Nembhard, and Uzumeri (2001) introduce a model that accounts for the recency of a task. Jaber and Sikström (2004) provide a numerical analysis that, for three models, identifies in which particular applications each model has the best statistical fit of the learning and forgetting measured in the application. While our work is amenable to other models of learning and forgetting, we focus this work on the exponential function used in Nembhard and Norman (2007, chap. 4). The function was introduced by Thomas and Nembhard (2005) as an exponential learning and forgetting function designed to achieve the performance of the three-parameter hyperbolic function while allowing for improved tractability in optimization models.

A number of authors have incorporated models of learning or of learning and forgetting into optimization models for workforce management. While the existing work covers a variety of different applications, a common theme is the challenge in solving large-sized instances. The success of exact solution approaches has been particularly limited. Nembhard and Norman (2007) introduce a task-assignment model that includes learning and forgetting. Computational results are presented for a two workers, four tasks, and 10 time periods example. For ease of exposition, we choose a model similar to what is presented in Nembhard and Norman (2007) to illustrate our reformulation technique. Heimerl and Kolisch (2010) also consider an assignment model with the addition of constraints that represent the firm's desired skill composition at the end of a specified period of time. Kim and Nembhard (2010) use a non-linear mixed integer program to test the effects of experimental factors on cross-training policy selection. While the number of tasks is difficult to determine, the largest problem size considered has three workers and 24 periods. For the situation in which each station has infinite buffers and a problem similar to that studied in this paper, Nembhard and Bentefouet (2012) identify the structure of the optimal policy for the case where the number of tasks and workers is the same. The result allows the authors to solve the problem up to 96 workers, 96 tasks, and 246 periods. Using the flow-line production scenario used in this paper as an application and for small problems of two to three workers and two to three tasks, Bentefouet and Nembhard (2013) identify structural properties that characterize the optimal solution. The results do not generalize.

In addition to the specific structural results, however, Bentefouet and Nembhard (2013) present a model that provides an upper bound on the production that can be achieved from the application discussed in this paper. The nature of the formulation of the upper bound problems allows the authors to linearize the learning functions in a way that is a special case of the reformulation presented in Section 3. As noted in Bentefouet and Nembhard (2013), solutions to their presented model do not generally offer implementable task assignment schedules. The techniques discussed in this paper do offer implementable task assignments. The authors of this paper are not aware of any papers that offer a mixed integer reformulation of a non-linear function with a discrete domain.

To overcome the non-linearities of the learning and forgetting functions, a number of authors consider approximation schemes. In the literature that follows, none of the authors addresses how well their schemes approximate the learning and forgetting functions nor do they discuss whether or not the approximation impacts solution quality relative to using the actual learning and forgetting functions. Nembhard and Bentefouet (2012) introduce a rectangular approximation for learning/forgetting functions. For a problem similar to the application used in this paper, the approximation can solve 10 worker, 10 task, and 40 period problems. Corominas, Olivella, and Pastor (2010) introduce a piece-wise linear transformation of a convex learning function for a task assignment problem that incorporates learning resulting from experience with related tasks. The largest problem solved in that paper has five tasks, four workers, and 20 time periods. Olivella, Corominas, and Pastor (2013) combines piecewise linearization with constraint relaxation. Sayin and Karabati (2007) also use a piecewise linear approximation of the learning function, but in a problem in which they first try to maximize utility and then skill improvement. In addition to the linearization, the problem is solved for only a single period with each period's solution implemented as part of a simulation. Problems with at most 18 workers and four different skills are solved. Gutjahr, Katzensteiner, Reiter, Stummer, and Denk (2008) consider maximizing a weighted average of economic gains and skill development for a project selection problem. As part of the selection process, assignments to selected projects are optimized. Gutjahr et al. (2008) introduce a first-order approximation of the non-linear learning curve and are resultantly able to solve the approximate model to optimality for an example with 14 candidate projects requiring a mix of 40 skills (analogous to tasks in our discussion) with 28 workers over 24 periods.

A number of heuristic approaches to task assignment with learning can be found in the literature. Yan and Wang (2011) consider a model with the same learning function considered in this paper. They introduce a genetic algorithm and solve a problem with six