

Sequential event prediction

Benjamin Letham · Cynthia Rudin · David Madigan

Received: 8 November 2011 / Accepted: 11 April 2013 / Published online: 8 June 2013
© The Author(s) 2013

Abstract In *sequential event prediction*, we are given a “sequence database” of past event sequences to learn from, and we aim to predict the next event within a current event sequence. We focus on applications where the *set* of the past events has predictive power and not the specific *order* of those past events. Such applications arise in recommender systems, equipment maintenance, medical informatics, and in other domains. Our formalization of sequential event prediction draws on ideas from supervised ranking. We show how specific choices within this approach lead to different sequential event prediction problems and algorithms. In recommender system applications, the observed sequence of events depends on user choices, which may be influenced by the recommendations, which are themselves tailored to the user’s choices. This leads to sequential event prediction algorithms involving a non-convex optimization problem. We apply our approach to an online grocery store recommender system, email recipient recommendation, and a novel application in the health event prediction domain.

Keywords Sequential event prediction · Supervised ranking · Recommender systems

1 Introduction

Sequential event prediction refers to a wide class of problems in which a set of initially hidden events are sequentially revealed. The goal is to use the set of revealed events, but

Editors: Eyke Hüllermeier and Johannes Fürnkranz.

B. Letham (✉)

Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, USA
e-mail: bletham@mit.edu

C. Rudin

MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA, USA
e-mail: rudin@mit.edu

D. Madigan

Department of Statistics, Columbia University, New York, NY, USA
e-mail: madigan@stat.columbia.edu