

Mining and Forecasting of Big Time-series Data

Yasushi Sakurai
Kumamoto University
yasushi@cs.kumamoto-u.ac.jp

Yasuko Matsubara
Kumamoto University
yasuko@cs.kumamoto-u.ac.jp

Christos Faloutsos
Carnegie Mellon University
christos@cs.cmu.edu

ABSTRACT

Given a large collection of time series, such as web-click logs, electric medical records and motion capture sensors, how can we efficiently and effectively find typical patterns? How can we statistically summarize all the sequences, and achieve a meaningful segmentation? What are the major tools for forecasting and outlier detection? Time-series data analysis is becoming of increasingly high importance, thanks to the decreasing cost of hardware and the increasing on-line processing capability.

The objective of this tutorial is to provide a concise and intuitive overview of the most important tools that can help us find patterns in large-scale time-series sequences. We review the state of the art in four related fields: (1) similarity search and pattern discovery, (2) linear modeling and summarization, (3) non-linear modeling and forecasting, and (4) the extension of time-series mining and tensor analysis. The emphasis of the tutorial is to provide the intuition behind these powerful tools, which is usually lost in the technical literature, as well as to introduce case studies that illustrate their practical use.

Categories and Subject Descriptors

H.2.8 [Database management]: Database applications—*Data mining*

General Terms

Algorithms, Experimentation, Theory

Keywords

Time-series; Tensors; Pattern discovery; Forecasting;

1. INTRODUCTION

Time series data occur naturally in countless domains including medical analysis [17, 33], financial analysis [56], online text [14], sensor network monitoring [22, 41] and social activity mining [26, 31, 32]. The increasing volume of online, time-stamped activity represents a vital new opportunity for data scientists and analysts to measure the collective behavior of social, economic, and other important evolutions [21, 13]. In many applications the logging rate

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGMOD'15, May 31–June 4, 2015, Melbourne, Victoria, Australia.
Copyright © 2015 ACM 978-1-4503-2758-9/15/05 ...\$15.00.
<http://dx.doi.org/10.1145/2723372.2731081>.

has increased greatly with the progress made on hardware and storage technology. In such a situation, the most fundamental requirements are the efficient and effective mining of “big” time-series sequences.

Time-series data analysis is a well-known topic that has attracted huge interest in various research communities (e.g., theory, databases, data mining, networking) for a few decades. However, recent studies has revealed some new directions for research on time-series analysis, which include:

- **Large-scale tensor analysis:** Many time-series data can be modeled as tensors, and tensor analysis is an important data mining tool that has various applications including sensor streams, hyperlinks, medical records and social networks over time.
- **Non-linear modeling:** Non-linear models are widely used in a variety of areas, such as epidemiology, biology, physics and economics. In the database and mining fields, analyses of social media and online user behavior have attracted considerable interest, and recent studies have focused on non-linear time-series analysis to understand the dynamics of social networks.
- **Automatic mining:** We also emphasize the importance of fully-automatic mining. Most of the existing time-series tools require parameter tuning, and they are very sensitive to these parameters. In fact, faced with “big data”, fully automatic mining is even more important: otherwise, the user would have to try several parameter tuning steps, each of which would take too long (e.g., hours, or days).

This tutorial provides a concise and intuitive overview of the most important tools that we can use to help us understand and find patterns in large-scale time evolving sequences. We will provide a comprehensive overview and the above new directions for time-series analysis, and deal specifically with the following key topics: (1) similarity search and pattern discovery, (2) linear modeling and summarization, (3) non-linear modeling and forecasting, and (4) the extension of time-series mining and tensor analysis.

Who should attend. The target audience is data mining and data management researchers who wish to get up to speed with the major tools used in time sequence analysis. Also, practitioners who want a concise, intuitive overview of the state of the art.

Prerequisites. None. The emphasis is on the intuition behind all these mathematical tools.

Related tutorials. Related tutorials have been presented, e.g., (a) Indexing and Mining Streams, by Christos Faloutsos, SIGMOD 2004, (b) Indexing and Mining Time Sequences, by Christos Faloutsos and Lei Li, SIGKDD 2010, and (c) Mining Shape and Time Series Databases with Symbolic Representations, by Eamonn Keogh,

SIGKDD 2007. However, these tutorials look at different approaches and tools. Moreover, they do not focus on large-scale evolving tensors, or non-linear dynamical systems.

2. CONTENT AND OUTLINE

Our tutorial is structured as follows:

1. Similarity search and pattern discovery (60 minutes)
 - (a) Why we need similarity search
 - i. Indexing
 - ii. Fast searching
 - (b) Distance function and similarity measure
 - i. Euclidean distance
 - ii. Dynamic time warping
 - iii. Correlation
 - (c) Feature extraction
 - i. Discrete Fourier transform, wavelets
 - ii. Singular value decomposition, independent component analysis
 - iii. Multi-dimensional scaling
 - (d) Streaming pattern discovery
 - i. Component analysis
 - ii. Correlation monitoring
 - iii. Time warping over streams
2. Linear modeling and summarization (30 minutes)
 - (a) Linear modeling and forecasting
 - i. Main idea behind linear forecasting
 - ii. AR methodology and multivariate regression
 - iii. Recursive least square
 - iv. Streaming algorithms
 - (b) Sequence summarization
 - i. Linear dynamic systems
 - ii. Probabilistic models
 - iii. Automatic mining of co-evolving sequences
3. Non-linear modeling and forecasting (30 minutes)
 - (a) Non-linear forecasting
 - i. Lag-plots
 - ii. Fractal dimension and power-law
 - (b) Non-linear dynamical systems
 - i. Main idea: non-linear equations
 - ii. Non-linear epidemic models
 - iii. Gray-box non-linear mining
 - iv. Non-linear dynamical systems for online activities
 - v. Information diffusion in social networks
4. Extension of time-series mining - tensor analysis (60 minutes)
 - (a) Tensor decomposition
 - i. Basic approaches
 - ii. Decompositions of higher-order tensors
 - (b) Mining of complex time-stamped tensors
 - i. Complex time-stamped events and big sparse tensors
 - ii. Feature extraction from sparse tensors

- iii. Forecasting of complex time-stamped events
- (c) New directions of tensor analysis
 - i. Non-linear modeling for tensors
 - ii. Automatic non-linear analysis

1. Similarity search and pattern discovery. In this first part of the tutorial, we explain the most common and fundamental tools of time series data mining. More specifically, we demonstrate some traditional approaches applied to time series data mining including similarity search e.g., euclidean distance, dynamic time warping and correlation [43, 46, 42, 34]; feature extraction including singular value decomposition (SVD), independent component analysis (ICA) [38, 44], segmentation [17]; multi-dimensional scaling [10]. We also answer the following important question: What can we do in the highly likely case that the users need real-time, on-line data processing while in practice they require high accuracy? Specifically, we introduce several mining algorithms for online data streams, including component analysis [40, 41], correlation monitoring [56, 45] and time warping over streams [45, 49].

2. Linear modeling and summarization. In this part, we explain the main ideas behind linear modeling and forecasting. For example, auto regression and moving averaging models have been studied for many years in statistics and finance [4], and have been applied to time-series data mining [6, 16, 23]. We introduce AR methodology and several important tools including MUSCLES [55] and AWSOM [39]. We also introduce linear dynamical systems (LDS), Kalman filters (KF) and their variants [15, 25, 24, 48]. As regards probabilistic time-series analysis, hidden Markov models (HMMs) have been used in various applications including speech recognition [52], and sensor monitoring [22, 32, 12]. For example, Wang et al. [51] presented a pattern-based hidden Markov model (pHMM) for time-series segmentation and clustering, while [28] developed AutoPlait, a fully-automatic mining algorithm for co-evolving time sequences. Given a large collection of co-evolving multiple sequences, which contains an unknown number of patterns of different durations, AutoPlait automatically identifies all distinct patterns and spots the time position of each variation.

3. Non-linear modeling and forecasting. In this part, we introduce several advanced techniques, and focus specifically on non-linear time series analysis. We start by explaining non-linear forecasting e.g., lag-plots [7], which is based on nearest-neighbor search. We also explain some fundamental concepts such as fractal dimension and power law [36, 2, 27]. We then review the most common non-linear equations, including the logistic function (LF) [5], the susceptible-infected (SI) model [1], the independent cascade (IC) model [9], the so-called “Bass” model [3], the Lotka-Volterra (LV) model [35] and other non-linear equations [37]. We explain the importance of non-linear equations and the concept of gray-box non-linear mining. In this part, we also review recent work on understanding the non-linear time evolution of online user activities. Analyses of epidemics, blogs, social media, propagation and the cascades they create have attracted much interest. We answer several important topics such as how popularity of “memes” changes over time [21]; how to find temporal patterns in information diffusion process through online media, e.g., blogs, hashtags [54, 53], and YouTube [8, 11]; how to describe rising and falling patterns of information propagation (e.g., memes, hashtags and keyword search volume) using non-linear dynamical systems [31, 29].

4. Extension of time-series mining - tensor analysis. The goal in this part is to present large-scale studies of complex time-stamped events and big sparse tensors. We first introduce some basic ap-

proaches including Tucker, PARAFAC, and higher-order SVD (HOSVD) [18, 19, 47, 20]. Complex time-stamped events can be represented as a tensor with several dimensions. For example, given a set of time-stamped event entries of the form { object, actor, timestamp} (e.g., web-clicks: {URL, userID, timestamp}), we can treat them as a 3rd order tensor. Here, one subtle, but important issue is that the complex time-stamped tensor is *very sparse*, which derails all typical time-series mining and forecasting tools. We introduce a scalable algorithm, TriMine [30] to deal with this issue. TriMine has the ability to find meaningful patterns in complex time-stamped tensors, and forecast future events, e.g., estimate the number of clicks from user “Smith” to URL “CNN.com” for the next 30 days.

Finally, we show new directions for tensor analysis, namely, automatic and non-linear analysis for big time-series tensors. Specifically, we introduce a unifying analytical model, FUNNEL [33], for mining and forecasting large-scale epidemiological data (e.g., the Project Tycho [50]) as well as an efficient fitting algorithm, which solves the problem. We also discuss the importance of fully-automatic mining for time-series tensor analysis. There are many fascinating and useful tools for time-series analysis. However, most existing methods require parameter settings and fine tuning, such as the number of coefficients, and the reconstruction error thresholds, and they are very sensitive to these parameters. The ideal method should look for arbitrary patterns and require no initial human intervention to guide it.

3. PRESENTERS - BIOS

Yasushi Sakurai is a Professor at Kumamoto University. He obtained his B.E. degree from Doshisha University in 1991, and his M.E. and Ph.D. degrees from Nara Institute of Science and Technology in 1996 and 1999, respectively. In 1998, he joined NTT Laboratories, and became a Senior Research Scientist in 2005. He was a Visiting Researcher at Carnegie Mellon University during 2004–2005. He received two KDD best paper awards in 2008 and 2010. His research interests include time-series analysis, web mining, and sensor data processing.

Yasuko Matsubara is an Assistant Professor in the Department of Computer Science and Electrical Engineering at Kumamoto University, Japan. She obtained her BS and MS degrees from Ochanomizu University in 2007 and 2009 respectively, and her PhD from Kyoto University in 2012. She was a Visiting Researcher at Carnegie Mellon University during 2011–2012 and 2013–2014. Her research interests include time-series data mining and non-linear dynamic systems.

Christos Faloutsos is a Professor at Carnegie Mellon University. He has received the Presidential Young Investigator Award by the National Science Foundation (1989), the Research Contributions Award in ICDM 2006, the SIGKDD Innovations Award (2010), twenty “best paper” awards (including two “test of time” awards), and four teaching awards. Five of his advisees have attracted KDD or SCS dissertation awards. He is an ACM Fellow, he has served as a member of the executive committee of SIGKDD; he has published over 300 refereed articles, 17 book chapters and two monographs. He holds eight patents and he has given over 35 tutorials and over 15 invited distinguished lectures. His research interests include data mining for graphs and streams, fractals, database performance, and indexing for multimedia and bio-informatics data.

Acknowledgement

This work was supported by JSPS KAKENHI Grant-in-Aid for Scientific Research Number 26730060, 24500138, 26280112, 25-7946.

This material is based upon work supported by the National Science Foundation under Grants No. CNS-1314632 and IIS-1408924; and by the Army Research Laboratory (ARL) under Cooperative Agreement Number W911NF-09-2-0053; and by a Google Focused Research Award. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, ARL, or other funding parties. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

4. REFERENCES

- [1] R. M. Anderson and R. M. May. *Infectious Diseases of Humans Dynamics and Control*. Oxford University Press, 1992.
- [2] A. L. Barabasi. The origin of bursts and heavy tails in human dynamics. *Nature*, 435, 2005.
- [3] F. M. Bass. A new product growth for model consumer durables. *Management Science*, 15(5):215–227, 1969.
- [4] G. E. Box, G. M. Jenkins, and G. C. Reinsel. *Time Series Analysis: Forecasting and Control*. Forecasting and Control Series. Prentice Hall, Englewood Cliffs, NJ, 3rd edition, 1994.
- [5] F. Brauer and C. Castillo-Chavez. *Mathematical models in population biology and epidemiology*, volume 40. Springer Verlag, New York, 2001.
- [6] P. J. Brockwell and R. A. Davis. *Time Series: Theory and Methods*. Springer-Verlag New York, Inc., New York, NY, USA, 1987.
- [7] D. Chakrabarti and C. Faloutsos. F4: large-scale automated forecasting using fractals. In *CIKM*, pages 2–9, 2002.
- [8] R. Crane and D. Sornette. Robust dynamic classes revealed by measuring the response function of a social system. In *PNAS*, 2008.
- [9] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.
- [10] C. Faloutsos and K.-I. Lin. Fastmap: A fast algorithm for indexing, data-mining and visualization of traditional and multimedia datasets. In *SIGMOD*, pages 163–174, 1995.
- [11] F. Figueiredo, J. M. Almeida, Y. Matsubara, B. Ribeiro, and C. Faloutsos. Revisit behavior in social media: The phoenix-r model and discoveries. In *PKDD*, pages 386–401, 2014.
- [12] Y. Fujiwara, Y. Sakurai, and M. Yamamuro. Spiral: efficient and exact model identification for hidden markov models. In *KDD*, pages 247–255, 2008.
- [13] J. Ginsberg, M. Mohebbi, R. Patel, L. Brammer, M. Smolinski, and L. Brilliant. Detecting influenza epidemics using search engine query data. *Nature*, 457:1012–1014, 2009.
- [14] M. D. Hoffman, D. M. Blei, and F. R. Bach. Online learning for latent dirichlet allocation. In *NIPS*, pages 856–864, 2010.
- [15] A. Jain, E. Y. Chang, and Y.-F. Wang. Adaptive stream resource management using kalman filters. In *SIGMOD*, pages 11–22, 2004.
- [16] K. Kalpakis, D. Gada, and V. Puttagunta. Distance measures for effective clustering of arima time-series. In *ICDM 2001: Proceeding of 2001 IEEE International Conference on Data Mining*, pages 273–280, 2001.
- [17] E. J. Keogh, S. Chu, D. Hart, and M. J. Pazzani. An online algorithm for segmenting time series. In *ICDM*, pages 289–296, 2001.

- [18] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.
- [19] T. G. Kolda, B. W. Bader, and J. P. Kenny. Higher-order web link analysis using multilinear algebra. In *ICDM*, pages 242–249, 2005.
- [20] L. D. Lathauwer, B. D. Moor, and J. Vandewalle. A multilinear singular value decomposition. *SIAM J. Matrix Anal. Appl.*, 21(4):1253–1278, 2000.
- [21] J. Leskovec, L. Backstrom, and J. M. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *KDD*, pages 497–506, 2009.
- [22] J. Letchner, C. Ré, M. Balazinska, and M. Philipose. Access methods for markovian streams. In *ICDE*, pages 246–257, 2009.
- [23] L. Li, C.-J. M. Liang, J. Liu, S. Nath, A. Terzis, and C. Faloutsos. Thermocast: A cyber-physical forecasting model for data centers. In *KDD*, 2011.
- [24] L. Li, J. McCann, N. Pollard, and C. Faloutsos. Dynammo: Mining and summarization of coevolving sequences with missing values. In *KDD*, 2009.
- [25] L. Li, B. A. Prakash, and C. Faloutsos. Parsimonious linear fingerprinting for time series. *PVLDB*, 3(1):385–396, 2010.
- [26] M. Mathioudakis, N. Koudas, and P. Marbach. Early online identification of attention gathering items in social media. In *WSDM*, pages 301–310, 2010.
- [27] Y. Matsubara, L. Li, E. E. Papalexakis, D. Lo, Y. Sakurai, and C. Faloutsos. F-trail: Finding patterns in taxi trajectories. In *PAKDD*, pages 86–98, 2013.
- [28] Y. Matsubara, Y. Sakurai, and C. Faloutsos. Autoplait: automatic mining of co-evolving time sequences. In *SIGMOD*, pages 193–204, 2014.
- [29] Y. Matsubara, Y. Sakurai, and C. Faloutsos. The web as a jungle: Non-linear dynamical systems for co-evolving online activities. In *WWW*, 2015.
- [30] Y. Matsubara, Y. Sakurai, C. Faloutsos, T. Iwata, and M. Yoshikawa. Fast mining and forecasting of complex time-stamped events. In *KDD*, pages 271–279, 2012.
- [31] Y. Matsubara, Y. Sakurai, B. A. Prakash, L. Li, and C. Faloutsos. Rise and fall patterns of information diffusion: model and implications. In *KDD*, pages 6–14, 2012.
- [32] Y. Matsubara, Y. Sakurai, N. Ueda, and M. Yoshikawa. Fast and exact monitoring of co-evolving data streams. In *ICDM*, 2014.
- [33] Y. Matsubara, Y. Sakurai, W. G. van Panhuis, and C. Faloutsos. FUNNEL: automatic mining of spatially coevolving epidemics. In *KDD*, pages 105–114, 2014.
- [34] Y. Matsubara, Y. Sakurai, and M. Yoshikawa. Scalable algorithms for distribution search. In *ICDM*, pages 347–356, 2009.
- [35] R. M. May. Qualitative stability in model ecosystems. *Ecology*, 54(3):638–641, 1973.
- [36] M. McGlohon, J. Leskovec, C. Faloutsos, M. Hurst, and N. Glance. Finding patterns in blog shapes and blog evolution. In *International Conference on Weblogs and Social Media*, Boulder, Colo., March 2007.
- [37] M. Nowak. *Evolutionary Dynamics*. Harvard University Press, 2006.
- [38] J.-Y. Pan, H. Kitagawa, C. Faloutsos, and M. Hamamoto. Autosplit: Fast and scalable discovery of hidden variables in stream and multimedia databases. In *PAKDD*, May 26-28 2004.
- [39] S. Papadimitriou, A. Brockwell, and C. Faloutsos. Adaptive, hands-off stream mining. In *VLDB*, pages 560–571, 2003.
- [40] S. Papadimitriou, J. Sun, and C. Faloutsos. Streaming pattern discovery in multiple time-series. In *VLDB*, pages 697–708, 2005.
- [41] S. Papadimitriou and P. S. Yu. Optimal multi-scale patterns in time series streams. In *SIGMOD*, pages 647–658, 2006.
- [42] T. Rakthanmanon, B. J. L. Campana, A. Mueen, G. E. A. P. A. Batista, M. B. Westover, Q. Zhu, J. Zakaria, and E. J. Keogh. Searching and mining trillions of time series subsequences under dynamic time warping. In *KDD*, pages 262–270, 2012.
- [43] Y. Sakurai, C. Faloutsos, and M. Yamamuro. Stream monitoring under the time warping distance. In *Proceedings of the 23rd International Conference on Data Engineering, ICDE 2007, April 15-20, 2007, The Marmara Hotel, Istanbul, Turkey*, pages 1046–1055, 2007.
- [44] Y. Sakurai, L. Li, Y. Matsubara, and C. Faloutsos. Windmine: Fast and effective mining of web-click sequences. In *SDM*, pages 759–770, 2011.
- [45] Y. Sakurai, S. Papadimitriou, and C. Faloutsos. Braid: Stream mining through group lag correlations. In *SIGMOD*, pages 599–610, 2005.
- [46] Y. Sakurai, M. Yoshikawa, and C. Faloutsos. Ftw: Fast similarity search under the time warping distance. In *PODS*, pages 326–337, Baltimore, Maryland, June 2005.
- [47] J. Sun, D. Tao, and C. Faloutsos. Beyond streams and graphs: dynamic tensor analysis. In *KDD*, pages 374–383, 2006.
- [48] Y. Tao, C. Faloutsos, D. Papadias, and B. Liu. Prediction and indexing of moving objects with unknown motion patterns. In *SIGMOD*, pages 611–622, 2004.
- [49] M. Toyoda, Y. Sakurai, and Y. Ishikawa. Pattern discovery in data streams under the time warping distance. *VLDB J.*, 22(3):295–318, 2013.
- [50] W. G. van Panhuis, J. Grefenstette, S. Y. Jung, N. S. Chok, A. Cross, H. Eng, B. Y. Lee, V. Zadorozhny, S. Brown, D. Cummings, and D. S. Burke. Contagious diseases in the united states from 1888 to the present. *NEJM*, 369(22):2152–2158, 2013.
- [51] P. Wang, H. Wang, and W. Wang. Finding semantics in time series. In *SIGMOD Conference*, pages 385–396, 2011.
- [52] J. G. Wilpon, L. R. Rabiner, C. H. Lee, and E. R. Goldman. Automatic recognition of keywords in unconstrained speech using hidden Markov models. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(11):1870–1878, 1990.
- [53] J. Yang and J. Leskovec. Modeling information diffusion in implicit networks. In *ICDM*, pages 599–608, 2010.
- [54] J. Yang and J. Leskovec. Patterns of temporal variation in online media. In *WSDM*, pages 177–186, 2011.
- [55] B.-K. Yi, N. Sidiropoulos, T. Johnson, H. Jagadish, C. Faloutsos, and A. Biliris. Online data mining for co-evolving time sequences. *ICDE*, pages 13–22, 2000.
- [56] Y. Zhu and D. Shasha. Statstream: Statistical monitoring of thousands of data streams in real time. In *VLDB*, pages 358–369, 2002.