Machine Learning for Time Series

List of mini-projects

Laurent Oudre
laurent.oudre@ens-paris-saclay.fr

Master MVA
2023-2024

- Projects can be done in groups of two, but no more than that.
- Students are allowed to propose additional project (please ask for approval beforehand)
- The mini project consists in reading the paper, implement it in Python and launch experiments on real time series
- Report (PDF file, ≈ 5 pages) + source code (Jupyter Notebook) should be submitted to laurent.oudre@ens-paris-saclay.fr and charles@doffy.net
- **Session 1**
  - Deadline for report: December 18th (23:59)
  - Oral presentations: December 20th and 22th (precise times TBA)
- **Session 2**
  - Deadline for report: January 9th (23:59)
  - Oral presentations: January, 11th and 12th (precise times TBA)
- The oral presentation will have a duration of 10 min
- Final grade is 25% report, 25% source code, 25% oral presentation and 25% tutorial

Session 1: Pattern Recognition and Detection

Project 1.1 Cuturi, M., & Blondel, M. (2017, July). Soft-dtw: a differentiable loss function for time-series. In International conference on machine learning (pp. 894-903). PMLR.
A differentiable DTW linked to optimal transport

Project 1.2 Zhao, J., & Itti, L. (2018). shapeDTW: Shape dynamic time warping. Pattern Recognition, 74, 171-184.
A variant of the DTW that takes local behavior into account

Project 1.3 Le Guen, V., & Thome, N. (2019). Shape and time distortion loss for training deep time series forecasting models. Advances in neural information processing systems, 32.
How to construct a DTW-based loss for deep learning

Project 1.4 Rakthanmanon, T., & Keogh, E. (2013, May). Fast shapelets: A scalable algorithm for discovering time series shapelets. In proceedings of the 2013 SIAM International Conference on Data Mining (pp. 668-676). Society for Industrial and Applied Mathematics.
How to automatically extract shapes from time series by using symbolic signal representation.

Project 1.5 Linardi, M., Zhu, Y., Palpanas, T., & Keogh, E. (2020). Matrix profile goes MAD: variable-length motif and discord discovery in data series. Data Mining and Knowledge Discovery
How to extend the matrix profile approach to variable lengths motifs.

Project 1.6 Yeh, C. C. M., Kavantzas, N., & Keogh, E. (2017, November). Matrix profile vi: meaningful multidimensional motif discovery. In 2017 IEEE international conference on data mining (ICDM) (pp. 565-574). IEEE.
How to extend the matrix profile approach to multivariate time series

Project 1.7 Alaee, S., Kamgar, K., & Keogh, E. (2020). Matrix Profile XXII: Exact Discovery of Time Series Motifs under DTW. arXiv preprint arXiv:2009.07907.
How to find patterns using the DTW.

Project 1.8 Hills, J., Lines, J., Baranauskas, E., Mapp, J., & Bagnall, A. (2014). Classification of time series by shapelet transformation. Data Mining and Knowledge Discovery, 28(4), 851-881.
How to use patterns for time series classification

Project 1.9 Pilastre, B., Silva, G., Boussouf, L., d'Escrivan, S., Rodríguez, P., & Tourneret, J. Y. (2020, May). Anomaly Detection in Mixed Time-Series Using A Convolutional Sparse Representation With Application To Spacecraft Health Monitoring. In ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 3242-3246). IEEE.
How to use convolutional dictionary learning for detecting anomaly

Project 1.10 La Tour, T. D., Moreau, T., Jas, M., & Gramfort, A. (2018). Multivariate convolutional sparse coding for electromagnetic brain signals. In Advances in Neural Information Processing Systems (pp. 3292-3302).
How to use convolutional dictionary learning to study the brain

Session 2: Feature Extraction and Selection

Project 2.1 Schäfer, P. (2015). The BOSS is concerned with time series classification in the presence of noise. Data Mining and Knowledge Discovery, 29(6), 1505-1530.
How to use local symbolic features to classify time series

Project 2.2 Elsworth, S., & Güttel, S. (2020). Time series forecasting using LSTM networks: A symbolic approach. arXiv preprint arXiv:2003.05672.
How to use symbolic representations for prediction

Project 2.3 Le Nguyen, T., Gsponer, S., Ilie, I., O'Reilly, M., & Ifrim, G. (2019). Interpretable time series classification using linear models and multi-resolution multi-domain symbolic representations. Data mining and knowledge discovery, 33(4), 1183-1222.
How to use multiscale symbolic representations to classify time series

Project 2.4 Richman, J. S., & Moorman, J. R. (2000). Physiological time-series analysis using approximate entropy and sample entropy. American Journal of Physiology-Heart and Circulatory Physiology, 278(6), H2039-H2049.
How to extract information theory based features to study physiological time series.

Project 2.5 Gidea, M., & Katz, Y. (2018). Topological data analysis of financial time series: Landscapes of crashes. Physica A: Statistical Mechanics and its Applications, 491, 820-834.
How to extract topological features to study financial data.

Project 2.6 Madiraju, N. S., Sadat, S. M., Fisher, D., & Karimabadi, H. (2018). Deep temporal clustering: Fully unsupervised learning of time-domain features. arXiv preprint arXiv:1802.01059.
How to use deep learning to extract time-domain features

Project 2.7 He, X., Cai, D., & Niyogi, P. (2006). Laplacian score for feature selection. In Advances in neural information processing systems (pp. 507-514).
How to apply unsupervised feature selection.

Project 2.8 Li, J., Cheng, K., Wang, S., Morstatter, F., Trevino, R. P., Tang, J., & Liu, H. (2017). Feature selection: A data perspective. ACM Computing Surveys (CSUR), 50(6), 1-45.
How to apply a large number of feature selection methods (multitude of topics in this article including information theory!)

Project 2.9 Barandas, M., Folgado, D., Fernandes, L., Santos, S., Abreu, M., Bota, P., ... & Gamboa, H. (2020). TSFEL: Time series feature extraction library. SoftwareX, 11, 100456.

Project 2.10 Längkvist, M., Karlsson, L., & Loutfi, A. (2014). A review of unsupervised feature learning and deep learning for time-series modeling. Pattern Recognition Letters, 42, 11-24.

Session 3: Models and Representation Learning

Project 3.1 Mairal, J., Bach, F., Ponce, J., & Sapiro, G. (2009, June). Online dictionary learning for sparse coding. In Proceedings of the 26th annual international conference on machine learning (pp. 689-696).
How to learn a dictionary from streaming data

Project 3.2 Tzagkarakis, G., Caicedo-Llano, J., & Dionysopoulos, T. (2015). Sparse modeling of volatile financial time series via low-dimensional patterns over learned dictionaries. Algorithmic Finance, 4(3-4), 139-158.
How to model financial data with sparse dictionary representations.

Project 3.3 Ho, S. L., Xie, M., & Goh, T. N. (2002). A comparative study of neural network and Box-Jenkins ARIMA modeling in time series prediction. Computers & Industrial Engineering, 42(2-4), 371-375.
How to compare deep learning and standard Box-Jenkins models for prediction

Project 3.4 Yazdi, S. V., & Douzal-Chouakria, A. (2018). Time warp invariant kSVD: Sparse coding and dictionary learning for time series under time warp. Pattern Recognition Letters, 112, 1-8.
How to mix Dynamic Time Warping and dictionary learning

Project 3.5 Lyu, H., Strohmeier, C., Menz, G., & Needell, D. (2020). COVID-19 time-series prediction by joint dictionary learning and online NMF. arXiv preprint arXiv:2004.09112.
How to use matrix factorization and dictionary learning to perform prediction

Project 3.6 Zhang, W., Wang, Z., Yuan, J., & Hao, S. (2020). Discriminative Dictionary Learning for Time Series Classification. IEEE Access, 8, 185032-185044.
How to combine symbolic representation and dictionary learning for time series classification

Project 3.7 Varoquaux, G., Gramfort, A., Pedregosa, F., Michel, V., & Thirion, B. (2011, July). Multi-subject dictionary learning to segment an atlas of brain spontaneous activity. In Biennial International Conference on information processing in medical imaging (pp. 562-573). Springer, Berlin, Heidelberg.
How to use dictionary learning for segmentation

Project 3.8 Sawada, H., Ono, N., Kameoka, H., Kitamura, D., & Saruwatari, H. (2019). A review of blind source separation methods: two converging routes to ILRMA originating from ICA and NMF. APSIPA Transactions on Signal and Information Processing, 8.
How to use dictionary learning for source separation

Session 4: Data Enhancement and Preprocessings

Project 4.1 Flandrin, P., Goncalves, P., & Rilling, G. (2004, September). Detrending and denoising with empirical mode decompositions. In 2004 12th European Signal Processing Conference (pp. 1581-1584). IEEE.
How to use EMD for denoising and detrending.

Project 4.2 Rhif, M., Ben Abbes, A., Farah, I. R., Martínez, B., & Sang, Y. (2019). Wavelet transform application for/in non-stationary time-series analysis: a review. Applied Sciences, 9(7), 1345.
How to use wavelets to work on non-stationary time series.

Project 4.3 Bayer, F. M., Kozakevicius, A. J., & Cintra, R. J. (2019). An iterative wavelet threshold for signal denoising. Signal Processing, 162, 10-20.
How to use adaptive wavelet thresholding for denoising

Project 4.4 Moussallam, M., Gramfort, A., Daudet, L., & Richard, G. (2014). Blind denoising with random greedy pursuits. IEEE Signal Processing Letters, 21(11), 1341-1345.
How to use statistical considerations to set the parameters in greedy denoising approaches

Project 4.5 Aharon, M., Elad, M., & Bruckstein, A. (2006). K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. IEEE Transactions on signal processing, 54(11), 4311-4322.
How to learn an overcomplete dictionary with K-SVD

Project 4.6 de Cheveigné, A., & Arzounian, D. (2018). Robust detrending, rereferencing, outlier detection, and inpainting for multichannel data. Neuroimage, 172, 903-912.
How to combine detrending, outlier detection and removal for multichannel data

Project 4.7 Hassani, H., & Mahmoudvand, R. (2013). Multivariate singular spectrum analysis: A general view and new vector forecasting approach. International Journal of Energy and Statistics, 1(01), 55-83.
How to use SSA for forecasting time series

Project 4.8 Adler, A., Emiya, V., Jafari, M. G., Elad, M., Gribonval, R., & Plumbley, M. D. (2011). Audio inpainting. IEEE Transactions on Audio, Speech, and Language Processing, 20(3), 922-932..
How to use sparse representation to perform audio inpainting

Session 5: Change-Point and Anomaly Detection

Project 5.1 How to contribute to the `ruptures` package (see with C. Truong) https://centre-borelli.github.io/ruptures-docs/

Project 5.2 Truong, C., Oudre, L., & Vayatis, N. (2017). Penalty learning for changepoint detection. In 2017 25th European Signal Processing Conference (EUSIPCO) (pp. 1569-1573). IEEE.
How to learn the penalty for change point detection

Project 5.3 Fearnhead, P., & Rigaill, G. (2019). Changepoint detection in the presence of outliers. Journal of the American Statistical Association, 114(525), 169-183.
How to detect change-points in presence of outliers

Project 5.4 Kim, H., Kim, B., Chung, D., Yoon, J., & Ko, S.-K. (2022). SoccerCPD: formation and role change-point detection in soccer matches using spatiotemporal tracking data. Proceedings of the ACM Conference on Knowledge Discovery and Data Mining (SIGKDD), 3146–3156.
How to use change point detection to study soccer games

Project 5.5 Fearnhead, P., Maidstone, R., & Letchford, A. (2019). Detecting Changes in Slope With an L0 Penalty. Journal of Computational and Graphical Statistics, 28(2), 265–275.
How to introduce a sparsity penalty into change-point detection problems

Project 5.6 Runge, V., Hocking, T. D., Romano, G., Afghah, F., Fearnhead, P., & Rigaill, G. (2020). gfpop: an R Package for Univariate Graph-Constrained Change-Point Detection. ArXiv E-Prints ArXiv:2002.03646.
How to introduce graph constraints into change-point detection problems

Project 5.7 Jewell, S. W., Hocking, T. D., Fearnhead, P., & Witten, D. M. (2020). Fast nonconvex deconvolution of calcium imaging data. Biostatistics, 21(4), 709–726.
How to apply change-point detection to biology

Project 5.8 Chin, S. C., Ray, A., & Rajagopalan, V. (2005). Symbolic time series analysis for anomaly detection: A comparative evaluation. Signal Processing, 85(9), 1859-1868.
How to use symbolic representation for detecting anomalies

Project 5.9 Chandola, V., Banerjee, A., & Kumar, V. (2010). Anomaly detection for discrete sequences: A survey. IEEE transactions on knowledge and data engineering, 24(5), 823-839.
How to detect anomalies in discrete time series

Project 5.10 Boniol, P., Linardi, M., Roncallo, F., & Palpanas, T. (2020, April). Automated Anomaly Detection in Large Sequences. In 2020 IEEE 36th International Conference on Data Engineering (ICDE) (pp. 1834-1837). IEEE.
How to use detect anomalies in large time series

Project 5.11 Nakamura, T., Imamura, M., Mercer, R., & Keogh, E. MERLIN: Parameter-Free Discovery of Arbitrary Length Anomalies in Massive Time Series Archives. In Proc. 20th IEEE Intl. Conf. Data Mining.
How to detect anomalies with different lengths

Project 5.12 Tatbul, N., Lee, T. J., Zdonik, S., Alam, M., & Gottschlich, J. (2018). Precision and recall for time series. arXiv preprint arXiv:1803.03639.
How to assess event detection techniques

Project 5.13 Schmidl, S., Wenig, P., & Papenbrock, T. (2022). Anomaly detection in time series: a comprehensive evaluation. Proceedings of the VLDB Endowment, 15(9), 1779-1797.
Wonderful article with tons of references, implementations, etc...

Project 5.14 Boniol, P., Paparrizos, J., Kang, Y., Palpanas, T., Tsay, R. S., Elmore, A. J., & Franklin, M. J. (2022). Theseus: navigating the labyrinth of time-series anomaly detection. Proceedings of the VLDB Endowment, 15(12), 3702-3705.

Project 5.15 Wenig, P., Schmidl, S., & Papenbrock, T. (2022). TimeEval: a benchmarking toolkit for time series anomaly detection algorithms. Proceedings of the VLDB Endowment, 15(12), 3678-3681.

Session 6: Multivariate Time Series

Project 6.1 Wang, D., Zheng, Y., Lian, H., & Li, G. (2020). High-dimensional vector autoregressive time series modeling via tensor decomposition. Journal of the American Statistical Association, 1-42.
How to apply VAR models to high dimensional data

Project 6.2 Li, H. (2019). Multivariate time series clustering based on common principal component analysis. Neurocomputing, 349, 239-247.
How to use PCA to perform clustering on multivariate time series

Project 6.3 Chen, X., & Sun, L. (2020). Low-rank autoregressive tensor completion for multivariate time series forecasting. arXiv preprint arXiv:2006.10436.
How to use tensor structure to forecast multivariate time series

Project 6.4 Barthélemy, Q., Gouy-Pailler, C., Isaac, Y., Souloumiac, A., Larue, A., & Mars, J. I. (2013). Multivariate temporal dictionary learning for EEG. Journal of neuroscience methods, 215(1), 19-28.
How to apply multivariate dictionary learning for EEG data

Project 6.5 Cotter, S. F., Rao, B. D., Engan, K., & Kreutz-Delgado, K. (2005). Sparse solutions to linear inverse problems with multiple measurement vectors. IEEE Transactions on Signal Processing, 53(7), 2477-2488.
How use multivariate sparse coding for solving inverse problems

Project 6.6 Dong, X., Thanou, D., Rabbat, M., & Frossard, P. (2019). Learning graphs from data: A signal representation perspective. IEEE Signal Processing Magazine, 36(3), 44-63..
How to learn a graph from smooth graph signals

Project 6.7 Kumar, S., Ying, J., de Miranda Cardoso, J. V., & Palomar, D. P. (2020). A Unified Framework for Structured Graph Learning via Spectral Constraints. Journal of Machine Learning Research, 21(22), 1-60.
How to learn a graph with spectral constraints