
Bibliography

- [1] N. Abe, B. Zadrozny, and J. Langford. Outlier Detection by Active Learning. *ACM KDD Conference*, 2006.
- [2] N. R. Adam, V. P. Janeja, and V. Atluri. Neighborhood-based Detection of Anomalies in High-Dimensional Spatio-temporal Sensor Datasets. *ACM SAC Conference*, 2004.
- [3] A. Adler, M. Elad, Y. Hel-Or, and E. Rivlin. Sparse Coding with Anomaly Detection. *Journal of Signal Processing Systems*, 79(2), pp. 179–188, 2015.
- [4] C. C. Aggarwal and P. S. Yu. Outlier Detection in High Dimensional Data. *ACM SIGMOD Conference*, 2001.
- [5] C. C. Aggarwal, C. Procopiuc, J. Wolf, P. Yu, and J. Park. Fast Algorithms for Projected Clustering. *ACM SIGMOD Conference*, 1999.
- [6] C. Aggarwal, J. Han, J. Wang, and P. Yu. A Framework for Projected Clustering of High Dimensional Data Streams. *VLDB Conference*, 2004.
- [7] C. C. Aggarwal and P. S. Yu. Finding Generalized Projected Clusters in High Dimensional Spaces. *ACM SIGMOD Conference*, 2000.
- [8] C. C. Aggarwal. Re-designing Distance Functions and Distance-based Applications for High Dimensional Data. *ACM SIGMOD Record*, 2001.
- [9] C. C. Aggarwal. On Abnormality Detection in Spuriously Populated Data Streams. *SIAM Conference on Data Mining*, 2005.
- [10] C. C. Aggarwal. *Data Streams: Models and Algorithms*, Springer, 2007.
- [11] C. C. Aggarwal. *Data Classification: Algorithms and Applications*, CRC Press, 2014.
- [12] C. C. Aggarwal. *Social Network Data Analytics*, Springer, 2011.
- [13] C. C. Aggarwal. On Effective Classification of Strings with Wavelets, *ACM KDD Conference*, 2002.
- [14] C. C. Aggarwal and K. Subbian. Evolutionary Network Analysis: A Survey. *ACM Computing Surveys*, May 2014.

- [15] C. C. Aggarwal, N. Ta, J. Wang, J. Feng, and M. J. Zaki. Xproj: A Framework for Projected Structural Clustering of XML Documents. *ACM KDD Conference*, 2007.
- [16] C. C. Aggarwal and P. S. Yu. On String Classification in Data Streams. *ACM KDD Conference*, 2007.
- [17] C. C. Aggarwal, Y. Zhao, and P. S. Yu. Outlier Detection in Graph Streams. *ICDE Conference*, 2011.
- [18] C. C. Aggarwal. Active Learning: A Survey, *Data Classification: Algorithms and Applications*, CRC Press, 2014.
- [19] C. C. Aggarwal. A Framework for Diagnosing Changes in Evolving Data Streams. *ACM SIGMOD Conference*, 2003.
- [20] C. C. Aggarwal and P. S. Yu. Online Analysis of Community Evolution in Data Streams. *SIAM Conference on Data Mining*, 2005.
- [21] C. C. Aggarwal. On the Effects of Dimensionality Reduction on High Dimensional Similarity Search. *ACM PODS Conference*, 2001.
- [22] C. C. Aggarwal. Managing and Mining Sensor Data, *Springer*, 2013.
- [23] C. C. Aggarwal and C. K. Reddy. Data Clustering: Algorithms and Applications, *CRC Press*, 2013.
- [24] C. Aggarwal and C. Zhai. Managing and Mining Text Data. *Springer*, 2012.
- [25] C. C. Aggarwal, A. Hinneburg, and D. Keim. On the Surprising Behavior of Distance Metrics in High Dimensional Space. *ICDT Conference*, 2001.
- [26] C. C. Aggarwal and P. S. Yu. Outlier Detection with Uncertain Data. *SIAM Conference on Data Mining*, 2008.
- [27] C. C. Aggarwal, Y. Xie, and P. S. Yu. On Dynamic Data-Driven Selection of Sensor Streams. *ACM KDD Conference*, 2011.
- [28] C. C. Aggarwal, J. Han, J. Wang, and P. Yu. A Framework for Clustering Evolving Data Streams. *VLDB Conference*, 2003.
- [29] C. C. Aggarwal and P. Yu. On Clustering Massive Text and Categorical Data Streams. *Knowledge and Information Systems*, 24(2), pp. 171–196, 2010.
- [30] C. C. Aggarwal and K. Subbian. Event Detection in Social Streams. *SIAM Conference on Data Mining*, 2012.
- [31] C. C. Aggarwal. Outlier Ensembles: Position Paper. *ACM SIGKDD Explorations*, 14(2), pp. 49–58, December, 2012.
- [32] C. C. Aggarwal and S. Sathe. Theoretical Foundations and Algorithms for Outlier Ensembles. *ACM SIGKDD Explorations*, 17(1), June 2015.
- [33] C. C. Aggarwal. Data Mining: The Textbook. *Springer*, 2015.
- [34] C. C. Aggarwal. Recommender Systems: The Textbook. *Springer*, 2016.

- [35] C. C. Aggarwal and S. Sathe. *Outlier Ensembles: An Introduction*. Springer, 2017.
- [36] R. Agrawal, T. Imielinski, and A. Swami. Mining Association Rules Between Sets of Items in Large Databases. *ACM SIGMOD Conference*, 1993.
- [37] R. Agrawal and R. Srikant. Fast algorithms for finding Association Rules in Large Databases. *VLDB Conference*, 1994.
- [38] M. Agyemang, K. Barker, and R. Alhajj. A Comprehensive Survey of Numeric and Symbolic Outlier Mining Techniques. *Intelligent Data Analysis*, 10(6). pp. 521–538, 2006.
- [39] A. Ahmad and L. Dey. A Method to Compute Distance between two Categorical Values of Same Attribute in Unsupervised Learning for Categorical Data Set. *Pattern Recognition Letters*, 28(1), pp. 110–118, 2007.
- [40] R. Ahuja, J. Orlin, and T. Magnanti. *Network Flows: Theory, Algorithms and Applications*. Prentice Hall, 1993.
- [41] L. Akoglu, M. McGlohon, and C. Faloutsos. Oddball: Spotting Anomalies in Weighted Graphs. *PAKDD Conference*, 2010.
- [42] L. Akoglu, H. Tong, J. Vreeken, and C. Faloutsos. Fast and Reliable Anomaly Detection in Categorical Data. *CIKM Conference*, 2012.
- [43] L. Akoglu, H. Tong, and D. Koutra. Graph-Based Anomaly Detection and Description: A Survey. *Data Mining and Knowledge Discovery*, 29(3), pp. 626–688, 2014.
- [44] L. Akoglu, E. Muller, and J. Vreeken. ACM KDD Workshop on Outlier Detection and Description, 2013. <http://www.outlier-analytics.org/odd13kdd/>
- [45] E. Aleskerov, B. Freisleben, and B. Rao. CARDWATCH: A Neural Network based Database Mining System for Credit Card Fraud Detection. *IEEE Computational Intelligence for Financial Engineering*, pp. 220–226, 1997.
- [46] T. Al-Khateeb, M. Masud, L. Khan, C. Aggarwal, J. Han, and B. Thuraisingham. Recurring and Novel Class Detection using Class-based Ensemble, *ICDM Conference*, 2012.
- [47] J. Allan, R. Papka, and V. Lavrenko. On-line New Event Detection and Tracking. *ACM SIGIR Conference*, 1998.
- [48] J. Allan, V. Lavrenko, and H. Jin. First Story Detection in TDT is Hard. *ACM CIKM Conference*, 2000.
- [49] J. Allan, J. Carbonell, G. Doddington, J. Yamron, and Y. Yang. Topic Detecting and Tracking Pilot Study Final Report, *CMU Technical Report, Paper 341*, 1998.
- [50] F. Alonso, J. Caraca-Valente, A. Gonzalez, and C. Montes. Combining Expert Knowledge and Data Mining in a Medical Diagnosis Domain. *Expert Systems with Applications*, 23(4), pp. 367–375, 2002.
- [51] Y. Altun, I. Tsochantaridis, and T. Hofmann. Hidden Narkov Support Vector Machines. *ICML Conference*, 2003.

- [52] M. Amer. Enhancing One-Class Support Vector Machines for Unsupervised Anomaly Detection. *Masters Thesis*, German University in Cairo, 2013. http://www.madm.eu/_media/theses/ma_thesis_amer.pdf
- [53] J. An and S. Cho. Variational Autoencoder based Anomaly Detection using Reconstruction Probability. *Technical Report*, SNU Data Mining Center, 2015. <http://dm.snu.ac.kr/static/docs/TR/SNUDM-TR-2015-03.pdf>
- [54] M. R. Anderberg. Cluster Analysis for Applications, *Academic Press*, New York, 1973.
- [55] D. Anderson, T. Lunt, H. Javitz, A. Tamaru, and A. Valdes. Detecting Unusual Program Behavior using the Statistical Components of NIDES, *Technical Report, SRI-CSL-95-06*, Computer Science Laboratory, SRI International, 1995.
- [56] D. Anderson, T. Frivold, A. Tamaru, and A. Valdes. Next-generation Intrusion Detection Expert System (nides), Software Users Manual, Beta-update Release. *Technical Report SRI-CSL-95-07*, Computer Science Laboratory, SRI International, 1994.
- [57] S. Ando. Clustering Needles in a Haystack: An Information Theoretic Analysis of Minority and Outlier Detection. *ICDM Conference*, 2007.
- [58] F. Angiulli and C. Pizzuti. Fast Outlier Detection in High Dimensional Spaces. *European Conference on Principles of Knowledge Discovery and Data Mining*, 2002.
- [59] F. Angiulli, F. Fassetti, and L. Palopoli. Finding Outlying Properties of Exceptional Objects. *ACM Transactions on Database Systems*, 34(1), 2009.
- [60] F. Angiulli and F. Fassetti. Detecting Distance-based Outliers in Streams of Data. *ACM CIKM Conference*, 2007.
- [61] M. Araujo, S. Papadimitriou, S. Gunnemann, C. Faloutsos, P. Basu, A. Swami, E. Papalexakis, and D. Koutra. Com2: Fast Automatic Discovery of Temporal (Comet) Communities. *PAKDD Conference*, 2014.
- [62] A. Arning, R. Agrawal, and P. Raghavan. A Linear Method for Deviation Detection in Large Databases. *ACM KDD Conference*, 1996.
- [63] I. Assent, P. Kranen, C. Beldauf, and T. Seidl. AnyOut: Anytime Outlier Detection in Streaming Data. *DASFAA Conference*, 2012.
- [64] I. Assent, R. Krieger, E. Muller, and T. Seidl. Subspace Outlier Mining in Large Multimedia Databases. *Parallel Universes and Local Patterns*, 2007.
- [65] A. Auer Jr. Correlation of Land Use and Cover with Meteorological Anomalies. *Journal of Applied Meteorology*, 17(5) pp. 636–643, 1978.
- [66] S. Babbar and S. Chawla. On Bayesian Network and Outlier Detection. *COMAD Conference*, 2010. Detailed Ph.D. thesis at <http://priyipati.library.usyd.edu.au/handle/2123/9371>
- [67] L. Backstrom, D. Huttenlocher, J. Kleinberg, and X. Lan. Group Formation in Large Social Networks: Membership, Growth, and Evolution. *KDD Conference*, 2006.
- [68] P. Baldi. Autoencoders, Unsupervised Learning, and Deep Architectures. *Unsupervised and Transfer Learning Challenges in Machine Learning*, Volume 7, 43, 2012.

- [69] D. Ballard. Generalizing the Hough Transform to Detect Arbitrary Shapes. *Pattern Recognition*, 11(2), pp. 111–122, 1981.
- [70] D. Barbara, Y. Li, J. Couto, J.-L. Lin, and S. Jajodia. Bootstrapping a Data Mining Intrusion Detection System. *Symposium on Applied Computing*, 2003.
- [71] D. Barbara, J. Couto, S. Jajodia, and N. Wu. ADAM: A Testbed for Exploring the Use of Data Mining in Intrusion Detection. *SIGMOD Record*, 30(4), pp. 15–24, 2001.
- [72] D. Barbara, J. Couto, S. Jajodia, and N. Wu. Detecting Novel Network Intrusions using Bayes Estimators. *SIAM Conference on Data Mining*, 2001.
- [73] R. Baragona and F. Battaglia. Outlier Detection in Multivariate Time Series by Independent Component Analysis. *Neural Computation*, 19(1), pp. 1962–1984, 2007.
- [74] V. Barnett and T. Lewis. *Outliers in Statistical Data*. Wiley, 1994.
- [75] S. Bay and M. Schwabacher. Mining Distance-based Outliers in Near Linear Time with Randomization and a Simple Pruning Rule. *ACM KDD Conference*, 2003.
- [76] S. Bay, K. Saito, N. Ueda, and P. Langley. A Framework for Discovering Anomalous Regimes in Multivariate Time-series Data with Local Models. *Technical report, Center for the Study of Language and Information*, Stanford University, 2004.
- [77] R. Beckman and R. Cook. Outliers, *Technometrics*, 25(2), pp. 119–149, 1983.
- [78] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger. The R*-tree: An Efficient and Robust Access method for Points and Rectangles. *ACM SIGMOD Conference*, 1990.
- [79] Y. Bengio, J. Paiement, P. Vincent, O. Delalleau, N. Le Roux, and M. Ouimet. Out-of-sample Extensions for LLE, Isomap, MDS, Eigenmaps, and Spectral Clustering. *Neural Information Processing Systems*, 16, pp. 177–184, 2006.
- [80] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy Layer-wise Training of Deep Networks. *Neural Information Processing Systems*, 19., pp. 153–160, 2007.
- [81] M. Berlingerio, F. Bonchi, B. Bringmann, and A. Gionis. Mining Graph Evolution Rules. *ECML/PKDD Conference*, 2009.
- [82] K. Beyer, J. Goldstein, R. Ramakrishnan, and U. Shaft. When is “nearest neighbor” meaningful? *International Conference on Database Theory*, 1999.
- [83] K. Bhaduri, B. Matthews, and C. Giannella. Algorithms for Speeding up Distance-based Outlier Detection. *ACM KDD Conference*, 2011.
- [84] C. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.
- [85] E. Blanzieri and A. Bryl. A Survey of Learning-based Techniques of Email Spam Filtering. *Artificial Intelligence Review*, 29(1), pp. 63–92, 2008.
- [86] M. Bosc, F. Heitz, J.-P. Armspach, I. Namer, D. Gounot, and L. Rumbach. Automatic Change Detection in Multimodal Serial MRI: Application to Multiple Sclerosis Lesion Evolution. *NeuroImage*, 20(2), 2003, Pages 643–656

- [87] Y. Bilberman. A Context Similarity Measure. *ECML Conference*, 1994.
- [88] P. Billingsley. Probability and Measure, Second Edition. *Wiley*, 1986.
- [89] D. Birant and A. Kut. Detecting Spatio-temporal Outliers in Large Databases. *Journal of Computing and Information Technology*, 14(4), pp. 291–297, 2006.
- [90] D. Blei and J. Lafferty. Dynamic Topic Models. *ICML Conference*, 2006.
- [91] D. Blei, A. Ng, and M. Jordan. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3: pp. 993–1022, 2003.
- [92] C. Bohm, K. Haegler, N. Muller, and C. Plant. Coco: Coding Cost for Parameter Free Outlier Detection. *ACM KDD Conference*, 2009.
- [93] S. Boriah, V. Chandola, and V. Kumar. Similarity Measures for Categorical Data: A Comparative Evaluation. *SIAM Conference on Data Mining*, 2008.
- [94] J. Branch, B. Szymanski, C. Giannella, R. Wolff, and H. Kargupta. In-Network Outlier Detection in Wireless Sensor Networks. *ICDCS Conference*, 2006.
- [95] T. Brants, F. Chen, and A. Farahat. A System for New Event Detection. *ACM SIGIR Conference*, 2003.
- [96] M. Breunig, H.-P. Kriegel, R. Ng, and J. Sander. LOF: Identifying Density-based Local Outliers. *ACM SIGMOD Conference*, 2000.
- [97] M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander. OPTICS-OF: identifying Local Outliers. *PKDD Conference*, 1999.
- [98] L. Brieman. Bagging Predictors. *Machine Learning*, 24: pp. 123–140, 1996.
- [99] L. Brieman and A. Cutler. Random Forests Manual v4.0, *Technical Report*, UC Berkeley, 2003. Available online at: https://www.stat.berkeley.edu/~breiman/Using_random_forests_v4.0.pdf
- [100] M. R. Brito, E. L. Chavez, A. J. Quiroz, and J. E. Yukich. Connectivity of the Mutual k -Nearest Neighbor Graph in Clustering and Outlier Detection. *Statistics and Probability Letters*, 35(1), pp. 33–42, 1997.
- [101] R. G. Brown and P. Hwang. Introduction to Random Signals and Applied Kalman Filtering. *John Wiley and Sons*, 1997.
- [102] Y. Bu, L. Chen, A. W.-C. Fu, and D. Liu. Efficient Anomaly Monitoring over Moving Object Trajectory Streams. *ACM KDD Conference*, 2009.
- [103] S. Budalakoti, A. Srivastava, R. Akella, and E. Turkov. Anomaly Detection in Large Sets of High-dimensional Symbol Sequences. *NASA Ames Research Center*, Technical Report NASA TM-2006-214553, 2006.
- [104] S. Budalakoti, A. Srivastava, and M. Otey. Anomaly Detection and Diagnosis Algorithms for Discrete Symbol Sequences with Applications to Airline Safety. *IEEE International Conference on Systems, Man, and Cybernetics*, 37(6), 2007.
- [105] P. Buhlmann. Bagging, Subbagging and Bragging for Improving some Prediction Algorithms. *Recent advances and trends in nonparametric statistics*, Elsevier, 2003.

- [106] P. Buhlmann and B. Yu. Analyzing Bagging. *Annals of Statistics*, pp. 927–961, 2002.
- [107] A. Buja and W. Stuetzle. Observations on Bagging. *Statistica Sinica*, 16(2), 323, 2006.
- [108] S. Burdakis and A. Deligiannakis. Detecting Outliers in Sensor Networks using the Geometric Approach. *ICDE Conference*, 2012.
- [109] T. Burnaby. On a Method for Character Weighting a Similarity Coefficient: Employing the Concept of Information. *Mathematical Geology*, 2(1), 25–38, 1970.
- [110] S. Byers and A. Raftery. Nearest-Neighbor Clutter Removal for Estimating Features in Spatial Point Processes. *JASIS*, 93, pp. 577–584, June 1998.
- [111] I. Cadez, D. Heckerman, C. Meek, P. Smyth, and S. White. Visualization of Navigation Patterns on a Web Site using Model-based Clustering. *ACM SIGMOD Conference*, 2000.
- [112] P. H. Calamai. Projected Gradient Methods for Linearly Constrained Problems. *Mathematical Programming*, 39: pp, 93–116, 1987.
- [113] C. Campbell and K. P. Bennett. A Linear-Programming Approach to Novel Class Detection. *Neural Information Processing Systems*, 2000.
- [114] G. O. Campos, A. Zimek, J. Sander, R. J. G. B. Campello, B. Micenkova, E. Schubert, I. Assent, and M. E. Houle. On the Evaluation of Unsupervised Outlier Detection: Measures, Datasets, and an Empirical Study. *DMKD Journal*, 30(4), pp. 891–927, 2016.
- [115] M. J. Canty. Image Analysis, Classification and Change Detection in Remote Sensing: with Algorithms for ENVI/IDL. *CRC Press*, 2006.
- [116] C. Caroni. Outlier Detection by Robust Principal Component Analysis. *Communications in Statistics – Simulation and Computation*, 29: pp. 129–151, 2000.
- [117] L. E. Carr and R. L. Elsberry. Monsoonal Interactions Leading to Sudden Tropical Cyclone Track Changes. *Monthly Weather Review*, 123(2), pp. 265–290, Feb. 1995.
- [118] P. Castro, D. Zhang, C. Chen, S. Li, and G. Pan. From Taxi GPS Traces to Social and Community Dynamics: A Survey. *ACM Computing Surveys (CSUR)*, 46(2), 17, 2013.
- [119] D. Chakrabarti and C. Faloutsos. Evolutionary Clustering. *ACM KDD Conference*, 2006.
- [120] D. Chakrabarti. AutoPart: Parameter-Free Graph Partitioning and Outlier Detection. *PKDD Conference*, 2004.
- [121] K. Chakrabarti and S. Mehrotra. Local Dimensionality Reduction: A New Approach to Indexing High Dimensional Spaces. *VLDB Conference Proceedings*, 2000.
- [122] S. Chakrabarti, S. Sarawagi, and B. Dom. Mining Surprising Patterns using Temporal Description Length. *VLDB Conference*, 1998.
- [123] C.-H. Chan and G. Pang. Fabric Defect Detection by Fourier Analysis. *IEEE Transactions on Industry Applications*, 36(5), 2000.

- [124] P. K. Chan and S. J. Stolfo. Toward Scalable Learning with Non-uniform Class and Cost Distributions: A Case Study in Credit Card Fraud Detection. *KDD Conference*, pp. 164–168, 1998.
- [125] V. Chandola, A. Banerjee, and V. Kumar. Anomaly Detection: A Survey. *ACM Computing Surveys*, 41(3), 2009.
- [126] V. Chandola, A. Banerjee, and V. Kumar. Anomaly Detection for Discrete Sequences: A Survey. *IEEE Transactions on Knowledge and Data Engineering*, 24(5): pp. 823–839, 2012.
- [127] V. Chandola, V. Mithal, and V. Kumar. A Comparative Evaluation of Anomaly Detection Techniques for Sequence Data. *ICDM Conference*, 2008.
- [128] C. Chang and C. J. Lin. LIBSVM: a Library for Support Vector Machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3), 27, 2011.
- [129] I. Chang, G. C. Tiao, and C. Chen. Estimation of Time Series Parameters in the Presence of Outliers. *Technometrics*, 30(2), pp. 193–204, 1988.
- [130] O. Chapelle. Training a Support Vector Machine in the Primal. *Neural Computation*, 19(5), pp. 1155–1178, 2007.
- [131] A. Chaudhary, A. S. Szalay, and A. W. Moore. Very Fast Outlier Detection in Large Multidimensional Data Sets. *DMKD Workshop*, 2002.
- [132] N. V. Chawla, N. Japkowicz, and A. Kotcz. Editorial: Special Issue on Learning from Imbalanced Data Sets. *ACM SIGKDD Explorations Newsletter*, 6(1), 1–6, 2004.
- [133] N. V. Chawla, K. W. Bower, L. O. Hall, and W. P. Kegelmeyer. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research (JAIR)*, 16, pp. 321–356, 2002.
- [134] N. V. Chawla, A. Lazarevic, L. Hall, and K. Bowyer. SMOTEBoost: Improving Prediction of the Minority Class in Boosting. *PKDD*, pp. 107–119, 2003.
- [135] N. V. Chawla, D. A. Cieslak, L. O. Hall, and A. Joshi. Automatically Counteracting Imbalance and its Empirical Relationship to Cost. *Data Mining and Knowledge Discovery*, 17(2), pp. 225–252, 2008.
- [136] C. Chen and L.-M. Liu. Joint Estimation of Model Parameters and Outlier Effects in Time Series. *Journal of the American Statistical Association*, 88(421), pp. 284–297, March 1993.
- [137] C. Chen, D. Zhang, P. Castro, N. Li, L. Sun, and S. Li. Real-time Detection of Anomalous Taxi Trajectories from GPS Traces. *Mobile and Ubiquitous Systems: Computing, Networking, and Services*, pp. 63–74, 2011.
- [138] D. Chen, C.-T. Lu, Y. Chen, and D. Kou. On Detecting Spatial Outliers. *Geoinformatica*, 12: pp. 455–475, 2008.
- [139] Y. Chen and L. Tu. Density-based Clustering for Real Time Stream Data. *ACM KDD Conference*, 2007.

- [140] H. Cheng, P.-N. Tan, C. Potter, and S. Klooster. Detection and Characterization of Anomalies in Multivariate Time Series. *SIAM Conference on Data Mining*, 2009.
- [141] T. Cheng and Z. Li. A Hybrid Approach to Detect Spatialtemporal Outliers. *International Conference on Geoinformatics*, 2004.
- [142] T. Cheng and Z. Li. A Multiscale Approach for Spatio-temporal Outlier Detection. *Transactions in GIS*, 10(2), pp. 253–263, March 2006.
- [143] Y. Chi, X. Song, D. Zhou, K. Hino, and B. Tseng. Evolutionary Spectral Clustering by Incorporating Temporal Smoothness. *ACM KDD Conference*, 2007.
- [144] A. Chiu and A. Fu. Enhancements on Local Outlier Detection. *Database Engineering and Applications Symposium*, 2003.
- [145] C. Chow and D. Yeung. Parzen-Window Network Intrusion Detectors. *International Conference on Pattern Recognition*, 4, 2002.
- [146] M. Chow, R. Sharpe, and J. Hung. On the Application and Design of Artificial Neural Networks for Motor Fault Detection. *IEEE Transactions on Industrial Electronics*, 40(2), 1993.
- [147] V. Ciesielski and V. Ha. Texture Detection using Neural Networks Trained on Examples of One Class. *LNCS*, Vol. 5866, pp. 140–149, *Springer*, 2009.
- [148] W. Cohen. Fast Effective Rule Induction. *ICML Conference*, 1995.
- [149] D. Cohn, R. Atlas, and N. Ladner. Improving Generalization with Active Learning. *Machine Learning*, 15, pp. 201–221.
- [150] R. Cooley, B. Mobashar, and J. Srivastava. Data Preparation for Mining World Wide Web Browsing Patterns. *Knowledge and Information Systems*, 1, pp. 5–32, 1999.
- [151] D. Cook and L. Holder. Graph-Based Data Mining. *IEEE Intelligent Systems*, 15(2), pp. 32–41, 2000.
- [152] G. Cormack. Email Spam Filtering: A Systematic Review. *Foundations and Trends in Information Retrieval*, 1(4), pp. 335–455, 2007.
- [153] X. Dang, B. Micenkova, I. Assent, and R. Ng. Outlier Detection with Space Transformation and Spectral Analysis. *SIAM Conference on Data Mining*, pp. 225–233, 2013.
- [154] G. Das and H. Mannila. Context-based Similarity Measures for Categorical Databases. *PKDD Conference*, 2000.
- [155] K. Das, J. Schneider, and D. Neill. Anomaly Pattern Detection in Categorical Data Sets. *ACM KDD Conference*, 2008.
- [156] S. Das, B. Matthews, A. Srivastava, and N. Oza. Multiple Kernel Learning for Heterogeneous Anomaly Detection: Algorithm and Aviation Safety Case Study. *ACM KDD Conference*, 2010.
- [157] D. Dasgupta and S. Forrest. Novelty Detection in Time Series using Ideas from Immunology. *International Conference on Intelligent Systems*, 1996.

- [158] D. Dasgupta and F. Nino. A Comparison of Negative and Positive Selection Algorithms in Novel Pattern Detection. *IEEE Conference on Systems, Man, and Cybernetics*, 1, pp. 125–130, 2000.
- [159] T. Dasu, S. Krishnan, S. Venkatasubramanian, and K. Yi. An Information-Theoretic Approach to Detecting Change in Multi-dimensional Data Streams. *Symposium on the Interface of Computer Science, Statistics, and Applications*, 2006.
- [160] H. Dau, V. Ciesielski, and A. Song. Anomaly Detection using Replicator Neural Networks Trained on Examples of One Class. In *Simulated Evolution and Learning*, pp. 311–322, Springer, 2014.
- [161] N. Delannay, C. Archambeau, and M. Verleysen. Improving the Robustness to Outliers of Mixtures of Probabilistic PCAs. *PAKDD Conference*, 2008.
- [162] S. T. Deerwester, S. T. Dumais, G. Furnas, and R. Harshman. Indexing by Latent Semantic Analysis, *JASIS*, 1990.
- [163] K. A. De Jong. Analysis of the Behavior of a Class of Genetic Adaptive Systems. *Ph.D. Dissertation, University of Michigan*, Ann Arbor, MI, 1975.
- [164] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society*, B, vol. 39(1), pp. 1–38, 1977.
- [165] D. Denning. An Intrusion Detection Model. *IEEE Transactions of Software Engineering*, 13(2), pp. 222–232.
- [166] R. Derrig. Insurance Fraud. *Journal of Risk and Insurance*, 69(3), pp. 271–287, 2002.
- [167] M. Desforges, P. Jacob, and J. Cooper. Applications of Probability Density Estimation to the Detection of Abnormal Conditions in Engineering. *Proceedings of Institute of Mechanical Engineers*, Vol. 212, pp. 687–703, 1998.
- [168] M. Deshpande and G. Karypis. Evaluation of Techniques for Classifying Biological Sequences. *PAKDD Conference*, 2002.
- [169] M. Deshpande and G. Karypis. Selective Markov Models for Predicting Web Page Accesses. *ACM Transactions on Internet Technology*, 4(2), pp. 163–184, 2004.
- [170] T. Dietterich. Ensemble Methods in Machine Learning. *First International Workshop on Multiple Classifier Systems*, 2000.
- [171] H. Ding, G. Trajcevski, P. Scheuermann, X. Wang, and E. Keogh. Querying and Mining of Time Series Data: Experimental Comparison of Representations and Distance Measures. *PVLDB*, 1(2), pp. 1542–1552, 2008.
- [172] S. Donoho. Early Detection of Insider Trading in Option Markets. *ACM KDD Conference*, 2004.
- [173] C. Drummond and R. Holte. C4.5, Class Imbalance, and Cost Sensitivity: Why Undersampling beats Oversampling. *ICML Workshop on Learning from Imbalanced Data Sets*, 2003.

- [174] C. Drummond and R. Holte. Explicitly Representing Expected Cost: An Alternative to ROC representation. *ACM KDD Conference*, pp. 198–207, 2001.
- [175] P. Domingos. MetaCost: A General Framework for Making Classifiers Cost-Sensitive, *ACM KDD Conference*, 1999.
- [176] R. Duda, P. Hart, and D. Stork. *Pattern Classification*. Wiley, 2001.
- [177] H. Dutta, C. Giannella, K. Borne, and H. Kargupta. Distributed top- k Outlier Detection in Astronomy Catalogs using the Demac System. *SIAM Conference on Data Mining*, 2007.
- [178] J. Dutta, B. Banerjee, and C. Reddy. RODS: Rarity based Outlier Detection in a Sparse Coding Framework. *IEEE Transactions on Knowledge and Data Engineering*, 28(2), p. 483–495, 2016.
- [179] W. Eberle and L. B. Holder. Mining for Structural Anomalies in Graph-based Data. *DMIN*, 2007.
- [180] F. Y. Edgeworth. On Discordant Observations. *Philosophical Magazine*, 23(5), pp. 364–375, 1887.
- [181] M. Elahi, K. Li, W. Nisar, X. Lv, and H. Wang. Efficient Clustering-Based Outlier Detection Algorithm for Dynamic Data Stream. *FSKD Conference*, 2008.
- [182] C. Elkan. The Foundations of Cost-Sensitive Learning. *IJCAI*, 2001.
- [183] C. Elkan and K. Noto. Learning Classifiers from only Positive and Unlabeled Data. *ACM KDD Conference*, 2008.
- [184] A. Emmott, S. Das, T. Dietterich, A. Fern, and W. Wong. Systematic Construction of Anomaly Detection Benchmarks from Real Data. *arXiv:1503.01158*, 2015.
<https://arxiv.org/abs/1503.01158>
- [185] D. Endler. Intrusion detection: Applying Machine Learning to Solaris Audit Data. *Annual Computer Security Applications Conference*, 1998.
- [186] D. Erhan, Y. Bengio, A. Courville, P. Manzagol, P. Vincent, and S. Bengio. Why does Unsupervised Pre-training help Deep Learning? *Journal of Machine Learning Research*, 11, 625-660, 2010.
- [187] E. Eskin. Anomaly Detection over Noisy Data using Learned Probability Distributions. *ICML Conference*, 2000.
- [188] E. Eskin, A. Arnold, M. Prerau, L. Portnoy, and S. Stolfo. A Geometric Framework for Unsupervised Anomaly Detection. *Applications of Data Mining in Computer Security*. Kluwer, 2002.
- [189] E. Eskin, W. Lee, and S. Stolfo, Modeling System Call for Intrusion Detection using Dynamic Window Sizes. *DISCEX*, 2001.
- [190] H. Fan, O. Zaiane, A. Foss, and J. Wu. A Nonparametric Outlier Detection for Efficiently Discovering top- n Outliers from Engineering Data. *PAKDD Conference*, 2006.

- [191] W. Fan, S. Stolfo, J. Zhang, and P. Chan. AdaCost: Misclassification Cost Sensitive Boosting. *ICML Conference*, 1999.
- [192] T. Fawcett. ROC Graphs: Notes and Practical Considerations for Researchers. *Technical Report HPL-2003-4*, Palo Alto, CA: HP Laboratories, 2003.
- [193] T. Fawcett and F. Provost. Activity Monitoring: Noticing Interesting Changes in Behavior. *ACM KDD Conference*, 1999.
- [194] W. Feng, C. Zhang, W. Zhang, J. Han, J. Wang, C. Aggarwal, and J. Huang. STREAMCUBE: Hierarchical Spatio-temporal Hashtag Clustering for Event Exploration over the Twitter Stream. *IEEE ICDE Conference*, pp. 1561–1572, 2015.
- [195] M. Fernandez-Delgado, E. Cernadas, S. Barro, and D. Amorim. Do we Need Hundreds of Classifiers to Solve Real World Classification Problems? *The Journal of Machine Learning Research*, 15(1), pp. 3133–3181, 2014.
- [196] D. Fetterly, M. Manasse, and M. Najork. Spam, Damn Spam, and Statistics: using Statistical Analysis to Locate Spam Web Pages. *WebDB*, 2004.
- [197] S. Forrest, C. Warrender, and B. Pearlmutter. Detecting Intrusions using System Calls: Alternate Data Models. *IEEE ISRSP*, 1999.
- [198] S. Forrest, S. Hofmeyr, A. Somayaji, and T. A. Longstaff. A Sense of Self for Unix Processes. *ISRSP*, 1996.
- [199] S. Forrest, P. D’Haeseleer, and P. Helman. An Immunological Approach to Change Detection: Algorithms, Analysis and Implications. *IEEE Symposium on Security and Privacy*, 1996.
- [200] S. Forrest, F. Esponda, and P. Helman. A Formal Framework for Positive and Negative Detection Schemes. *IEEE Transactions on Systems, Man and Cybernetics, Part B*, pp. 357–373, 2004.
- [201] S. Forrest, A. Perelson, L. Allen, and R. Cherukuri. Self-Nonself Discrimination in a Computer. *IEEE Symposium on Security and Privacy*, 1994.
- [202] A. Fox. Outliers in Time Series. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(3), pp. 350–363, 1972.
- [203] A. Frank and A. Asuncion. UCI Machine Learning Repository, Irvine, CA: University of California, School of Information and Computer Science, 2010. <http://archive.ics.uci.edu/ml>
- [204] C. Franke and M. Gertz. ORDEN: Outlier Region Detection and Exploration in Sensor Networks. *ACM SIGMOD Conference*, 2009.
- [205] A. Fu, O. Leung, E. Keogh, and J. Lin. Finding Time Series Discords based on Haar Transform. *Advanced Data Mining and Applications*, 2006.
- [206] R. Fujumaki, T. Yairi, and K. Machida. An Approach to Spacecraft Anomaly Detection Problem using Kernel Feature Space. *ACM KDD Conference*, 2005.
- [207] R. Fujamaki. Anomaly Detection Support Vector Machine and Its Application to Fault Diagnosis. *ICDM Conference*, 2008.

- [208] P. Galeano, D. Pea, and R. S. Tsay. Outlier detection in Multivariate Time Series via Projection Pursuit. *Statistics and Econometrics Working Papers WS044221*, Universidad Carlos III, 2004.
- [209] P. Gambaryan. A Mathematical Model of Taxonomy. *Izvest. Akad. Nauk Armen, SSR*, 17(12), pp. 47–53, 1964.
- [210] V. Ganti, J. Gehrke, and R. Ramakrishnan. CACTUS: Clustering Categorical Data using Summaries. *ACM KDD Conference*, 1999.
- [211] B. Gao, H.-Y. Ma, and Y.-H. Yang. HMMs (Hidden Markov Models) based on Anomaly Intrusion Detection Method. *International Conference on Machine Learning and Cybernetics*, 2002.
- [212] H. Gao, X. Wang, J. Tang, and H. Liu. Network Denoising in Social Media. *Technical Report, Arizona State University*, 2011.
- [213] J. Gao and P.-N. Tan. Converting Outlier Scores from Outlier Detection Algorithms into Probability Estimates. *ICDM Conference*, 2006.
- [214] J. Gao, F. Liang, W. Fan, C. Wang, Y. Sun, and J. Han. On Community Outliers and their Efficient Detection in Information Networks. *ACM KDD Conference*, pp. 813–822, 2010.
- [215] Y. Ge, H. Xiong, Z.-H. Zhou, H. Ozdemir, J. Yu, and K. Lee. Top-Eye: Top- k Evolving Trajectory Outlier Detection. *CIKM Conference*, 2010.
- [216] A. Ghosh, J. Wanken, and F. Charron. Detecting Anomalous and Unknown Intrusions against Programs. *Annual Computer Security Applications Conference*, 1998.
- [217] A. Ghosh, A. Schwartzbard, and M. Schatz. Learning Program Behavior Profiles for Intrusion Detection. *USENIX Workshop on Intrusion Detection and Network Monitoring*, pp. 51–62, 1999.
- [218] S. Ghosh and D. Reilly. Credit Card Fraud Detection with a Neural Network. *International Conference on System Sciences: Information Systems: Decision Support and Knowledge-Based Systems*, 3, pp. 621–630, 1994.
- [219] A. Ghoting, S. Parthasarathy, and M. Otey. Fast Mining of Distance-based Outliers in High Dimensional Spaces. *SIAM Conference on Data Mining*, 2006.
- [220] D. Gibson, J. Kleinberg, and P. Raghavan. Clustering Categorical Data: An Approach Based on Dynamical Systems. *The VLDB Journal*, 8(3), pp. 222–236, 2000.
- [221] M. Goldstein and S. Uchida. A Comparative Evaluation of Unsupervised Anomaly Detection Algorithms for Multivariate Data. *PLoS One*, 11(4), e0152173, 2016.
- [222] D. W. Goodall. A new similarity index based on probability. *Biometrics*, 22(4), pp. 882–907, 1966.
- [223] I. Goodfellow, Y. Bengio, and A. Courville. Deep Learning. *MIT Press*, 2016. <http://goodfeli.github.io/dlbook/>,
- [224] N. Gornitz, M. Kloft, K. Rieck, and U. Brefeld. Toward Supervised Anomaly Detection. *Journal of Artificial Intelligence Research*, 2013.

- [225] F. Grubbs. Procedures for Detecting Outlying Observations in Samples. *Technometrics*, 11(1), pp. 1–21, 1969.
- [226] S. Guha, R. Rastogi, and K. Shim. ROCK: A Robust Clustering Algorithm for Categorical Attributes. *Information Systems*, 25(5), pp. 345–366, 2000.
- [227] S. Guha, N. Mishra, G. Roy, and O. Schrijver. Robust Random Cut Forest Based Anomaly Detection On Streams. *ICML Conference*, pp. 2712–2721, 2016.
- [228] S. Gunnemann, N. Gunnemann, and C. Faloutsos. Detecting Anomalies in Dynamic Rating Data: A Robust Probabilistic Model for Rating Evolution. *ACM KDD Conference*, 2014.
- [229] D. Gunopulos and G. Das. Time-series Similarity Measures, and Time Series Indexing, *ACM SIGMOD Conference*, 2001.
- [230] S. Gunter, N. N. Schraudolph, and S. V. N. Vishwanathan. Fast Iterative Kernel Principal Component Analysis. *Journal of Machine Learning Research*, 8, pp 1893–1918, 2007.
- [231] M. Gupta, J. Gao, C. Aggarwal, and J. Han. Outlier Detection for Temporal Data. *Morgan and Claypool*, 2014.
- [232] M. Gupta, J. Gao, C. Aggarwal, and J. Han. Outlier Detection for Temporal Data: A Survey. *IEEE Transactions on Knowledge and Data Engineering*, 26(9), pp. 1, 2014.
- [233] M. Gupta, C. Aggarwal, J. Han, and Y. Sun. Evolutionary Clustering and Analysis of Bibliographic Networks. *ASONAM Conference*, 2011.
- [234] M. Gupta, C. Aggarwal, and J. Han. Finding Top- k Shortest Path Distance Changes in an Evolutionary Network. *SSDBM Conference*, 2011.
- [235] M. Gupta, J. Gao, Y. Sun, and J. Han. Community Trend Outlier Detection Using Soft Temporal Pattern Mining. *ECML/PKDD Conference*, 2012.
- [236] M. Gupta, J. Gao, Y. Sun, and J. Han. Integrating Community Matching and Outlier Detection for Mining Evolutionary Community Outliers. *KDD Conference*, 2012.
- [237] M. Gupta, A. Mallya, S. Roy, J. Cho, and J. Han. Local Learning for Mining Outlier Subgraphs from Network Datasets. *SIAM Conference on Data Mining*, pp. 73–81, 2014.
- [238] M. Gupta, J. Gao, X. Yan, H. Cam, and J. Han. Top- k Interesting Subgraph Discovery in Information Networks. *International Conference on Data Engineering (ICDE)*, 2014.
- [239] D. Gusfield. Algorithms for Strings, Trees and Sequences. *Cambridge University Press*, 1997.
- [240] D. Guthrie, L. Guthrie, and Y. Wilks. An Unsupervised Approach for the Detection of Outliers in Corpora. *LREC*, 2008.
- [241] S. Guttormsson, R. Marks, M. El-Sharkawi, and I. Kerszenbaum. Elliptical Novelty Grouping for Online Short-turn Detection of Excited Running Rotors. *IEEE Transactions on Energy Conversion*, 14(1), pp. 16–22, 1999.

- [242] R. Gwadera, M. Atallah, and W. Szpankowski. Markov Models for Identification of Significant Episodes. *SIAM Conference on Data Mining*, 2005.
- [243] R. Gwadera, M. Atallah, and W. Szpankowski. Detection of Significant Sets of Episodes in Event Sequences. *IEEE ICDM Conference*, 2004.
- [244] R. Gwadera, M. Atallah, and W. Szpankowski. Reliable Detection of Episodes in Event Sequences. *Knowledge and Information Systems*, 7(4), pp. 415–437, 2005.
- [245] F. Hampel. A General Qualitative Definition of Robustness. *Annals of Mathematics and Statistics*, 43, pp. 1887–1896, 1971.
- [246] J. Hanley and B. McNeil. The Meaning and Use of the Area under a Receiver Operating Characteristic (ROC) Curve. *Radiology*, 143, pp. 29–36, 1982.
- [247] J. Haslett, R. Brandley, P. Craig, A. Unwin, and G. Wills. Dynamic Graphics for Exploring Spatial Data With Application to Locating Global and Local Anomalies. *The American Statistician*, 45: pp. 234–242, 1991.
- [248] V. Hautamaki, I. Karkkainen, and P. Franti. Outlier Detection Using k -Nearest Neighbor Graph. *International Conference on Pattern Recognition*, 2004.
- [249] D. Hawkins. Identification of Outliers, *Chapman and Hall*, 1980.
- [250] S. Hawkins, H. He, G. Williams, and R. Baxter. Outlier Detection using Replicator Neural Networks. *International Conference on Data Warehousing and Knowledge Discovery*, pp. 170–180, Springer, 2002.
- [251] G. G. Hazel. Multivariate Gaussian MRF for Multispectral Scene Segmentation and Anomaly Detection. *GeoRS*, 38(3), pp. 1199–1211, 2000.
- [252] J. He, and J. Carbonell. Nearest-Neighbor-Based Active Learning for Rare Category Detection. *CMU Computer Science Department*, Paper 281, 2007.
<http://repository.cmu.edu/compsci/281>
- [253] Z. He, S. Deng, and X. Xu. Outlier Detection Integrating Semantic Knowledge. *Web Age Information Management (WAIM)*, 2002.
- [254] Z. He, X. Xu, J. Huang, and S. Deng. FP-Outlier: Frequent Pattern-based Outlier Detection. *COMSIS*, 2(1), 2005.
- [255] Z. He, X. Xu, and S. Deng. Discovering Cluster-based Local Outliers. *Pattern Recognition Letters*, Vol 24(9–10), pp. 1641–1650, 2003.
- [256] Z. He, X. Xu, and S. Deng. An Optimization Model for Outlier Detection in Categorical Data. *International Conference on Intelligent Computing*, 2005.
- [257] Z. He, S. Deng, X. Xu, and J. Huang. A Fast Greedy Algorithm for Outlier Mining. *PAKDD Conference*, 2006.
- [258] Z. He, S. Deng, and X. Xu. A Unified Subspace Outlier Ensemble Framework for Outlier Detection. *Advances in Web-Age Information Management*, pp. 632–637, 2005.
- [259] R. Hecht-Nielsen. Replicator Neural Networks for Universal Optimal Source Coding. *Science*, 269, 1860–1863, 1995.

- [260] P. Helman and J. Bhangoo. A Statistically-based System for Prioritizing Information Exploration under Uncertainty. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 27(4), pp. 449–466, 1997.
- [261] M. Henrion, D. Hand, A. Gandy, and D. Mortlock. CASOS: A Subspace Method for Anomaly Detection in High Dimensional Astronomical Databases. *Statistical Analysis and Data Mining*, 6(1), pp. 53–72, 2013.
- [262] S. Hido, Y. Tsuboi, H. Kashima, M. Sugiyama, and T. Kanamori. Statistical Outlier Detection using Direct Density Ratio Estimation. *Knowledge and Information Systems*, 26(2), pp. 309–336, 2011.
- [263] A. Hinneburg, C. Aggarwal, and D. Keim. What is the Nearest Neighbor in High-dimensional Spaces?, *VLDB Conference*, 2000.
- [264] G. Hinton and R. Salakhutdinov. Reducing the Dimensionality of Data with Neural Networks. *Science*, 313, (5786), pp. 504–507, 2006.
- [265] H. O. Hirschfeld. A Connection between Correlation and Contingency. *Proc. Cambridge Philosophical Society*, 31, pp. 520–524, 1935.
- [266] T. K. Ho. Random Decision Forests. *Third International Conference on Document Analysis and Recognition*, 1995. Extended version appears in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8), pp. 832–844, 1998.
- [267] S.-S. Ho. A Martingale Framework for Concept Change Detection in Time-Varying Data Streams. *ICML Conference*, 2005.
- [268] V. Hodge and J. Austin. A Survey of Outlier Detection Methodologies, *Artificial Intelligence Review*, 22(2), pp. 85–126, 2004.
- [269] J. Hodges. Efficiency in Normal Samples and Tolerance of Extreme Values for Some Estimates of Location. *Fifth Berkeley Symposium on Mathematics, Statistics and Probability*, 1, pp. 163–168, 1967.
- [270] H. Hoffmann. Kernel PCA for Novelty Detection. *Pattern Recognition*, 40(3), pp. 863–874, 2007.
- [271] T. Hofmann. Probabilistic Latent Semantic Indexing. *ACM SIGIR Conference*, 1999.
- [272] S. Hofmeyr, S. Forrest, and A. Somayaji. Intrusion Detection using Sequences of System Calls. *Journal of Computer Security*, 6(3), pp. 151–180, 1998.
- [273] J. H. Holland. Adaptation in Natural and Artificial Systems. *University of Michigan Press*, Ann Arbor, MI, 1975.
- [274] G. Hollier and J. Austin. Novelty Detection for Strain-gauge Degradation using Maximally Correlated Components. *European Symposium on Artificial Neural Networks*, 2002.
- [275] J. Hollmen and V. Tresp. Call-based Fraud Detection in Mobile Communication Networks using a Hierarchical Regime-switching Model. *Neural Information Processing Systems*, pp. 889–895, 1998.

- [276] P. Horn, L. Feng, Y. Li, and A. J. Pesce. Effect of Outliers and Nonhealthy Individuals on Reference Interval Estimation. *Clinical Chemistry*, 47(12), pp. 2137–2145, 2001.
- [277] R. Hu, C. Aggarwal, S. Ma, and J. Huai. An Embedding Approach to Anomaly Detection, *IEEE ICDE Conference*, 2016.
- [278] L. Huang, M. I. Jordan, A. Joseph, M. Garofalakis, and N. Taft. In-network PCA and Anomaly Detection. *Neural Information Processing Systems*, 2006.
- [279] J. W. Hunt and T. G. Szlymski. A Fast Algorithm for Computing Longest Common Subsequences. *Communications of the ACM*, 20(5), pp. 350–353, 1977.
- [280] T. Ide and H. Kashima. Eigenspace-based Anomaly Detection in Computer Systems. *ACM KDD Conference*, 2004.
- [281] P. Indyk, R. Motwani, P. Raghavan, and S. Vempala. Locality-preserving Hashing in Multidimensional Spaces. *ACM STOC Conference*, 1997.
- [282] D. Jackson and Y. Chen. Robust Principal Component Analysis and Outlier Detection with Ecological Data. *Environmetrics*, 15, pp. 129–139, 2004.
- [283] A. Jain and R. Dubes. Algorithms for Clustering Data. *Prentice Hall*, 1988.
- [284] H. Jagadish, N. Koudas, and S. Muthukrishnan. Mining Deviants in a Time-Series Database. *VLDB Conference*, 1999.
- [285] V. Janeja and V. Atluri. Random Walks to Identify Anomalous Free-form Spatial Scan Windows. *IEEE Transactions on Knowledge and Data Engineering*, 20(10), 2008.
- [286] P. Janeja and V. Atluri. Spatial Outlier Detection in Heterogeneous Neighborhoods. *Intelligent Data Analysis*, 13(1), 2008.
- [287] R. Jarvis and E. Patrick. Clustering Using a Similarity Measure based on Shared Near Neighbors. *IEEE Transactions on Computers*, 100(11), pp. 1025–1034, 1973.
- [288] H. Javitz and A. Valdez. The SRI IDES Statistical Anomaly Detector. *IEEE Symposium on Security and Privacy*, 1991.
- [289] Y. Jeong, M. Jeong, and O. Omiaomu. Weighted Dynamic Time Warping for Time Series Classification, *Pattern Recognition*, 44, pp. 2231–2240, 2010.
- [290] B. Jiang and J. Pei. Outlier Detection on Uncertain Data: Objects, Instances, and Inferences. *ICDE Conference*, 2011.
- [291] M. F. Jiang, S. S. Tseng, and C. M. Su. Two-phase Clustering Process for Outliers Detection. *Pattern Recognition Letters*, 22, 6–7, pp. 691–700, 2001.
- [292] R. Jiang, H. Fei, and J. Huan. Anomaly Localization for Network Data Streams with Graph Joint Sparse PCA. *ACM KDD Conference*, 2011.
- [293] W. Jin, A. Tung, and J. Han. Mining Top- n Local Outliers in Large Databases. *ACM KDD Conference*, 2001.
- [294] W. Jin, A. Tung, J. Han, and W. Wang. Ranking outliers using symmetric neighborhood relationship. *PAKDD Conference*, 2006.

- [295] T. Johnson, I. Kwok, and R. Ng. Fast Computation of 2-dimensional Depth Contours. *ACM KDD Conference*, 1998.
- [296] I. Jolliffe. *Principal Component Analysis*. Springer, 2002.
- [297] A. Ng, M. Jordan, and Y. Weiss. On Spectral Clustering Analysis and an Algorithm. *Neural Information Processing Systems*, pp. 849–856, 2001.
- [298] M. Joshi, R. Agarwal, and V. Kumar. Mining Needles in a Haystack: Classifying Rare Classes via Two-Phase Rule Induction, *ACM SIGMOD Conference*, 2001.
- [299] M. Joshi, V. Kumar, and R. Agarwal. Evaluating Boosting Algorithms to Classify Rare Classes: Comparison and Improvements. *ICDM Conference*, pp. 257–264, 2001.
- [300] M. Joshi and R. Agarwal. PNRule: A Framework for Learning Classifier Models in Data Mining (A Case Study in Network Intrusion Detection). *SIAM Conference on Data Mining*, 2001.
- [301] M. Joshi, R. Agarwal, and V. Kumar. Predicting Rare Classes: Can Boosting Make Any Weak Learner Strong? *ACM KDD Conference*, 2002.
- [302] M. Joshi. On Evaluating Performance of Classifiers on Rare Classes. *ICDM Conference*, 2003.
- [303] P. Juszczak and R. P. W. Duin. Uncertainty Sampling Methods for One-class Classifiers. *ICML Workshop on Learning from Imbalanced Data Sets*, 2003.
- [304] J. Kang, S. Shekhar, C. Wennen, and P. Novak. Discovering Flow Anomalies: A SWEET Approach. *ICDM Conference*, 2008.
- [305] G. Karakoulas and J. Shawe-Taylor. Optimising Classifiers for Imbalanced Training Sets. *Neural Information Processing Systems*, 1998.
- [306] S. Kasiviswanathan, P. Melville, and A. Banerjee. Emerging Topic Detection using Dictionary Learning. *CIKM Conference*, 2011.
- [307] L. Kaufman and P. Rousseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis*. Wiley-Interscience, 1990.
- [308] F. Keller, E. Muller, K. Bohm. HiCS: High-Contrast Subspaces for Density-based Outlier Ranking. *IEEE ICDE Conference*, 2012.
- [309] M. Kendall and J. Ord. *Time-series*. Hodder Arnold, 1990.
- [310] E. Keogh, S. Lonardi, and B. Y.-C. Chiu. Finding Surprising Patterns in a Time Series Database in Linear Time and Space. *ACM KDD Conference*, 2002.
- [311] E. Keogh, J. Lin, and A. Fu. HOT SAX: Finding the Most Unusual Time Series Subsequence: Algorithms and Applications. *ICDM Conference*, 2005.
- [312] E. Keogh, S. Lonardi, and C. Ratanamahatana. Towards Parameter-Free Data Mining. *ACM KDD Conference*, 2004.
- [313] S. Khan and M. Madden. One-class Classification: Taxonomy of Study and Review of Techniques. *Knowledge Engineering Review*, 29(03), 345–374, 2014.

- [314] N. L. D. Khoa and S. Chawla. Incremental Commute Time using Random Walks and Online Anomaly Detection. *ECML/PKDD Conference*, 2016.
- [315] D. Kifer, S. Ben-David, and J. Gehrke. Detecting Change in Data Streams. *VLDB Conference*, 2004.
- [316] J. Kim and C. Scott. Robust Kernel Density Estimation. *Journal of Machine Learning Research*, 13, pp. 2529–2565, 2012. <http://www.jmlr.org/papers/volume13/kim12b/kim12b.pdf>
- [317] E. Knorr and R. Ng. Algorithms for Mining Distance-based Outliers in Large Datasets. *VLDB Conference*, 1998.
- [318] E. Knorr and R. Ng. Finding Intensional Knowledge of Distance-Based Outliers. *VLDB Conference*, 1999.
- [319] E. M. Knorr, R. T. Ng, and V. Tucakov. Distance-based Outliers: Algorithms and Applications. *VLDB Journal*, 8(3), pp. 237–253, February 2000.
- [320] J. Koh, M.-L. Lee, W. Hsu, and W. Ang. Correlation-based Attribute Outlier Detection in XML. *ICDE Conference*, 2008.
- [321] G. Kollios, D. Gunopulos, N. Koudas, and S. Berchtold. Efficient Biased Sampling for Approximate Clustering and Outlier Detection in Large Data Sets. *IEEE Transactions on Knowledge and Data Engineering*, 15(5), pp. 1170–1187, 2003.
- [322] M. Kontaki, A. Gounaris, A. Papadopoulos, K. Tsihlias, and Y. Manolopoulos. Continuous Monitoring of Distance-based Outliers over Data Streams. *ICDE Conference*, 2011.
- [323] K. Kontonasis and T. Bie. An Information-Theoretic Approach to Finding Noisy Tiles in Binary Databases. *SIAM Conference on Data Mining*, 2003.
- [324] Y. Kou, C. T. Lu, and D. Chen. Spatial Weighted Outlier Detection. *SIAM Conference on Data Mining*, 2006.
- [325] H.-P. Kriegel, M. Schubert, and A. Zimek. Angle-based Outlier Detection in High-Dimensional Data. *ACM KDD Conference*, 2008.
- [326] H.-P. Kriegel, P. Kroger, and A. Zimek. Outlier Detection Techniques. *Conference Tutorial at SIAM Data Mining Conference*, 2010. <http://www.siam.org/meetings/sdm10/tutorial3.pdf>
- [327] H.-P. Kriegel, P. Kroger, E. Schubert, and A. Zimek. Outlier Detection in Axis-Parallel Subspaces of High Dimensional Data. *PAKDD Conference*, 2009.
- [328] H.-P. Kriegel, P. Kroger, E. Schubert, and A. Zimek. Outlier Detection in Arbitrarily Oriented Subspaces. *ICDM Conference*, 2012.
- [329] C. Kruegel and G. Vigna. Anomaly-Detection of Web-based Attacks, *CCS*, 2005.
- [330] C. Kruegel, D. Mutz, W. Robertson, and F. Valeur. Bayesian Event Classification for Intrusion Detection. *Computer Security Applications Conference*, 2003.

- [331] C. Kruegel, T. Toth, and E. Kirda. Service Specific Anomaly Detection for Network Intrusion Detection. *ACM Symposium on Applied computing*, 2002.
- [332] M. Kubat and S. Matwin. Addressing the Curse of Imbalanced Training Sets: One Sided Selection. *ICML Conference*, 1997.
- [333] L. Kuncheva. Change Detection in Streaming Multivariate Data using Likelihood Detectors. *IEEE Transactions on Knowledge and Data Engineering*, 25(5), pp. 1175–1180, 2013.
- [334] A. Lakhina, M. Crovella, and C. Diot. Mining Anomalies using Traffic Feature Distributions. *ACM SIGCOMM Conference*, pp. 217–228, 2005.
- [335] A. Lakhina, M. Crovella, and C. Diot. Diagnosing Network-wide Traffic Anomalies. *ACM SIGCOMM Conference*, pp. 219–230, 2004.
- [336] G. Lanckriet, L. Ghaoui, and M. Jordan. Robust Novelty Detection with Single Class MPM. *Neural Information Processing Systems*, 2002.
- [337] J. Lafferty, A. McCallum, and F. Pereira. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. *ICML Conference*, 2001.
- [338] T. Lane and C. Brodley. Temporal Sequence Learning and Data Reduction for Anomaly Detection. *ACM Transactions on Information and Security*, 2(3), pp. 295–331, 1999.
- [339] T. Lane and C. Brodley. An Application of Machine Learning to Anomaly Detection. *NIST-NCSC National Information Systems Security Conference*, 1997.
- [340] T. Lane and C. Brodley. Sequence matching and learning in anomaly detection for computer security. *AI Approaches to Fraud Detection and Risk Management*, pp. 43–49, 1997.
- [341] R. Lasaponara. On the Use of Principal Component Analysis (PCA) for Evaluating Interannual Vegetation Anomalies from SPOT/VEGETATION NDVI Temporal Series. *Ecological Modeling*, 194(4), pp. 429–434, 2006.
- [342] L. Latecki, A. Lazarevic, and D. Pokrajac. Outlier Detection with Kernel Density Functions. *Machine Learning and Data Mining in Pattern Recognition*, pp. 61–75, 2007.
- [343] J. Laurikkala, M. Juhola, and E. Kentala. Informal Identification of Outliers in Medical Data. *Fifth International Workshop on Intelligent Data Analysis in Medicine and Pharmacology*, pp. 20–24, 2000.
- [344] A. Lazarevic and V. Kumar. Feature Bagging for Outlier Detection. *ACM KDD Conference*, 2005.
- [345] A. Lazarevic, L. Ertöz, V. Kumar, A. Ozgur, and J. Srivastava. A Comparative Study of Anomaly Detection Schemes in Network Intrusion Detection. *SIAM Conference on Data Mining*, 2003.
- [346] S. Q. Le and T. B. Ho. An Association-based Dissimilarity Measure for Categorical Data. *Pattern Recognition Letters*, 26(16), pp. 2549–2557, 2005.

- [347] J.-G. Lee, J. Han, and X. Li. Trajectory Outlier Detection: A Partition-and-detect Framework. *ICDE Conference*, 2008.
- [348] W. Lee and B. Liu. Learning with Positive and Unlabeled Examples using Weighted Logistic Regression. *ICML Conference*, 2003.
- [349] W. Lee, S. Stolfo, and P. Chan. Learning Patterns from Unix Execution Traces for Intrusion Detection. *AAAI Workshop on AI methods in Fraud and Risk Management*, 1997.
- [350] W. Lee, S. Stolfo, and K. Mok. Adaptive Intrusion Detection: A Data Mining Approach. *Artificial Intelligence Review*, 14(6), pp. 533–567, 2000.
- [351] W. Lee and S. Stolfo. Data Mining Approaches for Intrusion Selection. *USENIX Security Symposium*, 1998.
- [352] W. Lee and D. Xiang. Information Theoretic Measures for Anomaly Detection. *IEEE Symposium on Security and Privacy*, 2001.
- [353] N. Lesh, M. J. Zaki, and M. Ogihara. Mining Features for Sequence Classification. *ACM KDD Conference*, 1999.
- [354] C. Leslie, E. Eskin, and W. Noble. The Spectrum Kernel: A String Kernel for SVM Protein Classification. *Pacific Symposium on Biocomputing*, pp. 566–575, 2002.
- [355] X. Li, J. Han, S. Kim, and H. Gonzalez. ROAM: Rule and Motif-based Anomaly Detection in Massive Moving Object Data Sets. *SIAM Conference on Data Mining*, 2007.
- [356] X. Li, B. Liu, and S. Ng. Negative Training Data can be Harmful to Text Classification. *EMNLP*, 2010.
- [357] X. Li, Z. Li, J. Han, and J.-G. Lee. Temporal Outlier Detection in Vehicle Traffic Data. *ICDE Conference*, 2009.
- [358] D. Lin. An Information-theoretic Definition of Similarity. *ICML Conference*, pp. 296–304, 1998.
- [359] Y. Lin and Y. Jeon. Random Forests and Adaptive Nearest Neighbors. *Journal of the American Statistical Association*, 101(474), pp. 578–590, 2006.
- [360] J. Lin, E. Keogh, A. Fu, and H. V. Herle. Approximations to Magic: Finding Unusual Medical Time Series. *Mining Medical Data (CBMS)*, 2005.
- [361] J. Lin, E. Keogh, S. Lonardi, and B. Y.-C. Chiu. A Symbolic Representation of Time Series, with Implications for Streaming Algorithms. *DMKD Workshop*, 2003.
- [362] B. Liu, W. S. Lee, P. Yu, and X. Li. Partially Supervised Text Classification. *ICML Conference*, 2002.
- [363] B. Liu, Y. Dai, X. Li, W. S. Lee, and P. Yu. Building Text Classifiers Using Positive and Unlabeled Examples. *ICDM Conference*, 2003.
- [364] B. Liu, W. Hsu, and Y. Ma. Integrating Classification and Association Rule Mining. *ACM KDD Conference*, pp. 80–86, 1998.

- [365] G. Liu, T. McDaniel, S. Falkow, and S. Karlin. Sequence Anomalies in the *cag7* Gene of the *Helicobacter Pylori* Pathogenicity Island. *National Academy of Sciences of the United States of America*, 96(12), pp. 7011–7016, 1999.
- [366] L. Liu and X. Fern. Constructing Training Sets for Outlier Detection. *SIAM Conference on Data Mining*, 2012.
- [367] F. T. Liu, K. M. Ting, and Z.-H. Zhou. Isolation Forest. *ICDM Conference*, 2008.
- [368] F. T. Liu, K. M. Ting, and Z. H. Zhou. Isolation-based Anomaly Detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1), 3, 2012.
- [369] F. T. Liu, K. M. Ting, and Z. H. Zhou. On Detecting Clustered Anomalies using SCiForest. *Machine Learning and Knowledge Discovery in Databases*, pp. 274–290, 2010.
- [370] S. Lin and D. Brown. An Outlier-based Data Association Method for Linking Criminal Incidents. *SIAM Conference On Data Mining*, 2003.
- [371] X.-Y. Liu, J. Wu, and Z.-H. Zhou. Exploratory Undersampling for Class-Imbalance Learning. *IEEE Transactions on Systems, Man and Cybernetics– Part B, Cybernetics*, 39(2), pp. 539–550, April 2009.
- [372] X. Liu, X. Wu, H. Wang, R. Zhang, J. Bailey, and K. Ramamohanarao. Mining Distribution Change in Stock Order Data Streams. *ICDE Conference*, 2010.
- [373] Z. Liu, W. Shi, D. Li, and Q. Qin. Partially Supervised Classification – based on Weighted Unlabeled Samples Support Vector Machine. *ADMA*, 2005.
- [374] X. Liu, P. Zhang, and D. Zeng. Sequence Matching for Suspicious activity Detection in Anti-money Laundering. *Lecture Notes in Computer Science*, Vol. 5075, pp. 50–61, 2008.
- [375] S. Loncarin. A Survey of Shape Analysis Techniques. *Pattern Recognition*, 31(5), pp. 983–1001, 1998.
- [376] C.-T. Lu, D. Chen, and Y. Kou. Algorithms for Spatial Outlier Detection. *ICDM Conference*, 2003.
- [377] A. Lung-Yut-Fong, C. Levy-Leduc, and O. Cappe. Distributed Detection/localization of Change-points in High-dimensional Network Traffic Data. *arXiv:0909.5524*, 2009.
- [378] U. von Luxburg. A Tutorial on Spectral Clustering. *Statistics and computing*, 17(4), pp. 395–416, 2007.
- [379] J. Ma and S. Perkins. Online Novelty Detection on Temporal Sequences. *ACM KDD Conference*, 2003.
- [380] J. Ma, L. Saul, S. Savage, and G. Volker. Learning to Detect Malicious URLs. *ACM Transactions on Intelligent Systems and Technology*, 2(3), Article 30, April 2011.
- [381] M. Mahoney and P. Chan. Learning Nonstationary Models of Normal Network Traffic for Detecting Novel Attacks. *ACM KDD Conference*, 2002.

- [382] M. Mahoney and P. Chan. Learning Rules for Anomaly Detection of Hostile Network Traffic. *ICDM Conference*, 2003.
- [383] F. Malliaros, V. Megalooikonomou, and C. Faloutsos. Fast Robustness Estimation in Large Social Graphs: Communities and Anomaly Detection. *SIAM Conference on Data Mining*, 2012.
- [384] L. M. Manevitz and M. Yousef. One-class SVMs for Document Classification. *Journal of Machine Learning Research*, 2: pp. 139–154, 2001.
- [385] L. Manevitz and M. Yousef. One-class Document Classification via Neural Networks. *Neurocomputing*, 70(7), pp. 1466–1481, 2007.
- [386] E. Manzoor, S. Milajerdi, and L. Akoglu. Fast Memory-Efficient Anomaly Detection in Streaming Heterogeneous Graphs. *ACM KDD Conference*, 2016.
- [387] C. Marceau. Characterizing the Behavior of a Program using Multiple-length n-grams. *Workshop on New Security Paradigms*, pp. 101–110, 2000.
- [388] M. Markou and S. Singh. Novelty detection: A Review, Part 1: Statistical Approaches. *Signal Processing*, 83(12), pp. 2481–2497, 2003.
- [389] M. Markou and S. Singh. Novelty Detection: A Review, Part 2: Neural Network-based Approaches. *Signal Processing*, 83(12), pp. 2481–2497, 2003.
- [390] M. Masud, Q. Chen, L. Khan, C. Aggarwal, J. Gao, J. Han, and B. Thuraisingham. Addressing Concept-Evolution in Concept-Drifting Data Streams. *ICDM Conference*, 2010.
- [391] M. Masud, T. Al-Khateeb, L. Khan, C. Aggarwal, J. Gao, J. Han, and B. Thuraisingham. Detecting Recurring and Novel Classes in Concept-Drifting Data Streams. *ICDM Conference*, 2011.
- [392] M. Masud, Q. Chen, L. Khan, C. Aggarwal, J. Gao, J. Han, A. Srivastava, and N. Oza. Classification and Adaptive Novel Class Detection of Feature-Evolving Data Streams. *IEEE Transactions on Knowledge and Data Engineering*, 25(7), pp. 1484–1497, 2013.
- [393] C. McDiarmid. On the Method of Bounded Differences. In *Surveys in Combinatorics*, pp. 148–188, *Cambridge University Press*, Cambridge, 1989.
- [394] P. Melville and R. Mooney. Diverse Ensembles for Active Learning. *ICML Conference*, 2004.
- [395] B. Micenkova, B. McWilliams, and I. Assent. Learning Outlier Ensembles: The Best of Both Worlds – Supervised and Unsupervised. *Outlier Detection and Description Workshop*, 2014. Extended version: <http://arxiv.org/abs/1507.08104>
- [396] C. Michael and A. Ghosh. Two State-based Approaches to Program-based Anomaly Detection. *Computer Security Applications Conference*, pp. 21, 2000.
- [397] B. Miller, N. Bliss, and P. Wolfe. Subgraph Detection using Eigenvector L1-Norms. *Neural Information Processing Systems*, 2010.
- [398] B. Miller, M. Beard, and N. Bliss. Eigenspace Analysis for Threat Detection in Social Networks. *International Conference on Information Fusion*, 2011.

- [399] D. Mladenic and M. Grobelnik. Feature Selection for Unbalanced Class Distribution and Naive Bayes. *ICML Conference*, 1999.
- [400] M. Mongiovi, P. Bogdanov, R. Ranca, A. Singh, E. Papalexakis, and C. Faloutsos. SIGSPOT: Mining Significant Anomalous Regions from Time-evolving Networks. *ACM SIGMOD Conference*, 2012.
- [401] F. Moosmann, B. Triggs, and F. Jurie. Fast Discriminative Visual Codebooks using Randomized Clustering Forests. *Neural Information Processing Systems*, pp. 985–992, 2006.
- [402] E. Muller, M. Schiffer, and T. Seidl. Statistical Selection of Relevant Subspace Projections for Outlier Ranking. *ICDE Conference*, pp. 434–445, 2011.
- [403] E. Muller, M. Schiffer, P. Gerwert, M. Hannen, T. Jansen, and T. Seidl. SOREX: Subspace Outlier Ranking Exploration Toolkit. *Joint ECML PKDD Conference*, 2010.
- [404] E. Muller, M. Schiffer, and T. Seidl. Adaptive Outlierness for Subspace Outlier Ranking. *CIKM Conference*, 2010.
- [405] E. Muller, F. Keller, S. Blanc, and K. Bohm. OutRules: A Framework for Outlier Descriptions in Multiple Context Spaces. *ECML/PKDD Conference*, 2012.
- [406] E. Muller, I. Assent, P. Iglesias, Y. Mülle, and K. Bohm. Outlier Analysis via Subspace Analysis in Multiple Views of the Data. *ICDM Conference*, 2012.
- [407] R. Motwani and P. Raghavan. Randomized Algorithms. *Cambridge University Press*, 1995.
- [408] A. Mueen, E. Keogh, and N. Young. Logical-Shapelets: An Expressive Primitive for Time Series Classification. *ACM KDD Conference*, 2011.
- [409] A. Naftel and S. Khalid. Classifying Spatiotemporal Object Trajectories using Unsupervised Learning in the Coefficient Feature Space. *Multimedia Systems*, 12(3), pp. 227–238, 2006.
- [410] K. Narita and H. Kitagawa. Outlier Detection for Transaction Databases using Association Rules. *WAIM*, 2008.
- [411] H. Nguyen, V. Gopalkrishnan, and I. Assent. An Unbiased Distance-based Outlier Detection Approach for High Dimensional Data. *DASFAA*, 2011.
- [412] H. Nguyen, H. Ang, and V. Gopalakrishnan. Mining Ensembles of Heterogeneous Detectors on Random Subspaces. *DASFAA*, 2010.
- [413] H. Nguyen, E. Muller, J. Vreeken, F. Keller, and K. Bohm. CMI: An Information-Theoretic Contrast Measure for Enhancing Subspace Cluster and Outlier Detection. *SIAM Conference on Data Mining (SDM)*, pp. 198–206, 2013.
- [414] V. Niennattrakul, E. Keogh, and C. Ratanamahatana. Data Editing Techniques to Allow the Applicability of Distance-based Outlier Detection in Streams. *ICDM Conference*, 2010.

- [415] H. Ning, W. Xu, Y. Chi, Y. Gong, and T. Huang. Incremental Spectral Clustering With Application to Monitoring of Evolving Blog Communities. *SIAM Conference on Data Mining*, 2007.
- [416] C. Noble, and D. Cook. Graph-based Anomaly Detection. *ACM KDD Conference*, 2003.
- [417] K. Noto, C. Brodley, and D. Slonim. FRaC: A Feature-Modeling Approach for Semi-Supervised and Unsupervised Anomaly Detection. *Data Mining and Knowledge Discovery*, 25(1), pp. 109–133, 2012.
- [418] P. Olmo Vaz de Melo, L. Akoglu, C. Faloutsos, and A. Loureiro. Surprising Patterns for the Call Duration Distribution of Mobile Phone Users. *ECML/PKDD Conference*, 2010.
- [419] G. Orair, C. Teixeira, W. Meira Jr, Y. Wang, and S. Parthasarathy. Distance-Based Outlier Detection: Consolidation and Renewed Bearing. *Proceedings of the VLDB Endowment*, 3(1–2), pp. 1469–1480, 2010.
- [420] M. Otey, S. Parthasarathy, A. Ghoting, G. Li, S. Narravula, and D. Panda. Towards NIC-based Intrusion Detection. *ACM KDD Conference*, 2003.
- [421] M. Otey, S. Parthasarathy, and A. Ghoting. Fast Distributed Outlier Detection in Mixed Attribute Data Sets. *Data Mining and Knowledge Discovery*, 12(2–3), pp. 203–228, 2006.
- [422] L. Ott, L. Pang, F. Ramos, and S. Chawla. On Integrated Clustering and Outlier Detection. *Neural Information Processing Systems*, pp. 1359–1367, 2014.
- [423] C. R. Palmer and C. Faloutsos. Electricity based External Similarity of Categorical Attributes. *PAKDD Conference*, 2003.
- [424] Y. Panatier. Variowin. Software For Spatial Data Analysis in 2D. *New York: Springer-Verlag*, 1996.
- [425] C. Papadimitriou, P. Raghavan, H. Tamakai, and S. Vempala. Latent Semantic Indexing: A Probabilistic Analysis. *ACM PODS Conference*, 1998.
- [426] S. Papadimitriou, H. Kitagawa, P. Gibbons, and C. Faloutsos. LOCI: Fast Outlier Detection using the Local Correlation Integral. *ICDE Conference*, 2003.
- [427] S. Papadimitriou, J. Sun, and C. Faloutsos. SPIRIT: Streaming Pattern Discovery in Multiple Time-Series. *VLDB Conference*, 2005.
- [428] L. Parra, G. Deco, and S. Andmiesbach. Statistical Independence and Novelty Detection with Information Preserving Nonlinear Maps. *Neural Computation*, 8(2), pp. 260–269, 1996.
- [429] H. Paulheim and R. Meusel. A Decomposition of the Outlier Detection Problem into a Set of Supervised Learning Problems. *Machine Learning*, 100(2–3), pp. 509–531, 2015.
- [430] Y. Pei, O. Zaiane, and Y. Gao. An Efficient Reference-based Approach to Outlier Detection in Large Datasets. *ICDM Conference*, 2006.

- [431] D. Pelleg and A. Moore. Active Learning for Anomaly and Rare Category Detection. *Neural Information Processing Systems*, 2004.
- [432] Z. Peng and F. Chu. Review Application of the Wavelet Transform in Machine Condition Monitoring and Fault Diagnostics. *Mechanical Systems and Signal Processing*, 18(2), pp. 199–221, March 2004.
- [433] B. Perozzi, L. Akoglu, P. Iglesias Sanchez, and E. Muller. Focused Clustering and Outlier Detection in Large Attributed Graphs. *ACM KDD Conference*, 2014.
- [434] B. Perozzi and L. Akoglu. Scalable Anomaly Ranking of Attributed Neighborhoods. *SIAM Conference on Data Mining*, 2016.
- [435] S. Petrovic, M. Osborne, and V. Lavrenko. Streaming First Story Detection with Application to Twitter. *ACL Conference*, pp. 181–189, 2010.
- [436] T. Pevny. Loda: Lightweight On-line Detector of Anomalies. *Machine Learning*, 102(2), pp. 275–304, 2016.
- [437] J. Pickands. Statistical Inference using Extreme Order Statistics. *The Annals of Statistics*, 3(1), pp. 119–131, 1975.
- [438] B. Pincombe. Anomaly Detection in Time Series of Graphs using ARMA Processes. *ASOR Bulletin*, 24(4): 2–10, 2005.
- [439] M. Pinsky. Introduction to Fourier Analysis and Wavelets. *American Mathematical Society*, 2009.
- [440] C. Phua, V. Lee, K. Smith, and R. Gayler. A Comprehensive Survey of Data Mining-based Fraud Detection Research, 2010.
<http://arxiv.org/abs/1009.6119>.
- [441] C. Phua, D. Alahakoon, and V. Lee. Minority Report in Fraud Detection: Classification of Skewed Data. *ACM SIGKDD Explorations Newsletter*, 6(1), pp. 50–59, 2004.
- [442] B. R. Preiss. Data Structures and Algorithms with Object Oriented Design Patterns in Java. *Wiley*, 1999.
- [443] D. Pokrajac, A. Lazarevic, and L. Latecki. Incremental Local Outlier Detection for Data Streams. *CIDM Conference*, 2007.
- [444] C. Potter, P. N. Tan, M. Steinbach, S. Klooster, V. Kumar, R. Myneni, and V. Genovese. Major Disturbance Events in Terrestrial Ecosystems detected using Global Satellite Data Sets. *Global Change Biology*, pp. 1005–1021, 2003.
- [445] A. Pires and C. Santos-Pereira. Using Clustering and Robust Estimators to Detect Outliers in Multivariate Data. *International Conference on Robust Statistics*, 2005.
- [446] L. Portnoy, E. Eskin, and S. Stolfo. Intrusion Detection with Unlabeled Data using Clustering. *ACM Workshop on Data Mining Applied to Security*, 2001.
- [447] B. Prakash, N. Valler, D. Andersen, M. Faloutsos, and C. Faloutsos. BGP-lens: Patterns and Anomalies in Internet Routing Updates. *ACM KDD Conference*, 2009.

- [448] M. Prastawa, E. Bullitt, S. Ho, and G. Gerig. A Brain Tumor Segmentation Framework based on Outlier Detection. *Medical Image Analysis*, 8, pp. 275–283, 2004.
- [449] C. Priebe, J. Conroy, D. Marchette, and Y. Park. Scan Statistics on Enron Graphs. *Computational and Mathematical Organizational Theory*, 11(3), pp. 229–247, 2005.
- [450] F. Provost, and T. Fawcett. Analysis and Visualization of Classifier Performance: Comparison under Imprecise Class and Cost Distributions. *KDD Conference*, 1997.
- [451] F. Provost, T. Fawcett, and R. Kohavi. The Case against Accuracy Estimation while Comparing Induction Algorithms. *ICML Conference*, 1998.
- [452] G. Qi, C. Aggarwal, and T. Huang. On Clustering Heterogeneous Social Media Objects with Outlier Links. *WSIAM Conference on Data Mining*, 2012.
- [453] J. Quinlan. Learning with Continuous Classes. *Australian Joint Conference on Artificial Intelligence*, 92, pp. 343–348, 1992.
- [454] L. R. Rabiner. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE*, 77(2), pp. 257–285, Feb. 1989.
- [455] M. Radovanovic, A. Nanopoulos, and M. Ivanovic. On the Existence of Obstinate Results in Vector Space Models. *ACM SIGIR Conference*, 2010.
- [456] S. Ramaswamy, R. Rastogi, and K. Shim. Efficient Algorithms for Mining Outliers from Large Data Sets. *ACM SIGMOD Conference*, pp. 427–438, 2000.
- [457] S. Ranshous, S. Shen, D. Koutra, S. Harenberg, C. Faloutsos, and N. Samatova. Anomaly Detection in Dynamic Networks: A Survey. *Wiley Interdisciplinary Reviews: Computational Statistics*, 7(3), pp. 223–247, 2015.
- [458] S. Ranshous, S. Harenberg, K. Sharma, and N. Samatova. A Scalable Approach for Outlier Detection in Edge Streams Using Sketch-based Approximations. *SIAM Conference on Data Mining*, 2016.
- [459] M. Ranzato, C. Poultney, S. Chopra, and Y. LeCun. Efficient Learning of Sparse Representations with an Energy-Based Model. *Neural Information Processing Systems*, 19 pp. 11371144, 2006
- [460] B. Raskutti and A. Kowalczyk. Extreme Rebalancing for SVMs: A Case Study. *SIGKDD Explorations*, 6(1): pp. 60–69, 2004.
- [461] S. Rayana and L. Akoglu. Less is More: Building Selective Anomaly Ensembles. *ACM Transactions on Knowledge Discovery and Data Mining*, 10(4), 42, 2016.
- [462] S. Roberts. Novelty Detection using Extreme Value Statistics. *IEEE Proceedings on Vision, Image and Signal Processing*, 146(3). pp. 124–129, 1999.
- [463] S. Roberts. Extreme Value Statistics for Novelty Detection in Biomedical Signal Processing. *International Conference on Advances in Medical Signal and Information Processing*. pp. 166–172, 2002.
- [464] J. Rogan, J. Miller, D. Stow, J. Franklin, L. Levien, and C. Fischer. Land-Cover Change Monitoring with Classification Trees Using Landsat TM and Ancillary Data. *Photogrammetric Engineering and Remote Sensing*, 69(7), pp. 793–804, 2003.

- [465] D. Ron, Y. Singer, and N. Tishby. The Power of Amnesia: Learning Probabilistic Automata with Variable Memory Length. *Machine Learning*, 25(2–3) pp. 117–149, 1996.
- [466] V. Roth. Kernel Fisher Discriminants for Outlier Detection. *Neural Computation*, 18(4), pp. 942–960, 2006.
- [467] P. Rousseeuw and A. Leroy. Robust Regression and Outlier Detection. *Wiley*, 2003.
- [468] I. Ruts and P. Rousseeuw. Computing Depth Contours of Bivariate Point Clouds. *Computational Statistics and Data Analysis*, 23, pp. 153–168, 1996.
- [469] M. Sahami, S. Dumais, D. Heckerman, and E. Horvitz. A Bayesian Approach to Filtering Junk E-mail. *AAAI Workshop on Learning for Text Categorization. Tech. Rep. WS-98-05*, 1998.
<http://robotics.stanford.edu/users/sahami/papers-dir/spam.pdf>
- [470] R. K. Sahoo, A. J. Oliner, I. Rish, M. Gupta, J. E. Moreira, S. Ma, R. Vilalta, and A. Sivasubramaniam. Critical Event Prediction for Proactive Management in Large-scale Computer Clusters. *ACM KDD Conference*, 2003.
- [471] M. Salehi, C. Leckie, M. Moshtaghi, and T. Vaithianathan. A Relevance Weighted Ensemble Model for Anomaly Detection in Switching Data Streams. *Advances in Knowledge Discovery and Data Mining*, pp. 461–473, 2014.
- [472] G. Salton and M. J. McGill. Introduction to Modern Information Retrieval. *McGraw Hill*, 1986.
- [473] P. Sanchez, E. Muller, F. Laforet, F. Keller, and K. Bohm. Statistical Selection of Congruent Subspaces for Mining Attributed Graphs. *ICDM Conference*, pp. 647–656, 2013.
- [474] S. Sarawagi, R. Agrawal, and N. Megiddo. Discovery-driven Exploration of OLAP Data Cubes. *EDBT Conference*, 1998.
- [475] S. Sathe and C. Aggarwal. LODS: Local Density Meets Spectral Outlier Detection. *SIAM Conference on Data Mining*, 2016.
- [476] S. Sathe and C. Aggarwal. Subspace Outlier Detection in Linear Time with Randomized Hashing. *ICDM Conference*, 2016.
- [477] G. Scarth, M. McIntyre, B. Wowk, and R. Somorjai. Detection of Novelty in Functional Images using Fuzzy Clustering. *Meeting of International Society for Magnetic Resonance in Medicine*, 1995.
- [478] R. Schapire and Y. Singer. Improved Boosting Algorithms using Confidence-rated Predictions. *Annual Conference on Computational Learning Theory*, 1998.
- [479] B. Scholkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the Support of a High-Dimensional Distribution. *Neural Computation*, 13(7), pp. 1443–1472, 2001.
- [480] B. Scholkopf, R. C. Williamson, A. J. Smola, J. Shawe-Taylor, and J. C. Platt. Support-vector Method for Novelty Detection. *Neural Information Processing Systems*, 2000.

- [481] B. Scholkopf, A. Smola, and K.-R. Muller. Nonlinear Component Analysis as a Kernel Eigenvalue Problem. *Neural Computation*, 10(5), pp. 1299–1319, 1998.
- [482] R. Schoen, T. Habetler, F. Kamran, and R. Bartfield. Motor Bearing Damage Detection using Stator Current Monitoring. *IEEE Transactions on Industry Applications*, 31(6), pp. 1275–1279, 1995.
- [483] E. Schubert, A. Zimek, and H.-P. Kriegel. Generalized Outlier Detection with Flexible Kernel Density Estimates. *SIAM Conference on Data Mining*, pp. 542–50, 2014.
- [484] K. Sequeira and M. Zaki. ADMIT: Anomaly-based Data Mining for Intrusions. *ACM KDD Conference*, 2002.
- [485] H. Seung, M. Opper, and H. Sompolinsky. Query by Committee. *ACM Workshop on Computational Learning Theory*, 1992.
- [486] B. Settles. Active Learning. *Morgan and Claypool*, 2012.
- [487] S. Shekhar, C. T. Lu, and P. Zhang. Detecting Graph-based Spatial Outliers: Algorithms and Applications. *ACM KDD Conference*, 2001.
- [488] S. Shekhar, C. T. Lu, and P. Zhang. A Unified Approach to Detecting Spatial Outliers. *Geoinformatica*, 7(2), pp. 139–166, 2003.
- [489] S. Shekhar and S. Chawla. A Tour of Spatial Databases. *Prentice Hall*, 2002.
- [490] S. Shekhar, C. T. Lu, and P. Zhang. Detecting Graph-based Spatial Outliers. *Intelligent Data Analysis*, 6, pp. 451–468, 2002.
- [491] T. Shi and S. Horvath. Unsupervised Learning with Random Forest Predictors. *Journal of Computational and Graphical Statistics*, 15(1), pp. 118–138, 2006.
- [492] P. Showbridge, M. Kraetzl, and D. Ray. Detection of Abnormal Change in Dynamic Networks. *International Conference on Information, Decision and Control*, pp. 557–562, 1999.
- [493] M.-L. Shyu, S.-C. Chen, K. Sarinnapakorn, and L. Chang. A Novel Anomaly Detection Scheme based on Principal Component Classifier. *ICDM Conference*, 2003.
- [494] A. Siebes, J. Vreeken, and M. van Leeuwen. Itemsets than Compress. *SIAM Conference on Data Mining*, 2006.
- [495] J. Silva and R. Willett. Detection of Anomalous Meetings in a Social Network. *SocialCom*, 2008.
- [496] B. W. Silverman. Density Estimation for Statistics and Data Analysis. *Chapman and Hall*, 1986.
- [497] K. Smets and J. Vreeken. The Odd One Out: Identifying and Characterising Anomalies. *SIAM Conference on Data Mining*, 2011.
- [498] E. S. Smirnov. On Exact Methods in Systematics. *Systematic Zoology*, 17(1), pp. 1–13, 1968.

- [499] R. Smith, A. Bivens, M. Embrechts, C. Palagiri, and B. Szymanski. Clustering Approaches for Anomaly Based Intrusion Detection. *Intelligent Engineering Systems through Artificial Neural Networks*, 2002.
- [500] P. Smyth. Clustering Sequences with Hidden Markov Models. *Neural Information Processing*, 1997.
- [501] P. Smyth. Markov Monitoring with Unknown States. *IEEE Journal on Selected Areas in Communications*, 12(9), pp. 1600-1612, 1994.
- [502] H. Solberg and A. Lahti. Detection of Outliers in Reference Distributions: Performance of Horn's Algorithm. *Clinical Chemistry*, 51(12), pp. 2326-2332, 2005.
- [503] X. Song, M. Wu, C. Jermaine, and S. Ranka. Conditional Anomaly Detection. *IEEE Transaction on Knowledge and Data Engineering*, 19(5), pp. 631-645, 2007.
- [504] X. Song, M. Wu, C. Jermaine, and S. Ranka. Statistical Change Detection for Multi-dimensional Data. *ACM KDD Conference*, 2007.
- [505] C. Spence, L. Parra, and P. Sajda. Detection, Synthesis and Compression in Mammographic Image Analysis with a Hierarchical Image Probability Model. *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, 2001.
- [506] A. Srivastava. Discovering System Health Anomalies using Data Mining Techniques. *Joint Army Navy NASA Airforce Conference on Propulsion*, 2005.
- [507] A. Srivastava. Enabling the Discovery of Recurring Anomalies in Aerospace Problem Reports using High-dimensional Clustering Techniques. *Aerospace Conference*, 2006.
- [508] A. Srivastava and B. Zane-Ulman. Discovering Recurring Anomalies in Text Reports regarding Complex Space Systems. *Aerospace Conference*, 2005.
- [509] A. Srivastava, A. Kundu, S. Sural, and A. Majumdar. Credit Card Fraud Detection using Hidden Markov Model. *IEEE Transactions on Dependable and Secure Computing*, 5(1), pp. 37-48, 2008.
- [510] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(1), pp. 1929-1958, 2014.
- [511] P. Srivastava, D. Desai, S. Nandi, and A. Lynn. HMM-ModE-Improved Classification using Profile Hidden Markov Models by Optimizing the Discrimination Threshold and Modifying Emission Probabilities with Negative Training Sequences. *BMC Bioinformatics*, 8 (104), 2007.
- [512] M. Stephens. Use of the Kolmogorov-Smirnov, Cramer-von Mises and Related Statistics without Extensive Tables. *Journal of the Royal Statistical Society. Series B*, pp. 115-122, 1970.
- [513] S. Stolfo, D. Fan, W. Lee, A.L. Prodromidis, and P. Chan. Credit Card Fraud Detection Using Meta-Learning: Issues and Initial Results. *AAAI Workshop AI Methods in Fraud and Risk Management*, pp. 83-90, 1997.

- [514] S. Stolfo, D. Fan, W. Lee, A. Prodromidis, and P. Chan. Cost-Based Modeling for Fraud and Intrusion Detection: Results from the JAM Project. *DARPA Information Survivability Conf. and Exposition*, 2, pp. 130–144, 2000.
- [515] G. Strang. An Introduction to Linear Algebra. *Wellesley Cambridge Press*, 2009.
- [516] A. Strehl and J. Ghosh. Cluster Ensembles– A Knowledge Reuse Framework for Combining Multiple Partitions. *Jour. of Machine Learning Research*, 3, pp. 583–617, 2003.
- [517] S. Subramaniam, T. Palpanas, D. Papadopoulos, V. Kalogeraki, and D. Gunopulos. Online Outlier Detection in Sensor Data using Non-parametric Models. *VLDB Conference*, 2006.
- [518] H. Sun, Y. Bao., F. Zhao, G. Yu, and D. Wang. CD-Trees: An Efficient Index Structure for Outlier Detection. *Web-Age Information Management (WAIM)*, pp. 600–609, 2004.
- [519] J. Sun, S. Papadimitriou, P. Yu, and C. Faloutsos. Graphscope: Parameter-free Mining of Large Time-Evolving Graphs. *ACM KDD Conference*, 2007.
- [520] J. Sun, D. Tao, and C. Faloutsos. Beyond Streams and Graphs: Dynamic Tensor Analysis. *ACM KDD Conference*, 2006.
- [521] J. Sun, H. Qu, D. Chakrabarti, and C. Faloutsos. Neighborhood Formation and Anomaly Detection in Bipartite Graphs. *ICDM Conference*, 2005.
- [522] J. Sun, Y. Xie, H. Zhang, and C. Faloutsos. Less is More: Compact Matrix Representation of Large Sparse Graphs. *SIAM Conference on Data Mining*, 2007.
- [523] P. Sun and S. Chawla. On Local Spatial Outliers. *IEEE ICDM Conference*, 2004.
- [524] P. Sun, S. Chawla, and B. Arunasalam. Mining for Outliers in Sequential Databases. *SIAM Conference on Data Mining*, 2006.
- [525] Y. Sun, Y. Yu, and J. Han. Ranking-based Clustering of Heterogeneous Information Networks with Star Network Schema. *ACM KDD Conference*, 2009.
- [526] C. Surace and K. Worden. A Novelty Detection Method to Diagnose Damage in Structures: An Application to an Offshore Platform. *International Conference of Offshore and Polar Engineering*, 4, pp. 64–70, 1998.
- [527] C. Surace, K. Worden, and G. Tomlinson. A Novelty Detection Approach to Diagnose Damage in a Cracked Beam. *Proceedings of SPIE*, Vol. 3089, pp. 947–953, 1997.
- [528] E. Suzuki, T. Watanabe, H. Yokoi, and K. Takabayashi. Detecting Interesting Exceptions from Medical Test Data with Visual Summarization. *ICDM Conference*, pp. 315–322, 2003.
- [529] P. Sykacek. Equivalent Error Bars for Neural Network Classifiers trained by Bayesian Inference. *ESANN*, 1997.
- [530] T. Takahashi, R. Tomioka, and K. Yamanishi. Discovering Emerging Topics in Social Streams via Link Anomaly Detection. *ICDM Conference*, 2011.
- [531] P.-N. Tan and V. Kumar. Discovery of Web Robot Sessions based on their Navigational Patterns. *Data Mining and Knowledge Discovery*, 6(1), pp. 9–35, 2002.

- [532] S. C. Tan, K. M. Ting, and T. F. Liu. Fast Anomaly Detection for Streaming Data. *IJCAI Conference*, 2011.
- [533] Y. Tao, X. Xiao, and S. Zhou. Mining Distance-based Outliers from Large Databases in any Metric Space. *ACM KDD Conference*, 2006.
- [534] J. Tang, Z. Chen, A. W.-C. Fu, and D. W. Cheung. Enhancing Effectiveness of Outlier Detections for Low Density Patterns. *PAKDD Conference*, 2002.
- [535] Y. Tang, Y.-Q. Zhang, N. V. Chawla, and S. Krasser. SVMs Modeling for Highly Imbalanced Classification. *IEEE Transactions on Systems, Man and Cybernetics- Part B: Cybernetics*, 39(1), pp. 281–288, 2009.
- [536] M. Taniguchi, M. Haft, J. Hollmen, and V. Tresp. Fraud Detection in Communications Networks using Neural and Probabilistic Methods. *IEEE International Conference in Acoustics, Speech and Signal Processing*, 2, pp. 1241–1444, 1998.
- [537] L. Tarassenko. Novelty detection for the Identification of Masses in Mammograms. *IEEE International Conference on Artificial Neural Networks*, 4, pp. 442–447, 1995.
- [538] D. Tax. One Class Classification: Concept-learning in the Absence of Counter-examples, *Doctoral Dissertation, University of Delft, Netherlands*, 2001. <http://prlab.tudelft.nl/sites/default/files/thesis.pdf>
- [539] D. Tax and R. Duin. Support Vector Data Description. *Machine learning*, 54(1), 45-66, 2004.
- [540] D. Tax and R. Duin. Combining One-Class Classifiers. *Multiple Classifier Systems*, pp. 299–308, 2001.
- [541] D. Tax and P. Juszczak. Kernel Whitening for One-Class Classification. *Pattern Recognition with Support Vector Machines*, pp. 40–52, 2002.
- [542] J. Tenenbaum, V. De Silva, and J. C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290, 5500, pp. 2319–2323, 2000.
- [543] P. Thompson, D. MacDonald, M. Mega, C. Holmes, A. Evans, and A. Toga. Detection and Mapping of Abnormal Brain Structure with a Probabilistic Atlas of Cortical Surfaces. *Journal of Computer Assisted Tomography*, 21(4), pp. 567–581, 1997.
- [544] M. Thottan and C. Ji. Anomaly Detection in IP Networks. *IEEE Transactions on Signal Processing*, 51(8), pp. 2191–2204, 2003.
- [545] S. Tian, S. Mu, and C. Yin. Sequence-similarity Kernels for SVMs to Detect Anomalies in System Calls. *Neurocomputing*, 70(4–6), pp. 859–866, 2007.
- [546] K. M. Ting. An Instance-Weighting Method to Induce Cost-sensitive Trees. *IEEE Transaction on Knowledge and Data Engineering*, 14: pp. 659–665, 2002.
- [547] K. M. Ting, G. T. Zhou, F. T. Liu, and S. C. Tan. Mass Estimation. *Machine Learning*, 90(1), pp. 127–160, 2013. [Short version in ACM KDD Conference, 2010]
- [548] K. M. Ting, Y. Zhu, M. Carman, and Y. Zhu. Overcoming Key Weaknesses of Distance-Based Neighbourhood Methods using a Data Dependent Dissimilarity Measure. *ACM KDD Conference*, 2016.

- [549] M. E. Tipping and C. M. Bishop. Probabilistic Principal Component Analysis. *Journal of the Royal Statistical Society*, B 61, pp. 611–622, 1999.
- [550] W. Tobler. Cellular Geography. In *Philosophy in Geography*, Dordrecht Reidel Publishing Company, pp. 379–386, 1979.
- [551] H. Tong and C.-Y. Lin. Non-Negative Residual Matrix Factorization with Application to Graph Anomaly Detection. *SIAM Conference on Data Mining*, 2011.
- [552] L. Toth and G. Gosztolya. Replicator Neural Networks for Outlier Modeling in Segmental Speech Recognition. *LNCS*, Vol. 3172, pp. 996–1001, 2004.
- [553] K. van Leemput, F. Maes, D. Vandermeulen, A. Colchester, and P. Suetens. Automated Segmentation of Multiple Sclerosis Lesions by Model Outlier Detection. *IEEE Transactions on Medical Imaging*, vol. 20, pp. 677–688, August 2001.
- [554] V. Venkatasubramanian, R. Rengaswamy, K. Yin, and S. N. Kavuri. A Review of Process Fault Detection and Diagnosis Part I: Quantitative Model-based Methods. *Computers and Chemical Engineering*, 27(3), pp. 293–311, 2003.
- [555] C. Vens and F. Costa. Random Forest based Feature Induction. *IEEE ICDM Conference*, pp. 744–753, 2011.
- [556] S. Viaene, R. Derrig, B. Baesens, and G. Dedene. A Comparison of State-of-the-Art Classification Techniques for Expert Automobile Insurance Claim Fraud Detection. *Journal of Risk and Insurance*, 69(3), pp. 373–421, 2002.
- [557] T. De Vries, S. Chawla, and M. Houle. Finding Local Anomalies in Very High Dimensional Space. *ICDM Conference*, 2010.
- [558] N. Wale, X. Ning, and G. Karypis. Trends in Chemical Data Mining. *Managing and Mining Graph Data*, Springer, 2010.
- [559] B. Wang, G. Xiao, H. Yu, and X. Yang. Distance-based Outlier Detection on Uncertain Data. *International Conference on Computer and Information Technology*, 2009.
- [560] B. Wang, X. Yang, G. Wang, and G. Yu. Outlier Detection over Sliding Windows for Probabilistic Data Streams. *Journal of Computer Science and Technology*, 25(3), pp. 389–400, 2010.
- [561] M. Wang, C. Zhang, and J. Yu. Native API-based Windows Anomaly Intrusion Detection Method using SVM. *International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing*, 2006.
- [562] X. Wang, C. Zhai, X. Hu, and R. Sproat. Mining Correlated Bursty Topic Patterns from Coordinated Text Streams. *ACM KDD Conference*, 2007.
- [563] L. Wei, W. Qian, A. Zhou, and W. Jin. HOT: Hypergraph-based Outlier Test for Categorical Data. *PAKDD Conference*, 2007.
- [564] L. Wei and E. Keogh. Semi-supervised Time Series Classification. *ACM KDD Conference*, 2006.
- [565] L. Wei, E. Keogh, and X. Xi. SAXually Explicit Images: Finding Unusual Shapes. *ICDM Conference*, 2006.

- [566] G. Weiss and F. Provost. Learning When Training Data are Costly: The Effect of Class Distribution on Tree Induction. *Journal of Artificial Intelligence Research*, 19: pp. 315–354, 2003.
- [567] G. Williams, R. Baxter, H. He, S. Hawkins, and L. Gu. A Comparative Study of RNN for Outlier Detection in Data Mining. *IEEE ICDM Conference*, 2002.
- [568] W.-K. Wong, A. Moore, G. Cooper, and M. Wagner. Rule-based Anomaly Pattern Detection for Detecting Disease Outbreaks. *National Conference on Artificial Intelligence*, 2002.
- [569] E. Wu, W. Liu, and S. Chawla. Spatio-temporal Outlier Detection in Precipitation Data. *Knowledge Discovery from Sensor Data, Springer, LNCS 5840*, 2008.
- [570] G. Wu and E. Y. Chang. Class-boundary Alignment for Imbalanced Dataset Learning. *ICML Workshop on Learning from Imbalanced Data Sets*, 2003.
- [571] K. Wu, K. Zhang, W. Fan, A. Edwards, and P. Yu. RS-Forest: A Rapid Density Estimator for Streaming Anomaly Detection. *ICDM Conference*, pp. 600–609, 2014.
- [572] M. Wu and C. Jermaine. Outlier Detection by Sampling with Accuracy Guarantees. *ACM KDD Conference*, 2006.
- [573] M. Wu, X. Song, C. Jermaine, S. Ranka, and J. Gums. A LRT Framework for Fast Spatial Anomaly Detection. *ACM KDD Conference*, 2009.
- [574] X. Xi, E. Keogh, C. Shelton, L. Wei, and C. Ratanamahatana. Fast Time Series Classification using Numerosity Reduction. *ICML Conference*, 2006.
- [575] Z. Xing, J. Pei, and E. Keogh. A Brief Survey on Sequence Classification. *ACM SIGKDD Explorations*, 12(1), 2010.
- [576] L. Xiong, X. Chen, and J. Schneider. Direct Robust Matrix Factorization for Anomaly Detection. *ICDM Conference*, 2011.
- [577] L. Xiong, B. Póczos, J. Schneider, A. Connolly, and J. VanderPlas. Hierarchical Probabilistic Models for Group Anomaly Detection, *Artificial Intelligence and Statistics*, 2011.
- [578] K. Yaminshi, J. Takeuchi, and G. Williams. Online Unsupervised Outlier Detection using Finite Mixtures with Discounted Learning Algorithms. *ACM KDD Conference*, 2000.
- [579] K. Yaminshi and J. Takeuchi. A Unified Framework for Detecting Outliers and Change Points from Time Series Data. *ACM KDD Conference*, 2002.
- [580] R. Yan, Y. Liu, R. Jin, and A. Hauptmann. On Predicting Rare Classes with SVM Ensembles in Scene Classification. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2003.
- [581] J. Yang and W. Wang. CLUSEQ: Efficient and Effective Sequence Clustering. *ICDE Conference*, 2003.
- [582] P. Yang and Q. Zhu. Finding Key Outlying Subspaces for Outlier Detection. *Knowledge-based Systems*, 24(2), pp. 269–274, 2011.

- [583] X. Yang, L. Latecki, and D. Pokrajac. Outlier Detection with Globally Optimal Exemplar-based GMM. *SIAM Conference on Data Mining*, 2009.
- [584] Y. Yang, J. Carbonell, R. Brown, T. Pierce, B. T. Archibald, and X. Liu. Learning Approaches for Detecting and Tracking News Events. *IEEE Intelligent Systems*, 14(4):32–43, 1999.
- [585] Y. Yang, T. Pierce, and J. Carbonell. A Study on Retrospective and On-line Event Detection. *ACM SIGIR Conference*, 1998.
- [586] Y. Yang, J. Zhang, J. Carbonell, and C. Jin. Topic-conditioned Novelty Detection. *ACM KDD Conference*, 2002.
- [587] D. Yankov, E. Keogh, and U. Rebbapragada. Disk-aware Discord Discovery: Finding Unusual Time Series in Terabyte Sized Data Sets. *ICDM Conference*, 2007.
- [588] N. Ye. A Markov Chain Model of Temporal Behavior for Anomaly Detection. *IEEE Information Assurance Workshop*, 2004.
- [589] N. Ye and Q. Chen. An Anomaly Detection Technique based on a Chi-square Statistic for Detecting Intrusions into Information Systems. *Quality and Reliability Engineering International*, 17, pp. 105–112, 2001.
- [590] L. Ye and E. Keogh. Time Series Shapelets: a New Primitive for Data Mining. *ACM KDD Conference*, 2009.
- [591] Q. Ye and W. Zhi. Outlier Detection in the Framework of Dimensionality Reduction. *International Journal of Pattern Recognition and Artificial Intelligence*, 29(4), 2015.
- [592] B.-K. Yi, N. D. Sidiropoulos, T. Johnson, H. Jagadish, C. Faloutsos, and A. Biliris. Online Data Mining for Co-evolving Time Sequences. *ICDE Conference*, 2000.
- [593] A. Ypma and R. Duin. Novelty Detection using Self-Organizing Maps. *Progress in Connectionist-Based Information Systems*, 2, pp. 1322–1325, 1997.
- [594] D. Yu, G. Sheikholeslami, and A. Zhang. Findout: Finding Outliers in Very Large Datasets. *Knowledge And Information Systems*, 4(4), pp. 387–412, 2002.
- [595] H. Yu, J. Han, and K. C.-C. Chang. PEBL: Web Page Classification without Negative Examples. *IEEE Transactions on Knowledge and Data Engineering*, 16(1), pp. 70–81, 2004.
- [596] J. X. Yu, W. Qian, H. Lu, and A. Zhou. Finding Centric Local Outliers in Categorical/Numeric Spaces. *Knowledge and Information Systems*, 9(3), pp. 309–338, 2006.
- [597] W. Yu, C. Aggarwal, S. Ma, and H. Wang. On Anomalous Hotspot Discovery in Graph Streams. *IEEE ICDM Conference*, 2013.
- [598] W. Yu, C. Aggarwal, and W. Wang. Temporally Factorized Network Modeling for Evolutionary Network Analysis. *WSDM Conference*, 2017.
- [599] B. Zadrozny and C. Elkan. Transforming Classifier Scores into Accurate Multiclass Probability Estimates. *ACM KDD Conference*, 2002.

- [600] B. Zadrozny, J. Langford, and N. Abe. Cost-Sensitive Learning by Cost-Proportionate Example Weighting, *ICDM Conference*, 2003.
- [601] B. Zadrozny and C. Elkan. Learning and Making Decisions when Costs and Probabilities are Unknown. *KDD Conference*, 2001.
- [602] D. Zhang and G. Lu. Review of Shape Representation and Description Techniques. *Pattern Recognition*, 37(1), pp. 1–19, 2004.
- [603] D. Zhang and W. S. Lee. A Simple Probabilistic Approach to Learning from Positive and Unlabeled Examples. *Annual UK Workshop on Computational Intelligence*, 2005.
- [604] J. Zhang, Q. Gao, and H. Wang. SPOT: A System for Detecting Projected Outliers from High-Dimensional Data Stream. *ICDE Conference*, 2008.
- [605] J. Zhang, M. Lou, T. W. Ling, and H. Wang. HOS-Miner: A System for Detecting Outlying Subspaces of High-dimensional Data. *VLDB Conference*, 2004.
- [606] J. Zhang, Q. Gao, and H. Wang. A Novel Method for Detecting Outlying Subspaces in High-dimensional Databases Using Genetic Algorithm. *ICDM Conference*, 2006.
- [607] J. Zhang and H. Wang. Detecting Outlying Subspaces for High-Dimensional Data: the New Task, Algorithms and Performance. *Knowledge and Information Systems*, 10(3), pp. 333–355, 2006.
- [608] J. Zhang, Z. Ghahramani, and Y. Yang. A Probabilistic Model for Online Document Clustering with Application to Novelty Detection. *Neural Information Processing Systems*, 2005.
- [609] J. Zhang and I. Mani. KNN Approach to Unbalanced Data Distributions: A Case Study Involving Information Extraction. *ICML Workshop on Learning from Imbalanced Datasets*, 2003.
- [610] K. Zhang, M. Hutter, and H. Jin. A New Local Distance-Based Outlier Detection Approach for Scattered Real-World Data. *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pp. 813–822, 2009.
- [611] T. Zhang, R. Ramakrishnan, and M. Livny. BIRCH: An Efficient Data Clustering Method for Very Large Databases. *ACM SIGMOD Conference*, 1996.
- [612] X. Zhang, P. Fan, and Z. Zhu. A New Anomaly Detection Method based on Hierarchical HMM. *International Conference on Parallel and Distributed Computing, Applications, and Technologies*, 2003.
- [613] Y. Zhang, J. Hong, and L. Cranor. CANTINA: A Content-based Approach to Detecting Phishing Web Sites. *WWW Conference*, 2007.
- [614] Y. Zhang, P. Meratnia, and P. Havinga. Outlier Detection for Wireless Sensor Networks: A Survey. *IEEE Communications Surveys and Tutorials*, 12(2), 2010.
- [615] J. Zhao, C.-T. Lu, and Y. Kou. Detecting Region Outliers in Meteorological Data. *ACM GIS Conference*, 2003.
- [616] Z. Zheng, X. Wu, and R. Srihari. Feature Selection for Text Categorization on Imbalanced Data. *SIGKDD Explorations*, 6(1), pp. 80–89, 2004.

- [617] Z.-H. Zhou. Ensemble Methods: Foundations and Algorithms. *Chapman and Hall/CRC Press*, 2012.
- [618] C. Zhu, H. Kitagawa, S. Papadimitriou, and C. Faloutsos. OBE: Outlier by Example. *PAKDD Conference*, 2004.
- [619] C. Zhu, H. Kitagawa, and C. Faloutsos. Example-based Robust Outlier Detection in High Dimensional Data Sets. *ICDM Conference*, 2005.
- [620] A. Zimek, A. Schubert, and H.-P. Kriegel. A Survey on Unsupervised Outlier Detection in High-dimensional Numerical Data. *Statistical Analysis and Data Mining*, 5(5), pp. 363–387, 2012.
- [621] A. Zimek, M. Gaudet, R. Campello, and J. Sander. Subsampling for Efficient and Effective Unsupervised Outlier Detection Ensembles. *KDD Conference*, 2013.
- [622] <http://www.itl.nist.gov/iad/mig/tests/tdt/tasks/fsd.html>
- [623] D. D. Lewis. Reuters-21578 Data Set.
<http://www.daviddlewis.com/resources/test-collections/reuters21578>.
- [624] <http://kdd.ics.uci.edu/databases/20newsgroups>
- [625] <http://www.informatik.uni-trier.de/~ley/db/>
- [626] <http://www.kdnuggets.com/software/deviation.html>
- [627] <http://www.kdnuggets.com/software/index.html>
- [628] <http://www.cs.waikato.ac.nz/ml/weka/>
- [629] http://scikit-learn.org/stable/auto_examples/svm/plot_oneclass.html
- [630] <http://scikit-learn.org/dev/modules/generated/sklearn.ensemble.IsolationForest.html>
- [631] <https://sourceforge.net/projects/iforest/>
- [632] http://scikit-learn.org/stable/modules/outlier_detection.html
- [633] <http://elki.dbs.ifi.lmu.de/>
- [634] <http://www.rdatamining.com/examples/outlier-detection>
- [635] http://www.cs.ucr.edu/~eamonn/time_series_data/
- [636] <http://www.cs.ucr.edu/~eamonn/SAX.htm>
- [637] <http://www-03.ibm.com/software/products/en/network-ips>
- [638] <http://www.ibm.com/software/analytics/spss>
- [639] <http://www-01.ibm.com/software/analytics/spss/products/statistics/>
- [640] <http://www.sas.com/>
- [641] <http://www.sas.com/software/security-intelligence/index.html>

[642] <http://www.oracle.com/technetwork/database/options/advanced-analytics/odm/index.html>

[643] <http://www.wizsoft.com>

Index

- Activation Function of Neural Network, 99
Active Learning, 26, 236
Adaboost, 230
AdaCost, 230
Adaptive Re-sampling, 228
Aggregate Change Points, 301
Aggregate Statistical Similarity, 256
ALSO, 239
Angle-based Outlier Detection, 49
Applications of Outlier Analysis, 399
Apriori, 166
AR Models, 276
Arbitrarily Oriented Subspaces, 170
ARIMA Model, 279
ARMA Model, 278
Astronomical Applications, 416
Autoencoder, 102
Autoregression Integrated Moving Average, 279
Autoregressive Models for Spatial Data, 352
Autoregressive Models for Time Series, 276
Autoregressive Moving Average, 278
Aviation Safety, 416
- Baum-Welch Algorithm, 335
Bayes Classifier, 225
Bayesian Network, 62
Behavioral Attribute, 345
Binary Data Outliers, 260
Binary Labels, 2
Biological Sequences, 311, 416
Boosting for Rare Class Classification, 230
Box and Whisker Diagrams, 45
Box Plots, 45
- Categorical Outlier Detection, 249
Categorization of Outlier Models, 10
CBLOF, 145
Cell-based Methods, 122
Central Limit Theorem, 42
Centroid Distance Signature, 356
Change Analysis, 273
Change Detection, 273
Chebychev Inequality, 38
Chernoff Bound (Lower Tail), 39
Chernoff Bound (Upper Tail), 41
Class Imbalance, 219, 221
CLUSEQ, 327
Clustering for Outlier Analysis, 112
COF, 145
Collective Outlier, 22
Combination Outliers, 320
Combination Outliers in Sequences, 312
Combining Novel and Rare Class Detection, 234
Complex Sequence Outliers, 338
Compression-based Dissimilarity Measure, 325
Conditional Outlier, 22
Contextual Attribute, 345
Contextual Outlier, 22
Contextual Similarity in Categorical Data, 257
Cosine Similarity, 268
Cost-Sensitive Learning, 223
Covariance Matrix Diagonalization, 75
Credit-Card Fraud, 404
- Data Cleaning, 415
Data Types, 21

- Data-Dependent Similarity, 128, 257
- Decision Trees, 226
- Dependent Variable Regression Analysis, 70
- Depth-based Outliers, 47
- Deviation-based Outliers, 48
- DFT, 288
- Differential Graph, 390
- Dimensionality Reduction, 13
- Discrete Attribute Outlier Detection, 249
- Discrete Fourier Transform, 288
- Discrete Sequence Outlier Detection, 24, 311
- Discrete Wavelet Transform, 288
- Disease Outbreaks, 416
- Distance Distribution-based Outliers, 51
- Distance-based Outliers, 118
- Distance-based Sequence Outliers, 324
- Distance-based Subspace Outlier Detection, 152
- Distribution Change, 304
- Distribution-based Outlier Modeling, 54
- Disturbance Events in Ecosystems, 416
- Dropout Training of Deep Networks, 106
- DWT, 288
- Dynamic Programming in HMM, 335
- Dynamic Time Warping, 293

- Early Discrete Sequence Anomalies, 340
- Early Termination Trick, 125
- Edit Distance, 325
- Egonet, 372
- EM Algorithm for Categorical Data, 250
- EM Algorithm for Continuous Data, 54
- ENetClus, 388
- Ensemble-based Outlier Detection, 19, 32
- Ensembles for Rare Classes, 229
- ERC-Forest, 144, 164, 182
- Eskin Measure, 256
- Evaluating Outlier Detection, 26
- Events in Social Streams, 412
- Evolutionary Network Analysis, 370
- Evolving Blogs, 413
- Evolving Social Networks, 413
- Explaining Sequence Anomalies, 334
- Exploratory Analysis, 418
- External Validity Measures, 27
- Extreme-Value Analysis, 11, 37, 46

- Failure Management of Computer Clusters, 416

- Fault Detection, 402
- Feature Bagging, 157, 158
- Feature Engineering, 235
- Feature Selection, 10
- Financial Applications, 404
- Financial Interaction Networks, 406
- Finite State Automaton, 316
- First Story Detection, 269, 411
- Flow Anomalies, 367
- FocusCO Algorithm, 395
- Forward Algorithm, 334
- Forward-backward Algorithm, 335
- FP-Outlier, 261
- Frequency-based Sequence Outliers, 327
- Frequent Pattern Mining Methods, 260

- GASP, 244
- Generalized Subspaces, 170
- Generative Models, 54
- Geometric Subsampling, 210
- Goodall Measure, 256
- Grammar Correction, 320
- Graph Evolution Rules, 396
- Graph Kernel, 372
- Graph Outliers, 369
- Graph-based Spatial Outliers, 351
- GraphScope, 390
- Grid-based Projected Outliers, 153

- Haar Wavelet, 288
- Half-Space Trees, 164, 182
- Harmonic k -Nearest Neighbor Detector, 120
- Heterogeneous Markov Random Field, 384
- HiCS, 164
- Hidden Markov Models, 313, 329
- Hidden Variables, 282
- High-Contrast Subspaces, 164
- High-Dimensional Outliers, 17, 149
- Histograms, 137, 160
- HMM, 329
- HMM Design Choices, 331
- Hoeffding Inequality, 41
- HOS-Miner, 157
- Host-based Intrusion Detection, 408
- HOTSAX, 293, 295
- Hotspot in Network, 385

- IBM SPSS, 421
- Image Anomalies, 416

- Incomplete Data-based Outliers, 96
- Independent Ensembles, 20
- Indexing for Distance-based Outliers, 126
- INFLO, 145
- Information Theoretic Measures, 48, 261, 307, 325, 339, 343, 381
- Information-Theoretic Models, 16
- Infrequently Recurring Classes, 306
- Inlier, 2
- Insider Trading Detection, 405
- Intensional Knowledge, 130
- Intensional Knowledge of Distance-based Outliers, 130
- Internal Validity Measures, 27
- Intrusion Detection, 305
- Intrusion Detection Applications, 407
- Inverse Document Frequency, 256, 268
- Inverse Occurrence Frequency, 256
- Isolation Forests, 129, 161, 203, 420
- ISOMAP, 86
- Isomorphism, 371
- Iterated Contextual Distance, 257

- KDD Nuggets, 421
- Kernel Density Estimation, 138
- Kernel PCA, 85
- Kernel Trick, 85
- Kernel Whitening, 84
- Kolmogorov Complexity, 17, 343
- Kurtosis, 10, 163

- Land Cover Anomalies, 414
- Latent Dirichlet Allocation, 268
- Latent Semantic Analysis, 21, 264
- Latent Semantic Indexing, 80
- Law Enforcement, 2
- LDA, 268
- LDOF, 133
- Leaky Bucket, 326
- LFC, 326
- Linear Models for Categorical Data, 254
- Linear Regression Models, 68
- Linkage Outliers, 5
- Linking Criminal Incidents, 416
- Local Correlation Integral, 135
- Local Outlier Factor, 132
- Local Spatial Outliers, 145, 350
- Local Subspace Selection, 166
- Locality Frame Count, 326
- LOCI, 135
- LOCI Plot, 136
- LODA, 141, 217
- LOF, 132
- LSA, 264
- LSA and PCA Relationship, 265

- MA Model, 278
- Mahalanobis Distance, 52, 114
- Malicious URL Detection, 416
- Market Basket Outliers, 260
- Markov Inequality, 37
- Markovian Models, 316
- Matrix Factorization, 95, 109, 374, 392
- MDEF, 135
- Medical Applications, 410
- Medical Imaging Diagnostics, 411
- Medical Sensor Diagnostics, 410
- MetaCost, 223
- Minimum Bounding Rectangles, 126
- Minimum Description Length, 339, 343, 381
- Mixed Attribute Outlier Detection, 249
- Mixture Modeling, 54
- Mobile Phone Fraud, 406
- Movement Pattern Outliers, 415
- Moving Average Model, 278
- Multidimensional Change Points, 301
- Multidimensional Spatial Outliers, 350
- Multiple Time Series Models, 279
- Multiple-Proclus, 159
- Multivariate Discrete Sequences, 338
- Multivariate Spatial Outliers, 351
- MUSCLES, 281

- NDCG, 30
- Nearest Neighbor Classifier, 225
- Neighborhood-based Spatial Outliers, 349
- NetClus, 388
- Network Intrusion Detection, 409
- Network Outliers, 25, 369
- Neural Networks, 98
- Noise Correction with PCA, 80
- Noise vs Anomaly, 3
- Non-negative Matrix Factorization, 395
- Normalization for PCA, 80
- Normalized Discounted Cumulative Gain, 30
- Novel Class Detection, 232
- Novelties in Temporal Transactions, 262
- Novelty Detection, 262, 273

- ODIN, 129
- One Class Learning, 88, 98, 233, 298
- One Class Neural Networks, 98
- One Class Novelty Detection, 88, 234, 298
- One Class SVM, 88, 233, 298
- One Class SVM for Novelty Detection, 88, 298
- Online Discrete Sequence Anomalies, 340
- Online Novelty Detection, 234, 269, 300, 305
- Open Source Software, 421
- Oracle Data Miner, 421
- Outlier Detection and Description, 10
- Outlier Ensembles, 18, 185
- Outlier Scores, 2
- Outliers from Small Graphs, 371
- OutRank, 159
- OUTRES, 167
- Overlap Measure, 255

- Parzen-Rosenblatt Density Estimator, 139
- PCA for Categorical Data, 254
- Perceptron, 99
- PLSA, 267
- Pocket Plots, 348
- Pooled Active Learning, 236
- Position Outliers, 312, 313
- Positive Unlabeled Classification, 231
- PPCA, 109
- PR Curve, 28
- Precision-Recall Curve, 28
- Pretraining Deep Networks, 104
- Primitive Sequence Anomaly, 322
- Principal Component Analysis, 13, 282
- Probabilistic and Statistical Models, 12
- Probabilistic Latent Semantic Analysis, 267
- Probabilistic Models for Categorical Data, 250
- Probabilistic Models for Mixed Data, 253
- Probabilistic Outlier Modeling, 54
- Probabilistic PCA, 109
- Probabilistic Suffix Trees, 318
- Projected Outlier Detection, 149, 151
- Projected Outliers, 153
- Proximity Models for Categorical Data, 255
- Proximity Models for Mixed Data, 259
- Proximity-based Classifiers, 225
- PST, 318
- PUC, 231

- Quality Control, 39, 401
- Quality Control Applications, 401
- Query by Committee, 238

- Random Forests, 161
- Random Subspace Ensembles, 153, 158
- Randomized Feature Weighting, 210
- Ranking Subspace Outliers, 153
- Rare Class Detection, 221
- Receiver Operating Characteristic, 28
- Recommender System Outliers, 96
- Regime Anomalies in Time Series, 308
- Regression Model, 7
- Regression Modeling with Dependent Variables, 70
- Relabeling, 223
- Replicator Neural Networks, 99, 102
- Reservoir Sampling, 386
- Reverse Nearest Neighbor, 129
- RIPPER, 316
- ROAM, 365
- ROC Curve, 28
- RODS, 109
- ROF, 144
- Rotated Bagging, 175
- Rotated Subspace Sampling, 175
- RS-Hash, 138, 141, 146, 160, 182, 217, 309, 420
- RS-Stream, 161, 309
- Rule-based Classifiers, 226
- Rule-based Models for Position Outliers, 315

- Sampling for Distance-based Outliers, 124
- SAS, 421
- SAS Security Intelligence, 421
- SAX, 291
- Sea Surface Temperature Anomalies, 413
- SELECTIVE MUSCLES, 281
- Semi-supervised Outlier Detection, 232
- Sequence Classification, 340
- Sequence Outlier Detection, 311
- Sequential Ensembles, 19
- Set-based Sequences, 339
- Shape Change Detection in Spatial Data, 361
- Shared Nearest Neighbor Similarity, 128
- Short Memory Property, 315
- Shortest Path Distance Changes, 392
- Similarity in Categorical Data, 255
- Similarity in Mixed Data, 259

- Simple Matching Coefficient, 324
- SLOM, 145, 350
- SMOTE, 229
- SMOTEBoost, 231
- SMT, 318
- Social Media Applications, 411
- Social Stream Evolution, 271, 396
- SOD, 169
- Software Resources, 421
- Space-filling Curves, 143
- Spam Filtering, 412
- Spam Link Detection, 412
- Sparse Coding, 109
- Sparse Markov Transducers, 318
- Spatial Heteroscedasticity, 348
- Spatial Autocorrelations, 348
- Spatial Outlier Detection, 345
- Spectral Methods, 14, 86, 108, 116, 378, 395
- SPIRIT, 282
- SPOT Algorithm, 301
- Stacking, 18, 236
- Statistical Tail Confidence Tests, 43
- Stock Market Anomalies, 405
- STORM Algorithm, 299
- Streaming Novel Class Detection, 306
- Streaming Novelty Detection, 234, 269, 300, 305
- Streaming Outlier Detection, 22, 182, 273, 298, 309
- Streaming Rare Class Detection, 305
- Strong Outliers, 4
- Structural Defect Detection, 403
- Structural Reservoir Sampling, 386
- Student's t-distribution, 43
- SUBDUE, 381
- Subgraph Outliers, 381
- Subsampling, 206
- Subspace Ensembles, 153
- Subspace Histograms, 160
- Subspace Methods for Transaction Data, 260
- Subspace Outlier Degree, 169
- Subspace Outlier Detection, 18, 149, 151
- Supervised Outlier Detection, 25, 219
- Supervised Sequence Outliers, 340
- Supervised Shape Anomalies in Time Series, 298
- Supervised Shape Discovery in Spatial Data, 360
- Support Vector Data Description, 93
- SVDD, 93
- SVM Classifier, 88, 226
- Symbolic Aggregate Approximation, 291
- Synthetic Over-sampling, 229
- Systems Diagnosis, 402
- t-Distribution, 43
- t-value Test, 43
- Tail Confidence Tests, 43
- Tail Inequalities, 37
- TARZAN, 328
- Temporal Description Length, 339
- Temporal Graph Outliers, 384
- Temporal Outlier Detection, 273
- Temporal Smoothness in Networks, 392
- Text Applications, 411
- Text Outliers, 262
- Time Series Outlier Detection, 22
- Top- n Local Outliers, 145
- Topic Detection and Tracking, 268
- Topic Modeling, 267
- Traffic Anomalies, 415
- Trajectory Outliers, 291, 363, 415
- Transaction Data Outliers, 260
- TRINITY, 421
- TROAD, 293, 363
- Tukey Box Plot, 46
- Unsupervised Regression Modeling, 74
- Unusual Shapes of Time-Series, 286
- Variable Subsampling, 207
- Variable Subsampling with Rotated Bagging, 209
- Variogram Clouds, 353
- Velocity Density Estimation, 302
- Viterbi Algorithm, 334
- VR, 209
- Wagging, 210
- Weak Outliers, 4
- Web Log Analytics, 406
- Weighting for Supervision, 225
- Whole Sequence Anomaly, 323
- Wilcoxon Test, 304
- WizRule, 421
- Z-value test, 6