

Emerging Trends in Machine Learning: Beyond Conventional Methods and Data

Luca Oneto¹, Nicolò Navarin², Michele Donini³, and Davide Anguita¹

1 - DIBRIS - University of Genova
Via Opera Pia 13, I-16145 Genova - Italy

2 - Department of Mathematics - University of Padua
Via Trieste, 63, I-35121 Padova - Italy

3 - IIT - Istituto Italiano di Tecnologia
Via Morego, 30, I-16163 Genova - Italy

Abstract. Recently, new promising theoretical results, techniques, and methodologies have attracted the attention of many researchers and have allowed to broaden the range of applications in which machine learning can be effectively applied in order to extract useful and actionable information from the huge amount of heterogeneous data produced everyday by an increasingly digital world. Examples of these methods and problems are: learning under privacy and anonymity constraints, learning from structured, semi-structured, multi-modal (heterogeneous) data, constructive machine learning, reliable machine learning, learning to learn, mixing deep and structured learning, semantics-enabled recommender systems, reproducibility and interpretability in machine learning, human-in-the-loop, adversarial learning. The focus of this special session is to attract both solid contributions or preliminary results which show the potentiality and the limitations of new ideas, refinements, or contaminations between the different fields of machine learning and other fields of research in solving real world problems. Both theoretical and practical results are welcome to our special session.

1 Introduction

Recently, new promising theoretical results, techniques, and methodologies have attracted the attention of many researchers and have allowed to broaden the range of applications in which machine learning can be effectively applied in order to extract useful and actionable information from the huge amount of heterogeneous data produced everyday by an increasingly digital world.

One important field of research is the privacy-preserving data mining that is being studied extensively, because of the wide proliferation of sensitive information on the internet [1]. A number of algorithmic techniques have been designed for privacy-preserving data mining. Different methods for randomization, k-anonymization, and distributed privacy-preserving data mining have been

proposed [2]. Other researchers discuss cases in which the output of data mining applications needs to be sanitized for privacy-preservation purposes. Computational and theoretical limits associated with privacy-preservation over high dimensional data sets have been also investigated [3, 4, 5]. It is important to highlight how this field of research has deep connections to the novel idea of fair machine learning, where the goal is to generate a model avoiding any discriminatory behavior with respect to a sensitive feature, e.g. gender, ethnic group and others [6, 7, 8].

Another relevant research field is the application of machine learning techniques to structured data [9, 10]. Indeed, in many application domains, it can be difficult to encode data in a fixed vectorial representation, while it can be more natural to represent it in a structured form. The simplest form of structured data are sequences, that can naturally encode text, protein sequences, event logs [11, 12]. Trees allow more flexibility with respect to sequences, and can encode more complex information. Trees can naturally encode, for instance, XML documents or parse trees of sentences (widely used in natural language processing) [13]. The more general form of structured data are graphs. Relevant domains where learning techniques for graph data has been successfully applied include computer vision, where each image can be represented by its segmentation graph [14]. Moreover, in Chemoinformatics chemical compounds can be represented as graphs, where each atom is a vertex and the edges represent the bonds between atoms [15]. Another class of techniques that involve the analysis of the nodes in huge graphs, motivated mainly by social networks analysis, have been introduced [16, 17].

Very recently, researchers are exploring a field referred to as constructive machine learning, where the output, in contrast with the classic classification or regression settings, is itself a structure. The main challenge of this task is to generate outputs that exhibits some of the properties of the inputs, without being too similar to the training instances [18, 19]. The ultimate goal of constructive machine learning is not to find a good model for predicting properties of the data, but instead to find one or more particular instances of the domain which are likely to exhibit desired properties. While traditional approaches choose these domain instances from a given set/databases of unlabeled domain instances, constructive machine learning is typically iterative and searches an infinite or exponentially large instance space [20].

Another important rising topic in machine learning is how to make it more reliable. Reliable machine learning has to deal with basically five important aspects: robustness, awareness, adaptation, value learning, and monitoring [21, 22, 23, 24, 25]. This implies to answer many possible questions touching on each of these categories. For example in order to create a robust system we have to make it robust to novel or potentially adversarial inputs and we have to handle

model misspecification or corrupted training data.

In order to create an aware system and make it aware of its environment and of its own limitations, we have to successfully identify “strange” inputs or situations and take appropriately conservative actions, we have to detect when changes in the environment have occurred that require re-training, and we have to detect that its model might be misspecified or poorly-calibrated. For what concerns adaptation, the machine learning system must detect and adapt itself to changes in its environment, especially large changes, and has to act properly when confronting radically new contexts. With respect to value learning we have to learn a value function that captures and balances all relevant considerations, we have to handle uncertainty about its value function, and we have to make sure that a system reflects the values of the humans who use it. Finally, respect to monitoring we have to monitor large-scale systems in order to judge if they are performing well and if things go wrong we have to fit it.

Another recent topic of research is the learning to learn or meta-learning [26, 27, 28]. The aim of meta-learning is a lifting of a standard machine learning problem: the task is to learn an algorithm capable of solving ground learning problems originated by a unknown distribution. A meta-dataset is thus a (possibly infinite) collection of datasets, or episodes, sampled from the unknown distribution, where each dataset is linked to a specific task. Specifically, the goal is to learn an algorithm capable of “producing” ground models. The meta-learner is viewed as a function which maps datasets to models (or weights), effectively making it a (non-standard, usually highly parametrized) learning algorithm. Since the ground models should exhibit good generalization performances on their specific task, each dataset can be split into training and validation sets, and the meta-learner can be trained to minimize the average validation error over tasks, which constitutes a natural outer objective in this setting [29].

In the last few years, several research groups are working on the problem of learning meaningful representations for complex data. Recent literature has shown the merits of having deep representations in the context of neural networks. The ideas are not new [30], but several proposals appeared in literature very recently [31, 32], sharing the goal of defining deep neural networks directly over structured data. An emerging challenge in kernel learning is the definition of similar multilayer [17], adaptive representations [33]. The research on this topic is still on its infancy, and the attention of the community to this topic is increasing.

Recommender systems are a well studied family of methods. These methods are used in different applications to help users exploring possibly interesting items. Usually, they exploit the preferences expressed in form of ratings. In the last years, the so-called content-based recommender systems has become very popular. They suggest the items that are similar to those the users previously

positively rated [34]. The first methods used simple retrieval models, for example the Vector Space Model. Due to the great interest from academia and companies (e.g. Google, Amazon and others), a new trend of research is arising. This trend studies the description of the items enriching the recommendation by using semantic knowledge. Consequently, content-based recommender systems evolved employing external sources of information, like ontologies, in order to improve their accuracy [34, 35]. Recent approaches are based, among others, on deep learning [36] or word embeddings [37]. As a future prospective, the researchers are combining in synergy semantic technologies with cognitive computing [38], in order to create new cognitive recommender systems.

Another big problem is reproducibility and interpretability of machine learning. Papers from the Machine Learning community are supposed to be a valuable asset. They can help to inform and inspire future research. They can be a useful educational tool for students. They can give guidance to applied researchers in industry. Reproducibility [39, 40, 41, 42], while not always possible in science (consider the study of a transient astrological phenomenon like a passing comet), is a powerful criteria for improving the quality of research. A result which is reproducible is more likely to be robust and meaningful and rules out many types of experimenter error (either fraud or accidental). Contemporary machine learning has a huge potential to improve products, processes and research but interpretability is a crucial problem [43, 44, 45]. In fact, machines usually don't give an explanation for their predictions, which hurts trust and creates a barrier for the adoption of machine learning. Many works try to find answers to important questions such as: can everyone trust the learned model? The model might perform well on the training data, but are the learned associations general enough to transfer to new data? Are there some oddities in the training data which the machine learning model dutifully picked up?

Human-in-the-loop leverages both human and machine intelligence to generate machine learning models [46, 47]. In the human-in-the-loop approach, people are involved in a virtuous circle where they train, tune, and test a particular algorithm. Human-in-the-loop is in clear opposition with respect to the so called Automatic Machine Learning, where the goal is to bring the humans out-of-the-loop. On the other hand, several real world applications contain missing data, noisy sources, unwanted information, and some problems in the domain could be hard to solve. This set of problems makes the application of automated methods difficult or, in certain cases, impossible. In fact, the quality of results from automatic approaches in these task might be questionable (e.g. methods applied on medical data [48]).

Finally a quite interesting and relatively novel research field is the adversarial machine learning, that lies at the intersection of machine learning and computer security [49]. It aims to enable the safe adoption of machine learning techniques

in adversarial settings. The problem arises from the fact that machine learning techniques were originally designed for stationary environments in which the training and test data are assumed to be generated from the same (although possibly unknown) distribution. In the presence of intelligent and adaptive adversaries, however, this working hypothesis is likely to be violated: a malicious adversary can carefully manipulate the input data exploiting specific vulnerabilities of learning algorithms to compromise the whole system security [50, 51]. Evasion attacks are the most prevalent type of attack that may be encountered in adversarial settings during system operation. In the evasion setting, malicious samples are modified at test time to evade detection; that is, to be misclassified as legitimate. No influence over the training data is assumed. When instead the data used for training purposes varies in time an attacker may poison the training data by injecting carefully designed samples to eventually compromise the whole learning process. This type of impairment is called poisoning attack.

This special session (SS) have attracted both solid contributions and preliminary results which show the potentiality and the limitations of new ideas, refinements, or contaminations between the different fields of machine learning and other fields of research in solving real world problems. Both theoretical and practical results have been submitted to our special session. Eight papers have been accepted and we will describe them in the next section.

2 Accepted Works

The first work accepted in our SS is entitled *Finding the most interpretable MDS rotation for sparse linear models based on external features* [52]. Authors state that one approach to interpreting multidimensional scaling (MDS) embeddings is to estimate a linear relationship between the MDS dimensions and a set of external features. However, because MDS only preserves distances between instances, the MDS embedding is invariant to rotation. As a result, the weights characterizing this linear relationship are arbitrary and difficult to interpret. Authors propose a procedure for selecting the most pertinent rotation for interpreting a two dimension MDS embedding.

The second work accepted in our SS is entitled *Mixture of Hidden Markov Models as Tree Encoder* [53]. The paper introduces a new probabilistic tree encoder based on a mixture of Bottom-up Hidden Tree Markov Models. The ability to recognize similar structures in data is experimentally assessed both in clusterization and classification tasks. The results of these preliminary experiments suggest that the model can be successfully used to compress the tree structure and label patterns in a vectorial representation.

The third work accepted in our SS is entitled *Differential private relevance learning* [54]. In this paper it is observed that digital information is collected

daily in growing volumes and mutual benefits drive the demand for the exchange and publication of data among parties. However, it is often unclear how to handle these data properly in the case that the data contains sensitive information. Differential privacy has become a powerful principle for privacy-preserving data analysis tasks in the last few years, since it entails a formal privacy guarantee for such settings. This is obtained by a separation of the utility of the database and the risk of an individual to lose his/her privacy. In this paper, authors introduced the Laplace mechanism and a stochastic gradient descent methodology which guarantee differential privacy [55]. Then, authors show how these paradigms can be incorporated into two popular machine learning algorithm, namely GLVQ and GMLVQ. Authors demonstrate the results of privacy-preserving LVQ based on three benchmarks.

The fourth work accepted in our SS is entitled *On aggregation in ranking median regression* [56]. In this work authors observed that the present era of personalized customer services and recommender systems, predicting the preferences of an individual/user over a set of items indexed by $[[n]] = \{1, \dots, n\}$, $n \geq 1$, based on its characteristics, modelled as a r.v. X say, is an ubiquitous issue. Though easy to state, this predictive problem referred to as ranking median regression (RMR in short) is very difficult to solve in practice. The major challenge lies in the fact that, here, the (discrete) output space is the symmetric group \mathfrak{S}_n , composed of all permutations of $[[n]]$, of explosive cardinality $n!$, and which is not a subset of a vector space. It is thus far from straightforward to build predictive rules taking their values in \mathfrak{S}_n , except by means of ranking aggregation techniques implemented at a local level, as proposed in [57] or [58]. However, such local learning techniques exhibit high instability and it is the main goal of this paper to investigate to which extent Kemeny ranking aggregation of randomized RMR rules may remedy this drawback. Beyond a theoretical analysis establishing its validity, the relevance of this novel ensemble learning technique is supported by experimental results.

The fifth work accepted in our SS is entitled *LANN-DSVD: A new privacy-preserving distributed algorithm for machine learning* [59]. In this work authors observed that in the Big Data era new challenges have arisen in machine learning related with the Volume (high number of samples or variables), the Velocity, etc. making many of the classic and brilliant methods not applicable. One main concern derives from Privacy issues when data is distributed and cannot be shared among locations. In their work, authors present LANN-DSVD, a non iterative algorithm for One-Layer Neural Networks that allows distributed learning guaranteeing privacy. Moreover, it is non iterative, parameter-free and provides incremental learning, thus making it very suitable to manage huge and/or continuous data. Results demonstrate its competitiveness both in efficiency and efficacy.

The sixth work accepted in our SS is entitled *Set point thresholds from topological data analysis and an outlier detector* [60]. In this work authors provide an algorithm for unsupervised or semi-supervised learning to determine, once the input settings are given, a very easily described zone of optimal execution settings for a production. A region is very easily described if anyone can determine whether a point is inside it and select a point on it with a certain range of choice. This can be applied both in production optimization and in predictive maintenance. Part of the method is based on a topological data analysis tool: Mapper. Authors also provide a method to detect outliers on new data.

The seventh work accepted in our SS is entitled *Vector Field Based Neural Networks* [61]. In this work authors state that a novel Neural Network architecture is proposed using the mathematically and physically rich idea of vector fields as hidden layers to perform nonlinear transformations in the data. The data points are interpreted as particles moving along a flow defined by the vector field which intuitively represents the desired movement to enable classification. The architecture moves the data points from their original configuration to a new one following the streamlines of the vector field with the objective of achieving a final configuration where classes are separable. An optimization problem is solved through gradient descent to learn this vector field.

The eighth work accepted in our SS is entitled *Temporal transfer learning for drift adaptation* [62]. Whereas detecting and adapting to concept drift has been well studied, predicting temporal drift of decision boundaries has received much less attention. This paper proposes a method for drift prediction, drift projection, and active-learning for adjusting the projected decision boundary so as to regain accuracy with minimal additional labeled samples. The method works with different underlying learning algorithms. Results on several data sets with translational and rotational drift and corresponding boundary projection show regained accuracy with significantly fewer labeled samples, even in the presence of noisy drift.

References

- [1] C. C. Aggarwal and S. Y. Philip. A general survey of privacy-preserving data mining models and algorithms. In *Privacy-preserving data mining*, 2008.
- [2] Y. A. A. S. Aldeen, M. Salleh, and M. A. Razzaque. A comprehensive review on privacy preserving data mining. *SpringerPlus*, 4(1):694, 2015.
- [3] A. Friedman and A. Schuster. Data mining with differential privacy. In *ACM SIGKDD international conference on Knowledge discovery and data mining*, 2010.
- [4] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- [5] L. Oneto, S. Ridella, and D. Anguita. Differential privacy and generalization: Sharper bounds with applications. *Pattern Recognition Letters*, 89:31–38, 2017.

- [6] C. Dwork, N. Immorlica, A. T. Kalai, and M. D. M. Leiserson. Decoupled classifiers for group-fair and efficient machine learning. In *Conference on Fairness, Accountability and Transparency*, 2018.
- [7] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, 2016.
- [8] M. Donini, S. Ben-David, M. Pontil, and J. Shawe-Taylor. An efficient method to impose fairness in linear models. In *NIPS Workshop on Prioritising Online Content*, 2017.
- [9] G. Bakir, T. Hofman, B. Scholkopf, A. J. Smola, B. Taskar, and S. V. N. Vishwanathan. *Predicting structured data*. MIT press, 2007.
- [10] G. Da San Martino and A. Sperduti. Mining Structured Data. *IEEE Computational Intelligence Magazine*, 5(1):42–49, 2010.
- [11] C. S. Leslie, E. Eskin, and W. S. Noble. The Spectrum Kernel: A String Kernel for SVM Protein Classification. In *Pacific Symposium on Biocomputing*, 2002.
- [12] N. Navarin, B. Vincenzi, M. Polato, and A. Sperduti. LSTM Networks for Data-Aware Remaining Time Prediction of Business Process Instances. In *IEEE Symposium on Deep Learning*, 2017.
- [13] A. Moschitti. Making Tree Kernels Practical for Natural Language Learning. In *Conference of the European Chapter of the Association for Computational Linguistics*, 2006.
- [14] F. Bach. Image Classification with Segmentation Graph Kernels e. In *Conference on Computer Vision and Pattern Recognition*, 2007.
- [15] C. C. Aggarwal and H. Wang. Manging and Mining Graph Data. In *Managing and Mining Graph Data*, 2010.
- [16] R. I. Kondor and J. Lafferty. Diffusion Kernels on Graphs and Other Discrete Input Spaces. In *Proceedings of the Nineteenth International Conference on Machine Learning*, 2002.
- [17] L. Oneto, N. Navarin, A. Sperduti, and D. Anguita. Multilayer Graph Node Kernels: Stacking While Maintaining Convexity. *Neural Processing Letters*, 2017.
- [18] D. Oglic, R. Garnett, and G. Thomas. Active Search in Intensionally Specified Structured Spaces. *Conference on Artificial Intelligence*, 2017.
- [19] F. Costa. Learning an efficient constructive sampler for graphs. *Artificial Intelligence*, 244:217–238, 2017.
- [20] R. Gómez-Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernández-Lobato, B. Sánchez-Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *American Chemical Society Central Science*, 2016.
- [21] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané. Concrete problems in ai safety. In *arXiv preprint arXiv:1606.06565*, 2016.
- [22] S. Athey and G. Imbens. A measure of robustness to misspecification. *American Economic Review*, 105(5):476–80, 2015.
- [23] M. Basseville. Detecting changes in signals and systems—a survey. *Automatica*, 24(3):309–326, 1988.
- [24] Y. Chow, A. Tamar, S. Mannor, and M. Pavone. Risk-sensitive and robust decision-making: a cvar optimization approach. In *Advances in Neural Information Processing Systems*, 2015.
- [25] A. Daniely, A. Gonen, and S. Shalev-Shwartz. Strongly adaptive online learning. In *International Conference on Machine Learning*, 2015.

- [26] M. Andrychowicz, M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, and N. de Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, 2016.
- [27] O. Wichrowska, N. Maheswaranathan, M. W. Hoffman, S. Gomez Colmenarejo, M. Denil, N. De Freitas, and J. Sohl-Dickstein. Learned optimizers that scale and generalize. In *International Conference on Machine Learning*, 2017.
- [28] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- [29] L. Franceschi, M. Donini, P. Frasconi, and M. Pontil. A bridge between hyperparameter optimization and learning-to-learn. In *NIPS workshop on Meta-learning*, 2017.
- [30] A. Sperduti and A. Starita. Supervised neural networks for the classification of structures. *IEEE Transactions on Neural Networks*, 8(3):714–735, 1997.
- [31] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph Attention Networks. In *International Conference on Learning Representations*, 2018.
- [32] T. N. Kipf and M. Welling. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*, 2017.
- [33] F. Aiolli, M. Donini, N. Navarin, and A. Sperduti. Multiple Graph-Kernel Learning. In *IEEE Symposium on Computational Intelligence and Data Mining*, Cape Town, 2015.
- [34] M. De Gemmis, P. Lops, C. Musto, F. Narducci, and G. Semeraro. Semantics-aware content-based recommender systems. In *Recommender Systems Handbook*, 2015.
- [35] M. Capelle, F. Hogenboom, A. Hogenboom, and F. Frasincar. Semantic news recommendation using wordnet and bing similarities. In *Annual ACM Symposium on Applied Computing*, 2013.
- [36] Tomas Mikolov, Quoc V Le, and Ilya Sutskever. Exploiting similarities among languages for machine translation. In *arXiv preprint arXiv:1309.4168*, 2013.
- [37] C. Musto, G. Semeraro, M. de Gemmis, and P. Lops. Learning word embeddings from wikipedia for content-based recommender systems. In *European Conference on Information Retrieval*, 2016.
- [38] D. S. Modha, R. Ananthanarayanan, S. K. Esser, A. Ndirango, A. J. Sherbondy, and R. Singh. Cognitive computing. *Communications of the ACM*, 54(8):62–71, 2011.
- [39] J. Vanschoren, J. N. Van Rijn, B. Bischl, and L. Torgo. Openml: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60, 2014.
- [40] H. Larochelle. Some opinions on reproducibility in ml. In *International Conference in Machine Learning*, 2017.
- [41] J. Langford. Reproducibility in machine learning. In *International Conference in Machine Learning*, 2017.
- [42] R. Williamson. Beyond reproducibility. In *International Conference in Machine Learning*, 2017.
- [43] A. Vellido, J. D. Martín-Guerrero, and P. J. G. Lisboa. Making machine learning models interpretable. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2012.
- [44] C. Rudin. Algorithms for interpretable machine learning. In *ACM SIGKDD international conference on Knowledge discovery and data mining*, 2014.
- [45] C. Molnar. *Interpretable Machine Learning. A Guide for Making Black Box Models Explainable*. in-press, 1998.

- [46] A. Thomaz, G. Hoffman, and M. Cakmak. Computational human-robot interaction. *Foundations and Trends in Robotics*, 4(2-3):105–223, 2016.
- [47] F. M. Zanzotto. Human-in-the-loop artificial intelligence. In *arXiv preprint arXiv:1710.08191*, 2017.
- [48] A. Holzinger. Interactive machine learning for health informatics: when do we need the human-in-the-loop? *Brain Informatics*, 3(2):119–131, 2016.
- [49] L. Huang, A. D. Joseph, B. Nelson, B. I. P. Rubinstein, and J. D. Tygar. Adversarial machine learning. In *ACM workshop on Security and artificial intelligence*, 2011.
- [50] M. Barreno, B. Nelson, R. Sears, A. D. Joseph, and J. D. Tygar. Can machine learning be secure? In *ACM Symposium on Information, computer and communications security*, 2006.
- [51] M. Barreno, B. Nelson, A. D. Joseph, and J. D. Tygar. The security of machine learning. *Machine Learning*, 81(2):121–148, 2010.
- [52] A. Bibal, R. Marion, and Fréney. Finding the most interpretable mds rotation for sparse linear models based on external features. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2018.
- [53] D. Bacciu and D. Castellana. Mixture of hidden markov models as tree encoder. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2018.
- [54] J. Brinkrolf, K. Berger, and B. Hammer. Differential private relevance learning. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2018.
- [55] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *ACM SIGSAC Conference on Computer and Communications Security*, 2016.
- [56] S. Cléménçon and A. Korba. On aggregation in ranking median regression. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2018.
- [57] L. H. Philip, W. M. Wan, and P. H. Lee. Decision tree modeling for ranking data. In *Preference learning*, 2010.
- [58] S. Cléménçon, A. Korba, and E. Sibony. Ranking median regression: Learning to order through local consensus. *arXiv preprint arXiv:1711.00070*, 2017.
- [59] O. Fontenla-Romero, B. Guijarro-Berdinās, B. Pérez-Sánchez, and M. Gómez-Casal. Lann-dsvd: A privacy-preserving distributed algorithm for machine learning. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2018.
- [60] A. Carrega. Set point thresholds from topological data analysis and an outlier detector. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2018.
- [61] D. Vieira, F. Rangel, F. Firmino, and J. Paixao. Vector field based neural networks. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2018.
- [62] D. Won, P. J. Jansen, and J. G. Carbonell. Temporal transfer learning for drift adaptation. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2018.