



Optimal transport for novelty and out-of-distribution detection

Master/Ecole d'Ingénieur internship

Expected starting: February/March 2021 (6 months)

Key-words Machine learning, optimal transport, out-of-distribution samples, novelty detection.

Context To be safe, a decision device learned from data requires a mechanism that adapts the decision according to whether or not there is a discrepancy between the distribution $p_{train}(X_{train}, Y_{train})$ of the training samples and the ones of test samples $p_{test}(X_{test}, Y_{test})$. In case of distribution shift, deep-based-approaches may be overconfident and tend to treat the given inputs as one of the previously seen situations leading to mislabelling. This brings to the scientific challenges of detecting out-of-distribution (OOD) samples (the test point x_0 is marginally sampled from $p_{test}(x_0) \neq p_{train}(x_0)$) or of recognizing that point x_0 belongs to an unseen class (new type of object occurs in the scenes) [1]. Moreover due to the multimodal nature of the inputs and sensors availability, the samples may not be embedded into the same space, and hence compromising the success of the detection task. We envision to leverage on the optimal transport theory [2] to implement algorithms dealing with out-of-distribution detection, with specific applications on road scene.

Optimal transport (OT) [2] has emerged as a powerful tool to compute distances (a.k.a. Wasserstein or earth mover's distances) between empirical distribution of data, thanks to new computational schemes that make the transport computation tractable [3]. It has wide applications in computer vision, statistics, imaging and has been recently introduced in the machine learning community to efficiently solve classification or transfer learning problems [4]. The advantage of OT is that it can compare possibly high dimensional empirical probability measures, taking into account the geometry of the underlying metric spaces and dealing with discrete measures.

Classical optimal transport problem seeks a transportation map that preserves the total mass between two probability distributions, requiring their mass to be the same. This may be too restrictive in certain applications such as color or shape matching, since the distributions may have arbitrary masses and/or that only a fraction of the total mass has to be transported. This happens also when datasets X_{train} and/or X_{test} are contaminated by outliers, in which we may want to discard them from the transportation plan: this is the partial OT problem [5]. Several algorithms have been devised to solve the problem, among them [6] solve the exact problem when given as input the total mass that has to be transported between the two empirical distributions.

Scientific objectives and expected achievements The objective of the internship is to study and implement OT-based strategies for dealing with OOD samples or when the datasets are contaminated by outliers. In many cases, the number of such samples are unknown and should be estimated from the data. To do so, one can rely on two-sample tests and their Wasserstein counterparts [7]; when there is a shift between p_{train} and p_{test} , or even when the 2 distributions do not lie on the same space, one can rather build on the Gromov-Wasserstein [8] based tests [9].

In more details, the aim of the internship is to study how the partial formulation of OT can be used in the OOD and outliers scenarii. Integration of two-sample tests within the OT formulation as a regularization term will be considered first. As such, we aim at estimating from the data the proportion of contaminated samples in the datasets, together with the optimal transport plan in a unified formulation, even when the 2 distributions live in incomparable spaces. Integration of partial-OT-based loss in deep-based approaches will serve as a playground to evaluate the method.

From an application point of view, a particular attention will be given on OOD detection for road scene. The intended methods will be evaluated on real-world datasets comprising of automotive images (such as nuScenes [10], KITTI[11]) or on autonomous car scene benchmark <https://github.com/OATML/oatomobile> in order to build robust system for road scene analysis. The developed methods will be challenged with some current position approaches and their applications [12, 13, 14].

Outcomes Outcomes of the internship will lead to publication in the machine learning community. This internship will be part of the AI chair RAIMO (A road toward safe artificial intelligence in mobility) funded by ANR, for which funding is available for the candidate to **enter a PhD program after the internship**. The applied part of the internship will lead to development in Python. The candidate will build upon the python toolbox for optimal transport (POT: <https://github.com/rflamary/POT>) developed by members of the team among others, and the developments proposed in the internship may be integrated to the toolbox. On this technical aspect, the intern will benefit from the expertise of ongoing collaborations with other academic partners on this subject.

Research environnement/Location The research will take place (depending on the candidate):

- either within the LITIS laboratory (<https://www.insa-rouen.fr/recherche/laboratoires/litis>) located at INSA Rouen, France.
- or within the OBELIX research group (www.irisa.fr/obelix) from IRISA located in the UBS (Université Bretagne Sud) campus in Vannes, France

The internship will be jointly supervised by Gilles Gasso (LITIS) and Laetitia Chapel (Obelix).

Candidate profile Applicants are expected to be graduated in computer science and/or machine learning and/or signal & image processing and/or applied mathematics/statistics, and show an excellent academic profile. Beyond, good programming skills are expected.

Application procedure Send a resume to Gilles Gasso (gilles.gasso@insa-rouen.fr) and Laetitia Chapel (laetitia.chapel@irisa.fr). Potential candidates will be contacted for interview. Feel free to contact us for any question.

References

- [1] A. Shafaei, M. Schmidt, and J. J. Little, “Does your model know the digit 6 is not a cat? a less biased evaluation of” outlier” detectors,” 2018.
- [2] C. Villani, *Optimal transport: old and new*. Springer Science, 2008.
- [3] M. Cuturi, “Sinkhorn distances: Lightspeed computation of optimal transport,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2292–2300.
- [4] N. Courty, R. Flamary, D. Tuia, and A. Rakotomamonjy, “Optimal transport for domain adaptation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016.
- [5] A. Figalli, “The optimal partial transport problem,” *Archive for rational mechanics and analysis*, vol. 195, 2010.

- [6] L. Chapel, M. Z. Alaya, and G. Gasso, “Partial optimal transport with applications on positive-unlabeled learning,” in *NeurIPS*, 2020.
- [7] A. Ramdas, N. G. Trillos, and M. Cuturi, “On wasserstein two-sample testing and related families of nonparametric tests,” *Entropy*, 2017.
- [8] G. Peyré, M. Cuturi, and J. Solomon, “Gromov-wasserstein averaging of kernel and distance matrices,” in *ICML*, vol. 16, 2016, pp. 2664–2672.
- [9] C. Bréchet, “A statistical test of isomorphism between metric-measure spaces using the distance-to-a-measure signature,” *Electronic Journal of Statistics*, 2019.
- [10] N. Sünderhauf, O. Brock, W. Scheirer, R. Hadsell, D. Fox, J. Leitner, B. Upcroft, P. Abbeel, W. Burgard, M. Milford *et al.*, “The limits and potentials of deep learning for robotics,” *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 405–420, 2018.
- [11] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” in *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2012, pp. 3354–3361.
- [12] A. Filos, P. Tigas, R. McAllister, N. Rhinehart, S. Levine, and Y. Gal, “Can autonomous vehicles identify, recover from, and adapt to distribution shifts?” in *International Conference on Machine Learning (ICML)*, 2020.
- [13] J. Nitsch, M. Itkina, R. Senanayake, J. Nieto, M. Schmidt, R. Siegwart, M. J. Kochenderfer, and C. Cadena, “Out-of-distribution detection for automotive perception,” *arXiv preprint arXiv:2011.01413*, 2020.
- [14] J. Ren, P. J. Liu, E. Fertig, J. Snoek, R. Poplin, M. Depristo, J. Dillon, and B. Lakshminarayanan, “Likelihood ratios for out-of-distribution detection,” in *Advances in Neural Information Processing Systems*, 2019, pp. 14 707–14 718.