

Fairness in Clustering

Brian Brubach,¹ Deeparnab Chakrabarty,² John P. Dickerson,³ Seyed Esmaeili,³ Matthäus Kleindessner,⁴ Marina Knittel,³ Jamie Morgenstern,⁵ Samira Samadi,⁶ Aravind Srinivasan,³ Leonidas Tsepenekas³

¹ Wellesley College, ² Dartmouth College, ³ University of Maryland, ⁴ Amazon, ⁵ University of Washington, ⁶ Max Planck Institute for Intelligent Systems (MPI-IS)
bb100@wellesley.edu, deeparnab@gmail.com, john@cs.umd.edu, esmaeili@cs.umd.edu, matkle@amazon.de, mknittel@cs.umd.edu, jamiemmt@cs.washington.edu, samira.samadi@tuebingen.mpg.de, asriniv1@umd.edu, ltsepene@umd.edu

Introductory Information

Keywords: Fairness; Clustering; Unsupervised learning

Duration: Half-day (4 hours)

Expected Participants: 150 (if in person); more (if virtual)

History of Tutorial: This is a new tutorial.

Goal of the Tutorial

Clustering problems are fundamental in areas ranging from Unsupervised Learning and Operations Research, to Data Mining and Optimization. Even more importantly, due to the numerous practical applications of this family of problems, algorithms developed for them are employed on a daily basis in settings affecting the lives of millions of individuals. Because of this, a flurry of research has emerged on guaranteeing that the aforementioned clustering problems will explicitly take into account the issue of fairness, thus providing socially just solutions. This research endeavour has led to a vast collection of results, where multiple and often unrelated or conflicting definitions of fairness are considered.

The goal of this tutorial is to introduce a wide audience interested in algorithmic fairness to the nascent research area of *fair clustering*. Specifically, we wish to present a variety of fairness notions used in the context of clustering, argue about the necessity of each of those through corresponding applications, discuss the relationships between different notions, sketch the algorithmic ideas that were developed in order to address the corresponding computational problems, and finally share our thoughts about the future of research in algorithmic fairness. By the end of the tutorial, the audience will have achieved a significant level of familiarity with multiple definitions of fairness in the unsupervised learning context, and we hope that researchers will use these ideas in contexts both within and adjacent to the clustering context, in both industrial and academic applications.

Brief Outline of the Tutorial

We will begin the tutorial by formally introducing standard variants of clustering problems. Moving on, we will discuss applications in which unfairness may naturally arise, and for each such application we will discuss the corresponding notion of fairness that was introduced in order to mitigate the aforementioned situation. Furthermore, we will sketch the

high-level algorithmic ideas that were developed in order to solve each corresponding fair clustering problem. However, we plan on putting more emphasis on presenting the actual mathematical definitions of fairness, and not on the algorithmic aspects of the results. The reason for this is that we want to expose the audience to as many notions of fairness as possible, while discussing the necessity of each of them and the relationships between them. The last section of the tutorial will involve a broader discussion of the future of research in algorithmic fairness, by proposing specific directions.

Prerequisite Knowledge

We will cater our tutorial to the modal *junior* participant (e.g., an early-stage PhD student or similar). That is, we will assume the audience has a basic CS and AI/ML background, but not necessarily deep clustering experience. Specifically, any participant who has had an undergraduate- or early graduate-level Algorithms and Machine Learning course will be able to follow the entirety of the tutorial. This is because, as mentioned earlier, the emphasis will be put into the definitions of fairness and their corresponding relationships, and relatively less on the algorithmic aspects of (traditional, constrained, and/or fair) clustering.

Content

The first part of the tutorial will consist of introducing classical clustering paradigms, together with appropriate applications for each of them. These paradigms define what we call the unfair clustering problem, where fairness is not part of the computational problem. The unfair clustering problems we will discuss are well-studied and longstanding, however we will only need their definitions and not algorithmic results on them. Specifically, in the first part of the tutorial we will present the formal definitions for:

- Metric k -clustering: k -center, k -median and k -means.
- Clustering with outliers.
- Hierarchical clustering.
- Correlation clustering.

In the second part of the tutorial we will start discussing specific notions of fairness for clustering, and we will see how these are incorporated in the classical/unfair models.

For each such notion we will present realistic applications demonstrating its importance. One of the biggest take-home messages of our tutorial will be that these different definitions of fairness are equally important and they are introduced in order to capture different and unique scenarios. In addition, besides simply enumerating different concepts of fairness, we will present a more structured taxonomy. The full list of the work we are going to cover can be found in the References section below.

The third part of the tutorial will revolve around some of the main algorithmic techniques that were developed in order to tackle fair clustering problems. Our description of those will be on a high-level, nonetheless we would still like to make the audience aware of the kind of approaches that have been used throughout the literature. Examples of techniques we want to cover include the widely-used two-step approach of (Bercea et al. 2019; Bera et al. 2019), the fairlet decomposition of (Chierichetti et al. 2017), and network-flow roundings such as those in (Bercea et al. 2019).

The final part of the tutorial will feature a discussion of where fair clustering and the broader field of algorithmic fairness is going. We will explore important questions, criticisms, and developing norms in this growing research area. Emphasis will be placed on connecting foundational research in fair clustering more closely to real applications and communities served. This discussion will conclude with exciting future directions and open problems.

References

- Abbasi, M.; Bhaskara, A.; and Venkatasubramanian, S. 2021. Fair Clustering via Equitable Group Representations. *FACCT '21*.
- Ahmadian, S.; Epasto, A.; Knittel, M.; Kumar, R.; Mahdian, M.; Moseley, B.; Pham, P.; Vassilvitskii, S.; and Wang, Y. 2020a. Fair hierarchical clustering. *Advances in Neural Information Processing Systems* 33: 21050–21060.
- Ahmadian, S.; Epasto, A.; Kumar, R.; and Mahdian, M. 2020b. Fair correlation clustering. In *International Conference on Artificial Intelligence and Statistics*, 4195–4205. PMLR.
- Anderson, N.; Bera, S. K.; Das, S.; and Liu, Y. 2020. Distributional Individual Fairness in Clustering. *CoRR* abs/2006.12589. URL <https://arxiv.org/abs/2006.12589>.
- Anegg, G.; Angelidakis, H.; Kurpisz, A.; and Zenklusen, R. 2020. A Technique for Obtaining True Approximations for k -Center with Covering Constraints. In Bienstock, D.; and Zambelli, G., eds., *Integer Programming and Combinatorial Optimization*, 52–65. Cham: Springer International Publishing.
- Bandyapadhyay, S.; Inamdar, T.; Pai, S.; and Varadarajan, K. 2019a. A Constant Approximation for Colorful k -Center. In *27th Annual European Symposium on Algorithms (ESA 2019)*, volume 144 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 12:1–12:14. ISBN 978-3-95977-124-5. ISSN 1868-8969.
- Bandyapadhyay, S.; Inamdar, T.; Pai, S.; and Varadarajan, K. 2019b. A Constant Approximation for Colorful k -Center. In *27th Annual European Symposium on Algorithms (ESA 2019)*, volume 144 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 12:1–12:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Bera, S.; Chakrabarty, D.; Flores, N.; and Negahbani, M. 2019. Fair Algorithms for Clustering. In *Advances in Neural Information Processing Systems* 32, 4954–4965.
- Bercea, I. O.; Groß, M.; Khuller, S.; Kumar, A.; Rösner, C.; Schmidt, D. R.; and Schmidt, M. 2019. On the Cost of Essentially Fair Clusterings. In *APPROX/RANDOM 2019*, volume 145, 18:1–18:22.
- Brubach, B.; Chakrabarti, D.; Dickerson, J. P.; Khuller, S.; Srinivasan, A.; and Tsepenekas, L. 2020. A Pairwise Fair and Community-preserving Approach to k -Center Clustering. In *ICML*.
- Brubach, B.; Chakrabarti, D.; Dickerson, J. P.; Srinivasan, A.; and Tsepenekas, L. 2021. Fairness, Semi-Supervised Learning, and More: A General Framework for Clustering with Stochastic Pairwise Constraints. In *AAAI*.
- Chakrabarti, D.; Esmaeili, S. A.; Dickerson, J. P.; Srinivasan, A.; and Tsepenekas, L. 2022. A New Notion of Individually Fair Clustering: α -Equitable k -Center. In *AISTATS*.
- Chen, X.; Fain, B.; Lyu, L.; and Munagala, K. 2019. Proportionally Fair Clustering. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 1032–1041. PMLR.
- Chierichetti, F.; Kumar, R.; Lattanzi, S.; and Vassilvitskii, S. 2017. Fair Clustering Through Fairlets. In *Advances in Neural Information Processing Systems* 30.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness Through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12.
- Esmaeili, S. A.; Brubach, B.; Srinivasan, A.; and Dickerson, J. P. 2021. Fair Clustering Under a Bounded Cost. In *NeurIPS*.
- Esmaeili, S. A.; Brubach, B.; Tsepenekas, L.; and Dickerson, J. P. 2020. Probabilistic Fair Clustering. In *NeurIPS*.
- Ghadiri, M.; Samadi, S.; and Vempala, S. 2021. Socially Fair k -Means Clustering. *FACCT '21*.
- Goyal, D.; and Jaiswal, R. 2021. Tight FPT Approximation for Socially Fair Clustering.
- Harris, D. G.; Pensyl, T.; Srinivasan, A.; and Trinh, K. 2019. A Lottery Model for Center-Type Problems With Outliers. *ACM Trans. Algorithms* 15(3). ISSN 1549-6325. doi:10.1145/3311953. URL <https://doi.org/10.1145/3311953>.
- Jia, X.; Sheth, K.; and Svensson, O. 2021. Fair colorful k -center clustering. *Mathematical Programming* 1–22.
- Jones, M.; Nguyen, H.; and Nguyen, T. 2020. Fair k -centers via maximum matching. In *International Conference on Machine Learning*, 4940–4949. PMLR.
- Jung, C.; Kannan, S.; and Lutz, N. 2019. A Center in Your Neighborhood: Fairness in Facility Location.

Kleindessner, M.; Awasthi, P.; and Morgenstern, J. 2019. Fair k -Center Clustering for Data Summarization. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 3448–3457. PMLR.

Kleindessner, M.; Awasthi, P.; and Morgenstern, J. 2020. A Notion of Individual Fairness for Clustering.

Kleindessner, M.; Samadi, S.; Awasthi, P.; and Morgenstern, J. 2019. Guarantees for spectral clustering with fairness constraints. In *International Conference on Machine Learning*, 3458–3467. PMLR.

Mahabadi, S.; and Vakilian, A. 2020. Individual Fairness for k -Clustering. In *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research, 6586–6596. PMLR.

Makarychev, Y.; and Vakilian, A. 2021. Approximation Algorithms for Socially Fair Clustering. In *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, 3246–3264. PMLR.

Micha, E.; and Shah, N. 2020. Proportionally Fair Clustering Revisited. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, volume 168 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 85:1–85:16. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. ISBN 978-3-95977-138-2. ISSN 1868-8969.

Negahbani, M.; and Chakrabarty, D. 2021. Better Algorithms for Individually Fair k -Clustering. *Advances in Neural Information Processing Systems* 34.

Rösner, C.; and Schmidt, M. 2018. Privacy Preserving Clustering with Constraints. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 96:1–96:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-076-7. ISSN 1868-8969.

Thejaswi, S.; Ordozgoiti, B.; and Gionis, A. 2021. Diversity-Aware k -median: Clustering with Fair Center Representation. In *Machine Learning and Knowledge Discovery in Databases. Research Track*. Springer International Publishing.

Vakilian, A.; and Yalçiner, M. 2021. Improved approximation algorithms for individually fair clustering. *arXiv preprint arXiv:2106.14043* .