Deciding the Future of Refugees: Rolling the Dice or Algorithmic Location Assignment?

Clara Strasser Ceballos[1], Christoph Kern[1,2]

[1]Institute of Statistics, Ludwig-Maximilian University of München, Ludwigstr. 33, 80539 München, Germany
[2]Munich Center for Machine Learning (MCML), Oettingenstraße 67, 80538 München, Germany

Abstract
Upon arrival in Germany, refugees are distributed among the 16 federal states. This distribution decision is based on a fixed formula consisting of two components: tax revenue and the population size of the federal state. Research suggests that optimal refugee-location matching enhances refugee integration into the labor market. However, the current mechanism fails to align refugees' characteristics with their assigned locations, resulting in a missed opportunity to leverage synergies. To this end, we use comprehensive refugee data in Germany and exploit an existing machine learning matching tool to assign refugees to states algorithmically. Our findings reveal potential improvements in refugee employment, depending on the modeling setup. Our study provides two key contributions. First, we evaluate the effectiveness of an algorithmic matching tool within Germany. Second, we investigate the fairness implications of such an algorithmic decision-making tool by evaluating the impact of different train data setups on group-specific model performance.

Keywords
refugee assignment, geographic matching, labour market integration, subgroup performance

1. Introduction

Between 2014 and 2016, Germany faced one of the largest refugee influxes since World War II [1]. Almost one million people, including nearly half a million Syrians, sought protection during this period [2]. Persecution, unrest, and conflict around the world make the influx of people seeking protection a recurring challenge [3].
The enduring nature of wars and conflicts that force people to seek refuge in Germany emphasizes the urgent need for successful refugee integration. Integration is not uniformly conceptualized [4], but can be understood as both a process and final aim of mutual adaptation between refugees and host society members [5, 6]. Integration comprises four central dimensions: structural, cultural, social, and emotional [7]. The structural dimension includes the labor market participation of refugees. The latter is crucial for successful integration as it promotes financial independence and facilitates interaction with members of the host society [5]. Thus, our study focuses primarily on the structural dimension, but can be applied to other dimensions. Research suggests that the success of refugees' economic integration may depend on the location to which they are assigned, as specific locations may be better suited to certain

EWAF'24: European Workshop on Algorithmic Fairness, July 01–03, 2024, Mainz, Germany
Clara.StrasserCeballos@stat.uni-muenchen.de (C. S. Ceballos); Christoph.Kern@stat.uni-muenchen.de (C. Kern)
0009-0008-5330-3220 (C. S. Ceballos); 0000-0001-7363-4299 (C. Kern)

characteristics of refugees [8, 9]. In Switzerland, for instance, French-speaking refugees face better employment opportunities in French-speaking cantons [8]. The potential existence of an optimal location-refugee match underlines the importance of addressing this issue in the German context.

Thus far, the allocation of refugees and asylum seekers to federal states in Germany has been guided by a distribution key called "Koenigsstein key" [10, 11]. In line with this key, refugees and asylum seekers are distributed among the federal states according to two state characteristics - tax revenue and population size [11]. Thus, instead of exploiting historical synergies between refugees and federal states for allocation, refugees are randomly assigned [12]. Additionally, once allocated and granted a temporary or permanent residence permit, refugees and asylum seekers are restricted to residing in the assigned federal state for up to three years [9]. This prompts us to question whether an allocation method informed by integration outcomes of past refugee-location matching could improve employment integration of new arriving refugees.

To this end, we employ a matching tool called GeoMatch. The tool was designed and developed by the Immigration Policy Lab [8, 13]. GeoMatch draws on supervised machine learning to predict refugee integration outcomes in potential assignment locations [8]. Further, the tool applies an optimal matching approach to strategically assign refugees to locations where the probability of a desired integration outcome is maximized [8]. The tool has been piloted by the Swiss State Secretariat for Migration (SEM) in Switzerland since 2020 and by the Lutheran Immigration and Refugee Service (LIRS) in the U.S. since 2023 [14]. In our study, we use and critically assess the GeoMatch tool for the first time in the German context. For this purpose, we draw on data from the German Socio-Economic Panel (SOEP) and collect comprehensive information on the socio-demographic profile of refugees and migrants, including their assigned federal state and employment outcomes in Germany [15, 16, 17].

1.1. Contribution

The contribution of our study is twofold. First, we use the matching tool, GeoMatch, for the first time in the German context, and evaluate its effectiveness in improving refugees' employment opportunities. While the study presenting GeoMatch was based on data from Switzerland and the U.S. [8], we draw on data from Germany. Second, we investigate the impact of different training setups on model performance across groups and downstream assignments. Specifically, we consider two training data configurations: one consisting exclusively of data on refugees and asylum seekers, and a second covering refugees and migrants. This approach allows us to mimic scenarios where information about new incoming refugee populations is scarce, and to asses model performance when the train data cover different sub-populations.

1.2. Related Work

Our study is related to literature investigating algorithmic refugee allocation. In particular, this includes the study that introduces the GeoMatch tool [8], as well as studies that build on the matching tool and extend it in several aspects. For instance, by incorporating refugee preferences [18]; by extending the tool's functionality to serve as a decision-making tool for economic immigrants [19]; by addressing the operational burden on resettlement agencies and incorporating time allocation balancing [20]; by exploring the fairness of the assignments [21]. It is worth noting that these studies are based on data from Switzerland and the U.S. Furthermore, our study is related to the literature discussing the impact of distributional shifts and training setups on model predictions [22, 23]. Finally, our study links to research that studies (sub)group-specific performance of algorithmic decision-making systems [24, 25, 26].

2. Methods

2.1. Data

This study draws on two German longitudinal studies on refugees and migrants from the German Socio-Economic Panel (SOEP): the IAB-BAMF-SOEP and the IAB-SOEP-MIG [16, 17]. The IAB-BAMF-SOEP is a longitudinal survey of refugees and asylum seekers in Germany [16, 27]. Responsible for the study are the Institute for Employment Research (IAB), the Research Centre of the Federal Office for Migration and Refugees (BAMF-FZ), and the SOEP. The survey is conducted annually since 2016. The study collects representative information on the refugees and asylum seekers who arrived in Germany between January 2013 and September 2022 by drawing random samples from the Central Register of Foreigners (AZR). The IAB-SOEP-MIG sample is a longitudinal survey of migrants in Germany [17, 28]. Responsible for the study are the IAB and the SOEP. The study is conducted annually since 2013 and collects representative information on the people who immigrated to Germany since 1995.
We construct the following variables from both data sets. First, our binary outcome variable indicating whether the person was employed within the first year(s) of arrival in Germany. Second, our predictor variables, including country of origin, German level in speaking, writing, and reading before arrival, vocational training before arrival, education level before arrival, and further demographic characteristics. Finally, we use the information on the first assigned federal state of residence recorded in both studies.

2.2. Analytical Strategy

First, we set up the data. Subsequently, we employ the matching tool comprising three key stages: modeling, mapping, and matching [8].

Data Setup The training data contains historical data and is used to fit the prediction model. In our study, we explore two training scenarios, each containing information for different sub-populations: (1) only refugees and asylum seekers, (2) both refugees and migrants. The training data includes information solely for individuals who arrived in Germany before 2016. Whereas we generate two different training data sets, we only create one test data set. The test data set contains information only on refugees who arrived in Germany from 2016 onwards.

Modeling In the first stage of the algorithm, we fit prediction models using the training data. The process unfolds as follows. First, we divide the selected training data set into subsets for each assigned federal state, i.e., one subset contains, e.g., all data of refugees and asylum seekers assigned to the federal state of Hesse upon their arrival. Second, we fit a model on each training subset. We rely on gradient-boosted trees for model fitting as they offer several benefits and align with established practices in prior studies [8, 18, 29, 30]. Third, we generate employment predictions for each individual in the test data set for each federal state.

Mapping In the second stage of the algorithm, we transform the generated individual-level predicted probabilities of employment at each federal state in the test set to case-level predicted probabilities. This transformation is necessary because some individuals in the test data set may belong to a "case," e.g. if they are members of the same family and are therefore assigned to the same federal state. The case-level metric is the probability that at least one individual in the case finds employment at a given federal state.

Matching In the final stage of the algorithm, we assign each case to the federal state where an optimality criterion is satisfied, considering existing constraints. While both the optimality criterion and the constraints are adjustable, we adopt the following approach in our analysis: We assign cases across states to maximize the global average employment probability. Based on the test data, we determine the capacity of cases that each state can accommodate. We use this as a capacity constraint. This algorithm stage produces the following results: for each individual in the test data set, in addition to the actual state assignment and employment outcome, we obtain the optimal state assignment and the predicted probability of employment in that state. This comprehensive information allows us to evaluate the (sub)group effectiveness of the GeoMatch tool in different scenarios.

3. Preliminary Results

The results of our ongoing statistical research can be divided into two parts. The first part covers the potential impact results of the matching tool in the German context. The preliminary results suggest an improvement in global average employment, if refugees were matched algorithmically rather than by the default allocation. In Figure 1, we observe a relative increase in average employment of 135% two, 123% three and 113% four years after arrival when refugees are algorithmically assigned. This results are obtained when using training data covering only refugees. We observe similar results for training scenario 2. High employment gains are also observed in the study presenting the GeoMatch tool, where a relative increase in average employment of 40% to 70% is observed for the Swiss and US contexts [8]. The second part presents the model performance results, evaluated by various metrics like ROC-AUC and PR-AUC, for each train set configuration. We show how the model performance changes when the training data covers different sub-populations. This includes, for instance, determining whether the inclusion of additional data, such as data from migrants, improves the performance of the 16 fitted models.

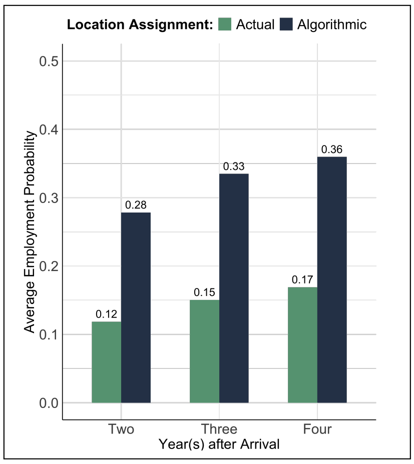

Figure 1: Employment Gains

References

[1] Federal Office for Migration and Refugees (BAMF), Figures on Asylum, 2024. URL: https://www.bamf.de/EN/Themen/Statistik/Asylzahlen/asylzahlen.html.

[2] Statistisches Bundesamt (DESTATIS), Persons Seeking Protection by Protection Status, Year of First Entry and Countries of Origin (Citizenship), 2023. URL: https://www.destatis.de/EN/Themes/Society-Environment/Population/Migration-Integration/Tables/protection-year-off-first-entry-status.html.

[3] UNHCR (The UN Refugee Agency), Global Trends, 2024. URL: https://www.unhcr.org/global-trends.

[4] F. Adam, S. Föbker, D. Imani, C. Pfaffenbach, G. Weiss, C.-C. Wiegandt, “Lost in Transition”? Integration of Refugees into the Local Housing Market in Germany, Journal of Urban Affairs 43 (2021) 831–850. doi:10.1080/07352166.2018.1562302.

[5] A. Ager, A. Strang, Understanding Integration: A Conceptual Framework, Journal of Refugee Studies 21 (2008) 166–191. doi:10.1093/jrs/fen016.

[6] J. Franzke, J. M. Ruano de la Fuente, New Challenges in Local Migrant Integration Policy in Europe, in: J. Franzke, J. M. Ruano de la Fuente (Eds.), Local Integration of Migrants Policy : European Experiences and Challenges, Palgrave Studies in Sub-National Governance, Springer International Publishing, Cham, 2021, pp. 1–9. doi:10.1007/978-3-030-50979-8_1.

[7] H. Esser, Assimilation, Ethnische Schichtung oder Selektive Akkulturation? Neuere Theorien der Eingliederung von Migranten und das Modell der Intergenerationalen Integration, Kölner Zeitschrift für Soziologie und Sozialpsychologie (Special Issue) 60 (2008) 81–107.

[8] K. Bansak, J. Ferwerda, J. Hainmueller, A. Dillon, D. Hangartner, D. Lawrence, J. Weinstein, Improving Refugee Integration through Data-Driven Algorithmic Assignment, Science 359 (2018) 325–329. doi:10.1126/science.aao4408.

[9] C. G. Aksoy, P. Poutvaara, F. Schikora, First Time Around: Local Conditions and Multi-Dimensional Integration of Refugees, Journal of Urban Economics 137 (2023) 103588. doi:10.1016/j.jue.2023.103588.

[10] C. G. Aksoy, P. Poutvaara, F. Schikora, S. Sardoschau, P. Jaschke, T. Kuchler, H. Vorländer, T. Heidland, M. Lücke, A. Steinmayr, et al., Schulterschluss oder Sackgasse–gelingt Europa eine wirksame Asyl-und Integrationspolitik?, ifo Schnelldienst 76 (2023) 03–28.

[11] Federal Office for Migration and Refugees (BAMF), Initial Distribution of Asylum Seekers (EASY), 2022. URL: https://www.bamf.de/EN/Themen/AsylFluechtlingsschutz/AblaufAsylverfahrens/Erstverteilung/erstverteilung-node.html.

[12] J. Tjaden, C. Spörlein, How Much Do “Local Policies” Matter for Refugee Integration? An Analytical Model and Evidence from a Highly Decentralized Country, International Migration Review (2023). doi:10.1177/01979183231205561.

[13] Immigration Policy Lab (IPL), GeoMatch, 2024. URL: https://immigrationlab.org/geomatch/.

[14] Global Compact on Refugees, GeoMatch: Connecting People to Places Using Artificial Intelligence, 2024. URL: https://globalcompactrefugees.org/good-practices/geomatch-connecting-people-places-using-artificial-intelligence.

[15] J. Goebel, M. M. Grabka, S. Liebig, M. Kroh, D. Richter, C. Schröder, J. Schupp, The German Socio-Economic Panel (SOEP), Jahrbücher für Nationalökonomie und Statistik 239 (2019)

345–360. doi:10.5684/soep.core.v38eu.
[16] H. Brücker, N. R. Rother, J. Schupp, IAB-BAMF-SOEP Befragung von Geflüchteten 2016. Studiendesign, Feldergebnisse sowie Analysen zu schulischer wie beruflicher Qualifikation, Sprachkenntnissen sowie kognitiven Potenzialen, IAB Forschungsbericht 13 (2017). doi:10.5684/soep.core.v38eu.
[17] H. Brücker, M. Kroh, S. Bartsch, J. Goebel, S. Kühne, E. Liebau, P. Trübswetter, I. Tucci, J. Schupp, The new IAB-SOEP Migration Sample: An Introduction into the Methodology and the Contents, SOEP Survey Paper 216 (2014). doi:10.5684/soep.core.v38eu.
[18] A. Acharya, K. Bansak, J. Hainmueller, Combining Outcome-Based and Preference-Based Matching: A Constrained Priority Mechanism, Political Analysis 30 (2022) 89–112. doi:10.1017/pan.2020.48.
[19] J. Ferwerda, N. Adams-Cohen, K. Bansak, J. Fei, D. Lawrence, J. M. Weinstein, J. Hainmueller, Leveraging the Power of Place: A Data-Driven Decision Helper to Improve the Location Decisions of Economic Immigrants, Papers 2007.13902, arXiv.org, 2020. doi:10.48550/arXiv.2007.13902.
[20] K. Bansak, E. Paulson, Outcome-Driven Dynamic Refugee Assignment with Allocation Balancing, in: Proceedings of the 23rd ACM Conference on Economics and Computation, EC ’22, Association for Computing Machinery, New York, NY, USA, 2022, p. 1182–1183. doi:10.1145/3490486.3538340.
[21] D. Freund, T. Lykouris, E. Paulson, B. Sturt, W. Weng, Group Fairness in Dynamic Refugee Assignment, in: Proceedings of the 24th ACM Conference on Economics and Computation, EC ’23, Association for Computing Machinery, New York, NY, USA, 2023, p. 701. doi:10.1145/3580507.3597758.
[22] J. G. Moreno-Torres, T. Raeder, R. Alaiz-Rodríguez, N. V. Chawla, F. Herrera, A Unifying View on Dataset Shift in Classification, Pattern Recognition 45 (2012) 521–530. doi:10.1016/j.patcog.2011.06.019.
[23] K. Bansak, E. Paulson, D. Rothenhäusler, Learning under Random Distributional Shifts, 2023. doi:10.48550/arXiv.2306.02948.
[24] M. P. Kim, A. Ghorbani, J. Zou, Multiaccuracy: Black-Box Post-Processing for Fairness in Classification, in: Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, AIES ’19, Association for Computing Machinery, New York, NY, USA, 2019, pp. 247–254. doi:10.1145/3306618.3314287.
[25] S. Desiere, L. Struyven, Using Artificial Intelligence to classify Jobseekers: The Accuracy-Equity Trade-off, Journal of Social Policy 50 (2021) 367–385. doi:10.1017/S0047279420000203.
[26] M. v. Zahn, O. Hinz, S. Feuerriegel, Locating Disparities in Machine Learning, in: 2023 IEEE International Conference on Big Data (BigData), 2023, pp. 1883–1894. doi:10.1109/BigData59044.2023.10386485.
[27] DIW, IAB-BAMF-SOEP-Befragung Geflüchteter 2021, Daten der Jahre 2016-2021, 2023. doi:10.5684/soep.iab-bamf-soep-mig.2021.
[28] DIW, IAB-SOEP Migrationsstichproben 2021, Daten der Jahre 2013-2021, 2023. doi:10.5684/soep.iab-soep-mig.2021.
[29] J. H. Friedman, Greedy Function Approximation: A Gradient Boosting Machine, The Annals of Statistics 29 (2001) 1189–1232. Publisher: Institute of Mathematical Statistics.

[30] T. Hastie, R. Tibshirani, J. Friedman, The Elements of Statistical Learning, Springer Series in Statistics, Springer, New York, NY, 2009. URL: http://link.springer.com/10.1007/978-0-387-84858-7. doi:10.1007/978-0-387-84858-7.