

Distributed Calibration of Agent-based Models

Aditi Garg
Amazon Web Services
Boston, MA

Ayush Chopra
Massachusetts Institute of Technology
Cambridge, MA

ABSTRACT

Agent-based models (ABMs) simulate complex systems by modeling the interactions between individual agents. Calibrating ABMs to real-world data is critical for their practical utility, but is hindered by the fact that granular data is often siloed across institutions due to privacy concerns. We propose a new protocol for distributed calibration of ABMs that allows institutions to collaborate on model calibration without sharing raw data. The protocol splits the calibration neural network (CalibNN) between the data clients and a central server. Each client generates embeddings from their local data and transmits them to the server, which merges the embeddings to calibrate the ABM. Gradients are propagated back to the clients to update their local models. On preliminary experiments simulating the COVID-19 pandemic, we find the distributed protocol achieves calibration accuracy on par with centralized calibration using pooled data. This demonstrates the potential to leverage sensitive data to improve ABMs while preserving privacy.

KEYWORDS

Agent-based Modeling, Distributed Machine Learning, Calibration

ACM Reference Format:

Aditi Garg and Ayush Chopra. 2024. Distributed Calibration of Agent-based Models. In *Proceedings of Workshop on Epidemiology meets Data Mining and Knowledge discovery (epiDAMIK at KDD 2024)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Agent-based models (ABM) comprise a collection of agents that act and interact within a computation world. They enable simulation of interventions by modeling the interplay between individual behaviors and environmental dynamics and hence are valuable for addressing policy questions. ABMs have helped study a multitude of complex systems across epidemiology [7, 21, 24], economics [5, 6], and disaster response [18, 19]. During the COVID-19 pandemic, ABMs helped measure the effectiveness of lockdowns [32], evaluate immunization protocols [33] and prioritize testing schedules [25]. Their utility for practical decision making requires capturing realistic environment dynamics, integrating with real-world data streams while efficiently simulating million-size population.

Their widespread adoption has been hindered by two concerns: a) computational resources required to simulate, calibrate and analyze

an ABM; b) need for microdata to generate underlying populations. Recent advances in deep learning have addressed some of the computational challenges associated with ABMs. First, vectorized programming has helped ABM scale to million-size populations [10, 13]; and differentiable programming has enabled using gradient-based learning to calibrate structural parameters [11, 13], conduct sensitivity analysis in zero-shot using autograd [31] and compose ABMs with deep neural network to ingest heterogeneous data sources for guiding the calibration [11]. As a result, it is now possible to conduct simulation, calibration and analysis of country-scale ABMs on commodity hardware. Unfortunately, improvements in the computational efficiency of ABMs is of little value if the quality of the underlying population microdata is poor.

Conventional ABMs rely on data generated using sparse summary statistics derived from real-world observations. Privacy considerations, not data sparsity, are the cause for this limited granularity as data remains siloed across diverse institutions [9]. Consider the case of epidemic modeling, which requires leveraging multimodal sources to capture feedback loop between individual behavior and disease dynamics. This include access to demographic data from the US census to define population; mobility traces from SafeGraph and Google Mobility to capture interaction patterns; response on digital surveys (such as facebook) to understand compliance and risk behavior; health data (from CDC) to capture individual susceptibility; insurance and employment data (from bureau of labor statistics) to contextualize socio-economic considerations. Assimilating these dynamic data streams into the ABM calibration process has become essential for address observational gaps and generate real-time insights.

However, such information is usually sensitive and personally identifiable, and hence securely siloed under regulatory oversight. Previous attempts to release such demographic, mobility and health data for scientific research have resulted in leaks that exposed agents' personal information [1, 14, 23]. Hence, the data is shared in form of sparse summary statistics generated using privacy methods to protect individual information on an aggregate level, as used for US census [8] and Google mobility data [3]. In practice, these privacy consideration come at utility loss which inhibits real-world utility of the released data. As ABMs continue to scale towards one-to-one representations of real-world systems, there remains a fundamental limitation in their modeling potential as long as data access is constrained. Alleviating this challenge is the focus of our research.

In this paper, we introduce a mechanism for distribution simulation that can allow the institutions to collaborate but without sharing their individual data. We leverage the composability of differentiable ABMs to calibrate simulation parameters while keeping the data decentralized. We build upon existing work in collaborative machine learning. The simulation model, initialized using synthetic population, is key on a centralized server. For each of the datasets,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

epiDAMIK at KDD 2024, August 26, 2024, Barcelona, Spain

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06

<https://doi.org/XXXXXXXX.XXXXXXX>

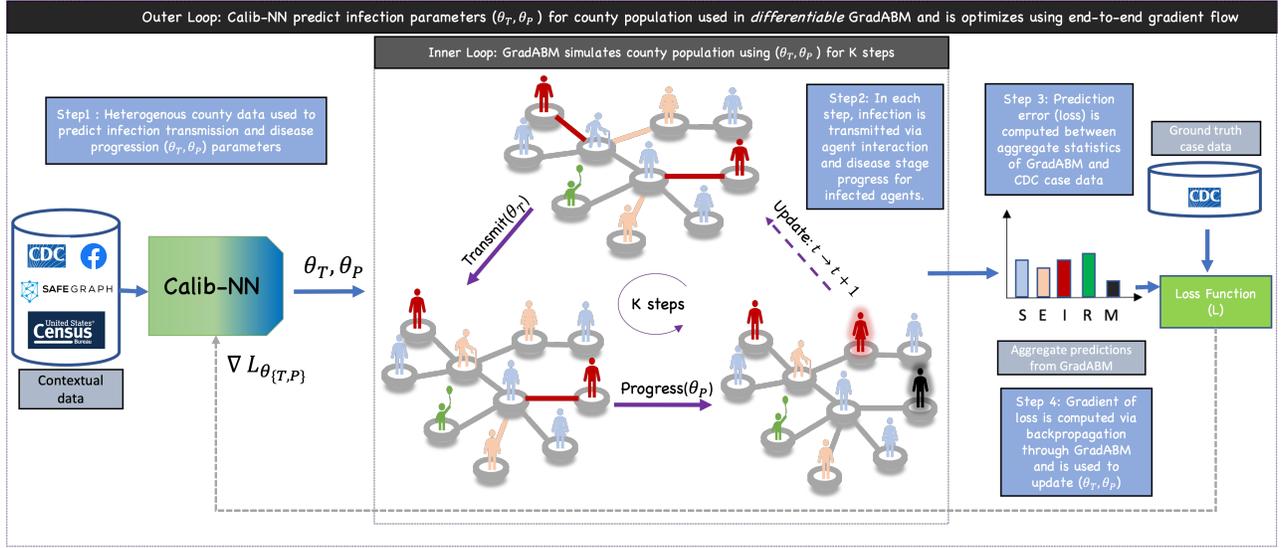


Figure 1: Background - Centralized Calibration of ABM (source [11]). The protocol consists of four steps: i) heterogeneous macro-level population data (CDC, census, behavioral, survey) is input to a calibration model (CalibNN) to predict epidemiological parameters (θ_T, θ_P) . ii) (θ_T, θ_P) are used run K forward steps of the fully-differentiable epidemiological model (GradABM) which simulates micro-level infection transmission (Transmit) and disease progression (Progress) over individual contact networks. Disease statistics are aggregated (Aggregate) at end of K steps to obtain the macro-level simulation output (\hat{y}). iii) Error between predicted \hat{y} and real-world case statistics (y) is used to define a loss ($L(\hat{y}; y)$). iv) Gradient of this loss is computed by automatic differentiation through the micro-level GradABM to update weights of macro-level CalibNN using gradient descent. In practice, the contextual data may be siloed due privacy and logistical concerns. This work proposes a technique to execute calibration while keeping the contextual data decentralized.

client embeddings are generated with siloed data and communicated on the server. The server aggregates the client embeddings to generate simulation parameters and execute simulation. The simulated gradients are propagated back to the clients to update CalibNN. Periodically, the CalibNN across clients are synchronized via federated averaging. This method helps achieve effective calibration, without compromising sensitivity of data. Second, we also use this method for hierarchical calibration of multiple simulations, which achieves in improved performance across different ecosystems. While our analysis is focused on ABMs, the proposed mechanism can be generalized to any differentiable simulation.

2 BACKGROUND

First, we review differentiable ABM and how they can be composed with deep neural networks to ingest heterogeneous data for calibration. Second, we review how this can be bridged with federated deep learning to enable calibration on siloed data.

2.1 Differentiable Agent-based Modeling

Consider an ABM where $X_N = F(X_0; \theta)$ where F is a stochastic N -step simulator with input state X_0 , structural parameters θ and output state X_N . The ABM is composed of repeated simulation steps of function g s.t. $X_{t+1} = g(X_t; \theta)$. The ABM is differentiable if the gradient

$$\eta = \nabla_{\theta} \mathbb{E}[F(\theta)] \quad (1)$$

exists and can be computed. In such case, for a smooth objective function $Y = c(X_n)$, the partial derivatives dY/dX_0 and $dY/d\theta$ can be computed using autograd. g is typically a stochastic function (eg: individual agent's probability of infection) and research in differentiable ABM has proposed methods to differentiate through discrete and continuous stochasticity [4, 11, 32] and also built frameworks built to generalize these capabilities across domains [13].

Calibration of an ABM refers to the process of estimating a set of structural parameters $\hat{\theta}$, or a probability distribution over θ , such that $y = c(X_n)$ is consistent with real-world data. There is extensive literature on how to calibrate ABMs with techniques including approximate Bayesian computation [29], neural likelihood and posterior estimation [15], among others. When the ABM is differentiable, gradient-assisted calibration techniques can be used. This enables the ABM to be composed with neural networks into end-to-end differentiable pipelines for calibration[11] and is empirically shown to improve performance. In these protocols, the structural parameters for the ABM can be optimized by training weights of the calibration neural network. This allows to guide calibration with heterogeneous data sources; as well as jointly calibrate multiple simulators improving sample efficiency [11]. The centralized protocol is visualized in figure 1. Consider, a calibration neural network C (CalibNN) with weights ϕ and contextual data D . Then, the structural parameters are generated by $\theta = C(D; \phi)$ and the ABM is executed at $F(\theta)$. The calibration gradient is:

$$\eta_c = \nabla_{\phi} \mathbb{E}[F(C(D; \phi))] \quad (2)$$

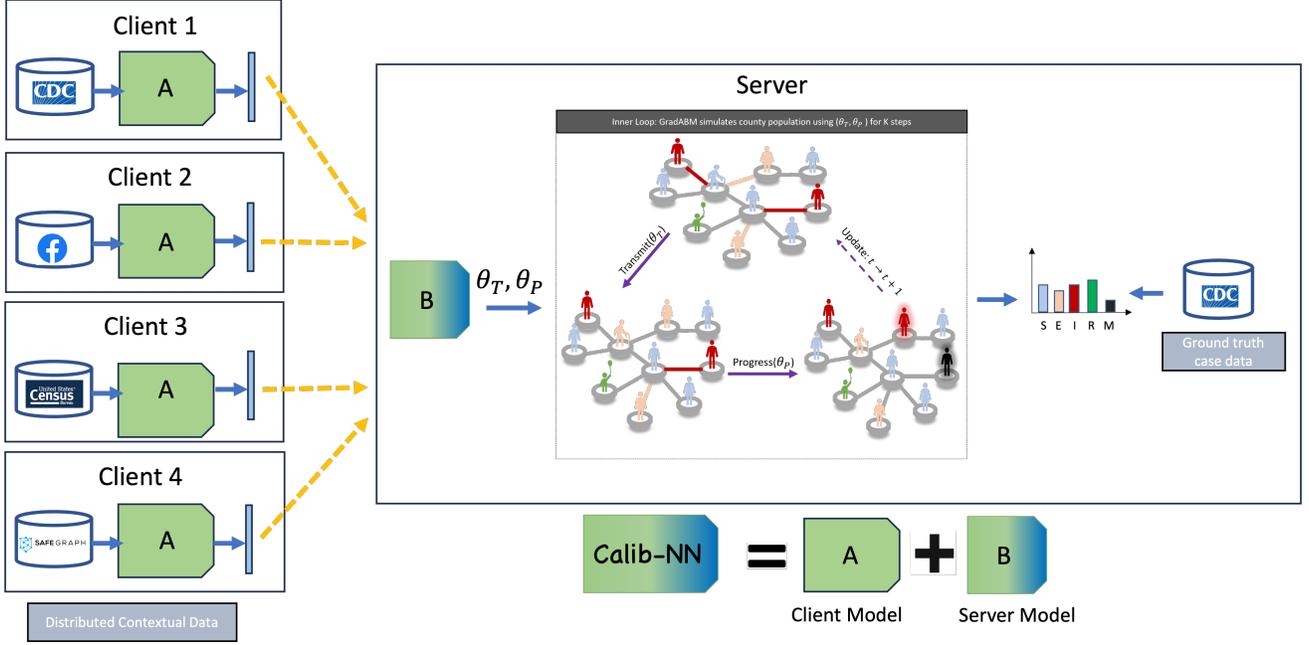


Figure 2: Proposed - Distributed Calibration of Agent-based Models. We execute calibrate with contextual data distributed across multiple clients. To achieve this, we split CalibNN between the clients and server to ensure calibration on multi-modal data, without centralizing siloed information. Each client uses its local data to generate embeddings (in yellow) which are transmitted to server and used to predict structural parameters (θ_T, θ_P) to execute the simulation. We describe the protocol in section 4. In principle, our framework can be extended to any differentiable simulation, even beyond agent-based models.

In practice, the contextual data D can be siloed due to privacy considerations. The goal of this work is to calibrate the ABM while keeping the data decentralized. To achieve this, we draw upon work in distributed machine learning which we briefly introduce below.

2.2 Distributed Machine Learning

This is a setting where many clients collaborative train a model under the orchestration of a central server, while keeping the training data decentralized. As strict regulations emerge for data capture and storage, such as GDPR [17] and CCPA [2], distributed deep learning is being used to enable privacy-aware learning across a wide range of applications [16, 20, 30]. The key insight is to decentralize the model training instead of centralizing the training data. In a general distributed deep learning pipeline, there are multiple rounds of *training* and *synchronization* steps where the model is *trained* with local client data in each round and updates made by multiple clients are *synchronized* by the server into a global model. Federated learning (FL) [22, 27, 28] and split learning (SL) [12, 34] are two paradigms which provide different designs for the training and synchronization algorithms. Briefly, FL executes the training entirely on local client devices and synchronizes the models by averaging gradients on the server; while SL distributes model training between client and server and synchronizes the model by constraining clients to update shared parameters on the server model. We refer the readers to [26] for a detailed review of these paradigms. In this paper, we generalize these protocols to execute calibration of differentiable ABMs. Specifically, we extend split learning design to

distribute the CalibNN between client and server to enable learning on the siloed data during calibration.

3 METHOD: DISTRIBUTED AGENT-BASED MODELING

We consider a scenario where the contextual dataset d is composed of multiple components $((d_1, d_2, d_3, d_4))$ which are siloed across distinct entities (called *data clients*). As motivated before, this is plausible, given the heterogeneous data requirements for ABM calibration. We consider an ABM F with structural parameters θ and initial state X_0 . The protocol is visualized in figure 2 and described below. The configuration is composed of 4 clients and 1 server. The calibration network C is split as C_a and C_b , following the split learning design [12], which are executed at client and server respectively. At the beginning, the server initializes a calibration parameters ϕ and transmits the model weights ϕ_a to all the data clients. Each episode executes as follows:

- First, each of the data client ($i = 1$ to 4) generate embeddings $E_i = C_a(d_i; \phi_a)$ and transmits the embeddings to the server.
- Second, the server merges the embeddings ($E = ([E_1, E_2, E_3, E_4])$) and generates the structural parameters $\theta = C_b(E; \phi_b)$.
- Third, the server uses θ to execute ABM simulation to generate aggregate predictions $\hat{y} = F(X_0; \theta)$ and apply a loss function $L = L(\hat{y}, y)$.
- Fourth, the gradients of the loss function $\nabla_{\phi_a} L$ are used to calibrate the structural parameters by: i) optimizing ϕ_b on

server using $\nabla_{\phi_b} L$ and ii) transmitting gradient ($\nabla_{\phi_b} L$) to clients to update ϕ_a on each client.

- Fifth, the models (ϕ_a) at each client are periodically synchronized by the server, using federated averaging [28], to ensure consistent convergence and optimize over non-iid distributions.

We execute this process for multiple episode and the final calibrated model can be used for simulation. We empirically validate the effectiveness of the protocol in the next section.

4 EXPERIMENTS

To illustrate the effectiveness of our framework, we conduct experiments for COVID-19 counties of the state of Massachusetts, USA. We follow experimental protocol from [11] and consider an epidemiological ABM and calibrate R_0 (θ_T in figure 1 and 2) to ground-truth data on death statistics (obtained from the CDC). The contextual data (d) has 4 signals:

- Mobility Signals - This record people visiting points of interest (POIs) in various regions and are obtained from Google Community Surveys
- Symptomatic surveys - This includes statistics on COVID-like illness obtained from voluntary surveys conducted by Facebook
- Symptom search data - This includes records of searches related to symptoms for multiple conditions and syndromes across the US and different states collected by Google
- Deaths and Hospitalizations - This includes data of number of hospitalizations and deaths obtained from the department of Health & Human Services and the CDC

We vertically partition the training dataset with each of the 4 signal stored in separate clients; which is more consistent with practical considerations. In practice, the same calibration protocol can be used when the data is horizontally partitioned. The architecture of the Calibration network and the simulation design follows from [11].

We compare calibration performance with the centralized protocol (fig 1) and the proposed distributed mechanism (fig 2). We evaluate the performance using standard metrics like normal deviation(ND), root mean squared error (RMSE) and mean absolute error(MAE). Results in Table 1 shows that we distributed mechanism can preserve the performance of centralized baseline. This is an encouraging observation since it highlights the potential to use more granular signal for calibration, in the future, while ensuring data privacy.

5 CONCLUSION

Granular and heterogeneous data is critical for calibrating agent-based models to real-world dynamics. However, much of this data is siloed across institutions due to valid privacy concerns. We have introduced a protocol to enable distributed calibration of ABMs without centralizing raw data. By splitting the calibration network between data clients and a server, and communicating only embeddings and gradients, it allows leveraging multi-modal data while keeping it decentralized. Our experiments on a COVID-19 ABM demonstrate the distributed protocol can match the accuracy of centralized calibration using pooled data. This is a promising result,

Calibration data	ND	RMSE	MAE
No data	2.79 \pm 0.65	160.55 \pm 23.59	69.16 \pm 10.84
Centralized [11]	1.15 \pm 0.24	53.86 \pm 14.61	28.43 \pm 6.39
Distributed (Proposed)	1.23 \pm 0.35	54.74 \pm 10.61	30.21 \pm 8.45

Table 1: We compare the proposed distributed calibration mechanism with two baselines: a) centralized calibration where all contextual data for calibNN is available in one-place, b) centralized calibration without using any contextual data (use gradient-based learning with CalibNN, as equation 1). We observe that the proposed method achieves consistent performance with the centralized data mechanism, while preserving privacy. We also observe that using contextual data is consistently better than not using the data. This work is promising in ability to ingest more granular data for calibration, while protecting privacy.

as it provides a path to improve ABMs with sensitive datasets that would otherwise be inaccessible. While we focused on ABMS, the core distributed learning techniques are quite general and could be extended to other types of differentiable simulations. There are several important directions for future work. First, conducting more extensive experiments across a range of ABM applications. Second, enhancing the protocol with additional privacy safeguards such as differential privacy. Finally, exploring alternative architectures and learning algorithms to further improve efficiency and robustness. Enabling privacy-preserving use of sensitive data to enhance modeling and simulation could have a profound impact across domains from epidemiology to economics. We believe bridging techniques from multi-agent systems and collaborative learning is a fruitful direction to realize this potential. Our work provides an initial step in this direction.

REFERENCES

- [1] 2021. Indonesia Probes Suspected Data Breach on COVID-19 App. *Reuters* (Aug. 2021).
- [2] Accountability Act. 1996. Health insurance portability and accountability act of 1996. *Public law* 104 (1996), 191.
- [3] Ahmet Aktay, Shailesh Bavadekar, Gwen Cossoul, John Davis, Damien Desfontaines, Alex Fabrikant, Evgeniy Gabrilovich, Krishna Gadepalli, Bryant Gipson, Miguel Guevara, et al. 2020. Google COVID-19 Community Mobility Reports: Anonymization Process Description (Version 1.1). *arXiv preprint arXiv:2004.04145* (2020). arXiv:2004.04145
- [4] Gaurav Arya, Moritz Schauer, Frank Schäfer, and Christopher Rackauckas. 2022. Automatic Differentiation of Programs with Discrete Randomness. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 10435–10447.
- [5] R. Axtell and J. Farmer. 2022. Agent Based Modeling in Economics and Finance: Past, Present, and Future. *Journal of Economic Literature* (2022).
- [6] Robert L. Axtell. 2016. 120 Million Agents Self-Organize into 6 Million Firms: A Model of the U.S. Private Sector. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems (AAMAS '16)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 806–816.
- [7] Joseph Aylett-Bullock, Carolina Cuesta-Lazaro, Arnau Quera-Bofarull, Miguel Icaza-Lizaola, Aidan Sedgewick, Henry Truong, Aoife Curran, Edward Elliott, Tristan Caulfield, Kevin Fong, Ian Vernon, Julian Williams, Richard Bower, and Frank Krauss. 2021. June: Open-Source Individual-Based Epidemiology Simulation. *Royal Society Open Science* 8, 7 (July 2021), 210506. <https://doi.org/10.1098/rsos.210506>
- [8] US Census Bureau. [n. d.]. Why the Census Bureau Chose Differential Privacy. <https://www.census.gov/library/publications/2023/decennial/c2020br-03.html>.

- [9] Ayush Chopra, Arnau Quera-Bofarull, Nurullah Giray-Kuru, Michael Wooldridge, and Ramesh Raskar. 2024. Private Agent-Based Modeling. (2024).
- [10] Ayush Chopra, Ramesh Raskar, Jayakumar Subramanian, Balaji Krishnamurthy, Esma S Gel, Santiago Romero-Brufau, Kalyan S Pasupathy, and Thomas C Kingsley. 2021. DeepABM: scalable and efficient agent-based simulations via geometric learning frameworks—a case study for COVID-19 spread and interventions. In *2021 Winter Simulation Conference (WSC)*. IEEE, 1–12.
- [11] Ayush Chopra, Alexander Rodríguez, Jayakumar Subramanian, Arnau Quera-Bofarull, Balaji Krishnamurthy, B. Aditya Prakash, and Ramesh Raskar. 2023. Differentiable Agent-based Epidemiology. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1848–1857.
- [12] Ayush Chopra, Surya Kant Sahu, Abhishek Singh, Abhinav Java, Praneeth Vepakomma, Vivek Sharma, and Ramesh Raskar. 2021. Adasplit: Adaptive trade-offs for resource-constrained distributed deep learning. *arXiv preprint arXiv:2112.01637* (2021).
- [13] Ayush Chopra, Jayakumar Subramanian, Balaji Krishnamurthy, and Ramesh Raskar. 2024. A Framework for Learning in Agent-Based Models. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems.
- [14] Joseph Cox. 2019. T-Mobile ‘Put My Life in Danger’ Says Woman Stalked With Black Market Location Data.
- [15] Joel Dyer, Patrick Cannon, J. Doyne Farmer, and Sebastian M. Schmon. 2024. Black-Box Bayesian Inference for Agent-Based Models. *Journal of Economic Dynamics and Control* 161 (April 2024), 104827. <https://doi.org/10.1016/j.jedc.2024.104827>
- [16] Alessandro Epasto, Mohammad Mahdian, Vahab Mirrokni, and Peilin Zhong. 2022. Massively parallel and dynamic algorithms for minimum size clustering. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1613–1660.
- [17] General Data Protection Regulation GDPR. 2018. General data protection regulation. URL: [https://gdpr-info.eu/\[accessed 2020-11-21\]](https://gdpr-info.eu/[accessed 2020-11-21]) (2018).
- [18] Jiaqi Ge. 2017. Endogenous rise and collapse of housing price: an agent-based model of the housing market. *Computers, Environment and Urban Systems* 62 (2017), 182–198.
- [19] Saman Ghaffarian, Debraj Roy, Tatiana Filatova, and Norman Kerle. 2021. Agent-based modelling of post-disaster recovery with remote sensing data. *International Journal of Disaster Risk Reduction* 60 (2021), 102285.
- [20] Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. 2018. Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604* (2018).
- [21] Robert Hinch, William J. M. Probert, Anel Nurtaý, Michelle Kendall, Chris Wymant, Matthew Hall, Katrina Lythgoe, Ana Bulas Cruz, Lele Zhao, Andrea Stewart, Luca Ferretti, Daniel Montero, James Warren, Nicole Mather, Matthew Abueg, Neo Wu, Olivier Legat, Katie Bentley, Thomas Mead, Kelvin Van-Vuuren, Dylan Feldner-Busztin, Tommaso Ristori, Anthony Finkelstein, David G. Bonsall, Lucie Abeler-Dörner, and Christophe Fraser. 2021. OpenABM-Covid19—An Agent-Based Model for Non-Pharmaceutical Interventions against COVID-19 Including Contact Tracing. *PLOS Computational Biology* 17, 7 (July 2021), e1009146. <https://doi.org/10.1371/journal.pcbi.1009146>
- [22] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J. Reddi, Sebastian U. Stich, and Ananda Theertha Suresh. 2021. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. *arXiv:1910.06378*
- [23] Bennett Cyphers and Jason Kelley. 2021. Illinois Bought Invasive Phone Location Data From Banned Broker Safegraph. <https://www.eff.org/deeplinks/2021/08/illinois-bought-invasive-phone-location-data-banned-broker-safegraph>.
- [24] Cliff C. Kerr, Robyn M. Stuart, Dina Mistry, Romesh G. Abeysuriya, Katherine Rosenfeld, Gregory R. Hart, Rafael C. Núñez, Jamie A. Cohen, Prashanth Selvaraj, Brittany Hagedorn, Lauren George, Michał Jastrzębski, Amanda S. Izzo, Greer Fowler, Anna Palmer, Dominic Delpport, Nick Scott, Sherrie L. Kelly, Caroline S. Bennette, Bradley G. Wagner, Stewart T. Chang, Assaf P. Oron, Edward A. Wenger, Jasmina Panovska-Griffiths, Michael Famulare, and Daniel J. Klein. 2021. Covasim: An Agent-Based Model of COVID-19 Dynamics and Interventions. *PLOS Computational Biology* 17, 7 (July 2021), e1009149. <https://doi.org/10.1371/journal.pcbi.1009149>
- [25] Daniel B Larremore, Bryan Wilder, Evan Lester, Soraya Shehata, James M Burke, James A Hay, Milind Tambe, Michael J Mina, and Roy Parker. 2021. Test sensitivity is secondary to frequency and turnaround time for COVID-19 screening. *Science advances* 7, 1 (2021), eabd5393.
- [26] Qinbin Li, Zeyi Wen, Zhaomin Wu, Sixu Hu, Naibo Wang, Yuan Li, Xu Liu, and Bingsheng He. 2021. A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering* 35, 4 (2021), 3347–3366.
- [27] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated Optimization in Heterogeneous Networks. *arXiv:1812.06127*
- [28] H. Brendan McMahan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas. 2016. Federated Learning of Deep Networks using Model Averaging. *CoRR abs/1602.05629* (2016).
- [29] Donovan Platt. 2020. A Comparison of Economic Agent-Based Model Calibration Methods. *Journal of Economic Dynamics and Control* 113 (April 2020), 103859. <https://doi.org/10.1016/j.jedc.2020.103859>
- [30] Jason Posner, Lewis Tseng, Moayad Aloqaily, and Yaser Jararweh. 2021. Federated learning in vehicular networks: Opportunities and solutions. *IEEE Network* 35, 2 (2021), 152–159.
- [31] Arnau Quera-Bofarull, Ayush Chopra, Joseph Aylett-Bullock, Carolina Cuesta-Lazaro, Anisoara Calinescu, Ramesh Raskar, and Michael Wooldridge. 2023. Don't Simulate Twice: One-shot Sensitivity Analyses via Automatic Differentiation. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (Aamas '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1867–1876.
- [32] Arnau Quera-Bofarull, Ayush Chopra, Anisoara Calinescu, Michael Wooldridge, and Joel Dyer. 2023. Bayesian Calibration of Differentiable Agent-Based Models. <https://doi.org/10.48550/arXiv.2305.15340> arXiv:2305.15340 [cs, stat]
- [33] Santiago Romero-Brufau, Ayush Chopra, Alex J Ryu, Esma Gel, Ramesh Raskar, Walter Kremers, Karen S Anderson, Jayakumar Subramanian, Balaji Krishnamurthy, Abhishek Singh, et al. 2021. Public health impact of delaying second dose of BNT162b2 or mRNA-1273 covid-19 vaccine: simulation agent based modeling study. *bmj* 373 (2021).
- [34] Praneeth Vepakomma, Otkrist Gupta, Tristan Swedish, and Ramesh Raskar. 2018. Split learning for health: Distributed deep learning without sharing raw patient data. *arXiv:1812.00564* (2018).