

Decentralized Sharing and Valuation of Fleet Robotic Data

Anonymous Author(s)

Affiliation

Address

email

Abstract:

We propose a decentralized learning framework for networked robots to trade, price, and discover valuable machine learning (ML) training data. Today’s robotic fleets, such as self-driving vehicles, can gather terabytes of rich video and LIDAR data [1] in diverse, geo-distributed environments. Often, robots in one city or home might observe training data that is commonplace for them but is actually a valuable, out-of-distribution (OoD) dataset to train robust ML models at robots elsewhere. However, simply sharing all this diverse data in cloud databases is infeasible due to limits on privacy and network bandwidth. Inspired by decentralized file-sharing protocols like BitTorrent, we propose a novel system where each robot is provisioned with a learnable privacy filter and sharing model. Importantly, this sharing model attempts to predict and prioritize which sensory percepts are of high value to *other* robotic peers using a decentralized voting and feedback mechanism. Our scheme naturally raises timely questions on data privacy and valuation as multiple companies start to deploy robots in our homes, hospitals, and roads.

Keywords: Decentralized Fleet Learning, Swarm Robotics, Data Valuation

1 Introduction

Imagine a fleet of robots that are deployed at diverse hospitals and nursing homes across the world. These robots should seamlessly interact with nurses and patients that speak different languages, have diverse preferences on human-robot interaction, and rapidly adapt to new hospital layouts and visual appearances of medical equipment. To do so, these robots can continually re-train their perception, natural language processing, and control policies using hundreds of gigabytes of rich video, audio, and LIDAR sensory streams they observe locally. Moreover, these robots should ideally share their rare, valuable training data with their peers in diverse locations in order to train more robust ML models. However, how should we incentivize heterogeneous robots (often owned by different companies) to share valuable training data while protecting human privacy and proprietary ML models?

Designing private, efficient data-sharing protocols is of prime importance today. Already, the Toyota Research Institute is advocating large-scale fleet learning in homes [2], Diligent Robotics is deploying robot swarms to assist nurses with fighting COVID-19 [3], and platforms like SafetyPool allow multiple self-driving car companies to crowd-source hazardous driving scenarios [4].

Despite this potential, our key observation is that today’s mechanisms for data sharing are largely insufficient since they are *centralized and data-agnostic*. In other words, we typically share *all* data with trusted cloud servers regardless of whether it is useful to other robots, a redundant concept in our training dataset, or private. As such, today’s mechanisms can lead to privacy concerns and waste valuable network bandwidth, cloud storage, and cloud training resources. In this paper, we argue that individual robots should have sovereignty over their data and sharing should be **selective, private, and decentralized**.

The State-of-the-Art: Historically, robotic fleet learning has been addressed by related work in cloud robotics [5, 6, 7, 8], active learning [9, 10, 11, 12], and crowd-sourced platforms such as RoboTurk [13] and RoboEarth [14]. However, these mechanisms rely on trusted cloud servers to

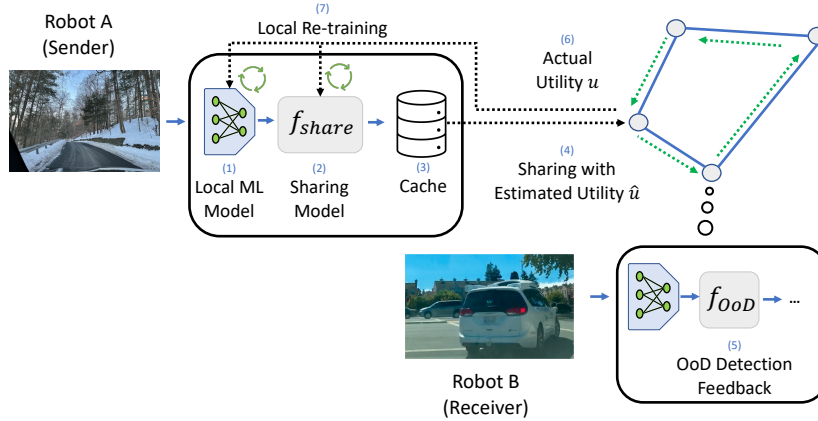


Figure 1: We collected dashcam videos from 3 diverse rural and metropolitan areas thousands of miles apart. For example, a car in a rural snowy area should have a specialized vision model for snowy terrain. While it might find a new image of snow commonplace, that very same image could be OoD and valuable to a peer thousands of miles away in a sunny environment to build a more robust model. As such, our system would *proactively* share this data using a global utility (GU) model (steps 1-3) and receive feedback from networked peers (steps 4-7).

41 aggregate data and do not address privacy or how to fairly assess the contribution of each training
 42 example to overall model accuracy. Moreover, today’s approaches to privacy in ML, such as differ-
 43 ential privacy (DP) [15, 16] and federated learning (FL) [17, 18], do not readily apply to robots that
 44 operate amongst humans. For example, DP approaches would add excessive noise to anonymize
 45 high-dimensional robotic visual data, which severely limits its utility for training ML models. Like-
 46 wise, FL allows mobile devices to train ML models on local private data and simply share gradient
 47 updates with a central cloud server. However, FL does not apply when a robot sees a novel image it
 48 cannot classify locally and must share this image with a peer to receive a ground-truth label. Finally,
 49 our proposal is inspired by classical work on consensus in multi-agent systems [19], which pre-dates
 50 our focus on trading data for ML models.

51 **Contributions:** Given the shortcomings of today’s methods, our principal contribution is to propose
 52 a communication-efficient, private, and proactive data sharing system. By proactive, we mean that
 53 data sharing can potentially allow robots to improve their models on otherwise unseen data and thus
 54 avoid failure before they actually encounter such scenarios. We now describe our proposal in detail.

55 2 System Architecture and Preliminary Results

56 Figure 1 depicts our modular system architecture. Each robot in a peer-to-peer network acts both
 57 as a sender as well as a receiver of training data. For clarity, Figure 1 illustrates a simple sharing
 58 scenario where a sender robot (left) shares presumably valuable data to receivers (right), which
 59 in turn provide feedback to the sender to adjust the sharing process. Crucially, to limit network
 60 bandwidth, training, and data annotation costs, each robot is limited to share only 1-5% of local data
 61 with its peers. Data sharing and model re-training occur in periodic rounds, such as the end of a
 62 day. While our architecture is broadly applicable to many ML models, we consider a scenario in the
 63 self-driving domain. Here, a robot in a rural snowy area shares data with a vehicle in a sunny city to
 64 build a robust ML model across multiple weather conditions. We now describe the information flow
 65 in one round of data sharing by referencing the numbered blocks in Figure 1.

66 1. Local ML Model.

67 Each robot starts with a generic ML model pre-trained on anonymized public datasets like ImageNet.
 68 However, each robot can specialize this model based on its own local data with assistance of trusted,
 69 local data annotators. In the context of robotic perception, the specialized vision model maps a new
 70 image observation X into a predicted class or object detection \hat{y} and embedding vector e .

71 **2. Global Utility (GU) Sharing Model.** The crux of our approach is to provision each robot
72 with a global utility (GU) model which learns the value (e.g. utility) of a new image X to other
73 robotic peers. One example of a utility score is the proportion of other robots in a fleet that consider
74 image X to be out-of-distribution with respect to their current vision model, which could indicate
75 a valuable training scenario. To estimate the utility of a specific image X , the GU model naturally
76 takes in the vision model’s embedding e and predictions \hat{y} . Crucially, the GU model has learnable
77 parameters θ_{share} that allow it to predict a utility score \hat{u} per image. While the utility might be
78 uniform initially, our next steps show how to estimate the true utility for each image using prioritized
79 *network feedback*.

80 **3-4. Cache, Privacy Filter, and Network Transmission.** After images are ranked by their utility,
81 the next step is to prioritize them for sharing while respecting bandwidth limits. As such, each robot
82 only caches the top 1 – 10% of images with the highest utilities \hat{u} in its local storage. Then, these
83 cached images can be blurred or altogether dropped to restrict sensitive human faces or other private
84 features. Finally, each sender can now relay the cached and anonymized images to other peers.

85 **5-6. Network Feedback Based on Out-of-Distribution Detection.** The next key step is to have
86 other robots evaluate the utility of shared data in order to eventually tune the GU model. Each shared
87 image X has a default *true* utility counter u , initialized to zero, to indicate how many robots found
88 image X interesting. Each robot then uses an OoD detector [20, 21, 22, 23, 24, 25] to determine
89 if shared image X displays novel features far from its local training distribution. If so, this robot
90 increments the utility u before relaying image X and the current utility to new robots. Finally, the
91 utility u is normalized to represent the fraction of peers that found the data interesting. This feedback
92 is then relayed back to the sender robot to re-train its GU model.

93 **7. Closing the Loop: Re-training Sharing and Local ML Models.** Crucially, each sender robot
94 has now been relayed a dataset of its original shared images X , their embeddings and predictions
95 e, \hat{y} , estimated utility \hat{u} , and true utility u from peers. As such, we can re-train the GU model to
96 update parameters θ'_{share} to improve the sharing process. Moreover, each robot can re-train its vision
97 model with new data it received with possible help from trusted local data annotators. Ideally, this
98 process will converge when each local vision model is more robust to previously OoD data.

99 **Preliminary Evaluation** We now describe a simple toy illustrative example of data sharing. We
100 consider a scenario with three networked ‘robots’, each of which has a pre-trained vision model to
101 detect handwritten digits from the benchmark MNIST dataset. Crucially, to simulate the presence of
102 rare, valuable training data, one robot also observes data from the KMNIST [26] variant of MNIST
103 while another observes data from FashionMNIST [27], which both exhibit distinct visual properties
104 from standard MNIST. Our goal is to train a GU model that learns to share KMNIST and FashionMNIST
105 data between appropriate robots so that all robots have access to all classes of training
106 data. Our evaluation metric is the final accuracy of each robot’s vision model on all three variants of
107 MNIST to demonstrate robust performance on data each robot was not initially exposed to.

108 We implemented the data sharing protocol illustrated in Figure 1. First, each robot passes a new
109 image into its local vision model, whose embedding and predictions are passed to the GU sharing
110 model. Since the GU sharing model is initially un-trained, it first randomly caches images to share
111 with other peers. Importantly, each robot can only share 10% of its training data per round due to
112 network bandwidth limits. As data is shared, each peer provides a utility score u based on whether
113 a shared image was valuable (e.g. OoD). Concretely, this OoD detector uses recently-developed
114 Bayesian uncertainty estimates to provide a scalar OoD signal to indicate if a new image is from the
115 local vision model’s training distribution. After the utility score for shared images is relayed back
116 to the sender, the sender re-trains its GU sharing model to prioritize relevant content.

117 **Results:** Figure 2 illustrates strong experimental results for our toy example of data sharing. We
118 compare our protocol based on a GU sharing model (green) with upper and lower bounds of oracle
119 and random sampling in blue and red respectively. In particular, the un-realizable oracle approach
120 has a perfect GU model that shares OoD data with perfect accuracy. In stark contrast, the random
121 policy is a lower-bound that wastes precious network bandwidth by randomly sharing images as
122 opposed to prioritizing them based on network feedback. Our key evaluation metric is the *overall*
123 *multi-class* accuracy (y-axis) across data sharing and training rounds on the x-axis.

124 Clearly, our decentralized sharing protocol (green) significantly improves each robot’s vision model
125 by sharing previously OoD training data with peers that need it most. Moreover, our scheme is




Figure 2: **Toy Experiment Results:** Initially, a robot’s local vision model does not have access to KMNIST and FashionMNIST data and hence performs poorly on these OoD concepts in round 0. Clearly, our learned sharing model (green) allows a robot to quickly accrue valuable training data from its peers, which allows the local vision model to efficiently be re-trained to match the upper-bound accuracy of an oracle data collection strategy (blue).

126 efficient and adaptive – the GU model learns to correctly cache 88.18 % of rare KMNIST data and
 127 88.96% of rare FashionMNIST data, which is significantly higher than random sharing.

128 3 Future Research Questions on Data Valuation and Privacy

129 We now describe open research questions that cross-cut ML, privacy research, and cloud robotics.

130 **Game-Theoretic Approaches to Data Valuation** In order to incentivize data sharing, each robot
 131 should be paid a fair price for the data it shares. To do so, we can adapt the recently-proposed Data
 132 Shapley [28] metric, which is inspired by cooperative game theory. Specifically, this metric consid-
 133 ers each training example to be a player in a cooperative game, which is to collectively train a large
 134 ML model. The payout or value to each player (i.e., data point) is calculated based on the marginal
 135 accuracy gain when that specific data example is used for training as opposed to when it is omit-
 136 ted, averaged over all possible subsets of training data where a specific example can appear. While
 137 promising, today’s Data Shapley metric only scales to small, centralized datasets and a homogenous
 138 ML model. However, our network feedback scheme is a natural approach to value data based on
 139 the collective score assigned by peer robots’ OoD detectors. Future research can reconcile our data
 140 valuation metric and extend Data Shapley to decentralized, heterogenous ML models.

141 **Privacy-Utility Trade-offs in Robotics** Today’s federated learning [18] approaches make a strong
 142 assumption that ground-truth labels are present on mobile clients, which allows them to train models
 143 locally on private data and simply share anonymized gradient updates. However, such a restrictive
 144 assumption is infeasible in robotics - if a robot sees a novel scene it cannot classify locally, it must
 145 necessarily share the raw training data to ask for a label. Future research should investigate how to
 146 blur sensitive information before sharing, such as in Google Street View [29], but for a broad set of
 147 private classes. Moreover, we believe (but must certify) that our system guards against malicious
 148 actors that falsify their utility for data, since this will cause them to see less relevant data over time.

149 **Formal Convergence Analysis** Future work should establish whether our data sharing protocol
 150 converges to a steady state where rare training data is evenly distributed across robotic peers. Our
 151 preliminary theoretical analysis estimates the number of rare images cached by each robot based on
 152 the precision and recall of the GU model and OoD detectors for each round of data sharing.

153 **A Data-Addressable Network?** Inspired by internet protocols, we propose a Data Discovery Ser-
 154 vice where robots can quickly identify new peers joining the network that are likely to have valuable
 155 training data for a task of interest. Likewise, future research should explore whether OoD detectors
 156 and our voting scheme can scalably run on commodity networks and resource-constrained robots.

157 Overall, robots are poised to leave the confines of the factory floor to enter our homes, hospitals,
 158 and airspace. Given the diversity of data owners and interactions with humans, it is increasingly
 159 important to consider privacy and data valuation as key desiderata in robotic fleet learning. We hope
 160 this paper spurs discussion on these topics and invite other researchers to collaborate and share their
 161 perspectives.

References

- 162
- 163 [1] Data is the new oil in the future of automated driving. [https://newsroom.intel.com/](https://newsroom.intel.com/editorials/krzanich-the-future-of-automated-driving/#gs.LoDUaZ4b)
164 [editorials/krzanich-the-future-of-automated-driving/#gs.LoDUaZ4b](https://newsroom.intel.com/editorials/krzanich-the-future-of-automated-driving/#gs.LoDUaZ4b), 2016.
165 [Online; accessed 10-June-2021].
- 166 [2] T. R. Institute. Toyota research institute showcases latest robotics research aimed at
167 amplifying human ability in the home, 2020. URL [https://www.tri.global/news/](https://www.tri.global/news/virtual-robotics-event/)
168 [virtual-robotics-event/](https://www.tri.global/news/virtual-robotics-event/).
- 169 [3] S. Balasubramanian. Robots have become an essential part of the war against
170 covid-19, 2021. URL [https://www.forbes.com/sites/saibala/2021/01/26/](https://www.forbes.com/sites/saibala/2021/01/26/robots-have-become-an-essential-part-of-the-war-against-covid-19/?sh=19265225ef31)
171 [robots-have-become-an-essential-part-of-the-war-against-covid-19/?sh=](https://www.forbes.com/sites/saibala/2021/01/26/robots-have-become-an-essential-part-of-the-war-against-covid-19/?sh=19265225ef31)
172 [19265225ef31](https://www.forbes.com/sites/saibala/2021/01/26/robots-have-become-an-essential-part-of-the-war-against-covid-19/?sh=19265225ef31).
- 173 [4] Safety pool scenario database. <https://www.safetypool.ai/database>, 2020. [Online;
174 accessed 10-June-2021].
- 175 [5] Cloud robotics and automation. <http://goldberg.berkeley.edu/cloud-robotics/>,
176 2019. [Online; accessed 03-June-2021].
- 177 [6] B. Kehoe, S. Patil, P. Abbeel, and K. Goldberg. A survey of research on cloud robotics and
178 automation. *IEEE Trans. Automation Science and Engineering*, 12(2):398–409, 2015.
- 179 [7] J. Kuffner. Cloud-enabled robots in: Ieee-ras international conference on humanoid robots.
180 *Piscataway, NJ: IEEE*, 2010.
- 181 [8] S. Chinchali, E. Pergament, M. Nakanoya, E. Cidon, E. Zhang, D. Bharadia, M. Pavone, and
182 S. Katti. Harvestnet: Mining valuable training data from high-volume robot sensory streams.
183 In *2020 International Symposium on Experimental Robotics (ISER)*, Valetta, Malta, 2020.
- 184 [9] D. A. Cohn, Z. Ghahramani, and M. I. Jordan. Active learning with statistical models. *JOUR-*
185 *NAL OF ARTIFICIAL INTELLIGENCE RESEARCH*, 4:129–145, 1996.
- 186 [10] Y. Gal, R. Islam, and Z. Ghahramani. Deep Bayesian active learning with image data. In
187 D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Ma-*
188 *chine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1183–1192.
189 PMLR, 06–11 Aug 2017. URL <http://proceedings.mlr.press/v70/gal17a.html>.
- 190 [11] B. Settles. Active learning literature survey. Technical report, University of Wisconsin-
191 Madison Department of Computer Sciences, 2009.
- 192 [12] S. Tong. *Active learning: theory and applications*, volume 1. Stanford University USA, 2001.
- 193 [13] A. Mandlekar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta,
194 E. Orbay, S. Savarese, and L. Fei-Fei. Roboturk: A crowdsourcing platform for robotic skill
195 learning through imitation. In *Conference on Robot Learning*, 2018.
- 196 [14] M. Waibel, M. Beetz, J. Civera, R. D’Andrea, J. Elfring, D. Glvez-Lpez, K. Hussermann,
197 R. Janssen, J. Montiel, A. Perzylo, B. Schiele, M. Tenorth, O. Zweigle, and M. Molengraft.
198 Roboearth - a world wide web for robots. *Robotics and Automation Magazine, IEEE*, 18:69 –
199 82, 07 2011. doi:10.1109/MRA.2011.941632.
- 200 [15] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang.
201 Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference*
202 *on Computer and Communications Security, CCS ’16*, page 308318, New York, NY, USA,
203 2016. Association for Computing Machinery. ISBN 9781450341394. doi:10.1145/2976749.
204 2978318. URL <https://doi.org/10.1145/2976749.2978318>.
- 205 [16] C. Dwork. Differential privacy: A survey of results. In M. Agrawal, D. Du, Z. Duan, and A. Li,
206 editors, *Theory and Applications of Models of Computation*, pages 1–19, Berlin, Heidelberg,
207 2008. Springer Berlin Heidelberg. ISBN 978-3-540-79228-4.

- 208 [17] Federated learning: Collaborative machine learning without centralized training data. <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>, 2017.
209 [Online; accessed 15-June-2021].
210
- 211 [18] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon. Federated
212 learning: Strategies for improving communication efficiency. *CoRR*, abs/1610.05492, 2016.
213 URL <http://arxiv.org/abs/1610.05492>.
- 214 [19] R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-
215 agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
- 216 [20] J. Ren, P. J. Liu, E. Fertig, J. Snoek, R. Poplin, M. Depristo, J. Dillon, and
217 B. Lakshminarayanan. Likelihood ratios for out-of-distribution detection. In H. Wallach,
218 H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors,
219 *Advances in Neural Information Processing Systems*, volume 32. Curran As-
220 sociates, Inc., 2019. URL [https://proceedings.neurips.cc/paper/2019/file/
221 1e79596878b2320cac26dd792a6c51c9-Paper.pdf](https://proceedings.neurips.cc/paper/2019/file/1e79596878b2320cac26dd792a6c51c9-Paper.pdf).
- 222 [21] K. Lee, K. Lee, H. Lee, and J. Shin. A simple unified framework for detecting out-of-
223 distribution samples and adversarial attacks. In S. Bengio, H. Wallach, H. Larochelle, K. Grau-
224 man, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing
225 Systems*, volume 31. Curran Associates, Inc., 2018. URL [https://proceedings.neurips.
226 cc/paper/2018/file/abdeb6f575ac5c6676b747bca8d09cc2-Paper.pdf](https://proceedings.neurips.cc/paper/2018/file/abdeb6f575ac5c6676b747bca8d09cc2-Paper.pdf).
- 227 [22] S. Liang, Y. Li, and R. Srikant. Enhancing the reliability of out-of-distribution image detec-
228 tion in neural networks. In *6th International Conference on Learning Representations, ICLR
229 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. Open-
230 Review.net, 2018. URL <https://openreview.net/forum?id=H1VGkIxRZ>.
- 231 [23] S. Goyal, A. Raghunathan, M. Jain, H. V. Simhadri, and P. Jain. DROCC: Deep robust one-
232 class classification. In H. D. III and A. Singh, editors, *Proceedings of the 37th International
233 Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*,
234 pages 3711–3721. PMLR, 13–18 Jul 2020. URL [http://proceedings.mlr.press/v119/
235 goyal20c.html](http://proceedings.mlr.press/v119/goyal20c.html).
- 236 [24] D. Hendrycks and K. Gimpel. A baseline for detecting misclassified and out-of-distribution
237 examples in neural networks. In *5th International Conference on Learning Representations,
238 ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenRe-
239 view.net, 2017. URL <https://openreview.net/forum?id=Hkg4TI9xl>.
- 240 [25] C. Richter and N. Roy. Safe visual navigation via deep learning and novelty detection. 07
241 2017. doi:10.15607/RSS.2017.XIII.064.
- 242 [26] T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, and D. Ha. Deep learning
243 for classical japanese literature, 2018.
- 244 [27] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking
245 machine learning algorithms, 2017.
- 246 [28] A. Ghorbani and J. Zou. Data shapley: Equitable valuation of data for machine learning, 2019.
- 247 [29] Google-contributed street view imagery policy. [https://www.google.com/streetview/
248 policy/](https://www.google.com/streetview/policy/), 2021. [Online; accessed 16-June-2021].