# FAIRNESS, EFFICIENCY, AND UNCERTAINTY IN SOCIETAL RESOURCE ALLOCATION MODELS

Kate Donahue, www.katedonahue.me
Department of Computer Science, Cornell

## Research Summary

The key question my research studies is, "how can we use societal resources to do the most good"? These "resources" could be concrete, like doctors, or abstract, like data, and the objective of "good" could be concepts like fairness, privacy, or utilization (helping as many people as possible). For example, some of my past work has analyzed cases with limited resource allocation, studying the extent to which fairness may be achievable alongside other goals, such as utilization or stability of certain arrangements. In other papers, I model data as a resource, analyzing federated learning through the lens of cooperative game theory to explore the relationship between stability, optimality, and fairness of federating structures. Finally, in ongoing work, I view human time as a resource, modeling how human and AI collaboration can show when complementary performance may be achievable.

## How can we allocate resources to as many people as possible while ensuring fairness in access?

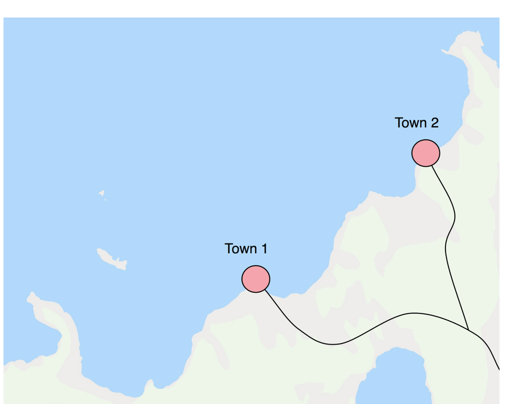


**Figure 1:** A stretch of coastline prone to unpredictable storms, causing levels of need in varying magnitude across towns.

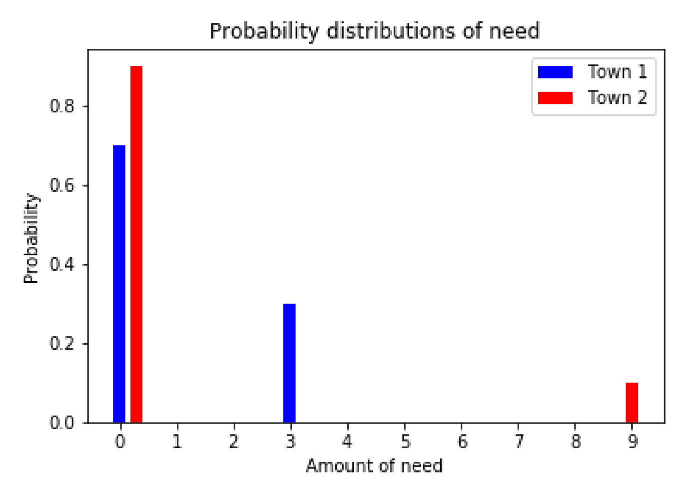


**Figure 2:** An example of distributions of need over the two towns. Note that need in Town 2 is larger and rarer than need in Town 1.

In [4] (which won Best CS paper at FAccT 2020), I study questions of how to allocate resources across multiple locations with stochastic levels of need. A community may wish to optimize for utilization (helping more people) or for fairness (equal probability of being helped). I analyze when these two goals might be in tension with each other- and when they can both be simultaneously achieved.

## Selected Publications


[1] Kate Donahue and Solon Barocas. "Better Together? How Externalities of Size Complicate Notions of Solidarity and Actuarial Fairness". In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 2021, pp. 185–195.

[2] Kate Donahue, Alexandra Chouldechova, and Krishnaram Kenthapadi. "Modeling Complementarity in Human-AI Collaboration". In: *FAccT 2022* (2022).

[3] Kate Donahue, Sreenivas Gollapudi, and Kostas Kollias. "'I pick you choose': Joint human-algorithm decision making in multi-armed bandits". Under review. 2022.

[4] Kate Donahue and Jon Kleinberg. "Fairness and utilization in allocating resources with uncertain demand". In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 2020, pp. 658–668.

[5] Kate Donahue and Jon Kleinberg. "Model-sharing Games: Analyzing Federated Learning Under Voluntary Participation". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 6. 2021, pp. 5303–5311.

[6] Kate Donahue and Jon Kleinberg. "Models of fairness in federated learning". In: *EAAMO 2022* (2022).

[7] Kate Donahue and Jon Kleinberg. "Optimality and Stability in Federated Learning: A Game-theoretic Approach". In: *NeurIPS 2021* (2021).


## How can we help agents with differing access to data build models collaboratively?

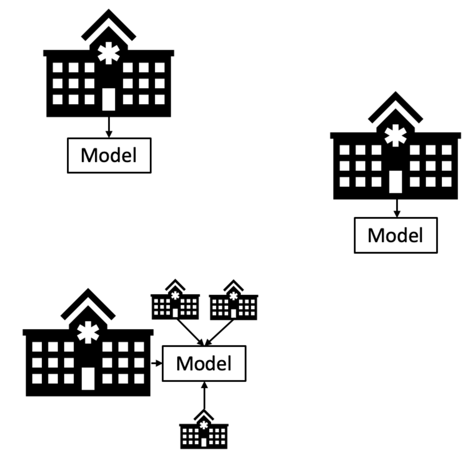


**Figure 3:** Agents of varying size (amount of data) federating in three coalitions. In federated learning, agents build a model with other agents, combining model parameters learned on their local data. For local learning, agents build a model based only on their local data.

In [5], I analyze federated learning, a distributed learning paradigm where agents combine models from local data to create a global model. If agents draw data from different distributions, however, the federated learning model may be biased and be sub-optimal for each agent. I use a lens of cooperative game theory to analyze which federating structures will be *stable* - where no agent wishes to move from its current federating coalition to another coalition (that would accept it).

## ... and how can we ensure such collaborative model-building is fair and societally optimal?

| Coalition $C$ | $err_s(C)$ | $err_l(C)$ | $\frac{err_s(\{n_s,n_\ell\})}{err_l(\{n_s,n_\ell\})}$ | $2 \cdot c + 1$ | $\frac{n_\ell}{n_s}$ | $err_w(C)$ |
|---|---|---|---|---|---|---|
| $\{\{n_s\},\{n_\ell\}\}$ | 1.67 | 0.5 | 3.33 | 5 | 3.33 | 20 |
| $\{n_s, n_\ell\}$ | 1.57 | 0.49 | 3.20 | 5 | 3.33 | 19.23 |

**Table 1:** Example of two-player game with $n_s = 6, n_\ell = 20$, with parameters $\mu_e = 10, \sigma^2 = 1$. Note that in this case, federating in the grand coalition $\{n_s, n_\ell\}$ is individually stable (and thus optimal, for weighted error). The grand coalition satisfies egalitarian fairness ($2 \cdot c + 1$ bound) and proportional fairness.

| Coalition $C$ | $err_s(C)$ | $err_l(C)$ | $\frac{err_s(\{n_s,n_\ell\})}{err_l(\{n_s,n_\ell\})}$ | $2 \cdot c + 1$ | $\frac{n_\ell}{n_s}$ | $err_w(C)$ |
|---|---|---|---|---|---|---|
| $\{\{n_s\},\{n_\ell\}\}$ | 1.67 | 0.250 | 6.68 | 9 | 6.67 | 20 |
| $\{n_s, n_\ell\}$ | 1.73 | 0.251 | 6.89 | 9 | 6.67 | 20.43 |

**Table 2:** Example of two-player game with $n_s = 6, n_\ell = 40$, with parameters $\mu_e = 10, \sigma^2 = 1$. Here, the grand coalition fails to be stable, so local learning minimizes weighted error. The grand coalition satisfies egalitarian fairness ($2 \cdot c + 1$ bound) but does not satisfy proportional fairness.

In follow-on work, I analyze the fairness and optimality of federating coalitions. In [6], I consider two competing notions of fairness in federated learning: egalitarian fairness (which aims to bound how dissimilar error rates can be) and proportional fairness (which aims to reward players for contributing more data). For egalitarian fairness, we obtain a tight multiplicative bound on how widely error rates can diverge between agents federating together. For proportional fairness, we show that sub-proportional error (relative to the number of data points contributed) is guaranteed for any individually rational federating coalition.

In [7], I consider the relationship between a federating structure's stability and its optimality (how low the overall average error is). After giving an efficient, constructive algorithm for calculating an optimal arrangement, I show that optimal arrangements are not always stable, but that the worst stable solution has a cost no more than 9 times that of an optimal arrangement.

## How can we combine human and algorithmic expertise to do better than either alone?

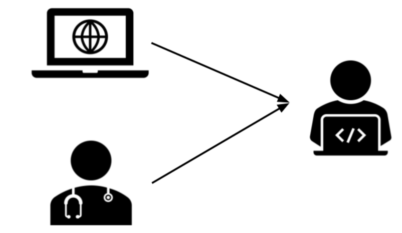

**Figure 4:** A model of human/AI collaboration: the performance of the combined system can be viewed as a function of the performance of the algorithm and the (unaided) human.

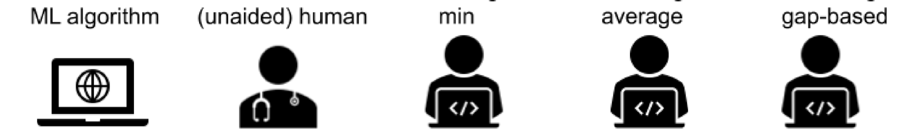

| | ML algorithm | (unaided) human | human using ML: min | human using ML: average | human using ML: gap-based |
|---|---|---|---|---|---|
| **Error regime 1 (50%)** | 10% | 3% | 3% | 6.5% | 6.25% |
| **Error regime 2 (50%)** | 4% | 12% | 4% | 8% | 7.7% |
| **Average error** | 7% | 7.5% | 3.5% | 7.25% | 6.9% |

**Figure 5:** A model of human/AI collaboration, for various examples of decision (combining) functions.

In [2], I analyze systems where both human and algorithmic expertise are incorporated. For example, a doctor making a medical prediction may rely on her own knowledge, as well as the prediction of an ML tool, in making a final decision. We analyze theoretically factors that influence *complementary* performance (lower error than either the algorithm or unaided human alone). In ongoing work [3], I extend this analysis to human-algorithm systems in multi-armed bandit settings.

## How can we allocate cost across dissimilar people in a way that is fair and incentivizes everyone to participate?

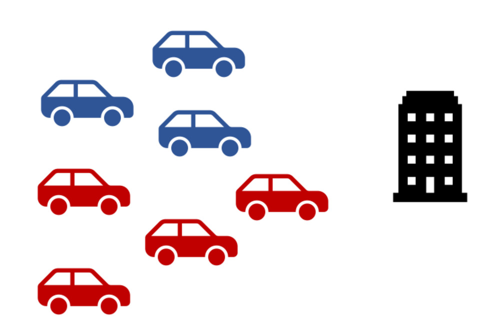

**Figure 6:** Consider a case of car insurance, where each person is either low risk or high risk .

| | Total | Alice's cost | Bob's cost |
|---|---|---|---|
| **Separate pools** | \$35,741 | \$32.52 | \$38.96 |
| **Pooled: even-split pricing** | \$31,878 | \$31.88 | \$31.88 |
| **Pooled: proportional pricing** | \$31,878 | \$28.36 | \$35.40 |

**Table 3:** Example with insolvency-based premiums. Alice has low risk (2%) and Bob has higher risk (2.5%)

In [1], I consider cost-sharing games where some players contribute more or less to overall cost: for example, car insurance with drivers of different levels of true risk. Two natural and competing notions of fairness might be to a) charge each individual the same price (solidarity) or b) charge each individual according to the cost that they bring to the pool (actuarial fairness). However, in many natural settings, cost-sharing games exhibit *externalities of size*, where larger groups are cheaper per person (all else being equal). We explore how this complicates traditional understandings of fairness, drawing on literature in cooperative game theory.