

# Exploring Geriatric Clinical Data and Mitigating Bias with Multi-Objective Synthetic Data Generation for Equitable Health Predictions

Jarren Briscoe, Chance DeSmet, Katherine Wuestney, Assefaw Gebremedhin, Roschelle Fritz, and Diane J. Cook

Washington State University, Department of Electrical Engineering and Computer Science  
255 E Main Street, Pullman, Washington, USA 99163

{jarren.briscoe,chance.desmet,katherine.wuestney,assefaw.gbremedhin,shelly.fritz,djcook}@wsu.edu

**Abstract** - In the data-centric landscape of modern healthcare, addressing bias in machine learning models is crucial for ensuring equitable health outcomes. When applied in clinical settings, biased predictions can exacerbate disparities in healthcare. This paper focuses on the domain of biomedical informatics and the challenge of mitigating bias in smart home datasets used for health monitoring. We assess existing bias metrics and a new metric, the Objective Fairness Index (OFI), to quantify bias related to sensitive attributes. To address these biases, we propose a novel method using a multi-objective generative adversarial network (GAN) that generates diverse synthetic data to improve data representation. This approach, validated on data from older adults managing chronic health conditions, demonstrates the potential to enhance both prediction accuracy and fairness in health outcomes.

**Keywords:** Bias metrics, clinical data, generative adversarial network, smart homes, synthetic data

© Copyright 2024 Authors - This is an Open Access article published under the Creative Commons Attribution License terms (<http://creativecommons.org/licenses/by/3.0>). Unrestricted use, distribution, and reproduction in any medium are permitted, provided the original work is properly cited.

## 1 INTRODUCTION

In bioinformatics and biomedicine, the potential of machine learning (ML) to revolutionize healthcare is exciting for many clinicians. Biomedical datasets are increasingly being used to inform clinical decision-making, contributing to the growing field of digital health, which leverages technology to monitor and improve health outcomes [1]–[5]. However, the adoption of these algorithms for critical decision making is still

limited. The black-box nature of algorithms frequently necessitates caution for clinicians [6].

Biomedical data often contain inherent biases due to factors such as demographic disparities in data collection, unequal access to healthcare, and historical health disparities. When not properly addressed, these biases can lead to skewed predictions and unequal health outcomes. For instance, a predictive model trained on biased data might disproportionately misclassify certain demographic groups, leading to suboptimal treatment recommendations for those groups. In the worst-case scenario, such biases could exacerbate existing health disparities, undermining the goal of equitable healthcare. Distrust in ML algorithms is heightened by publicized cases where machine learning algorithms yielded prejudicial inferences [7]. Unless ML algorithms are designed to avoid bias along a particular sensitive attribute, they will reflect the prejudices of the data used to train them.

In this paper, we use a new bias metric, called *Objective Fairness Index* (OFI) [26–27], that gives a novel quantitative perspective on bias consistent with legal precedent. Using prior metrics and OFI, we analyze bias in a clinical dataset containing smart home data. Then, we propose a multi-agent generative adversarial network tool, called HydraGAN, to create diverse synthetic data that mitigates bias due to lack of sample diversity. Finally, we provide insight into older adults by giving distributions of their activities, by gender and by age. We postulate that the new metric, OFI, is valuable because it considers all outcomes of a classification decision with context. The revelations behind geriatric activities show that OFI properly captures objective testing. Additionally, our proposed HydraGAN tool,

Date Received: 2024-02-21

Date Revised: 2024-09-20

Date Accepted: 2024-10-01

Date Published: 2024-11-04

which generates diverse synthetic data, has the potential to enhance the robustness of predictive models in bioinformatics, leading to more accurate and equitable health predictions.

Table 1: Study participants. F=female, M=male; HS=high school, B=bachelor, M=master, D=doctorate.

<b>Age</b>	89	83	88	75	95	89	81	63	92	79	83	88	90	76	93
<b>Gender</b>	F	M	F	M	F	F	F	F	F	M	F	F	F	F	M
<b>Education</b>	HS	B	HS	D	HS	B	HS	B	HS	D	B	B	B	B	M

## 2 RELATED WORK

There have been numerous attempts to reduce machine learning bias. Weighting data points can reduce bias by emphasizing a minority class, while transforming the feature representations may reduce correlation between sensitive attributes and other features [8]. A systematic search of models and hyperparameters can identify the combination yielding the lowest bias [9]. In the case of digital health, much of the bias is due to a lack of representation for underrepresented groups. Some prior research limit analysis to the data, while others examine how predictions and corresponding actions will affect target groups [10], [11]. Difficulties have been noted in aligning these measures with statistical requirements [12].

In this paper, we propose a multi-agent generative adversarial network (GAN) to generate sample data that improve diversity. Recent synthetic data creation for health applications relies more frequently on GANs [13]. Traditionally, GANs represent two-agent systems. However, the FairGAN architecture [14] balances the generator with two critics, one promoting data realism and the other supporting fairness. Our proposed approach extends these previous works by supporting an arbitrary number of agents, corresponding to a list of optimization criteria for the synthetic data. In this analysis, we harness the power of critics for pointwise realism, distribution realism, and distribution diversity.

## 3 CLINICIAN-IN-THE-LOOP SMART HOME STUDY

We collected continuous ambient sensor data for 22 older adults who are managing two or more chronic health conditions. Because 70% of the world's older adults are managing chronic conditions, the World Health Organization is asking for technology solutions to support these individuals [15]. The goal of this study is to design a clinician-in-the-loop (CIL) smart home that

identifies health condition exacerbations using clinician-guided machine learning techniques. Table 1 provides a summary of the study participants.

We installed a CASAS smart home in a box (SHiB) [16] in the home of each subject for one year. The CASAS SHiB sensors monitor movement, door use, ambient

light, and temperature. Nurses met weekly with each subject. Based on these interviews and nurse visual inspection of smart home data, changes in health status related to condition exacerbations were identified. In prior work, we extracted markers from these data that detect flare-ups in symptoms related to conditions such as congestive heart failure, diverticulitis, urinary tract infections, and Parkinson's disease [17].

## 4 BIAS IN ACTIVITY RECOGNITION

As a first step in analyzing the data, we created a machine learning approach to recognize activities in real time. In this process, a sliding window is moved over data and used as context for the learning algorithm to label the last sensor reading in the window. At least one month of data was manually labelled by research team members (inter-annotator agreement  $k=.80$ ). In prior work, we validated that eleven activities (bed-toilet transition, cook, eat, enter home, leave home, hygiene, relax, sleep, wash dishes, work, other) are recognized with accuracy=0.99 [18]. For this work, we categorize the eleven activities into binary classes of active or sedentary activities to train a multi-layer perceptron to predict the patient's activity status. This gives nurses and viewers a more understandable gist of a patient's activities and is easier to tell when something is wrong with a glance. We then use a multi-layer perceptron to predict the patient's activity status. In this paper, we focus our attention on binary activity recognition and analyze bias for this task. This machine learning task is a pivotal component of assessing health state and is central to many other mobile health technologies.

Table 2 summarizes popular bias metrics. As the table shows, many metrics do not reflect all desired properties. To meet this need, we introduce a bias metric called Objective Fairness Index (OFI). In our discussion,

we distinguish “benefit” from “bias”. While benefit references an advantage that is gained from an action (e.g., a ML prediction), bias reflects the distance between true and expected benefit. OFI is created based on legal precedents established by the US Supreme Court,

Next,  $\mathcal{B}$  quantifies the bias for a given group and is calculated as  $b - E[b]$ . Substituting and simplifying yields:

$$\mathcal{B} = (FP - FN)/n \in [-1,1] \quad (2)$$

Table 2: Popular bias metrics and their properties. Each metric contrasts group  $i$  with group  $j$ . A larger score indicates more bias toward group  $i$ .

$n$ =size of set,  $P$ =positive,  $N$ =negative,  $\hat{P}$ =predicted positive,  $\hat{N}$ =predicted negative,  $TP$ =true positive,  $FP$ =false positive,  $FN$  = false negative,  $TN$ =true negative

Metric	Formula	Directed	Symmetric	$n > 0$	Resilient	Bounded	Context
Disparate impact [19]	$\hat{P}_i/n_i \div \hat{P}_j/n_j$	✓	✗	✗	✓	✗	✗
Predictive parity [20]	$TP_i/P_i - TP_j/P_j$	✓	✓	✗	✗	✓	✗
Treatment equality [21]	$FP_i/FN - FP_j/FN_j$	✓	✓	✗	✗	✗	✗
FPR difference (FPRD) [22]	$FP_i/N_i - FP_j/N_j$	✓	✓	✗	✗	✓	✗
TPR difference (TPRD) [22]	$TP_i/P_i - TP_j/P_j$	✓	✓	✗	✗	✓	✗
Equalized odds [23]	$FPRD + TPRD$	✓	✓	✗	✗	✓	✗
Difference in conditional acceptance [22]	$P_i/\hat{P}_i - P_j/\hat{P}_j$	✓	✓	✗	✗	✗	✗
Difference in conditional rejection [22]	$N_i/\hat{N}_i - N_j/\hat{N}_j$	✓	✓	✗	✗	✗	✗
Difference in positive proportion & labels [22]	$\hat{P}_i/n_i - \hat{P}_j/n_j$	✓	✓	✓	✓	✓	✗
Objective Fairness Index	$\frac{FP_i - FN_i}{n_i} - \frac{FP_j - FN_j}{n_j}$	✓	✓	✓	✓	✓	✓

EU’s Charter of Fundamental Rights, the Canadian Human Rights Act, and the Constitution of India, indicating that a lower admittance rate of a discriminated group warrants further investigation.

The OFI metric is formalized in Equations 1-3. Using the legal precedents, we define benefit as a positive (+1) prediction, or desirable outcome, for each individual,  $q$ . We define  $b$  as the overall benefit for any group, estimated as the weighted mean benefit over individuals in the group. Next, we assign weights to the prediction classes, defaulting to a weight of 1 for the positive class and 0 for the negative class. Based on the ground truth values, we compute expected benefit  $E[b_q]$  as 1 for points with positive labels and -1 for negative points, resulting in Equation 1.

$$E[b] = \frac{1}{n} \sum_{\forall b_q} = \frac{1 \cdot P + 0 \cdot N}{n} = \frac{TP + FN}{n} \quad (1)$$

Where  $b$  is the benefit,  $b_q$  is an individual benefit.  $P$  is the count of positive labels,  $N$  is the count of negative labels,  $n = P + N$ ,  $TP$  is the count of true positives, and  $FN$  is the count of false negatives.

Note that  $FP$  is the count of false positives.

To find the bias for group  $i$  over group  $j$  (OFI), Equation 3 calculates the directional difference of bias for each group.

$$OFI = \mathcal{B}_i - \mathcal{B}_j = (FP_i - FN_i)/n_i - (FP_j - FN_j)/n_j \quad (3)$$

We note several benefits of OFI. While some metrics focus solely on benefit or harm ( $FP$  or  $FN$ ), OFI uses all confusion matrix cells. As a result, this metric includes more facets of the comparison between ground truth and predicted values. Additionally, this metric handles class imbalance without impacting calculations, creating metric resiliency. Furthermore, it is directed and symmetric. Importantly for our study, OFI can be adapted for binary classification, multi-class classification, regression, and ranking. In the class of multi-class classification, the calculation is based on an  $n$ -ary confusion matrix, summarized by mean values. For regression, OFI can replace  $FP$  and  $FN$  with error (e.g., mean absolute error, mean squared error) above and below the desired threshold. When applying to ranking, OFI considers the difference in the ranked positions.

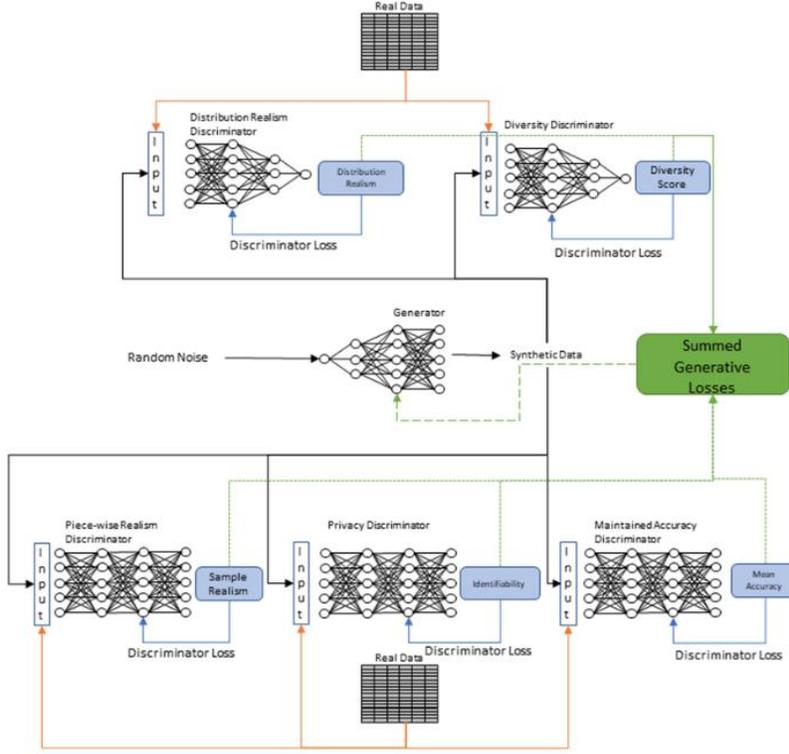


Figure 1: The HydraGAN multi-agent architecture.

Many clinical data are time series in nature (e.g., EHR entries, lab results, vital signs, or sensor readings). To evaluate bias in such data, we need to extend the bias metrics to apply to time series data. In time-series data, we consider each time step as an individual step that repeats benefit or harm, with a corresponding confusion matrix. In this scenario, bias is aggregated over individuals in a group and time steps in the series. For our experiments, we randomly sample 5% of all data to train, leaving 95% to validate and test. This process yields a larger undersampling effect and highlights potential bias.

## 5 MITIGATING BIAS WITH DIVERSE SYNTHETIC DATA

Because researchers recognize the surrogate role offered by synthetic data generators, they create methods to generate increasingly realistic data proxies. We consider the impact of creating realistic, diverse synthetic data for our dataset. What prior approaches lack is the ability to introduce multiple critics, each of which represents a distinct goal of the synthetic data. In some cases, emulating all characteristics of available real data is not the sole, or even desired, outcome. For example, the data may also need to achieve a diversity goal or obfuscate sensitive information. For this, we use

HydraGAN [24], a multi-agent generative adversarial network that performs multi-objective synthetic data generation.

We adopt a multi-agent GAN, called HydraGAN, that assigns a discriminator ( $D$ ) to each data goal (see Figure 1). Each of the critics separately critique individual or batches of synthetic data points. The generator's loss is the weighted sum of all critic scores. When the system converges (the weight changes for an epoch are below a threshold value), a Nash equilibrium is formed among the discriminator goals. Here, we focus on three discriminators that minimize loss for traditional pointwise data realism ( $D_p$ , Equation 4), distribution data realism ( $D_\tau$ , Equation 5), and data diversity (Equation 6).

$$\underset{x_r, x_g}{\text{minimize}} \sum_{i \in x_r, x_g} D_\rho(x_{g_i}) + (1 - D_\rho(x_{r_i})) \quad (4)$$

Here,  $D_p$  entices data realism.  $x_r$  and  $x_g$  represent batches of real and corresponding synthetic points.

$$\underset{x_r, x_q \in X}{\text{minimize}} D_\tau(x_g) + (1 - D_\tau(x_r)) \quad (5)$$

$D_\tau$  minimizes the difference between data distribution characteristics for the real ( $x_r$ ) and synthetic data ( $x_g$ ).

(6)

In this data diversity minimization,  $\alpha$  is the observed and  $\beta$  is the desired proportions for some feature  $f_i$ .

Diversity constraints may be designed to ensure equal representation among alternative groups. As an example, if 90% of a physical data collection represents one value for a sensitive feature (e.g., race) and 10% represents another, the diversity discriminator will move toward a more uniform distribution. Our architecture uses a combination of 1D convolutional layers, learnable positional encoding, and fully connected layers. Our regularization techniques include layer normalizations, instance normalizations, dropout layers, and Gaussian noise. We use the leaky ReLU activation function with a negative slope of 0.2. Each network uses the Adam optimizer with a learning rate of 0.0002,  $b_1=0.5$ , and  $b_2=0.9$ . We train with 75 epochs and conduct 100 steps per epoch. Each step contains a mini-batch of 64 time windows, each with a sequence length of 32.

The generator processes the sensitive attribute conditional and a noise vector, outputting normalized values in time windows. We partition synthetic features to give appropriate output activations. Since our real date-time features are represented with two sine and cosine pairs for the day of year and time of day, the synthetic time features are outputted with a sine activation. Each set of one-hot encoded features (sensor one, sensor two, activity) is then passed through a softmax activation.

## 6 EXPERIMENTAL RESULTS

We are interested in quantifying the bias contained in our clinical data using traditional metrics and our novel OFI metric. We then analyze bias for the newly generated dataset. While we focus on one dataset, the demographics in our study are similar to those found in many other clinical studies. A prevalent form of bias in clinical studies is sample selection bias. Many clinical study populations are largely devoid of diversity. As an example, Latinos and Asian Americans are disproportionately underrepresented in clinical trials

assessing cognitive decline, comprising only 1%-5% of research participants [25].

We focus on two sensitive attributes: age and gender. We minimize the difference between the observed proportions ( $\alpha_{f_i}$ ) and the desired proportions ( $\beta_{f_i}$ ) for each feature  $f_i$ . We assess bias for the older 25% of the sample in comparison with the group containing the younger 75% of the participants. Rather than evaluate bias for all the metrics listed in Table 2, we select a subset including Disparate Impact (DI), Difference in Conditional Acceptance (DCA), Difference in Proportionate Positives and Labels (DPPI), and Objective Fairness Index (OFI). This is a representative set: the remaining metrics yield similar results to these.

The selected bias metrics focus on a task, in this case activity recognition. To simplify analyses, we aggregate activity categories into two classes: active behavior (bed-toilet transition, cook, eat, enter home, leave home, hygiene, wash dishes) and sedentary behavior (relax, sleep, work, other). For DCA, DPPL, and OFI, a no-bias score is zero. For DI, the no-bias score is one. A closer value to the no-bias score indicates less bias. A positive value indicates a bias toward the sedentary categories, while a negative value indicates a bias toward the active categories.

### 6.1 Original dataset

Figure 2 depicts boxplots of quantified bias on the original dataset. Since Ages 90-95 receives high scores in DI and DPPL, we see they are more likely to be predicted as sedentary. Referencing OFI, which gives context for correct predictions, we see that these positive predictions are largely correct as the Ages 90-95 OFI score is near zero. This makes sense as our data reflects that these older patients tend to be sedentary more often than the younger class, as shown in Section 7. However, since OFI is positive, we see that there is still a slight bias for predicting the older age group as sedentary more often than they should.

### 6.2 Expanded dataset

We train the model using synthetic data generated by HydraGAN then quantify the bias by testing on real data. We summarize the bias results in Figure 3. Here, individual diversity is enforced by querying the generator with each individual's label. This individual diversity also improves the protected class's diversities. In total, 3,072,000 synthetic points are created to be

realistic and improve diversity for the underrepresented groups.

We see in Table 3 that the synthetic data mitigates bias overall (the bold cells indicate improvement). For OFI, DPPL, and DCA, a score of 0 indicates no bias. For DI, the ideal score is 1. Over many trials, these scores have a negligible standard deviation.

Table 4 shows that the synthetic data improves diversity. The improvement in uniformity is statistically significant ( $p < 0.001$ ) for the Kolmogorov-Smirnov (KS) statistic when comparing distributions for activities. Furthermore, our synthetic data's distribution over individuals reflects a distance of 0 from the uniform distribution (the desired result) due to the infinitely strong conditional passed to the generator. Reviewing the protected class's bias reductions in Table 3, we furthermore conclude that our synthetic data successfully mitigates bias for age and gender.

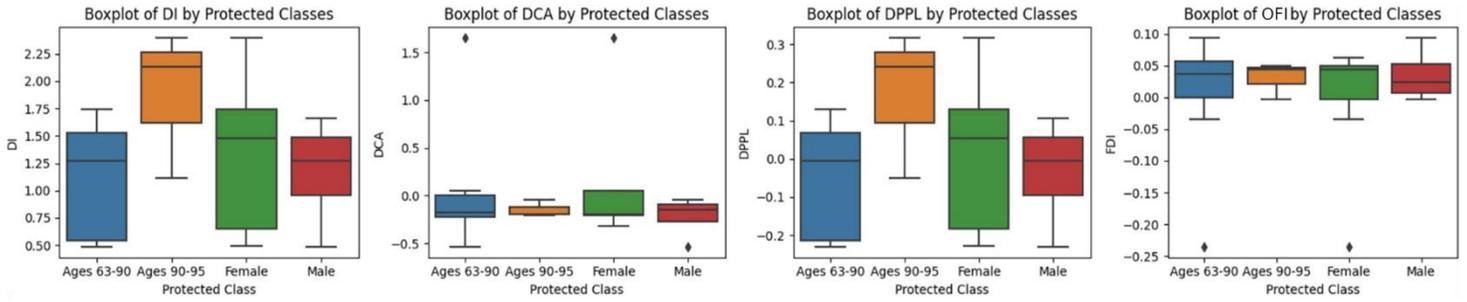


Figure 2: Boxplots of bias metrics applied to the synthetically expanded dataset.

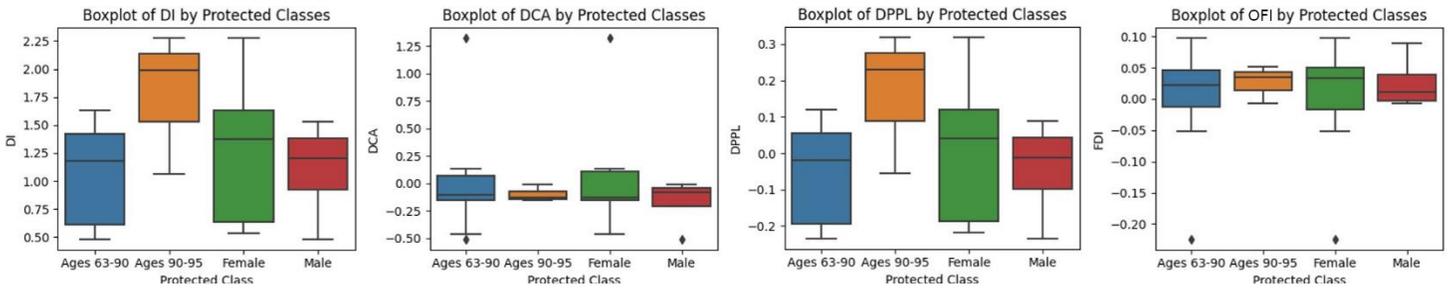


Figure 3: Boxplots of bias metrics applied to original dataset.

Table 3: Comparison of bias results for synthetic and real data for each possible sensitive attribute.

Metric	Data	Ages 63-90	Ages 90-95	Female	Male
OFI	Synthetic	0.002	0.026	-0.001	0.026
	Real	0.004	0.030	-0.001	0.033
	Improves	<b>0.002</b>	<b>0.005</b>	-0.000	<b>0.007</b>
DPPL	Synthetic	-0.056	0.165	0.012	-0.043
	Real	-0.055	0.170	0.012	-0.036
	Improves	-0.002	<b>0.005</b>	<b>0.000</b>	-0.007
DI	Synthetic	1.060	1.778	1.280	1.103
	Real	1.109	1.891	1.340	1.174
	Improves	<b>0.049</b>	<b>0.113</b>	<b>0.060</b>	<b>0.071</b>
DCA	Synthetic	0.003	-0.100	0.046	-0.172
	Real	0.012	-0.151	0.061	-0.219
	Improves	<b>0.009</b>	<b>0.051</b>	<b>0.014</b>	<b>0.048</b>

Table 4: Comparison of synthetic and real data to a uniform distribution using measures of KS (Kolmogorov-Smirnov) statistic, Kullback-Leibler (KL) divergence, and Jensen-Shannon (JS) distance.

Metric	Activities			Individuals		
	Synthetic	Real	Improvement	Synthetic	Real	Improvement
KS statistic	0.863	0.900	<b>0.027</b>	0.000	0.733	<b>0.733</b>
KL Divergence	1.822	2.302	<b>0.480</b>	0.000	0.252	<b>0.252</b>
JS Distance	0.652	0.725	<b>0.073</b>	0.000	0.250	<b>0.250</b>

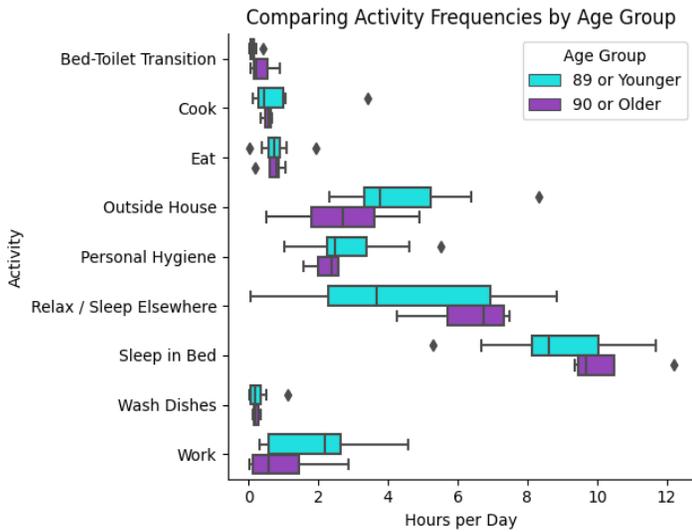


Figure 4: Comparing Activity Frequencies by Age Group

## 7 ACTIVITY DISTRIBUTION ANALYSES

In this section, we present a visual comparison of activity frequencies categorized by age and gender, offering further insights into the distribution of activities performed by older adults in our clinical dataset. These visualizations provide key context for understanding the behavioral patterns of individuals in these differing demographics and how these patterns may contribute to bias in our machine learning models. This analysis complements the bias metrics outlined earlier by visually capturing differences in time allocation to various activities.

### 7.1 Activity Frequency by Age Group

Figure 4 displays activity frequencies for participants segmented into two age groups: "89 or Younger" and "90 or Older." The x-axis represents the number of hours per day spent on specific activities, while the y-axis lists these activities. The activities include critical daily tasks such as bed-toilet transitions, cooking, eating, and personal hygiene, as well as sedentary activities like relaxation and sleep. Note that we update the activity "Relax" to be "Relax / Sleep Elsewhere" since these older adults sometimes fell asleep while relaxing.

We glean a few insights from Figure 4. Most notably, the older age group tends to spend more time on sedentary activities, particularly relaxing and sleeping, compared to the younger cohort. The younger group's wider interquartile ranges and whiskers indicate a

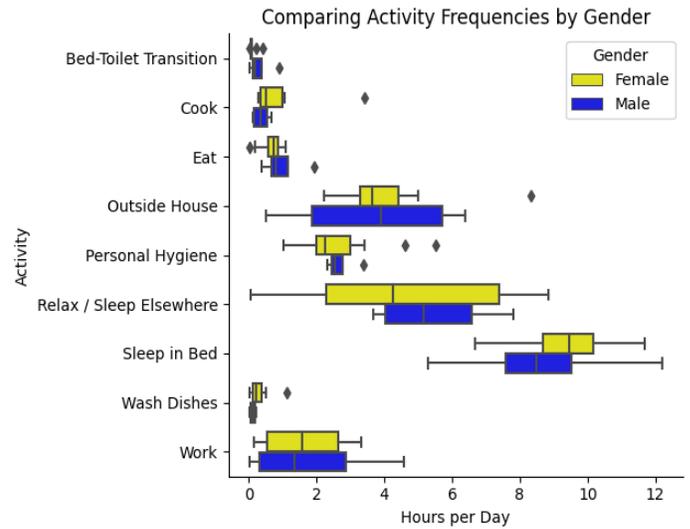


Figure 5: Comparing Activity Frequencies by Gender

higher variation in time sleeping and relaxing, which may reflect differences in health status or mobility. Subsequently, the 90 plus demographic is not as physically active as the younger group, with noticeable reductions in work and being outside.

Furthermore, we investigate necessary household activities. With food, the elder group about matches the younger group for time spent cooking, eating, and washing dishes. This suggests that even with age, these crucial tasks remain relatively consistent across individuals. However, the older ones tend to spend less time on personal hygiene and frequent the bathroom more often while in bed.

This visual analysis reveals that the older group generally spends more time on sedentary activities while engaging less in physical and socially active behaviors. These findings align with the bias analysis using our Objective Fairness Index (OFI) and DI in Figure 3. The elderly were predicted to be sedentary more often than others, hence the high DI score. However, OFI considered that objectively, the elderly were more sedentary. Hence, OFI does not display as much bias. Nevertheless, OFI's positive score still suggests the neural network has overconfidence in the elderly being sedentary.

### 7.2 Activity Frequency by Gender

Next, we illustrate activity frequencies split by gender, comparing male and female participants. The gender-based distribution highlights both similarities and differences in daily activity patterns.

In Figure 5 we see that females tend to sleep in bed more than males. Notably, men’s median relax time is greater than the women’s median time, yet women have a great variation in the amount that they relax. For physical activities, both genders work about the same amount of time in median and variation. Furthermore, females cluster around 3-4.5 hours outside while males have a wider IQR spread of between 1.75 and 6 hours.

Regarding household activities, males lead over females minorly in visits to the toilet from bed. However, females tend to cook and wash dishes for longer than males, while males eat longer than females.

These differences in activity frequency by gender are essential for understanding potential biases in our machine learning model’s predictions. For example, the observed variation in relaxation time between genders may lead to biased predictions if not appropriately accounted for in the data preprocessing and modeling stages.

The visual exploration of activity patterns across age and gender groups underscores the importance of controlling for demographic variables when developing predictive models for clinical applications. The observed disparities in activity frequency may inadvertently introduce bias in activity recognition models, particularly in predicting sedentary versus active behaviors. Our proposed HydraGAN synthetic data generator mitigates such biases by creating more diverse and realistic data distributions, ensuring fairer predictions for underrepresented groups. These visual insights into activity frequency further validate the need for multi-objective synthetic data generation to capture the heterogeneity of patient behavior in clinical settings.

## 8 DISCUSSION AND CONCLUSIONS

In this paper, we examine biases that may exist in a clinical dataset using smart home sensor data to model activities that are used for health assessment. Metrics of bias are varied yet do not consistently make use of all predicted outcomes. These outcomes lead to advantages for one group over another and so need to be considered in bias analyses. As a result, we not only use traditional metrics, but we also use a new metric based on legal precedent, Objective Fairness Index. As we show, the OFI metric considers all cells in the confusion matrix and compares predicted labels with ground truth labels,

leading to a more comprehensive analysis of bias and fairness.

The experimental results indicate that bias does exist in our data, even for a straightforward task such as activity labelling. Because activity recognition is used as a cornerstone for embedded and mobile technology strategies for health assessment and intervention, even this component necessitates unbiased reasoning and fair treatment of all groups. To potentially mitigate sample bias that results from a lack of diversity in the collected data, we introduce HydraGAN, a multi-agent synthetic data generator. Generating synthetic data with HydraGAN does reduce bias in the data based on multiple metrics.

This is an early analysis of the OFI metric and HydraGAN algorithm to analyze and lessen bias in clinical data. Our activity distribution visualizations give insight into geriatric activity, but further validation is needed to assess these contributions on a greater variety of clinical datasets and across additional protected attributes. We also note that HydraGAN can incorporate additional critics that consider metrics such as privacy preservation. Future work will analyze the role these optimization criteria can play in providing more trustworthy machine learning technologies for clinical data assessment and application.

## Acknowledgements

This work is supported in part by NINR grant R01NR016732.

## References

- [1] Q. Li, M. Jiang, and C. Ying, “An assistant decision-making method for rare diseases based on RNNs model,” in *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2022, pp. 2632–2639. doi: 10.1109/BIBM55620.2022.9995519.
- [2] X. Guo, Y. Qian, P. Tiwari, Q. Zou, and Y. Ding, “Kernel risk sensitive loss-based echo state networks for predicting therapeutic peptides with sparse learning,” in *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2022, pp. 6–11. doi: 10.1109/BIBM55620.2022.9994902.
- [3] Y. Lin, J. Jiang, Z. Ma, D. Chen, Y. Guan, X. Liu, H.

- You, J. Yang, and X. Cheng, "CGPG-GAN: An acne lesion inpainting model for boosting downstream diagnosis," in *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2022, pp. 1634–1638. doi: 10.1109/BIBM55620.2022.9995406.
- [4] D. Tan, J. Wang, R. Yao, J. Liu, J. Wu, S. Zhu, Y. Yang, S. Chen, and Y. Li, "CCA4CTA: A hybrid attention mechanism based convolutional network for analysing collateral circulation via multi-phase cranial CTA," in *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2022, pp. 1201–1206. doi: 10.1109/BIBM55620.2022.9995381.
- [5] R. Zemouri, N. Zerhouni, and D. Racoceanu, "Deep learning in the biomedical applications: Recent and future status," *Appl. Sci.*, vol. 9, no. 8, 2019, doi: 10.3390/app9081526.
- [6] H. Alami, P. Lehoux, Y. Auclair, and M. De, "Artificial intelligence and health technology assessment: Anticipating a new level of complexity," *J. Med. Internet Res.*, vol. 22, no. 7, p. e17707, 2020.
- [7] L. A. Celi, J. Cellini, M.-L. Charpignon, E. C. Dee, and F. DERNONCOURT, "Sources of bias in artificial intelligence that perpetuate healthcare disparities - A global review," *PLOS Digit. Heal.*, vol. 1, no. 3, p. e0000022, 2022.
- [8] D. Plecko and N. Meinshausen, "Fair data adaptation with quantile preservation," *J. Mach. Learn. Res.*, vol. 21, pp. 1–44, 2020.
- [9] A. Agarwal, M. Dudik, and Z. S. Wu, "Fair regression: Quantitative definitions and reduction-based algorithms," in *International Conference on Machine Learning*, 2019.
- [10] T. Speicher, H. Heidari, N. Grgic-Hlaca, and others, "A unified approach to quantifying algorithmic unfairness: Measuring individual & group unfairness via inequality indices," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2018, pp. 2239–2248.
- [11] G. Pleiss, M. Raghavan, F. Wu, and others, "On fairness and calibration," in *Advances in Neural Information Processing Systems*, 2017.
- [12] S. Corbett-Davies and S. Goel, "The measure and mismeasure of fairness: A critical review of fair machine learning," *arXiv Prepr. arXiv1808.00023*, 2018.
- [13] K. Baek and H. Shim, "Commonality in natural images rescues GANs: Pretraining GANs with generic and privacy-free synthetic data," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 7854–7876.
- [14] D. Xu, S. Yuan, L. Zhang, and X. Wu, "FairGAN: Fairness-aware generative adversarial networks," in *IEEE International Conference on Big Data*, 2018.
- [15] R. Yan, X. Liu, J. Dutcher, M. Tumminia, D. Villalba, S. Cohen, D. Creswell, K. Creswell, J. Mankoff, A. Dey, and A. Doryab, "A computational framework for modeling biobehavioral rhythms from mobile and wearable data streams," *ACM Trans. Intell. Syst. Technol.*, vol. 13, no. 3, p. 47, 2022.
- [16] D. J. Cook, A. Crandall, B. Thomas, and N. Krishnan, "CASAS: A smart home in a box," *IEEE Comput.*, vol. 46, no. 7, pp. 62–69, 2012.
- [17] S. Fritz, K. Wuestney, G. Dermody, and D. J. Cook, "Nurse-in-the-loop smart home detection of health events associated with diagnosed chronic conditions: A case-event series," *Int. J. Nurs. Stud. Adv.*, vol. 4, p. 100081, 2022.
- [18] S. Aminikhanghahi, T. Wang, and D. J. Cook, "Real-Time change point detection with application to smart home time series data," *IEEE Trans. Knowl. Data Eng.*, vol. 31, no. 5, pp. 1010–1023, 2019.
- [19] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian, "Certifying and removing disparate impact," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015, pp. 259–268.
- [20] R. Berk, H. Heidari, S. Jabbari, M. Kearns, and A. Roth, "Fairness in criminal justice risk assessments: The state of the art," *Sociol. Methods Res.*, vol. 50, no. 1, pp. 3–44, 2021.
- [21] N. Mehrabi, F. Forstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Comput. Surv.*, vol. 54, no.

6, pp. 1–35, 2021, doi: 10.1145/3457607.

- [22] S. Das, M. Donini, J. Gelman, K. Haas, M. Hardt, J. Katzman, K. Kenthapadi, P. Larroy, P. Yilmaz, and M. B. Zafar, “Fairness measures for machine learning in finance,” *J. Financ. Data Sci.*, 2021.
- [23] M. Hardt, E. Price, and N. Srebro, “Equality of Opportunity in Supervised Learning,” *CoRR*, vol. abs/1610.0, 2016, [Online]. Available: <http://arxiv.org/abs/1610.02413>
- [24] C. DeSmet and D. J. Cook, “HydraGAN: A cooperative agent model for multi-objective data generation,” *ACM Trans. Intell. Syst. Technol.*, 2024.
- [25] E. Arana-Chicas, F. Cartujano-Barrera, K. K. Rieth, K. K. Richter, E. F. Ellerbeck, L. S. Cox, K. D. Graves, F. J. Diaz, D. Catley, and A. P. Cupertino, “Effectiveness of recruitment strategies of Latino smokers: Secondary analysis of a mobile health smoking cessation randomized clinical trial,” *J. Med. Internet Res.*, vol. 24, no. 6, p. e34863, 2022.
- [26] J. Briscoe, C. DeSmet, K. Wuestney, A. Gebremedhin, R. Fritz, and D. J. Cook. 2024. Reducing Sample Selection Bias in Clinical Data through Generation of Multi-Objective Synthetic Data. In Proceedings of the 10th World Congress on Electrical Engineering and Computer Systems and Sciences (EECSS’24).
- [27] J. Briscoe and A. Gebremedhin. 2024. Facets of Disparate Impact: Evaluating Legally Consistent Bias in Machine Learning. In Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM ’24), October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA. <https://doi.org/10.1145/3627673.3679925>