

HW 6: Synthetic Data Release and Embedded Ethics

CS 208 Applied Privacy for Data Science, Spring 2025

Version 1.1: Due Fri, Mar. 14, 11:59pm.

Instructions: Submit a PDF file that contains both your written responses as well as your code to the assignment on Gradescope. Read the section "Collaboration & AI Policy" in the syllabus for our guidelines regarding the use of LLMs and other AI assistance on the assignments.

1. **DP Synthetic Data:** In this problem, you will create and analyze DP synthetic data. You will compare the results of running a regression on the synthetic data with your DP regression algorithm from HW4.

- (a) **Write a DP synthetic data function.** Write a function that takes a data bound b , a dataset $z = ((x_1, y_1), \dots, (x_n, y_n)) \in ([-b, b] \times [-b, b])^n$ a parameter $\varepsilon \geq 0$, a binning parameter k , and does the following: (a) bins the datapoints using a $k \times k$ grid, (b) constructs a ε -DP histogram h of the dataset, and (c) post-processes h to construct an ε -DP synthetic dataset $\tilde{z} = ((\tilde{x}_1, \tilde{y}_1), \dots, (\tilde{x}_m, \tilde{y}_m)) \in ([-b, b] \times [-b, b])^m$.
- (b) **Perform and evaluate linear regression on the synthetic data.** Given a DP synthetic data generator from Part 1a, we can perform DP linear regression by post-processing, running standard OLS regression on \tilde{z} to obtain an ε -DP slope $\tilde{\beta}$ minimizing $\sum_i (\tilde{y}_i - \tilde{\beta} \tilde{x}_i)^2$. Evaluate the resulting DP regression estimates exactly as in HW4 Problem 2d, using parameters $\varepsilon = .1$, $b = 1$, and $k = 20$, and many Monte Carlo trials of the following: For each $n = 100, 200, 300, \dots, 5000$, generate a dataset z , where the x_i 's are uniform in $[-1/2, 1/2]$ and $y_i = [x_i + \mathcal{N}(0, .02)]_{-1}^1$. Plot the bias and standard deviation of both the OLS estimate $\hat{\beta}$ and the DP estimate $\tilde{\beta}$ obtained by post-processing the DP histogram (note: you can use $m = n$ when generating the synthetic data). Your plot should have n on the x -axis, and bias and standard deviation on the y -axis on a scale from -1.0 to 1.0 . Try to run enough trials to obtain smooth curves.
- (c) **Compare the error.** Compare the bias and standard deviation of the above DP-histogram-based regression with the results obtained on HW4. Give an intuitive explanation for the differences you find.

2. **Applying Contextual Integrity:** Imagine a fictional technology company called Coachable. Coachable designs wearable fitness trackers for athletes. Coachable trackers collect data points about users blood flow and temperature in order to measure their resting heart rate, heart rate variability, and respiratory rate throughout the day and night. These measurements are used to calculate metrics on users sleep quality (including duration in bed, duration asleep, number of disturbances, length of time spent in different sleep stages, etc.), their level of physical and mental stress, their recovery rate, readiness for activity, and their overall cardiovascular health. In addition, users can log the following information in the Coachable journal to learn how different factors affect their training and performance:

- Alcohol and marijuana consumption
- Supplement use and dosage
- Caffeine consumption
- Medications and sleep aids
- Screen time and bedtime routines
- Air travel
- Stretching and other recovery modalities
- Nutrition and diet plans
- Menstruation and pregnancy
- Sexual activities

Coachable users receive detailed reports on how the behavior logged in their journal affects their athletic training, along with personalized training plans, lifestyle tips, and audio-guided workouts.

- (a) Using the data they collect, suppose Coachable is planning to compute various summary statistics about behaviors, traits of users, and athletic performance. They plan to use these statistics to micro-target ads to its users. For example, they may use these insights to target ads for products including, but not limited to, diet plans, supplements, workout equipment, and recreational activities. They also plan to make these statistics available to researchers, advertisers, and sports recruiters who are interested in the relationship between behaviors, traits of users, and athletic performance.

If we take the Coachable app to be playing a similar role as a human coach, explain how these additional data practices disrupt the informational norm(s) that operate in typical athlete-coach relationships. Explicitly identify the parameters of contextual integrity in your analysis.

- (b) Evaluate the disruptions you identified above. What are the context-specific values and goals of an athlete-coach relationship? How do these disruptions support or undermine these goals? Then, based on your evaluation, state whether you think Coachable should do anything differently with respect to these data practices.
- (c) Now imagine that Coachable plans to compute the summary statistics under differential privacy. How would your analysis and recommendations in Parts 2 and 3 change, if at all? How would deployment decisions, like how ϵ is set, impact your response?

3. **Sharing Project Ideas:** Recall that your final projects will be done in groups of 3-4 students. We would like you to form groups based on finding common interests with other students in the class, not just automatically grouping with people you already know. To that end, we will use the following process.

- By Tuesday 3/18 11:59pm: review the feedback you received on your hw4 project topic ideas, and add one or two of your ideas (possibly revised based on our feedback) to the following spreadsheet: Project Spreadsheet

- By Friday 3/21 11:59pm: add your name to the ‘Expression of Interest’ columns for at least two other topic ideas.

Looking ahead: by Monday 3/24 (immediately after Spring Break), you all will organize into groups (possibly with our help) and by Friday 3/28, you will submit a set of revised topic ideas with your group.

Collaborators

Please list all collaborators for this problem set. ChatGPT and other AI tools should be treated similarly to collaboration with your peers in the class. You may use these tools to help you understand the material and as part of your brainstorming process, but you should not be asking the tools to solve the homework problems for you. If you do use such tools, you must cite them and list the prompts you entered and responses obtained below.