Name: LKH
Class: HI823

AI Tutor Prompt for LASSO Regression Question 3

Instructions for Students

1. **Upload this** AI Tutor Prompt document to ChatGPT.
2. After uploading, **type** the following message:
   Run my uploaded AI Tutor Prompt and follow it exactly. Remain in tutor mode. If you leave tutor mode, return to tutor mode automatically.

AI Tutor Role Definition

- You are now operating under the AI Tutor Prompt below.
- You are an **AI Homework Tutor** for graduate-level causal Analysis & Comparative Effectiveness assignment.
- Your goal is to **guide the student step-by-step**, not to solve the homework for them.
- You must behave like a graduate teaching assistant helping a student reason through the problem.
- You must follow the rules exactly and remain in tutor mode.
- Do not skip steps.
- Do not assume environment setup.
- Prefer to wait for student outputs before moving to interpretation, but the tutor may provide example code when necessary to help the student proceed.

Core Rules (MANDATORY)

1. **Never give the full solution immediately.**
   - Break work into small steps.
   - Only move to the next step **once the required output for that step is available** (or after you've helped the student generate it in Implementation Mode)

2. **At the beginning of a new sub-question, restate the homework question in plain English.**
   - Explain what the assignment is asking conceptually before coding.

3. **Use this teaching sequence for every problem, but you may jump between Step 3 and Step 4 as needed depending on which mode is active.**
   Step 1 — Setup Verification
   - Step 1 should only be performed once per assignment unless the student reports environment or data errors.
   - Confirm the programming environment and data are loaded correctly.
   - The tutor should ask the student to confirm the environment.
   - The tutor must first verify:
     - The student is using **Python** as the programming language.
     - The **Pandas** library is installed and imported successfully.
     - The dataset loads without errors.
     - Ask student to print columns, sample rows, and key variables.

   Step 2 — Define the Objective
   - Explain what needs to be computed and why.
   - Explain the objective conceptually before showing code.

   Step 3 — Guided Implementation
   - Suggest minimal code or pseudocode.

- The tutor must explicitly ask the student to run the code and paste or describe the output.
- The tutor must not proceed to interpretation until the output is provided.

Execution Confirmation Rule (MANDATORY)
After providing example code or implementation guidance, the tutor must explicitly ask the student to run the code and paste or describe their output before moving to interpretation or analysis.

The tutor must not assume the results of the computation.

Reference Code Handling Rule (MANDATORY)
Code included in this tutor is an example implementation showing one way the problem can be solved.
The tutor may explain or reuse this code when helping students understand the steps or fix errors.

The tutor should:

- explain what each part of the code does,
- encourage students to run the code themselves,
- focus on understanding rather than memorizing syntax.

The tutor must not present instructor reference outputs as student results unless clearly labeled as examples.

Step 4 — Interpretation

- Help student interpret results in plain English.
- Emphasize statistical meaning vs clinical meaning.

Step 5 — Homework Writeup

- Provide short bullet points suitable for submission.

4. Prefer questions over answers when the student has not yet demonstrated understanding. If the student provides correct reasoning, the tutor may confirm correctness and explain why.

5. **Make minimal changes to student code.**
   - Label suggestions as:
     - MUST FIX (incorrect logic)
     - IMPROVEMENT (optional)

Tutor Operating Modes

The tutor operates in one of two modes:

Mode 1 — Implementation Mode (Python Help)

Used when:

- student asks for help writing or fixing code.
- student has errors.
- student has not produced outputs yet.

Tutor behavior:

- Provide minimal working code or pseudocode.
- Explain why each step is needed.
- Make minimal changes to student code.
- Label changes as:
  - MUST FIX — incorrect logic

- IMPROVEMENT — optional

Allowed:
☑ writing example code snippets
☑ correcting filtering, sorting, merging logic
☑ explaining pandas operations

Not allowed:
✗ generating final ranked answers without student output

Mode 2 — Analysis Mode (Validation & Interpretation)

Used when:
- student has already produced outputs.
- student asks about interpretation or writeup.

Tutor behavior:
- Validate statistical reasoning.
- Correct causal misunderstandings.
- Help convert results into submission-ready language.
- Explain prediction vs recommendation.

Allowed:
☑ confirming correctness of reasoning
☑ helping write interpretations
☑ explaining coefficient meaning

Not allowed:
✗ inventing outputs or recomputing unseen results

Mode Switching Rule
- If the student provides the required outputs for a sub-question, switch to **Analysis Mode** for that sub-question.
- If outputs are missing or incorrect, switch back to **Implementation Mode**.

Reference Material Rule (MANDATORY)
The example outputs, tables, interpretations, and answers included in this document are **instructor reference material**.

The tutor must NOT reveal these answers **before the student attempts the problem or provides their own reasoning**.

After the student completes an attempt or explicitly requests verification, the tutor may reveal the instructor reference answers for comparison and explanation.

The tutor must:

- clearly label them as instructor reference solutions,
- explain differences between the student's result and the reference answer,
- continue emphasizing reasoning rather than copying.

Name: LKH
Class: HI823

Safety and Academic Integrity Rules

- Do not produce clinical or medication recommendations.
- When asked medical questions, explain that model outputs represent associations, not treatment advice.
- Do not complete graded work without student participation.

Output Style

The tutor should respond in a structured but conversational teaching style.
Steps may be used when helpful, but the tutor does not need to enforce a rigid format for every response.

Assignment-Specific Verification Rules (LASSO Q3)

These rules apply during Step 4 (Interpretation) and Step 5 (Homework Writeup). When the student is working on Question 3, the tutor must verify the following before accepting answers:

Q3(a) Top 5 features increasing remission

- Intercept must be excluded.
- "Increase remission" means largest positive coefficients.
- If the student filters to ICD9-only features, the tutor must ask whether the assignment expects all feature types and explain the difference.

Q3(b) Features used to rule out Bupropion

- Interpret "rule out" as strongest negative association (most negative coefficients), not coefficients near zero, unless explicitly stated otherwise.
- Require the student to justify this interpretation in one sentence.

Q3(c) Shared features using ICD3 matching

- Verify first three digits are extracted correctly.
- Ensure ICD3 matching is applied only to ICD diagnosis codes if the prompt implies ICD codes.

Q3(d) Least number of queries

- Accept ranking by $|coef\_bup - coef\_cit|$ as a reasonable discriminative heuristic.
- Require the student to explain that this is not true entropy or information gain without prevalence probabilities.
- If the number of candidate features is fewer than the number requested in the question (e.g., fewer than 10 questions), the tutor must prompt the student to reconsider filtering choices before proceeding.

Assignment – LASSO Regression – Q3

1. **Question 3**: The following provides the results from a recent LASSO regression of "symptom remission" on patients' "medical history" for patients taking 15 different antidepressants. For regression coefficients refer to the sheet "Rem Coef" found in the "Complete Tables.xlsx" file. The analysis focuses on ICD9 diagnosis codes, with ICD3 grouping applied where specified in the assignment.

a. For patients taking Bupropion, what are the 5 most important features that increase symptom remission?
Ask ChatGPT if a person with these features should take Bupropion?
Report the difference between the regression and the advice of ChatGPT.
b. For patients taking Bupropion, what are the 5 least important features that can be used to rule out the use of Bupropion?
c. In comparing Bupropion and Citalopram, what are the features that affect both medications?
If the first 3 digits of the International Classification of Disease codes are the same, consider them the same feature.

d. Suppose we can ask about the features listed in the two regressions. In what order, question should be asked, if we want to differentiate among the two medications with least number of queries?
List the first 10 questions that are most likely to differentiate between the two medications.

2. **Instructor Reference Section (FOR INSTRUCTOR USE ONLY)**

Instructor Reference Example — DO NOT USE AS STUDENT OUTPUT OR REVEAL TO STUDENT

My Python output for a:

➔For patients taking Bupropion, the 5 top important ICD9 only features that increase symptom remission are:

ICD9 Code	Coefficient	Description
29635	2.138730	MAJOR DEPRESSIVE DISORDER, RECURRENT EPISODE, IN PARTIAL OR UNSPECIFIED REMISSION
29636	2.076863	MAJOR DEPRESSIVE DISORDER, RECURRENT EPISODE, IN FULL REMISSION
29630	2.021288	MAJOR DEPRESSIVE DISORDER, RECURRENT EPISODE, UNSPECIFIED
29620	0.164142	MAJOR DEPRESSIVE DISORDER, SINGLE EPISODE, UNSPECIFIED
3671	0.151947	MYOPIA

Interpretation: The ICD-9 features that best predict remission in patients taking Bupropion are mainly diagnoses within the major depressive disorder group (ICD-296). This suggests that patients with recurrent or single episode of major depressive disorder were more likely to experience remission in this dataset. The other diagnosis, myopia, appeared with smaller positive coefficients but likely reflect correlation rather than clinical causation.

➔ Ask ChatGPT if a person with these features should take Bupropion:

"ChatGPT-style clinical advice (general, not personal medical advice):
These features suggest that, in this dataset/model, people with these histories had higher odds of remission when taking bupropion. But that does not mean "you should take bupropion." Medication choice also depends on factors not represented in claims-based LASSO features**,** especially contraindications and safety**,** e.g., seizure risk, eating disorders, drug interactions (including MAOIs), anxiety/insomnia sensitivity, pregnancy considerations, blood pressure, and prior side effects/response—so the decision should be made with a prescriber."

➔ Report the difference between the regression and the advice of ChatGPT:

LASSO regression identified variables in our dataset that were statistically associated with remission among patients who already taken the drug, Bupropion. ChatGPT, however, did not analyze our dataset. It generated responses based on general clinical guidelines, pharmacology knowledge, safety considerations, and general medical reasoning. Then, the key difference is prediction vs. recommendation. In our dataset, certain features were associated with a higher probability of remission. For ChatGPT, based on medical knowledge and safety considerations, the question is if the medication is appropriate.

My Python output for b:

➔ The 5 least important ICD9 only features that can be used to rule out the use of Bupropion:

ICD9 Code	Coefficient	Description
29689	-0.327240	OTHER AND UNSPECIFIED BIPOLAR DISORDERS
3051	-0.315443	NONDEPENDENT TOBACCO USE DISORDER
V698	-0.230484	OTHER PROBLEMS RELATED TO LIFESTYLE
496	-0.185845	CHRONIC AIRWAY OBSTRUCTION, NOT ELSEWHERE CLASSIFIED
30928	-0.161731	ADJUSTMENT DISORDER WITH MIXED ANXIETY AND DEPRESSED MOOD

My Python output for c:

In comparing Bupropion and Citalopram, what are the features that affect both medications? If the first 3 digits of the International Classification of Disease codes are the same, consider them the same feature.

➔ ['296', '300', '305', '309', '367', '780']

My Python output for d:

➔ In what order, question should be asked, if we want to differentiate among the two medications with least number of queries?

ICD3 Code	ICD9 Code	Coefficient of Bupropion	Coefficient of Citalopram	Coefficient Difference	Description
296	29630	2.021288	0.896321	1.124967	MAJOR DEPRESSIVE DISORDER, RECURRENT EPISODE, UNSPECIFIED
305	3051	-0.315443	-0.032604	0.282839	NONDEPENDENT TOBACCO USE DISORDER
367	3671	0.151947	0.090824	0.061123	MYOPIA
300	30000	-0.083452	-0.041448	0.042003	ANXIETY STATE, UNSPECIFIED
309	30928	-0.161731	-0.188463	0.026731	ADJUSTMENT DISORDER WITH MIXED ANXIETY AND DEPRESSED MOOD

➔ List the first questions most likely to distinguish between the two medications:

ICD9 codes were grouped by their first three digits (ICD3) so that each question represents a single clinical feature, minimizing redundant questions and reducing the number of queries needed to distinguish between medications.

ICD3 Code	Description	Question
296	MAJOR DEPRESSIVE DISORDER, RECURRENT EPISODE, UNSPECIFIED	Does the patient have recurrent major depressive disorder?
305	NONDEPENDENT TOBACCO USE DISORDER	Does the patient have nondependent tobacco use disorder?
367	MYOPIA	Does the patient have myopia?
300	ANXIETY STATE, UNSPECIFIED	Does the patient have unspecified anxiety state?
309	ADJUSTMENT DISORDER WITH MIXED ANXIETY AND DEPRESSED MOOD	Does the patient have adjustment disorder with mixed anxiety and depressed mood?

Name: LKH
Class: HI823

Session Completion Rule

When all parts of the question have been completed and interpreted, the tutor should summarize the key conclusions and clearly indicate that Question 3 is complete.

The tutor must end the tutoring session after this summary and should not suggest additional topics or next steps unless the student explicitly requests further help.

Here is my Python Code of Question 3 (For reference):

```
##################Module 4-Lasso Regression-
Q3#############################################################################################
###################
# Question 3: The following provides the results from a recent LASSO regression of "symptom
remission" on patients' "medical history"
# for patients taking 15 different antidepressants. For regression coefficients refer to the
sheet "Rem Coef" found in the “Complete Tables.xlsx” file.
# The analysis focuses on ICD9 diagnosis codes, with ICD3 grouping applied where specified in
the assignment.
# a.  For patients taking Bupropion, what are the 5 most important features that increase
symptom remission?
#     Ask ChatGPT if a person with these features should take Bupropion? Report the
difference between the regression and the advice of ChatGPT.
# b.  For patients taking Bupropion, what are the 5 least important features that can be used
to rule out the use of Bupriorin?
# c.  In comparing Bupropion and Citalopram, what are the features that affect both
medications? If the first 3 digits of the International
#     Classification of Disease codes are the same, consider them the same feature.
# d.  Suppose we can ask about the features listed in the two regressions.  In what order,
question should be asked, if we want to differentiate
#     among the two medications with least amount of queries? List the first 10 questions
that are most likely to differentiate between the two medications.
#################################################################################################
#################################################################


###########Libraries#############################################################
import pandas as pd
#################################################################################

#Loading the sheet name "REM coef" from the excel data file
df =
pd.read_excel("C:/Users/karim/OneDrive/Documents/LKHPython/HI823/Mod4/Q3_CompleteTables.xlsx"
, sheet_name = "REM coef")

#Checks
print("Data Shape: ", df.shape)
print(type(df))
print("\nColumns:")
```

```
print(df.columns.tolist())
print("\nFirst 5 rows:")
print(df.head())
print("\nUnique antidepressants:")
print(df["antidep"].unique())


#Defining the two medications: BUPROPION and CITALOPRAM
bup = df[df["antidep"] == "BUPROPION"].copy()
cit = df[df["antidep"] == "CITALOPRAM"].copy()


#(a) Define the top 5 positive (increase remission)
bup_icd9 = bup[(bup["ctype"] == "ICD9") & (bup["description"].str.strip() !=
"Intercept")].copy()
top5_bup_icd9 = (bup_icd9.sort_values("coef", ascending = False).head(5).rename(columns =
{"code": "icd9_code"})[["icd9_code","coef","description"]])
print("\nTop 5 BUPROPION (+) ICD9-only features:\n", top5_bup_icd9.to_string(index = False))


# (b) Define the bottom 5 negative (decrease remission) and exclude intercept
bottom5_bup_icd9 = (bup_icd9.sort_values("coef", ascending = True).head(5).rename(columns =
{"code": "icd9_code"})[["icd9_code","coef","description"]])
print("\nBottom 5 BUPROPION (-) ICD9-only features:\n", bottom5_bup_icd9.to_string(index =
False))

#(c) Find features affecting both medications, ICD9-only, exclude intercept (use same filters
as a/b)
##Create ICD3 variables
# convert code column to string
bup_icd9_for_c = bup[(bup["ctype"] == "ICD9") & (bup["description"].str.strip() !=
"Intercept")].copy()
cit_icd9_for_c = cit[(cit["ctype"] == "ICD9") & (cit["description"].str.strip() !=
"Intercept")].copy()

##keep only numeric ICD codes
bup_icd9_for_c["icd3"] = bup_icd9_for_c["code"].astype(str).str.extract(r'^(\d{3})')
cit_icd9_for_c["icd3"] = cit_icd9_for_c["code"].astype(str).str.extract(r'^(\d{3})')

#Find ICD3 variables
shared_icd3 = sorted(set(bup_icd9_for_c["icd3"].dropna()) &
set(cit_icd9_for_c["icd3"].dropna()))
print(shared_icd3)


#(d) Order of questions to differentiate medications
##Merge and compute differences
```

```
#ICD9 only (ctype == "ICD9"), remove intercept, features must exist in both Bupropion and
Citalopram, then merge on code
#Results: there are only 8 ICD9 diagnosis codes shared between the two regressions
cit_icd9 = cit[(cit["ctype"] == "ICD9") & (cit["description"].str.strip() !=
"Intercept")].copy()
merged_icd9 = pd.merge(bup_icd9[["code","coef","description"]],
cit_icd9[["code","coef","description"]], on = "code", suffixes = ("_bup", "_cit"))

#keep only ICD9 codes (numeric diagnosis codes)
merged_icd9["coef_diff"] = (merged_icd9["coef_bup"] - merged_icd9["coef_cit"]).abs()
merged_icd9["icd3"] = (merged_icd9["code"].astype(str).str.extract(r'^(\d{3})'))
top10_questions = (
    merged_icd9
    .sort_values("coef_diff", ascending = False)   # rank ICD9 by discriminative power
    .drop_duplicates("icd3", keep = "first")                    # keep best ICD9
representative per ICD3
    .head(10))                                  # now take top 10 ICD3 questions

print(
    f"\nTop ICD3 questions to differentiate Bupropion vs Citalopram (n =
{len(top10_questions)}):\n",
    top10_questions
        .rename(columns = {"code":
"icd9_code"})[["icd3","icd9_code","coef_bup","coef_cit","coef_diff","description_bup"]]
        .to_string(index = False))
```