

A motivating application: Genome-Wide Association Studies (GWAS)

The human genome can be thought of as a word of length 3×10^9 in a 4-letter alphabet $ACGT$. Each person has two copies of this word, and the two copies are similar but not identical (similarly for copies from different individuals). A typical number cited is that two copies of the genome are about 99.9% identical, meaning several million letters different ($0.01\% = 3 \times 10^6$). A common representation of an individual genome (two copies) is as a *ternary* vector of length 3×10^9 : $x_j \in \{0, 1, 2\}$, $j = 1 \dots 3 \times 10^9$.

GWAS basic idea: for a number of people ($n = 1000$ traditionally, today $n = 10^5$ is common), measure their “entire” genome and also a property (phenotype) of interest that has a heritable component, such as height or whether they have diabetes. Then look for statistical connections between each point in the genome and the phenotype, that is test:

H_{0i} : point (or region) i has no statistical connection to the phenotype, $i = 1, \dots, 3 \times 10^9$.

If we reject some of the null hypotheses we learned something about the genetics of this phenotype. In this naive view we have $p = 3 \times 10^9$ tests with $n \approx 1000$ observations — very wide data.

A few more details that will be important in our further consideration of this problem:

1. The genomes of individual has a unique and interesting correlation structure called *linkage disequilibrium* (*LD*) that is related to how inheritance actually works: Points in the genome that are physically close tend to be co-inherited and therefore are highly correlated (that is, two chromosomes that have the same letter in point i will tend to have the same letter in point $i + k$ for small $|k|$). In contrast, far away points in the genome are uncorrelated under some assumptions, or more weakly correlated under others.
2. Measuring all 3×10^9 in the genome is called *sequencing*. However because of the combination of only 0.01% difference and LD, it is considered that the entire “common” diversity of the genome can be represented by measuring much fewer locations, typically about 10^6 — this is called *genotyping*.
3. Additional interesting aspects that will not be widely considered in this course (\Rightarrow Statistical Genetics course next semester):
 - Case-control sampling and its implications for statistical modeling and testing
 - The stochastic process of mutation and recombination and their spread in the population
 - Population structure and its implications for GWAS

This GWAS problem has been widely researched, and it has some important and interesting aspects related to core areas of our course. We will spend some time on several of these:

1. GWAS as an example of wide data with $n \ll p$, and implications for modeling it
2. Similar but slightly different view: GWAS as an example of multiple testing with high multiplicity, and the implications for how testing should be done
3. (TODAY) The privacy issues in releasing GWAS information for scientific research — how can we preserve the maximum useful information while preserving the privacy of study participants?

Privacy in big data

The basic problem: How to collect and publish data in a way that will be both:

- Useful for valid statistical analysis and scientific research
- “Safe” in terms of reasonably protecting the privacy of the individuals whose information was collected in the study

For most of the discussion we will assume standard tabular data n individuals, each with p pieces of information, as in GWAS.

We are looking for non-trivial privacy protection, that will be robust against:

- Smart statistical analysis by the “privacy attackers”
- Availability of additional outside information in helping to identify participants and violate privacy

Unsatisfactory but common solution 1: anonymization

It seems reasonable that releasing the information without the identifying information of the participants like name, address, etc. will protect their privacy. This has been proven to fail in GWAS: it is enough for someone to have a tiny part of someone’s genome to find whether that person is in the GWAS, and then know their entire genome.

Specific example: One of the most famous public genetic databases was called HapMap, which released increasingly detailed genomes of random individuals from 2005 onwards. The only personal details were sex, age and country/US state for each individual. In 2013 a Science paper demonstrated how the identities of some individuals in HapMap can be exposed:

- The male genome has a special small genetic pieces called the *Y-chromosome* that is directly inherited from father to son. It is widely used in relative search, and so millions of people have published some of their Y information online with their name and are actually finding relatives on their father’s side.
- If a relative of a HapMap individual published their Y information online, we can identify that the HapMap individual is their relative on their paternal inheritance line — so probably have the same last name.

- The combination of age, country/state and last name is sometime enough to uniquely identify a person in the phone book and other publicly available sources.

Using this approach, they were actually able to positively identify several HapMap participants, meaning they now have their non-anonymized genomes, with severe consequences.

Unsatisfactory but common solution 2: releasing summaries

Assume now our $n = 2000$ GWAS samples are made of $n_1 = 1000$ *cases* who have some disease, say Type-I diabetes, and $n_2 = 1000$ samples of healthy *controls*. All their genomes are measured at $p = 10^6$ locations. It is widely recognized that in addition to analyzing this dataset separately, releasing it to the scientific community is of great interest, for example to combine with other studies and increase power.

The summary approach amounts to releasing only two tables of size $3 \times p$, one summarizing the statistics of the cases genotypes and one summarizing the statistics of the controls.

Now assume we have a genome of a specific individual, but we don't know whether they are a case (sick) or control (healthy). Two questions arise:

- If we know that this individual was in the current GWAS, can we find out whether they are a case or control?
- If we don't know whether this individual is in the study, can we separate the three options: not in the study/case/control?

The surprising(?) result is that we can typically positively answer the two questions above: not only identify whether the individual is a case or control if in the study, but also whether they were in the study at all.

For simplicity let's now assume that the disease studied is not genetic at all (say HIV), and a slightly simpler genotype structure: binary and independent coin tosses, and of length 10^5 only, to make the problem tougher (why?). So the data is matrices $X_{1000 \times 10^5}^{(case)}$, $X_{1000 \times 10^5}^{(cont)}$ with $X_{ij} \sim \text{Ber}(0.5)$ all i.i.d. The released information is a pair of vectors $\hat{p}_{case}, \hat{p}_{cont} \in [0, 1]^{10^5}$.

Now assume we are given a "genome" $x \in \{0, 1\}^{10^5}$ and we want to see whether it is a case, control, or not in our study. Let's start from the simpler problem:

$$H_0 : x \text{ case} \Leftrightarrow x_j \sim \text{Ber}(\hat{p}_{case,j}) \quad \text{vs} \quad H_1 : x \text{ control} \Leftrightarrow x_j \sim \text{Ber}(\hat{p}_{cont,j}).$$

Of course we can switch the hypotheses. Note these are simple hypotheses (specify the entire distribution).

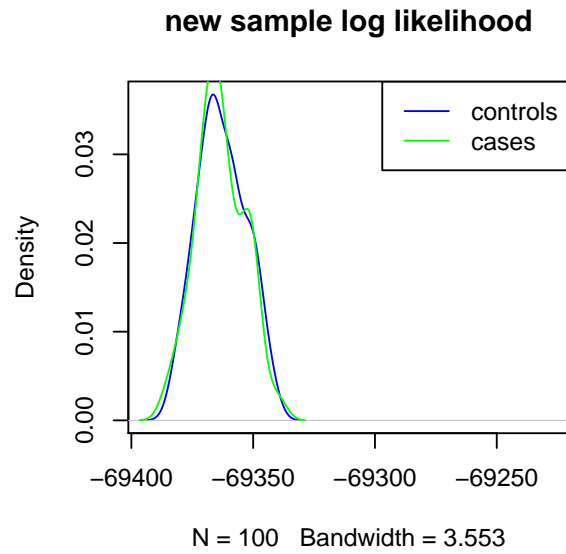
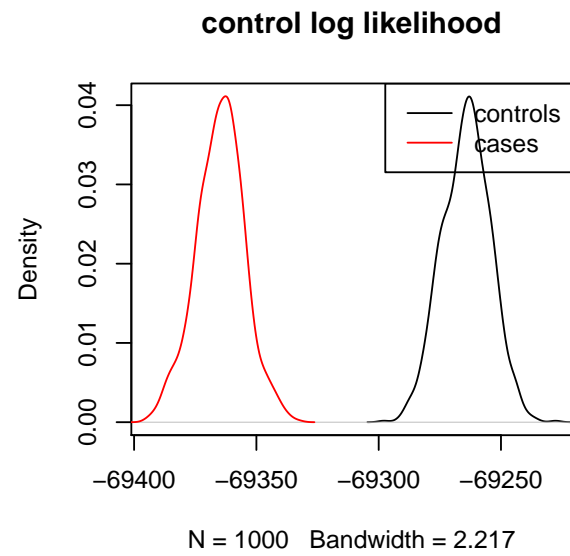
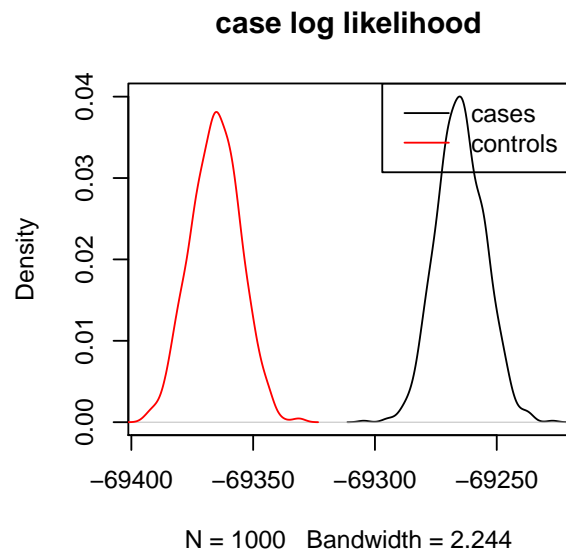
We can write the log-likelihood ratio for this problem:

$$\lambda = \left[\sum_j x_j \log(\hat{p}_{case,j}) + (1 - x_j) \log(1 - \hat{p}_{case,j}) \right] - \left[\sum_j x_j \log(\hat{p}_{cont,j}) + (1 - x_j) \log(1 - \hat{p}_{cont,j}) \right].$$

We can draw some data and present a histogram of the first sum (case likelihood) and second sum (control likelihood).

Next, we can consider the case of a third option that it is neither population, and not surprisingly the distribution of the two sums is the same and looks like the "wrong" distribution above.

The bottom line is, it turns out to be very easy to find out the hidden information about the disease status of the individual even if only summary statistics are released. Proving this rigorously will be part of the HW that will be given next week...



Example of non-privacy invasion

It is important to also understand that information can reveal information about a person without being an invasion of their privacy. For example, consider an imaginary study that shows that smoking hurts work productivity. After reading it the boss checks which of the employees smokes and fires them. So the fired employee is damaged by the results of the study that “taught” the manager that he is a bad employee. But assuming that the smoking status is not private, there was no violation of the employee’s privacy, since his personal information was not used in any way in the study.

Hence harming X through information is not the same as harming X’s privacy. There is privacy violation only if his information was used in the study and in that way it was exposed. Specifically this means “The study did not expose information on X” is not a good criterion for privacy.

The key to a proper definition: Limiting the information that can be exposed about *participants in our study* as a result of information we release. So we want to expose almost no additional information about the participants in our study, to any outside observer, whatever else they know about those participants or other people. Can we do this and still release information? It turns out we can, through the notion of Differential Privacy.