Kamel Lahouel, Johns Hopkins University

A statistical method for analyzing circulating tumor DNA

Kamel Lahouel

Summary

The setting :

- We have a certain number of mutations of interest for every patient.
- Every mutation has MAF measurements coming from different templates (usually 2 or 4 templates)
- For every mutation and template we have several MAF measurements corresponding to the number of wells observed.
- Number of wells for every mutation and template is 95 (digital).

General principles of the method :

1. For every mutation and template, score every well according to the observed MAF for that well.
2. For every mutation and template, combine the wells scores in order to get a score specific to every mutation and every template.
3. Combine the scores of all the mutations and templates to get a final score for the patient corresponding to the reported specificity.

Scoring a single well for a fixed mutation and template

This scoring is based on 2 evaluations :

- How unlikely is the MAF measurement with respect to the control distribution.
- Given an estimation of the Genome Equivalence number (GE), and assuming that we started with $\frac{1}{GE}$ mutants, decide if the MAF observed is plausible or too low for an initial signal of $\frac{1}{GE}$.

For the first evaluation:

- We learn the control (Null) distribution of the corresponding mutation from the control (”NL763”) templates.
- This distribution is modeled by a mixture of a point mass at 0 corresponding to the event of observing $MAF = 0$ and a continuous density modeling the positive MAFs.

- In general, the higher the mutation's background, the lower is the weight of the point mass at 0 and the heavier are the tails of the continuous density part.
- Important remark: If we don't observe an $MAF > 0$ for a particular mutation in the control templates, we do not put 100% of the mass on 0 in our corresponding control distribution. We still have a continuous part.
- The reason is the following: The estimated control distribution is a mix between a distribution using only the mutation specific MAFs of the controls and a global null using all mutations MAFs of the control templates.
- The more we observe MAFs under controls for the particular mutation the less weight we put on the global distribution and vice versa.

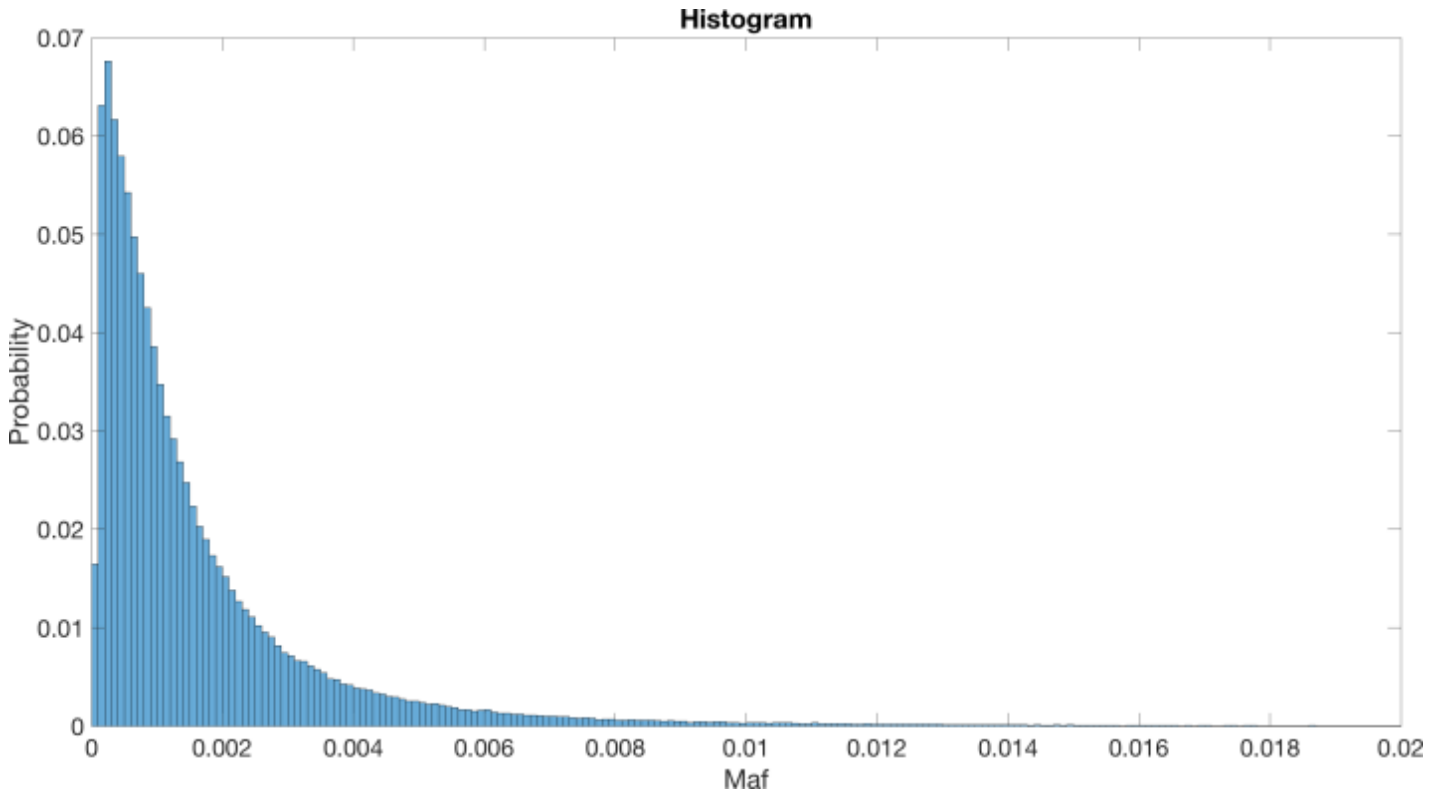

Figure: Example of continuous density part of the control distribution of a high background mutation. The right tails are heavy. The weight of the point mass at 0 (not represented here) is 0.92

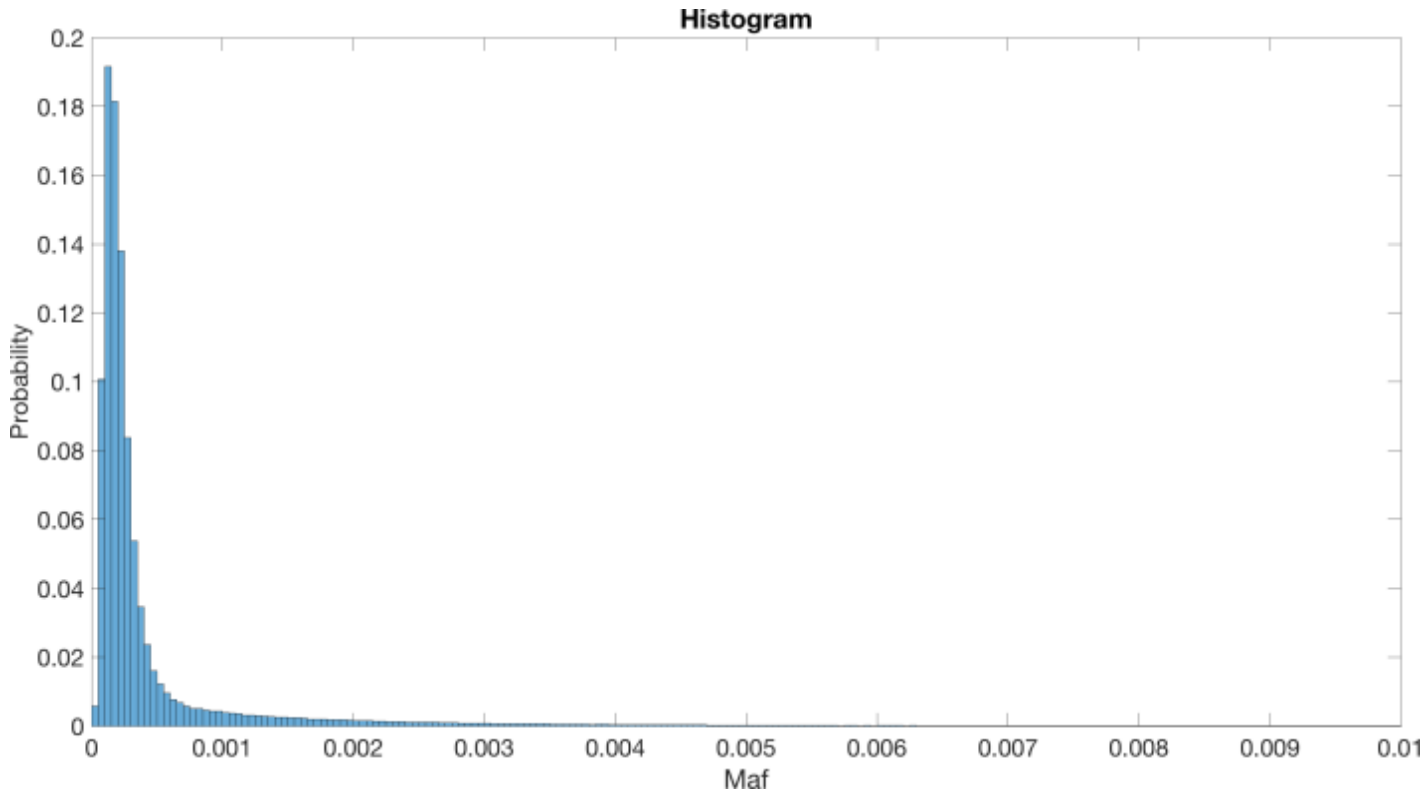

Figure: Example of continuous density part of the control distribution of a low background mutation. The right tails are not heavy. The weight of the point mass at 0 (not represented here) is 0.99

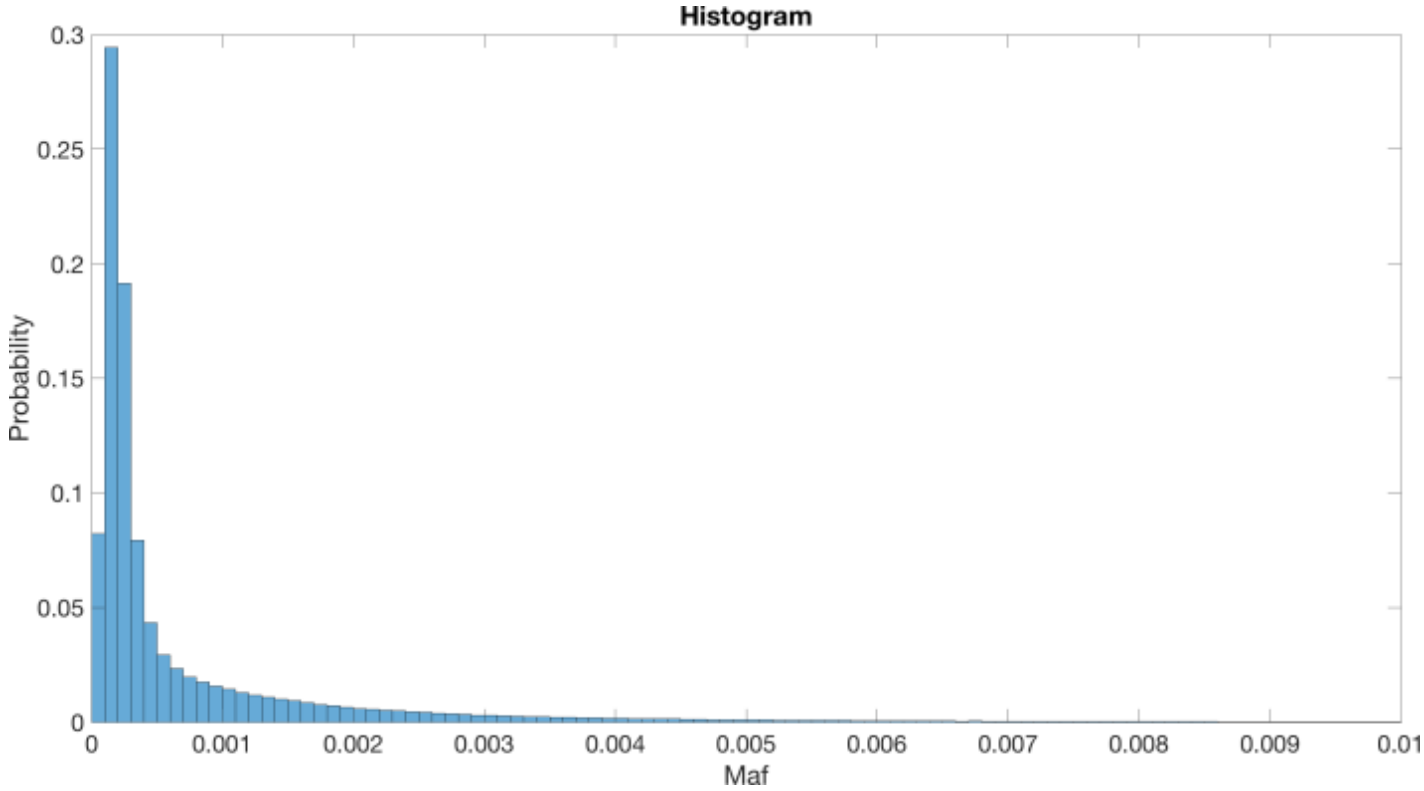

Figure: Example of continuous density part of the control distribution of a mutation where we did not observe any positive MAF under controls. Therefore, this density is nothing but the global null density. The right tails are heavier than the low background example but the weight of the point mass at 0 (not represented here) is 0.9985

- Once we have the control distribution for every mutation, we set a "tolerance level" p, that is 0.05,0.1 or 0.2. We compute the p-value of the observed MAF and correct it by the number of wells tested (multiply).
- We compute a preliminary score in the following way: if the corrected p-value is lower than the tolerance level, then the score is the the logarithm of the p-value. Otherwise, the score is 0. The lower is the score, the more significant.

- Example: Suppose that you observe a well with $MAF = 0.01$ where the mutation is the high background mutation illustrated previously and we set $p = 0.2$. The corrected p-value is 0.013 which is lower than 0.2. Therefore, the preliminary well score is $\log(0.013) = -4.3428$.
- Suppose now that you observe a well with $MAF = 0.002$ where the mutation is the high background mutation illustrated previously and we set $p = 0.2$. The corrected p-value is 0.24 which is higher than 0.2. Therefore, the preliminary well score is 0.

Second evaluation:

- We want to assess whether or not the observed MAF is plausible for an initial signal of $\frac{1}{GE}$
- For this, we use the "spike-in" experiments to learn the control distribution of the MAF s in the situation where we start with 1 mutant among a number of Genome equivalence GE.
- We then set a threshold t defined by:

$$P(MAF < t | MAF < \frac{1}{GE}) = 0.005.$$

- These thresholds are the same across all mutations and patients.
- The final score of a single well is then 0 if the corresponding MAF is lower than this threshold and equal to the preliminary score otherwise. This finalizes describing scoring a single well !

Table: Example of GEs and corresponding thresholds.

GEs	Thresholds
100	0.00147
200	0.00069
300	0.00044
400	0.00033
500	0.00024

Scoring a particular mutation in a particular template

- Now that we have a score for every well, we take the scores of all wells corresponding to the mutation and particular template of interest and sum them.
- To assess how significant is that sum: We compute the sums using the same procedure but for all mutations in the test templates (The "PLS" templates) that are not in the list of mutations we want to assess. The idea is that these give a conservative control distribution .
- The rank of our sum among these sums divided by the total number of mutations observed is the p-value and the score of the mutation for that particular template

- Now that we have a score for every well, we take the scores of all wells corresponding to the mutation and particular template of interest and sum them.
- To assess how significant is that sum: We compute the sums using the same procedure but for all mutations in the test templates (The "PLS" templates) that are not in the list of mutations we want to assess. The idea is that these give a conservative control distribution .
- The rank of our sum among these sums divided by the total number of mutations observed is the p-value and the score of the mutation for that particular template.

- For example, if for a mutation and template of interest, we observe 5 wells with $MAF > 0$ scored respectively $-1,\ 0,\ -1.5,\ -2,\ 0$. Then summing these scores gives -4.5.
- Assume now that we have 270 "control" mutations from the PLS templates and that only one value is lower that -4.5. The corresponding p-value of the mutation for that template would be $\frac{1}{270} = 0.0037$.

Scoring the patient

Now, we have a score for every mutation of interest and every template. For example assume we are assessing 2 mutations in 4 different templates:

Mutation \ Template	PLS01	PLS01A	PLS02	PLS02A
Mutation1	p_{11}	p_{12}	p_{13}	p_{14}
Mutation 2	p_{21}	p_{22}	p_{23}	p_{24}

- The final score is $\sum_{i=1}^{2}\sum_{j=1}^{4} -2\log(p_{ij})$. We compare that sum to a distribution of $\chi^2(16)$ (in general the degrees of freedom are number of templates times number of mutations assessed times 2) to generate a final p-value. One minus that p-value is the reported specificity of the patient !

5 minute Q&A

Chair/Co-Chair/NCI
feed Zoom Chat questions to presenter
and Track Time
NCI and Production Team
answer Chat questions not related to presentations
and use Slack