

# **Variant calling and filtering for INDELs**

Erik Garrison

SeqShop @ University of Michigan

# Overview

1. **Genesis of insertion/deletion (indel) polymorphism**
2. Standard approaches to detecting indels
3. Assembly-based indel detection
4. Haplotype-based indel detection
5. Primary filtering: Bayesian variant calling
6. Post-call filtering: SVM
7. Graph-based resequencing approaches

# **An INDEL**

A mutation that results from the gain or loss of sequence.

**AATTAGCCATTA**


**AATTA--CATTA**

# INDEL genesis

A number of processes are known to generate insertions and deletions in the process of DNA replication:


- Replication slippage
- Double-stranded break repair
- Structural variation (e.g. mobile element insertions, CNVs)

# DNA replication



# Polymerase *slippage*

## A) Slippage Event




(A) During replication, polymerase slippage and subsequent reattachment may cause a bubble to form in the new strand. Slippage is thought to occur in sections of DNA with repeated patterns of bases (such as CAG), represented here by matching colors. Then, DNA repair mechanisms realign the template strand with the new strand and the bubble is straightened out. The resulting double helix is thus expanded.

## B) No Slippage



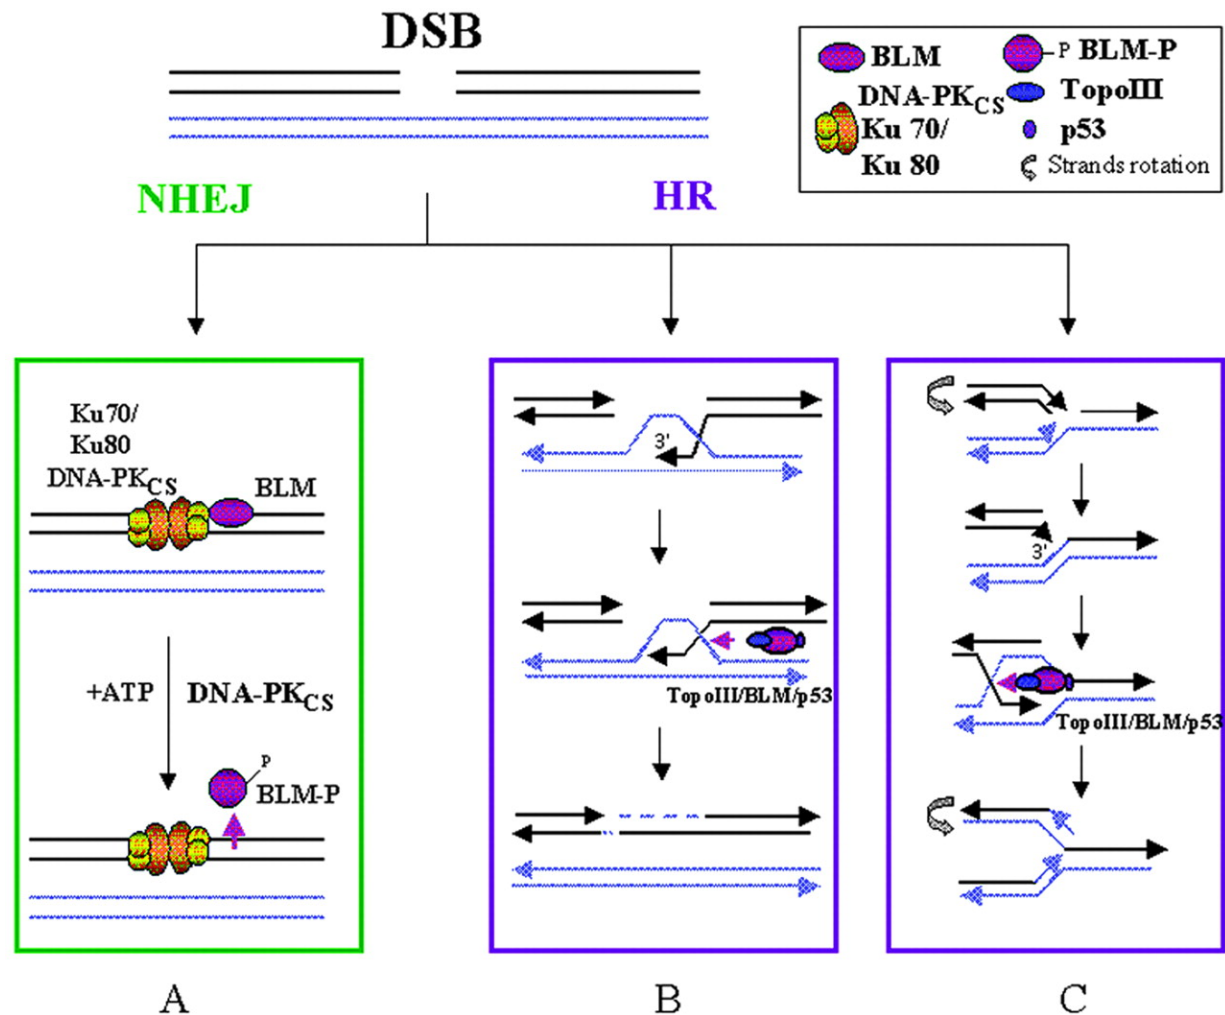
(B) Polymerase slippage, as theorized, cannot occur in DNA without repeating patterns of bases.

# Insertions and deletions via slippage



Energetic signatures of single base bulges:  
thermodynamic consequences and biological  
implications. Minetti CA, Remeta DP, Dickstein R,  
Breslauer KJ - Nucleic Acids Res. (2009)

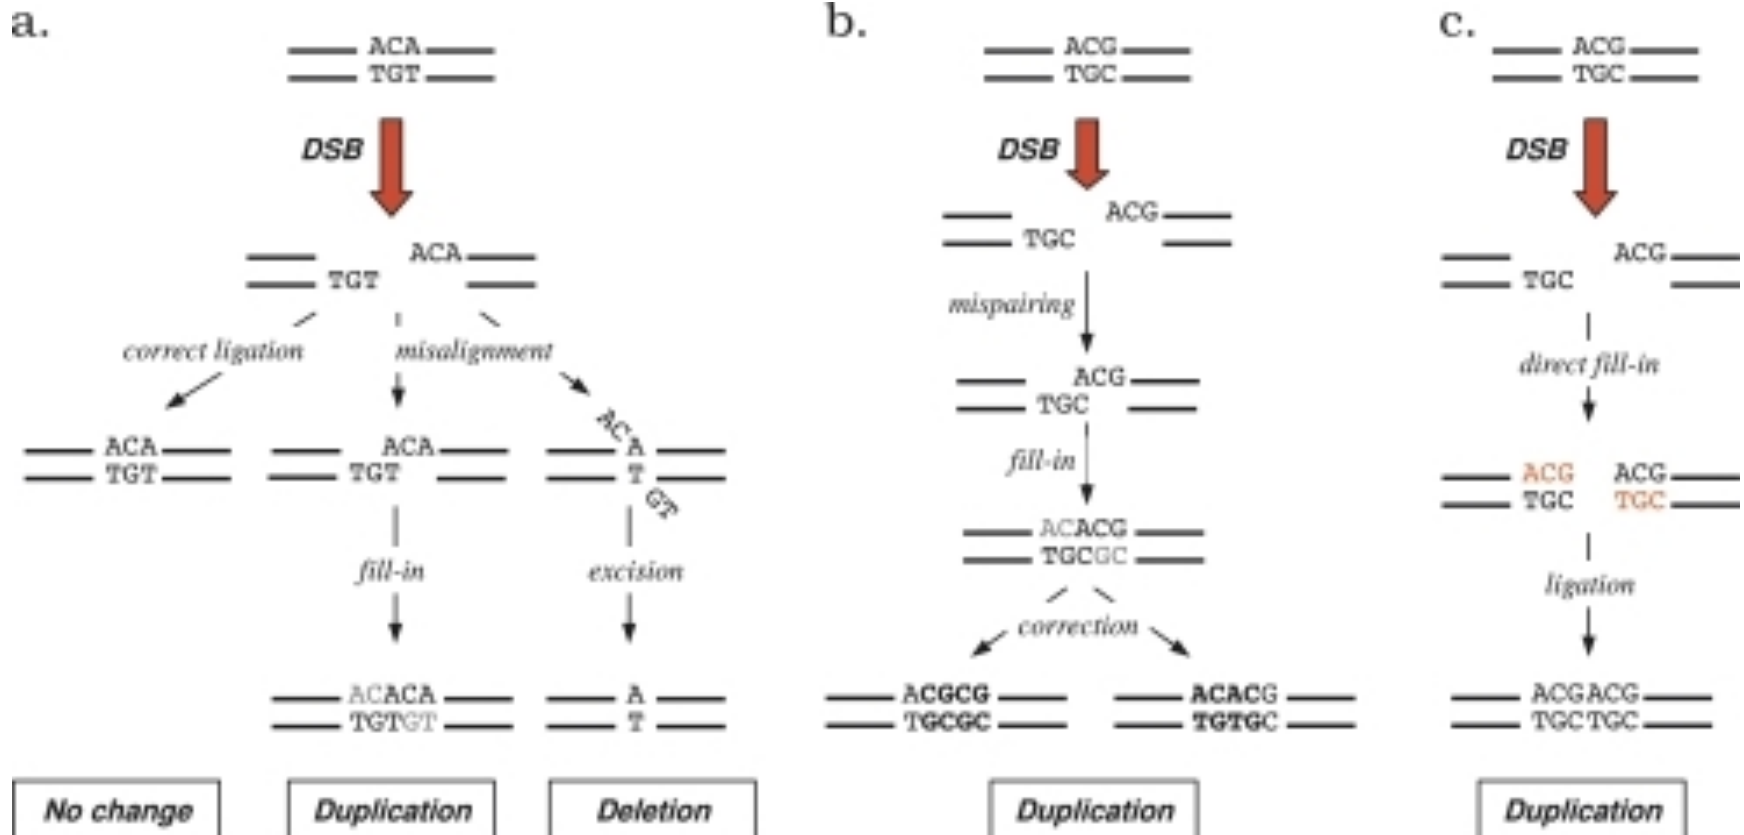
# Double-stranded break repair



Possible anti-recombinogenic role of Bloom's syndrome helicase in double-strand break processing. doi: [10.1093/nar/gkg834](https://doi.org/10.1093/nar/gkg834)




# NHEJ-derived indels



DNA Slippage Occurs at Microsatellite Loci without Minimal Threshold Length in Humans: A Comparative Genomic Approach. Leclercq S, Rivals E, Jarne P - Genome Biol Evol (2010)

# Structural variation (SV)



Transposable elements (in this case, an Alu) are sequences that can copy and paste themselves into genomic DNA, causing insertions.

Deletions can also be mediated by these sequences via other processes.

# Overview

1. Genesis of insertion/deletion (indel) polymorphism
2. **Standard approaches to detecting indels**
3. Assembly-based indel detection
4. Haplotype-based indel detection
5. Primary filtering: Bayesian variant calling
6. Post-call filtering: SVM
7. Graph-based resequencing approaches

# Calling INDEL variation


Can we quickly design a process to detect indels from alignment data?

What are the steps you'd do to find the indel between these two sequences?

CAAATAAGGTTTGGAGTTCATTATA  
CAAATAAGGTTTGGAAATTTTCTGGAGTTCATTATA

# Indel finder

We could start by finding the long matches in both sequences at the start and end:



# Indel finder

We can see this more easily like this:

CAAATAAGGTTTGGAGTTCTATTATA  
CAAATAAGGTTTGGAAATTTTCTGGAGTTCTATTATA

CAAATAAGGTTTGGAGTTCTATTATA  
CAAATAAGGTTTGGAAATTTTCTGGAGTTCTATTATA

CAAATAAGGTTTGGAGTTCTATTATA  
CAAATAAGGTTTGGAAATTTTCTGGAGTTCTATTATA

# Indel finder

The match structure implies that the sequence that doesn't match was inserted in one sequence, or lost from the other.

CAAATAAGGTT - - - - - TGGAGTTCTATTATA  
CAAATAAGGTTTGGAAATTTTCTGGAGTTCTATTATA

So that's easy enough....

# Something more complicated

These sequences are similar to the previous ones, but with different mutations between them.

CAAATAAGGAAATTTTCTGGAGTTCTATTATA  
CAAATAAGGTTTGTATCTAGGTTATTATA

They are still (kinda) homologous but it's not easy to see.



# Pairwise alignment

One solution, assuming a particular set of alignment parameters, has 3 indels and a SNP:

CAAATAAGGAAATTT - - - - TCTGGAGTTCTATTATA  
CAAATAAGG - - - TTTGCTATCT - - AGGT - TATTATA

But if we use a higher gap-open penalty, things look different:

CAAATAAGGAAATTT - - TCTGGAGTTCTATTATA  
CAAATAAGG - - - TTTGCTATCTAGGT - TATTATA

# Alignment as interpretation

Different parameterizations can yield different results.

Different results suggest “different” variation.

What kind of problems can this cause? (And how can we mitigate these issues?)

*First, let's review standard calling approaches.*


## Variation (VCF)

```
##source=mutatrix population genome simulator
##seed=1373972756
##reference=chrQ_fa
##phasing=true
##format=mutatrix-output -5 sample -p 2 -n 100 chrQ_fa
##filter="AC > 8"
##INFO=ID=Type,Number=A,Type=String,Descriptions="Type of each allele (snp, ins, del, mmp, complex)">
##INFO=ID=NA,Number=1,Type=Integer,Descriptions="Number of alternate alleles">
##INFO=ID=LEN,Number=A,Type=Integer,Descriptions="Length of each alternate allele">
##INFO=ID=HICROSAT,Number=0,Type=Flag,Descriptions="Generated at a sequence repeat loci">
##INFO=ID=IDGT,Number=1,Type=String,Descriptions="Genotype">
##INFO=ID=AC,Number=A,Type=Integer,Descriptions="Total number of alternate alleles in called genotypes">
##INFO=ID=AF,Number=A,Type=Float,Descriptions="Estimated allele frequency in the range (0,1]">
##INFO=ID=NS,Number=1,Type=Integer,Descriptions="Number of samples with data">
##INFO=ID=AN,Number=1,Type=Integer,Descriptions="Total number of alleles in called genotypes">
#CHROM POS ID REF ALT QUAL FILTER INFO FORMAT sample001 sample002
chrQ 12552 C T TC 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=ins
chrQ 6346 C T TC 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=ins
chrQ 7412 C T TC 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=ins
chrQ 7935 T C 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 8131 C C 99 AC2:AF=0.5;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 8606 AA TG 99 AC1:AF=0.25;AN=4;LEN=2;NA=1;NS=2;TYPE=mmp
chrQ 18926 T C 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 11921 G GTT 99 AC1:AF=0.25;AN=4;LEN=2;NA=1;NS=2;TYPE=ins
chrQ 12955 T G 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 13808 T TG 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=ins
chrQ 15271 A G 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=ins
chrQ 15467 A C 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 16486 C G 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 16563 T A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 16748 GTT G 99 AC1:AF=0.25;AN=4;LEN=2;NA=2;NS=2;TYPE=del
chrQ 17697 G C 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 18926 G A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 28750 G A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 21532 T C 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 22291 C T 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 23193 G A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 23954 CTAA TTA 99 AC1:AF=0.25;AN=4;LEN=4;NA=2;NS=2;TYPE=mmp
chrQ 24463 G A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 26189 G A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 29654 T A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 30862 T C 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 31790 A G 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 32792 T C 99 AC3:AF=0.75;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 33166 CC C 99 AC2:AF=0.5;AN=4;LEN=1;NA=1;NS=2;TYPE=del
chrQ 33403 T C 99 AC2:AF=0.5;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 33882 A G 99 AC2:AF=0.5;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 34450 C T 99 AC4:AF=1;AN=4;LEN=1;NA=1;NS=2;TYPE=snp GT
chrQ 34716 G A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 35484 G A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 36547 G A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 38015 T A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 38281 T C 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 40467 A G 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 48581 A G 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 48861 G T 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
chrQ 43268 G A 99 AC1:AF=0.25;AN=4;LEN=1;NA=1;NS=2;TYPE=snp
```

# Reads (FASTQ)

# Alignments to candidates

Reference




Reads

Variant observations

# The data exposed to the caller

Reference



Haplotype information is lost.

# INDELs have multiple representations and require normalization for standard calling

Left alignment allows us to ensure that our representation is consistent across alignments and also variant calls.

CGTATGATCTAGCGCGCTAGCTAGCTAGC  
CGTATGATCTA - - GCGCTAGCTAGCTAGC ← Left aligned

CGTATGATCTAGCGCGCTAGCTAGCTAGC  
CGTATGATCTAGC - - GCTAGCTAGCTAGC

CGTATGATCTAGCGCGCTAGCTAGCTAGC  
CGTATGATCTAGCGC - -TAGCTAGCTAGC

example: 1000G Phasel low coverage  
chr15:81551110, ref:CTCTC alt:ATATA

ref: TGTCACTCGCTCTCTCTCTCTCTCTATATATATATATTTGTGCAT  
alt: TGTCACTCGCTCTCTCTCTCTCTATATATATATATATATATTTGTGCAT

Interpreted as 3 SNPs

[illegible]

Interpreted as microsatellite expansion/contraction

example: 1000G Phasel low coverage  
chr20:708257, ref:AGC alt:CGA

ref: TATAGAGAGAGAGAGAGAGAGC GAGAGAGAGAGAGAGAGAGGGAGAGACGGAGTT  
alt: TATAGAGAGAGAGAGAGAGAGC GAGAGAGAGAGAGAGAGAGAGGGAGAGACGGAGTT

ref: TATAGAGAGAGAGAGAGAGAGC -- GAGAGAGAGAGAGAGAGAGGGAGAGACGGAGTT  
alt: TATAGAGAGAGAGAGAGAG -- CGAGAGAGAGAGAGAGAGAGGGAGAGACGGAGTT



# Overview

1. Genesis of insertion/deletion (indel) polymorphism
2. Standard approaches to detecting indels
3. **Assembly-based indel detection**
4. Haplotype-based indel detection
5. Primary filtering: Bayesian variant calling
6. Post-call filtering: SVM
7. Graph-based resequencing approaches

# ***Problem:* inconsistent indel representation makes alignment-based variant calling difficult**

If alleles are represented in multiple ways, then to detect them correctly with a single-position based approach we need:

1. An awesome normalization method
2. Perfectly consistent filtering (so we represent our entire context correctly in the calls)
3. Highly-accurate reads

# ***Solution: assembly and haplotype-driven detection***

We can shift our focus from the specific interpretation in the alignments:

- this is a SNP
- whereas this is a series of indels

... and instead focus on the underlying sequences.

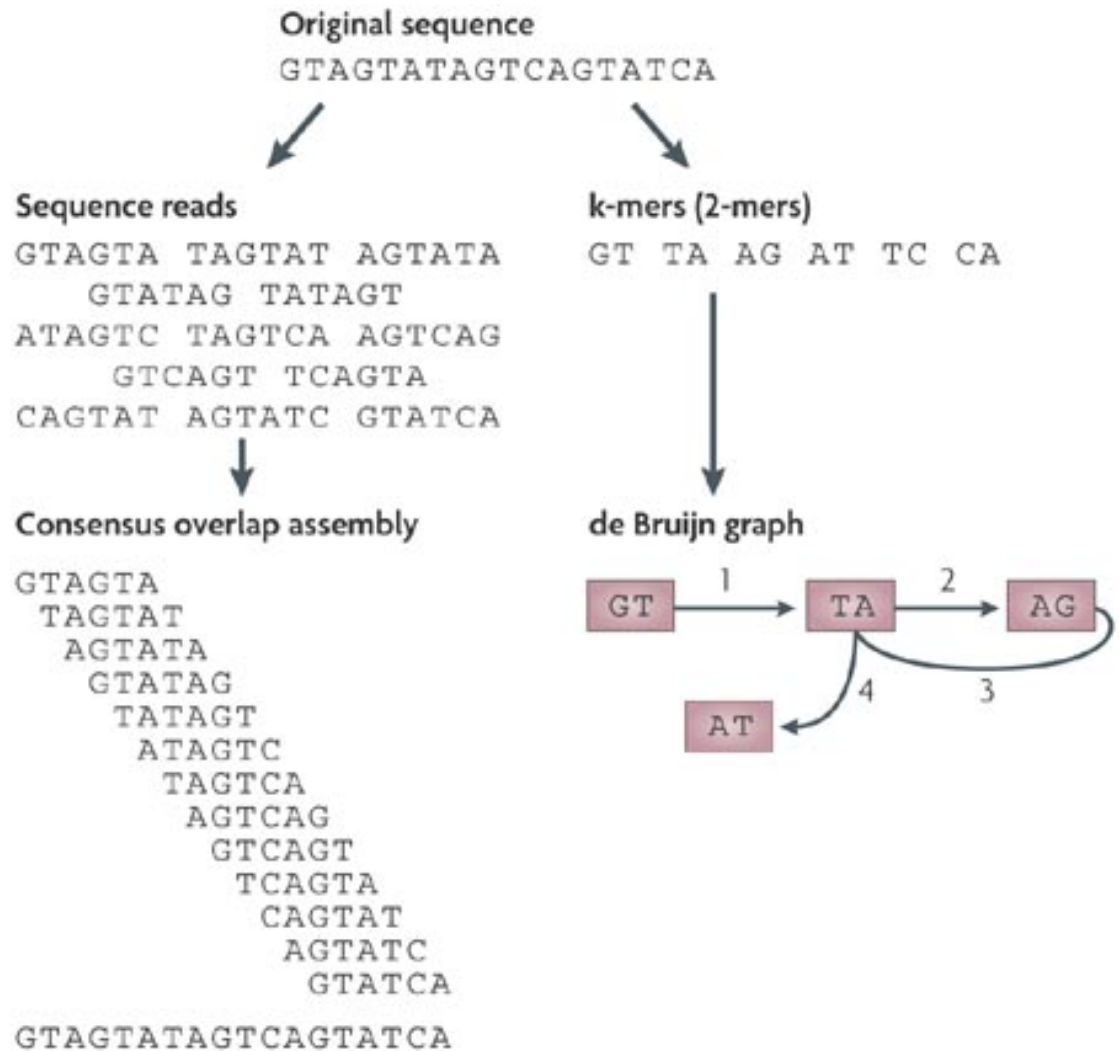
Basically, we use the alignments to localize reads, then process them again with assembly approaches to determine candidate alleles.

# Variant detection by assembly

Multiple methods have been developed by members of the 1000G analysis group:

- Global joint assembly
  - cortex
  - SGA (localized to 5 megabase chunks)
- Local assembly
  - Platypus (+cortex)
  - GATK HaplotypeCaller
- k-mer based detection
  - FreeBayes (anchored reference-free windows)

# Assembly

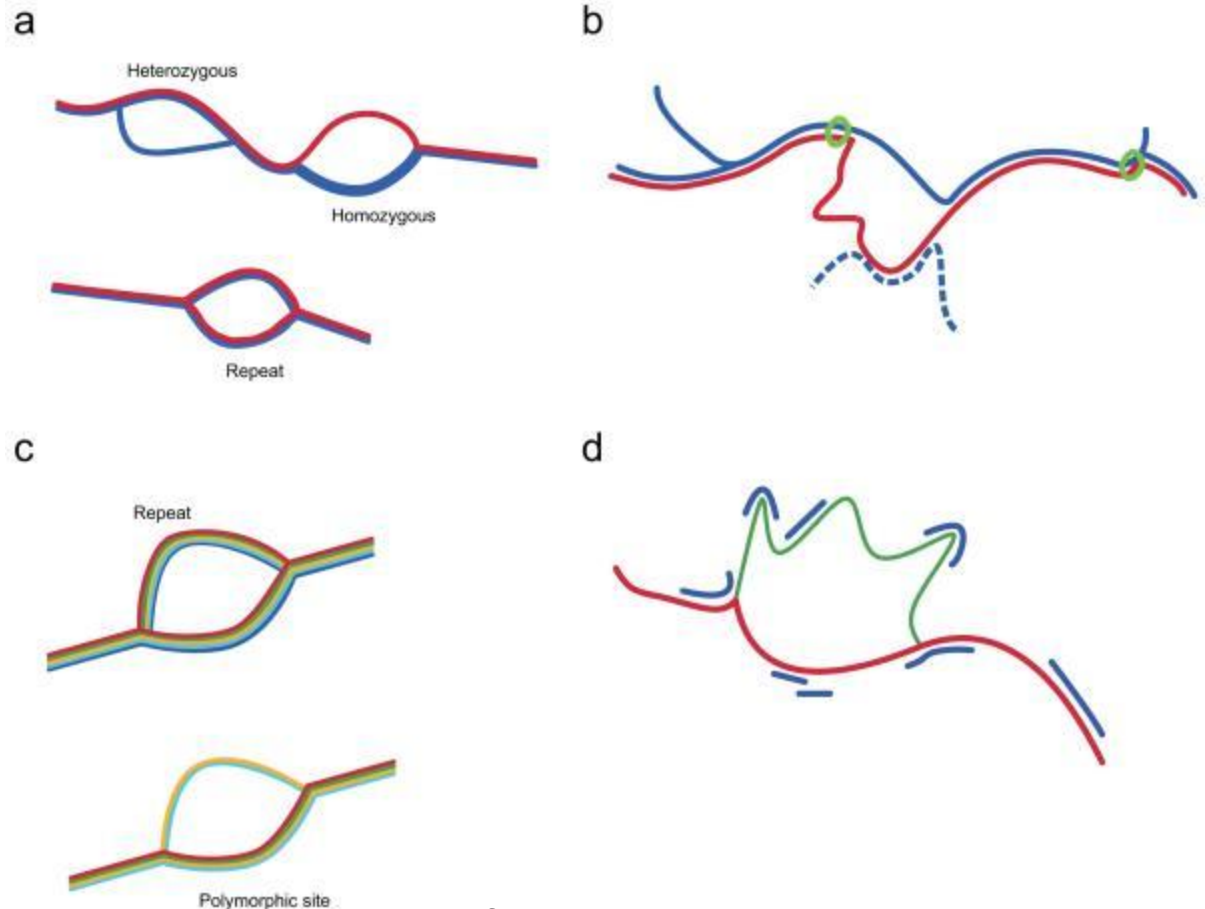


# Using colored graphs (Cortex)

Variants can be called using bubbles in deBruijn graphs.

Method is completely reference-free, except for reporting of variants. The reference is threaded through the colored graph.

Many samples can be called at the same time.




from Iqbal et. al., "De novo assembly and genotyping of variants using colored de Bruijn graphs." (2012)

# String graphs (SGA)

## Genome



## String Graph




A string graph has the advantage of using less memory to represent an assembly than a de Bruijn graph. In the 1000G, SGA is run on alignments localized to ~5mb chunks.

# Discovering alleles using graphs (GATK HaplotypeCaller)

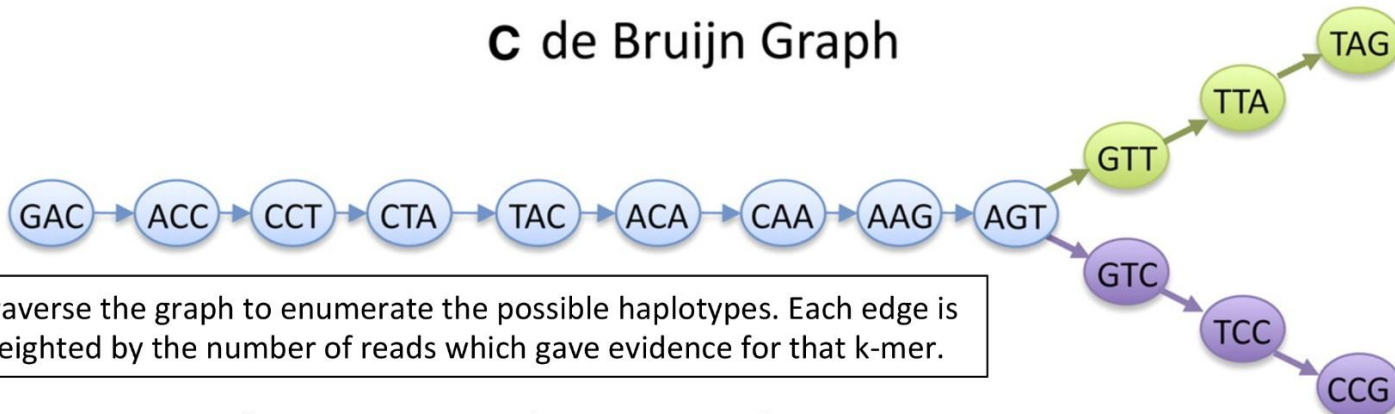
**A Read Layout**

R<sub>1</sub>: GACCTACA  
R<sub>2</sub>: ACCTACAA  
R<sub>3</sub>: CCTACAAG  
R<sub>4</sub>: CTACAAGT  
A: TACAAGTT  
B: ACAAGTTA  
C: CAAGTTAG  
X: TACAAGTC  
Y: ACAAGTCC  
Z: CAAGTCCG

**B Overlap Graph**



**C de Bruijn Graph**

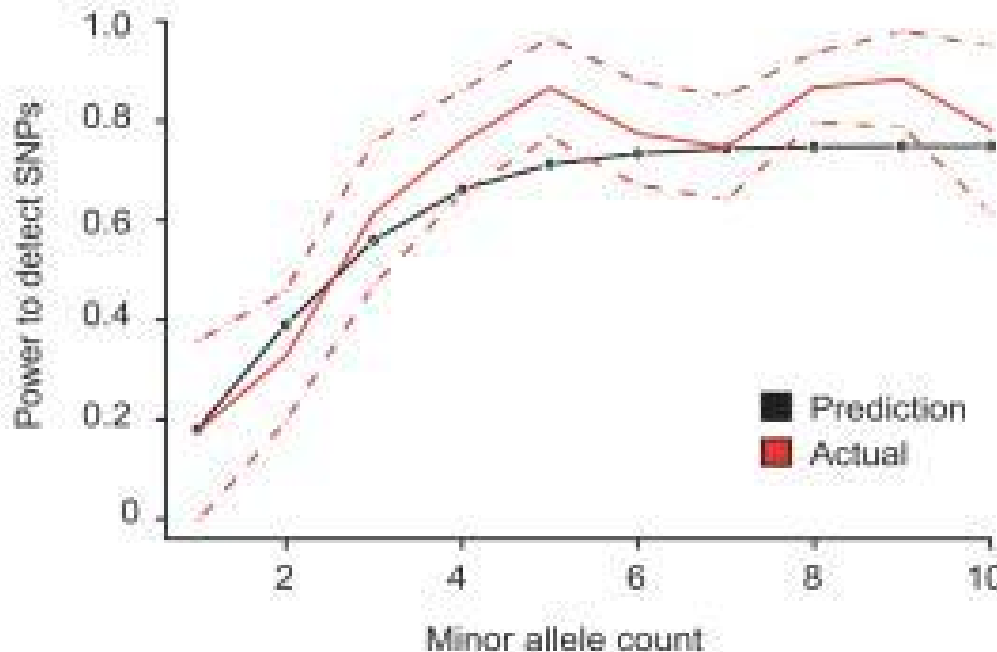


Traverse the graph to enumerate the possible haplotypes. Each edge is weighted by the number of reads which gave evidence for that k-mer.



# Why don't we just assemble?

Assembly-based calls tend to have high specificity, but sensitivity suffers.



The requirement of exact kmer matches means that errors disrupt coverage of alleles.


Existing assembly methods don't just detect point mutations--- they detect haplotypes.

from Iqbal et. al., "De novo assembly and genotyping of variants using colored de Bruijn graphs." (2012)

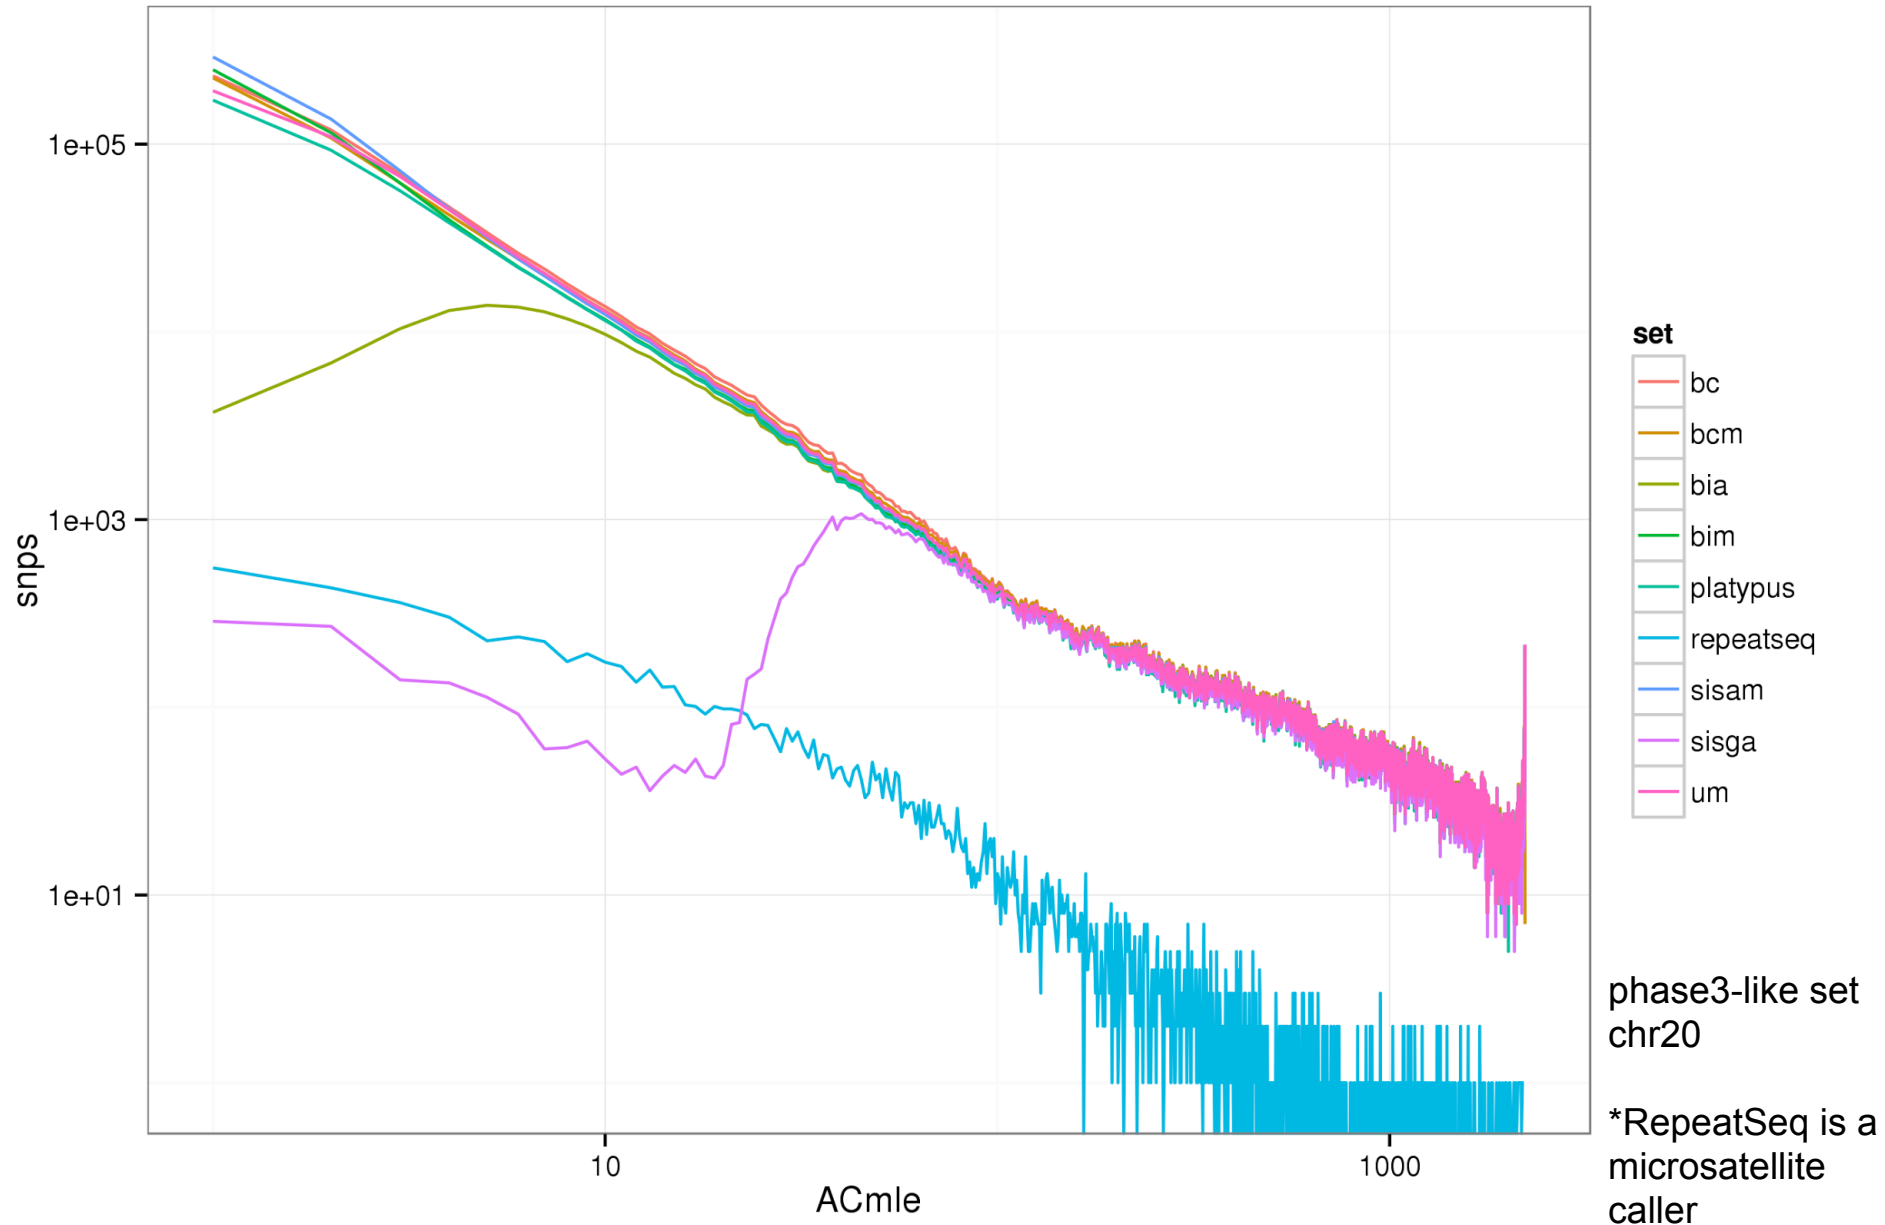
# Indel validation, 191 AFR samples

High-depth miSeq sequencing-based validation on 4 samples. Local assembly methods (BI2, BC, SI1)\* have higher specificity than baseline mapping-based calls (BI1), but lower sensitivity. Global assembly (OX2) yielded very low error, but also low sensitivity.


\*The local assembly-based method Platypus (OX1) had a genotyping bug which caused poor performance.



# Site-frequency spectrum, SNPs



# Site-frequency spectrum, indels



# Overview

1. Genesis of insertion/deletion (indel) polymorphism
2. Standard approaches to detecting indels
3. Assembly-based indel detection
- 4. Haplotype-based indel detection**
5. Primary filtering: Bayesian variant calling
6. Post-call filtering: SVM
7. Graph-based resequencing approaches

# Finding haplotype polymorphisms

Two  
reads

AGAACCCAGTGCTCTTTCTGCT  
AGAACCCAGTGTCTTTCTGCT

a SNP

AGAACCCAGTGCTCTTTCTGCT  
AGAACCCAGTGTCTTTCTGCT

Their  
alignmen  
t

Another read  
showing a SNP  
on the same  
haplotype as the  
first

AGAACCCAGTGCTCTATCTGCT

AGAACCCAGTGCTCTATCTGCT  
AGAACCCAGTGTCTTTCTGCT

A variant  
locus implied  
by  
alignments

# Direct detection of haplotypes



Ref  
Reads

	Variant Region	Variant Region
TACCGAT	CATTGGATCA	CGATTCC...GCATTGC
TACCGAT	CATTGGATCA	CGATTCC...GCATTGC
ACCGAT	TATTGCATCG	CGATTCC...GCATTGC
ACCGAT	CATTGGATCA	CGATTCC...GCATTGC
ACCGAT	TATTGGATCG	CGATTCC...GCATTGC
CCGAT	C-TTGGATCA	CGATTCC...GCATTGC
CCGAT	CATGGGATCA	CGATTCC...GCATTGC
...	...	...
		Variant Region
		AAAAAAAA-
		-AAAAAA-
		-AAAAAA-
		AAAAAA-A
		-AAAAAAA
		AAAAAAA-
		AAAAAAA A
		...

Observed Haplotypes

CATTGGATCA	x8
TATTGGATCG	x9
CTTGGATCA	x1
CATGGGATCA	x1
...	


(A) <sub>7</sub>	x10
(A) <sub>6</sub>	x7
(A) <sub>5</sub>	x1
(A) <sub>8</sub>	x1
...	



# Why haplotypes?

- Variants cluster.
- This has functional significance.
- Observing haplotypes lets us be more certain of the local structure of the genome.
- We can improve the detection process itself by using haplotypes rather than point mutations.
- We get the sensitivity of alignment-based approaches with the specificity of assembly-based ones.

# Sequence variants cluster



In ~1000 individuals,  $\frac{1}{2}$  of variants are within ~22bp of another variant.

Variance to mean ratio (VMR) = 1.4.

# The functional effect of variants depends on other nearby variants on the same haplotype

reference:   AGG   GAG   CTG  
              Arg   Glu   Leu

*OTOF* gene – mutations  
cause profound recessive  
deafness

apparent:   AGG   TAG   CTG  
              Arg   Ter   ---


Apparent nonsense variant,  
one YRI homozygote

actual:     AGG   TTG   CTG  
              Arg   Leu   Leu

Actually a block substitution  
that results in a missense  
substitution

(Daniel MacArthur)

# Importance of haplotype effects: frame-restoring indels



- Two apparent frameshift deletions in the *CASP8AP2* gene (one 17 bp, one 1 bp) on the same haplotype
- Overall effect is in-frame deletion of six amino acids

(Daniel MacArthur)

# Frame-restoring indels in 1000 Genomes Phase I exomes

chr6:117113761, GPRC6A (~10% AF in 1000G)


ref: A T T G T A A T T C T C A -- T A -- T T -- T G C C T T T G A A A G C  
alt: A T T G T A A T T C T C A G G T A A T T T C C T G C C T T T G A A A G C

chr6:32551935, HLA-DRB1 (~11% AF in 1000G)

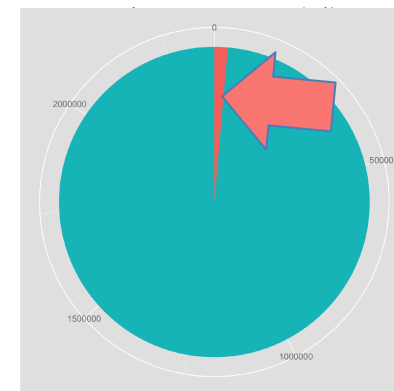

ref: C C A C C G C G G C C C G C G C C T G - C - T C C A G G A T G T C C  
Alt: C C A C C G C G G -- C G C G C C T G T C T T C C A G G A G G T C C

# Impact on genotyping chip design

- Biallelic SNPs detected during the 1000 Genomes Pilot project were used to design a genotyping microarray (Omni 2.5).
- When the 1000 Genomes samples were genotyped using the chip, 100k of the 2.5 million loci showed no polymorphism (monomorphs).





# Omni 2.5 1000G monomorphic loci, BC haplotype calls



## CLASS

- biallelic complex
- biallelic INDEL
- biallelic MNP
- biallelic SNP
- multiallelic complex
- multiallelic INDEL
- multiallelic INDEL, SNP, and MNP
- multiallelic mixed
- multiallelic SNP
- multiallelic SNP and MNP
- multiallelic SNP, MNP, and complex

# Omni 2.5 1000G polymorphic loci, BC haplotype calls

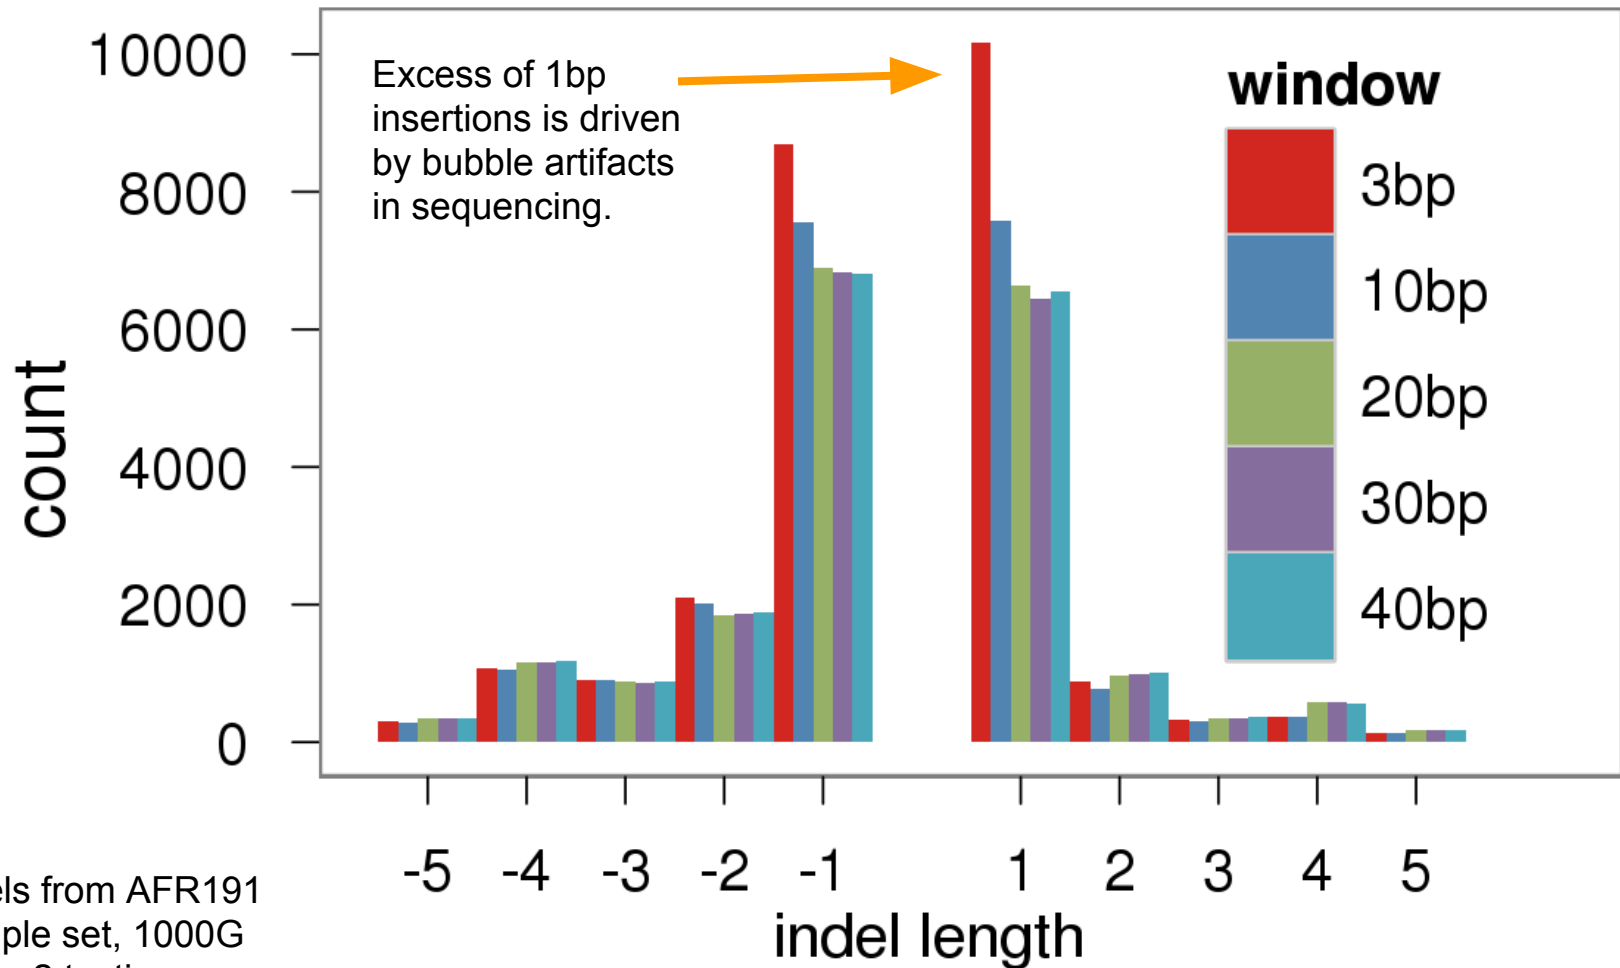


## CLASS

- biallelic complex
- biallelic INDEL
- biallelic MNP
- biallelic SNP
- multiallelic complex
- multiallelic INDEL
- multiallelic INDEL, SNP, and MNP
- multiallelic mixed
- multiallelic SNP
- multiallelic SNP and MNP
- multiallelic SNP, MNP, and complex




# Measuring haplotypes improves specificity




Indels from AFR191  
sample set, 1000G  
phase2 testing.

# 2bp MNPs and dinucleotide intermediates



# Direct detection of haplotypes can remove directional bias associated with alignment-based detection



# Overview

1. Genesis of insertion/deletion (indel) polymorphism
2. Standard approaches to detecting indels
3. Assembly-based indel detection
4. Haplotype-based indel detection
- 5. Primary filtering: Bayesian variant calling**
6. Post-call filtering: SVM
7. Graph-based resequencing approaches

# Filtering INDELs

As with SNPs, sequencing error rates are high.


So, we need to filter.

The standard filter of NGS is the Bayesian variant caller.


Combines population-based priors and data from many samples to make high-quality calls.

# Bayesian (visual) intuition

We have a universe of individuals.



A = samples with a  
variant at some locus



B = putative observations  
of variant at some locus

# probability(A|B)


We want to estimate the probability that we have a real polymorphism "A" given "|" that we observed variants in our alignments "B".

$$P(A|B) = \frac{|AB|}{|B|}$$



$$P(A|B) = \frac{P(AB)}{P(B)}$$

$$P(B|A) = \frac{P(AB)}{P(A)}$$

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$



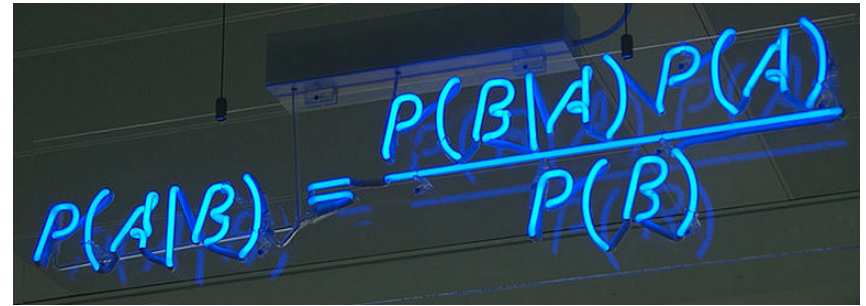
# In our case it's a bit more like this...



Observations (B) provide pretty good sensitivity, but poor specificity.



# The model



A photograph of a chalkboard with the formula for Bayes' theorem written in blue chalk. The formula is  $P(A|B) = \frac{P(B|A)P(A)}{P(B)}$ . The chalkboard has a dark background and the text is written in a clear, legible hand.

- Bayesian model estimates the probability of polymorphism at a locus given input data and the population mutation rate (~pairwise heterozygosity) and assumption of “neutrality” (random mating).
- Following Bayes theorem, the probability of a specific set of genotypes over some number of samples is:
  - **$P(G|R) = ( P(R|G) P(G) ) / P(R)$**
- Which in FreeBayes we extend to:
  - **$P(G,S|R) = ( P(R|G,S) P(G)P(S) ) / P(R)$**
  - **$G$**  = genotypes,  **$R$**  = reads,  **$S$**  = locus is well-characterized/mapped
  - **$P(R|G,S)$**  is our data likelihood,  **$P(G)$**  is our prior estimate of the genotypes,  **$P(S)$**  is our prior estimate of the mappability of the locus,  **$P(R)$**  is a normalizer.

# Handling non-biallelic/diploid cases


We compose our data likelihoods, **P(Reads|Genotype)** using a discrete multinomial sampling probability:

$$P(\text{reads}|\text{genotype}) = \binom{|\text{reads}|}{|\text{reads} = A|, |\text{reads} = B| \dots} \\ \times \prod_{\forall \text{alleles} \in \text{genotype}} \text{freq}(\text{allele} \in \text{genotype}) \\ \times \prod_{\forall \text{reads}} P(\text{correct}(\text{read}))$$

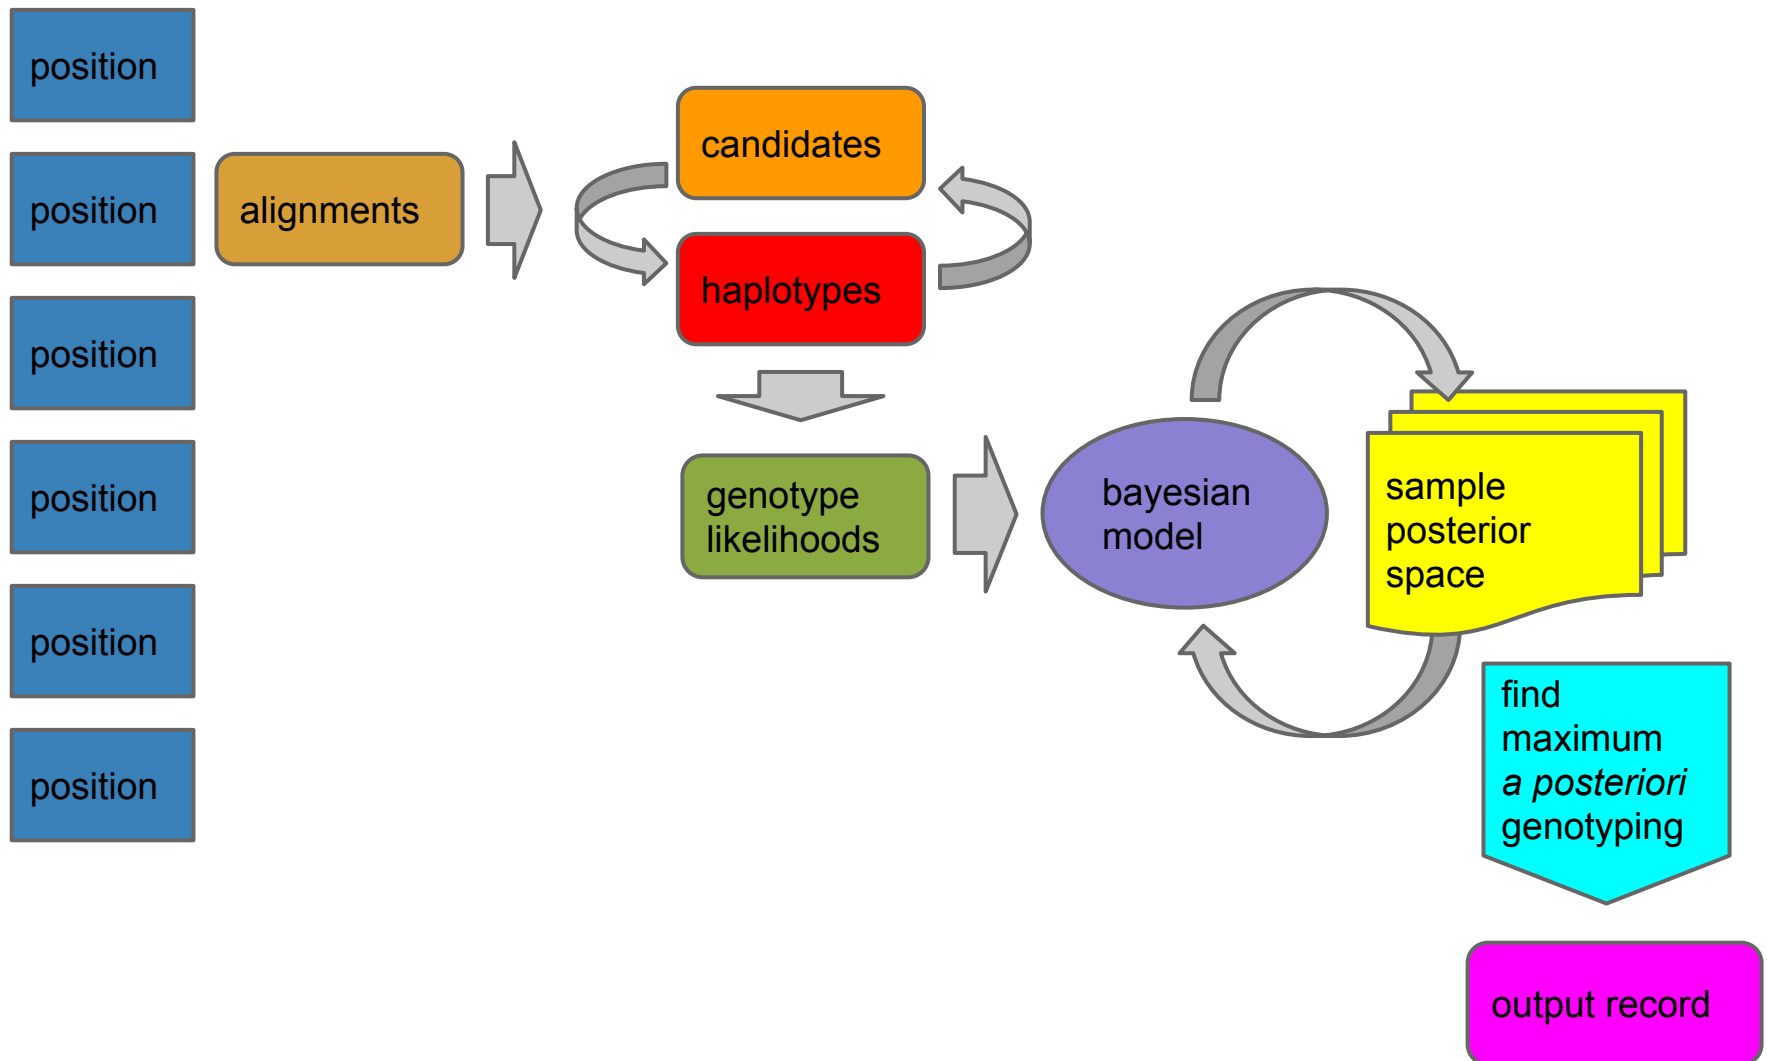
Our priors, **P(Genotypes)**, follow the Ewens Sampling Formula and the discrete sampling probability for genotypes.

# Are our locus and alleles sequenceable?

In WGS, biases in the way we observe an allele (placement, position, strand, cycle, or balance in heterozygotes) are often correlated with error. We include this in our posterior  $\mathbf{P}(\mathbf{G}, \mathbf{S} | \mathbf{R})$ , and to do so we need an estimator of  $\mathbf{P}(\mathbf{S})$ .



# The detection process



# Overview

1. Genesis of insertion/deletion (indel) polymorphism
2. Standard approaches to detecting indels
3. Assembly-based indel detection
4. Haplotype-based indel detection
5. Primary filtering: Bayesian variant calling
- 6. Post-call filtering: SVM**
7. Graph-based resequencing approaches

# SVM filtering

INDEL detection is hard.


*A priori* models can't capture all types of error.

It's especially difficult when we try to make a consensus set from lots of input variant callers.

We can use classifiers like Support Vector Machines (SVM) to further improve results.


# SVM classifier

Find a hyperplane (here a line in 2D) which separates observations.



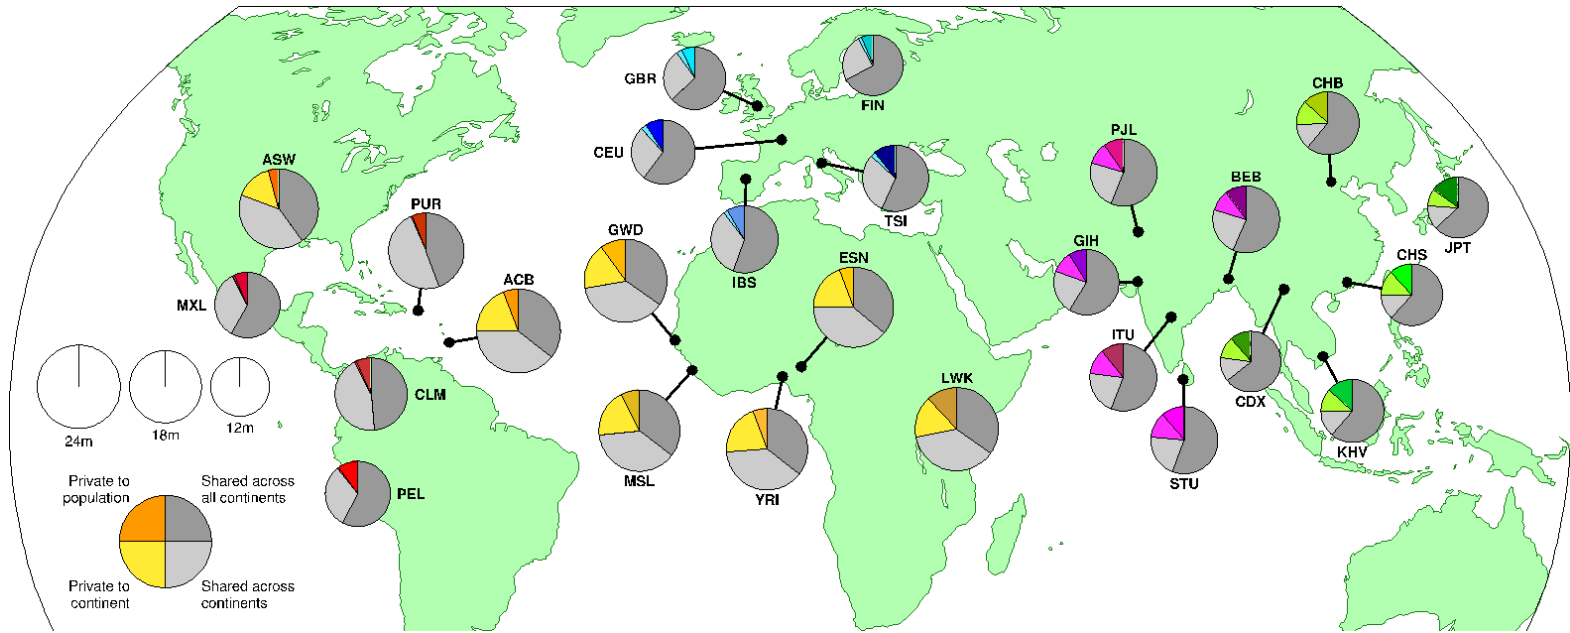
# SVM classifier

The best separating hyperplane is determined by maximum margin between groups we want to classify.



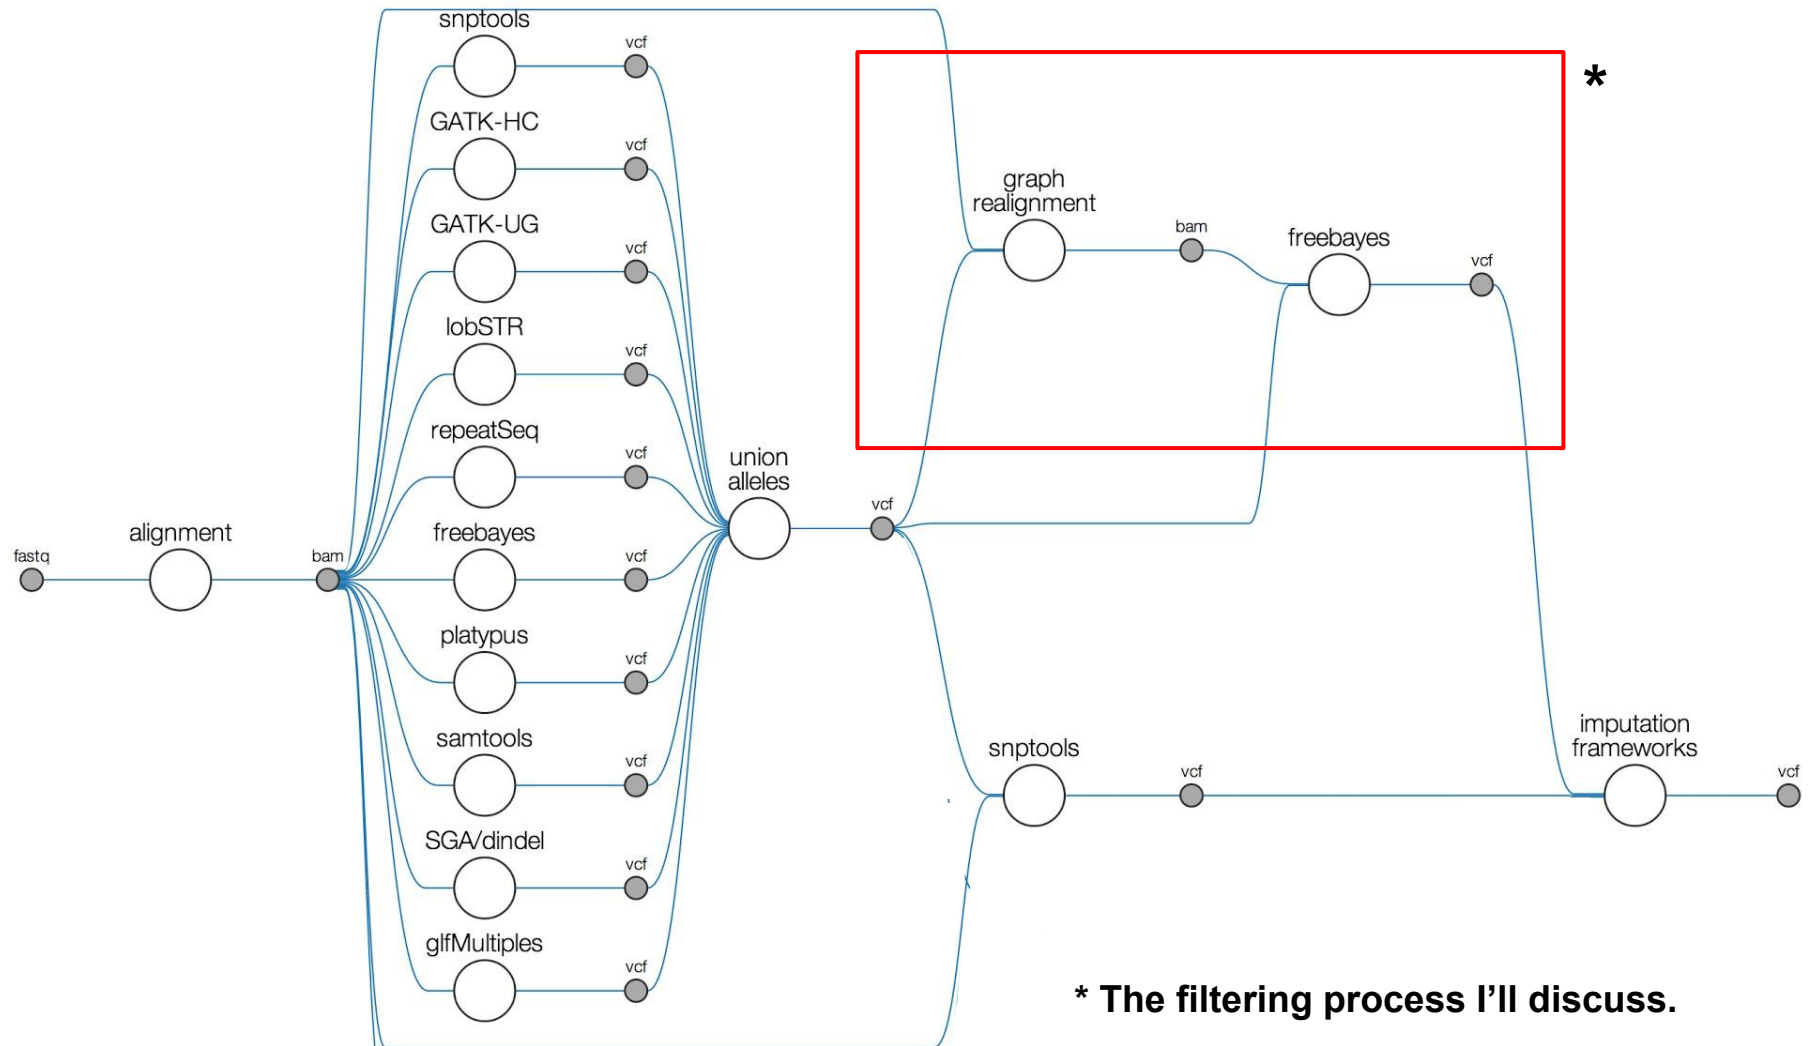


# SVM filtering in the 1000 Genomes



25 human populations X ~100 samples each.

# 1000G variant integration process



# SVM approach for INDEL filtering

Extract features that tend to vary with respect to call quality:

- call QUALity
- read depth
- sum of base qualities
- inbreeding coefficient (heterozygosity)
- entropy of sequence at locus
- mapping quality
- allele frequency in population
- read pairing rate
- etc.



# **SVM approach for INDEL filtering**

Now, use overlaps in validation samples or sites to determine likely errors and true calls.

Use this list + annotations of the calls to train an SVM model.

Apply the model to all the calls, filter, and measure validation rate of the whole set.

# Application of SVM to 1000G INDELs




Raw validation rates of indels in 1000G phase 3, “MVNCall” set.


*Tony Marcketta and Adam Auton*

# Application of SVM to 1000G INDELs

## Passing SVM



## Failing SVM




Filtering results, using SVM-based method.

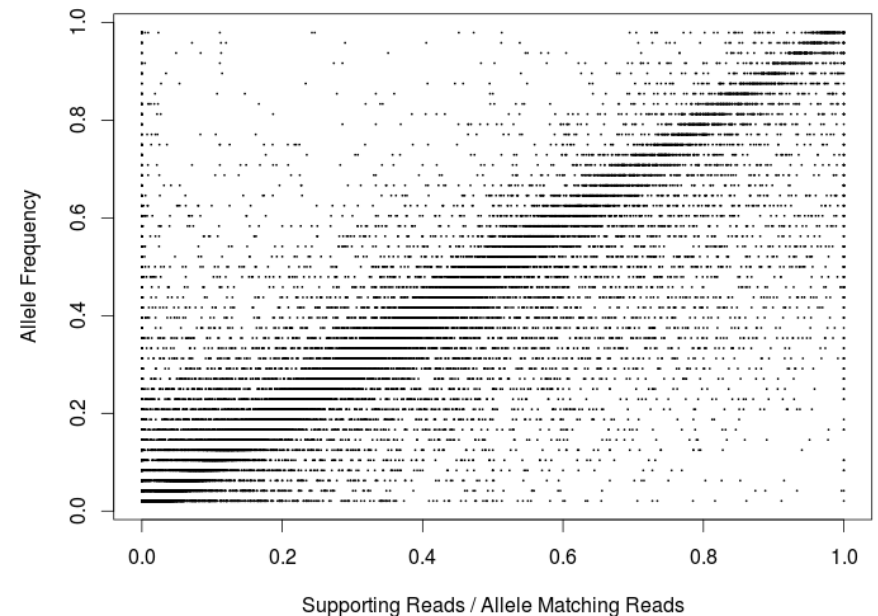
*Anthony Marcketta and Adam Auton*

# Application of SVM to 1000G INDELs

## Passing SVM



## Failing SVM



Correlation between allele frequency and observation counts.

*Anthony Marcketta and Adam Auton*


# Indel results from 1000G

Gold	Eval	R/R	R/A	A/A	All	NonRef
CGI SNPs	Phase3	0.9998	0.9930	0.9983	0.9994	0.9920
CGI Indels	Phase3	0.9990	0.9889	0.9923	0.9982	0.9805

Comparing the phase3 results to the genotypes for indels in the subset of samples for which we also had high-quality, high-coverage genomes from Complete Genomics.



# Genotype Accuracy by Allele Frequency



# Overview

1. Genesis of insertion/deletion (indel) polymorphism
2. Standard approaches to detecting indels
3. Assembly-based indel detection
4. Haplotype-based indel detection
5. Primary filtering: Bayesian variant calling
6. Post-call filtering: SVM
7. **Graph-based resequencing approaches**

**We know the variants,  
so why not use them in our analysis?**


We resequence new genomes and compare them to a single reference haplotype.

To determine anything more than short variants, we must do everything *de novo*.

*If we could merge sequence and variation, we could detect known alleles of arbitrary scale and divergence with minimal cost.*

# Pan-genomes as graphs


We can combine sequence and variation using  
*a variant graphs, or graph reference.*




\*This representation is directed (5' to 3'), and acyclic.

# Building the variant graph

#CHROM	POS	ID	REF	ALT
20	14370	.	T	TCGAA




#CHROM	POS	ID	REF	ALT
20	14371	.	A	C



#CHROM	POS	ID	REF	ALT
20	14370	.	TCGAA	T




# Local alignment against the graph



Christopher Lee, Catherine Grasso, Mark F. Sharlow. **Multiple sequence alignment using partial order graphs.** *Bioinformatics*, 2002.

# Local alignment against the graph

(A)



(B)


		T	T	G	G	A	T
	0	0	0	0	0	0	0
A	0	0	0	0	0	10	0
T	0	10	10	0	0	0	20
C	0	0	0	0	0	0	10
G	0	0	0	10	10	0	0
A	0	0	0	0	0	20	10
A	0	0	0	0	0	10	10

		C	G	A	A	T	T
	0	0	0	0	0	0	0
A	0	0	0	10	10	0	0
T	20	10	0	0	0	20	10
C	10	30	20	10	0	10	10
G	0	20	40	30	20	10	0
A	10	10	30	50	40	30	20
A	10	0	20	40	60	50	40

		A	T	G	G	G
	0	0	0	0	0	0
A	0	10	0	0	0	0
T	20	10	20	10	0	0
C	10	10	10	10	0	0
G	0	0	0	20	20	10
A	20	10	0	10	10	10
A	40	30	20	10	0	0

# “Striped” string/DAG alignment

We improved performance of our aligner >10-fold by generalizing Farrar’s striped Smith-Waterman algorithm to DAGs. **GSSW**




Data dependencies across DAG are limited to H and E vectors.

\*Implemented using SSE2 instruction set.

$$E_{i,j} = \max \begin{cases} E_{i,j-1} - G_{\text{ext}} \\ H_{i,j-1} - G_{\text{init}} \end{cases}$$


$$F_{i,j} = \max \begin{cases} F_{i-1,j} - G_{\text{ext}} \\ H_{i-1,j} - G_{\text{init}} \end{cases}$$

$$H_{i,j} = \max \begin{cases} 0 \\ E_{i,j} \\ F_{i,j} \\ H_{i-1,j-1} - W(q_i, d_j) \end{cases}$$






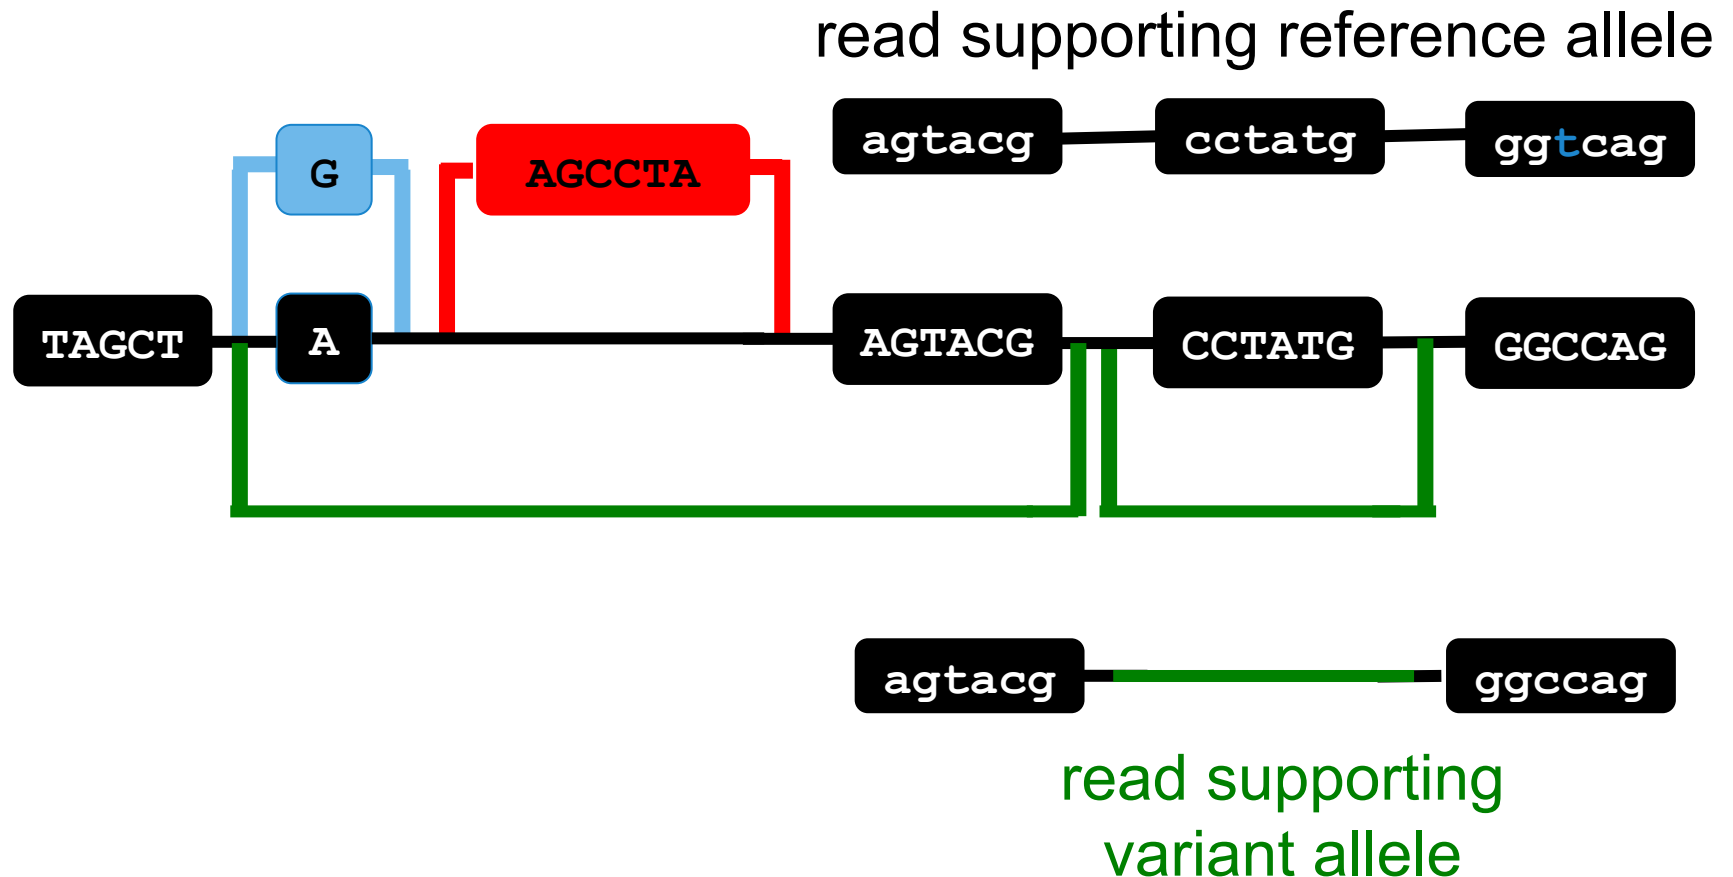
# Seeding graph-based alignments



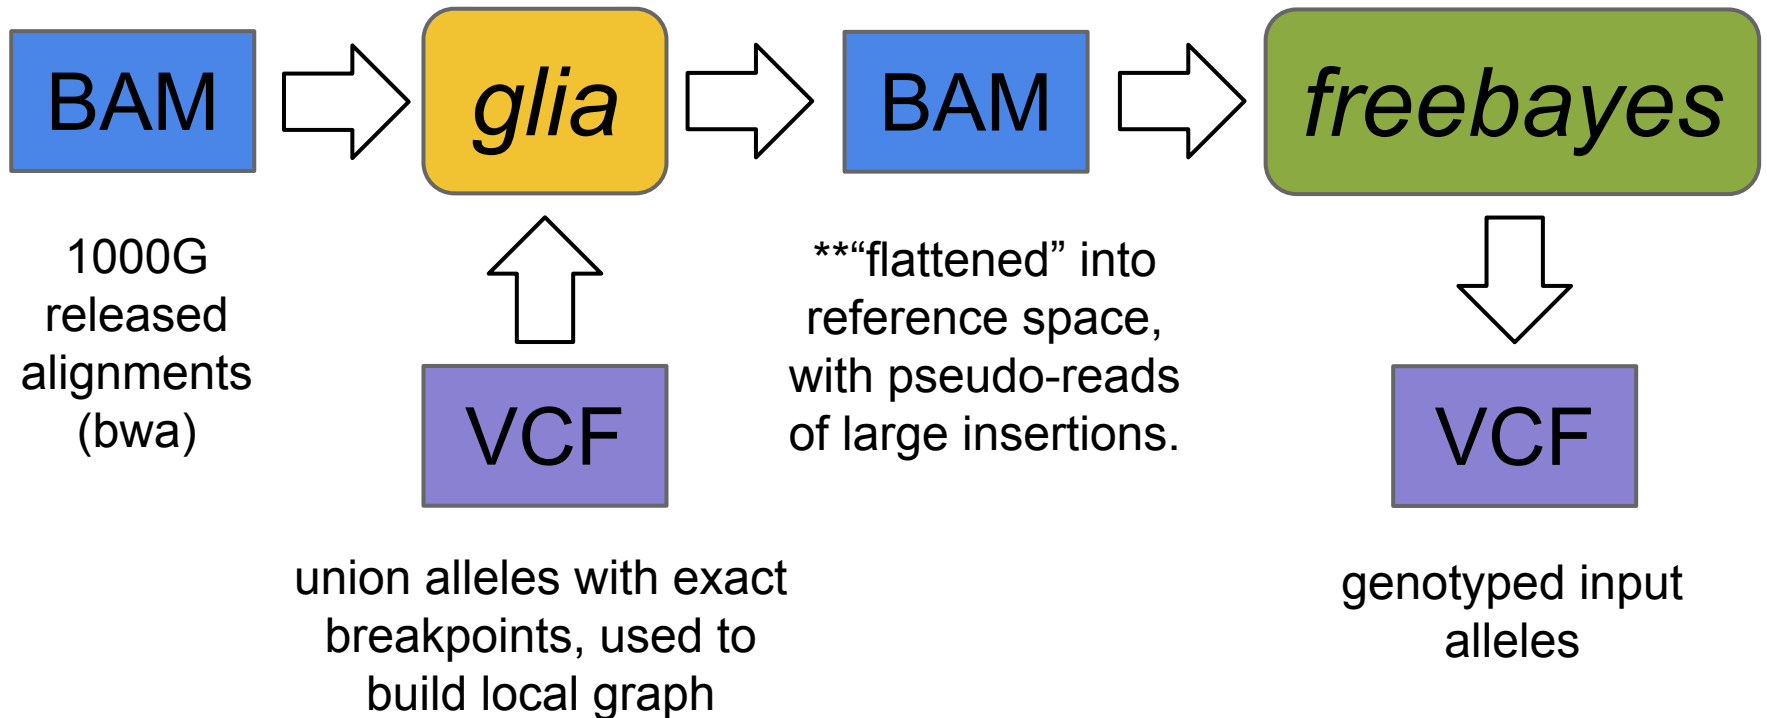
**Test imperfectly-mapped reads against graph.**



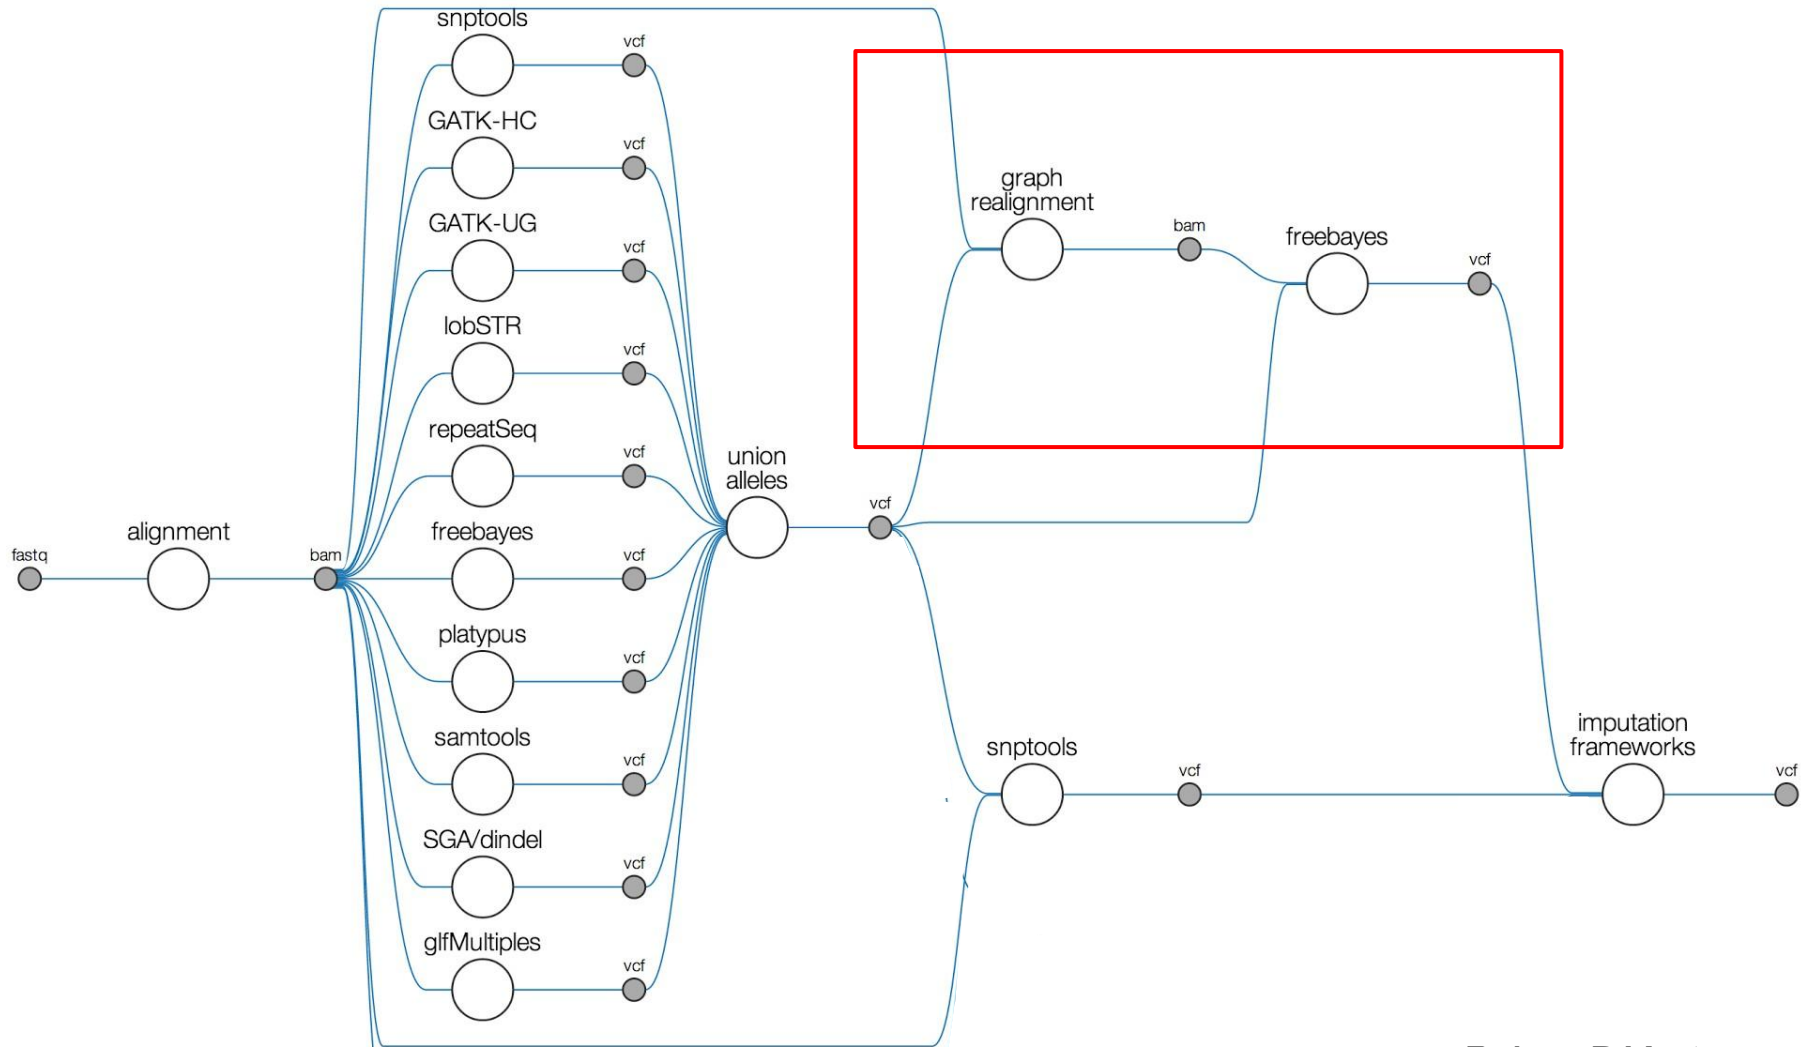
# Detecting variation on the graph




# Graph-based alignments with *glia*



# Application to 1000G variant integration



# Unifying calls from many methods

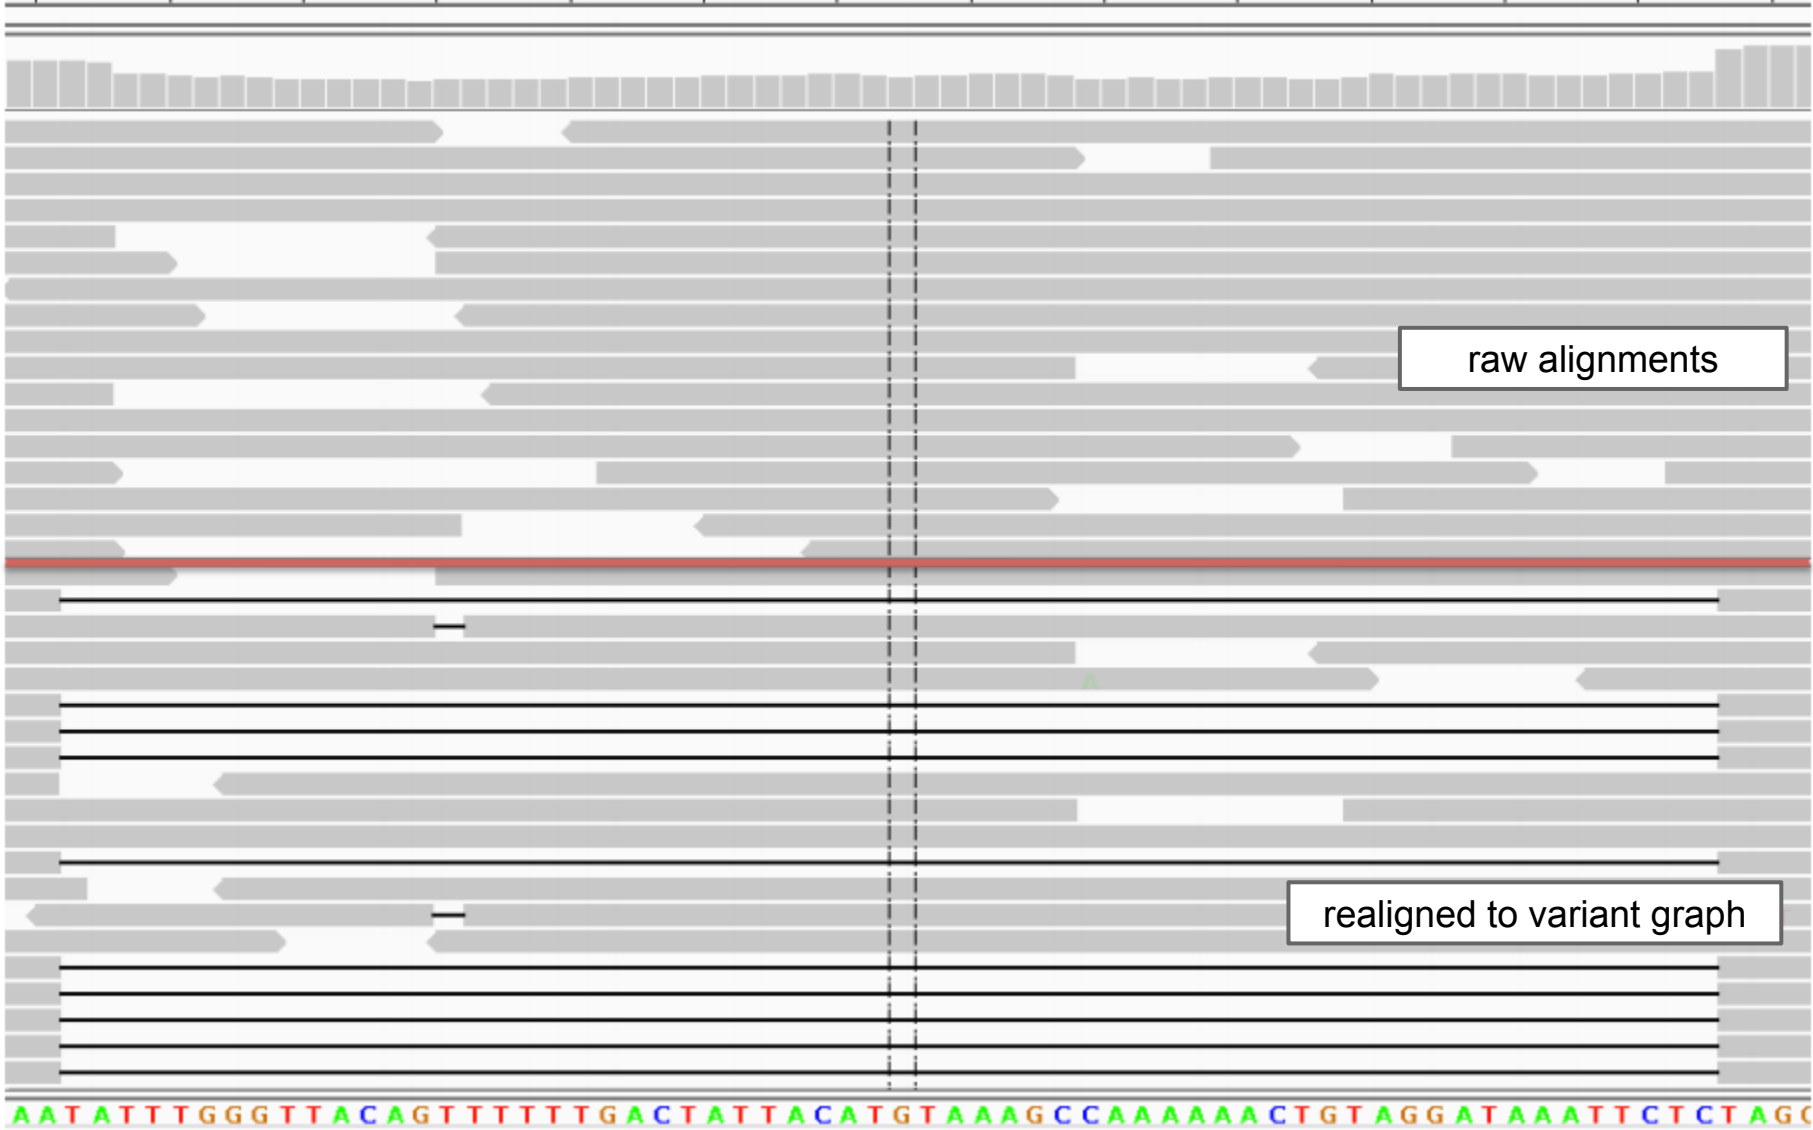


TATTTGGGTTACAGTTTTTTGACTATTACATGTAAAGCCAAAAAACTGTAGGATAAATTCTC


ACCCTTGAAGAA

TAGGATAAATG

TATTTGGGTTACAGTTTTTTGACTATTACATGTAAAGCCAAAAAACTGTAGGATAAATTCTC




# *glia* reduces reference bias



# *glia* reduces reference bias

Improvement in observation support



**Standard alignment is frustrated even by small variants!**

**Ratio between observations before and after  
realignment to graph of union variants**



# Improving genotype likelihoods

$$\text{Genotype Likelihood} = P(\text{data}|\text{genotype})$$

SET	GRP	N	RR	RA	AA	ALT	ALL
SVM indels	UM	6743	0.285	1.008	2.947	1.698	0.561
SVM indels	BC*	6743	<b>0.034</b>	<b>0.673</b>	<b>0.245</b>	<b>0.521</b>	<b>0.129</b>
SNPs	BCM	404270	0.029	1.373	0.445	1.093	0.111

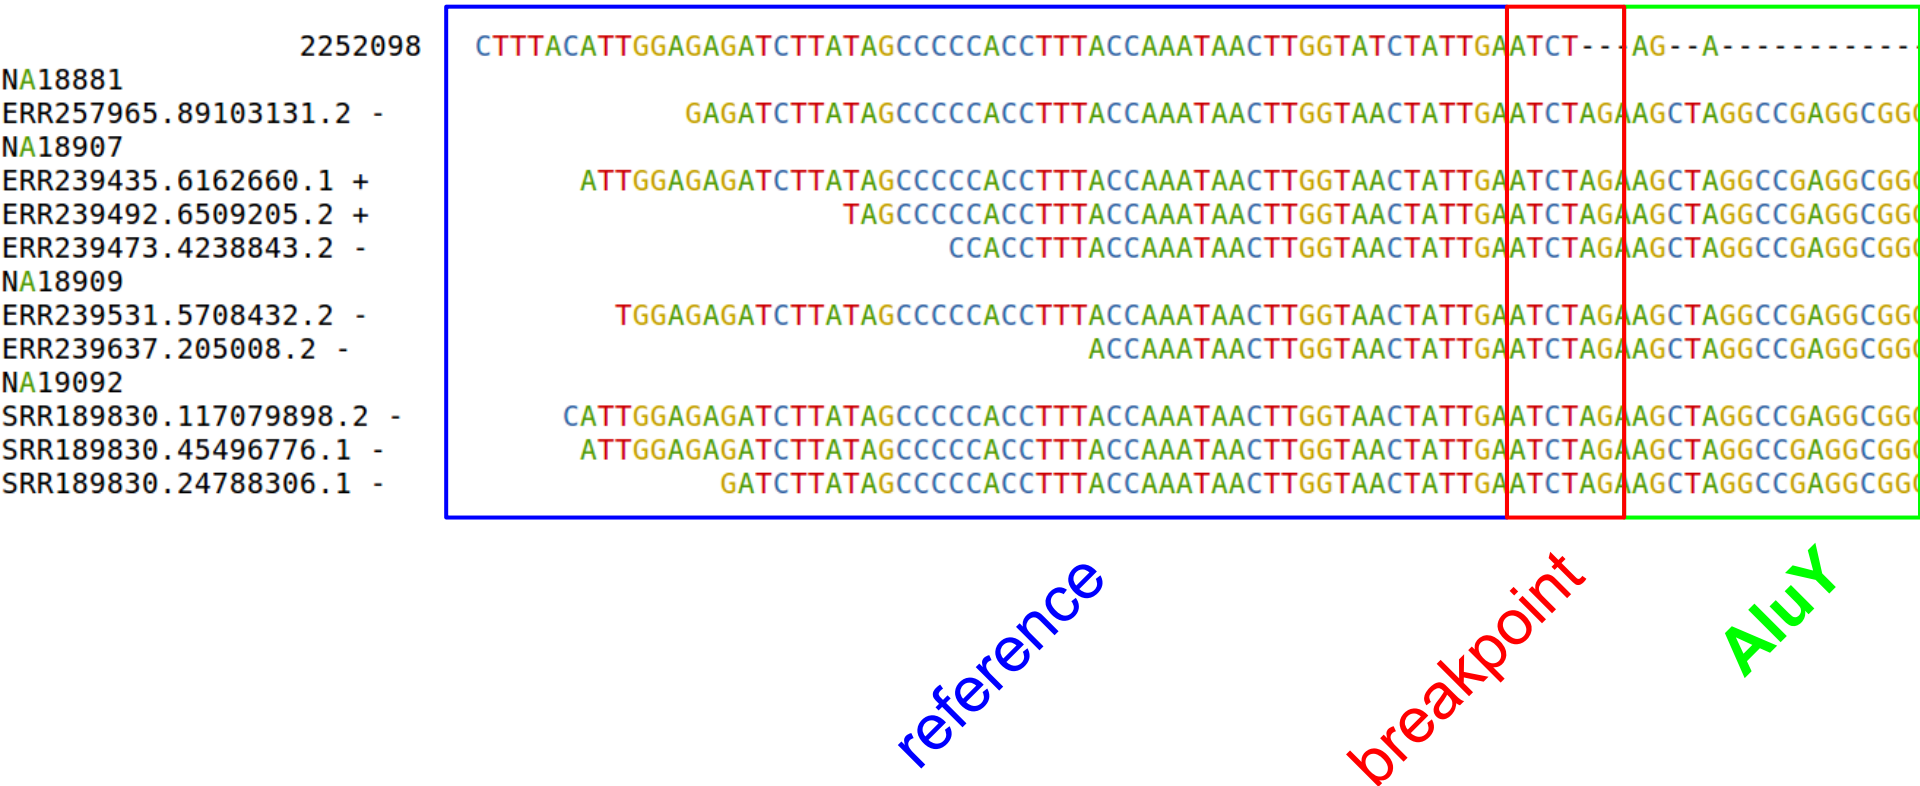
\* includes glia realignment

Imputation of variant calls on chr20 via SHAPEIT 2. Imputed results are tested against Complete Genomics samples in 1000 Genomes.

**We do as well for high-quality indels as SNPs!**

*Olivier Delaneau, Androniki Menelaou, Jonathan Marchini*

# Mobile element detection



Using *glia+freebayes* to re-genotype an AluY insertion at 20:2252139 in the YRI population. Insertion structure is estimated from split-read mappings.


# Alu genotyping efficiency

We re-call a substantial fraction of known (validated) Alus in 1000G low-coverage bwa alignments.


set	re-genotyped Alus	%
Pilot 2 data (source)	282	99.6
PCR-free NA12878	281	99.3
5x NA12878 (low-coverage)	173	61.1

Stewart et. al 2011. A Comprehensive Map of Mobile Element Insertion Polymorphisms in Humans. *PLoS Genetics*.

# Genotyping large deletions



# Performance using 1000G phase 3 SNPs and indels >1% frequency



**Deep-coverage 100bp Illumina data on NA12878** was downsampled to 5, 10, 20, 30, and 50-fold. Calling by both freebayes and freebayes+glia (realigning to 1000G variants >1% MAF), and comparing the results to the **Genome In a Bottle truth set** demonstrates marked improvement in sensitivity, particularly at low-coverage.

depth	snp AUC diff	indel AUC diff
5	6.02%	1.87%
10	1.07%	0.78%
20	0.26%	0.37%
30	0.08%	0.40%
50	0.02%	1.2%

# Questions?

...