Regularized sequence-context mutational trees capture variation in mutation rates across the human genome

Christopher J. Adams[1], Mitchell Conery[1], Benjamin J. Auerbach[1], Shane T. Jensen[2], Iain Mathieson[3], Benjamin F. Voight[3,4,5]

Author affiliations:

1. Genomics and Computational Biology Graduate Group, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, USA
2. Department of Statistics and Data Science, The Wharton School at the University of Pennsylvania, Philadelphia, USA
3. Department of Genetics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, USA
4. Department of Systems Pharmacology and Translational Therapeutics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, USA
5. Institute for Translational Medicine and Therapeutics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, USA

Correspondence to:

Benjamin F. Voight, PhD
Associate Professor of Systems Pharmacology and Translational Therapeutics
Associate Professor of Genetics
University of Pennsylvania - Perelman School of Medicine
3400 Civic Center Boulevard
10-126 Smilow Center for Translational Research
Philadelphia, PA 19104
Email: bvoight@pennmedicine.upenn.edu

ABSTRACT

Germline mutation is the mechanism by which genetic variation in a population is created. Inferences derived from mutation rate models are fundamental to many population genetics inference methods. Previous models have demonstrated that nucleotides flanking polymorphic sites – the local sequence context – explain variation in the probability that a site is polymorphic. However, limitations to these models exist as the size of the local sequence context window expands. These include a lack of robustness to data sparsity at typical sample sizes, lack of regularization to generate parsimonious models and lack of quantified uncertainty in estimated rates to facilitate comparison between models. To address these limitations, we developed Baymer, a regularized Bayesian hierarchical tree model that captures the heterogeneous effect of sequence contexts on polymorphism probabilities. Baymer implements an adaptive Metropolis-within-Gibbs Markov Chain Monte Carlo sampling scheme to estimate the posterior distributions of sequence-context based probabilities that a site is polymorphic. We show that Baymer accurately infers polymorphism probabilities and well-calibrated posterior distributions, robustly handles data sparsity, appropriately regularizes to return parsimonious models, and scales computationally at least up to 9-mer context windows. We demonstrate application of Baymer in three ways – first, identifying differences in polymorphism probabilities between continental populations in the 1000 Genomes Phase 3 dataset, second, in a sparse data setting to examine the use of polymorphism models as a proxy for *de novo* mutation probabilities as a function of variant age, sequence context window size, and demographic history, and third, comparing model concordance between different great ape species. We find a shared context-dependent mutation rate architecture underlying our models, enabling a transfer-learning inspired strategy for modeling germline mutations. In summary, Baymer is an accurate polymorphism probability estimation algorithm that automatically adapts to data sparsity at different sequence context levels, thereby making efficient use of the available data.

AUTHOR SUMMARY

Many biological questions rely on accurate estimates of where and how frequently mutations arise in populations. One factor that has been shown to predict the probability that a mutation occurs is the local DNA sequence surrounding a potential site for mutation. It has been shown that increasing the size of local DNA sequence immediately surrounding a site improves prediction of where, what type, and how frequently the site is mutated. However, current methods struggle to take full advantage of this trend as well as capturing how certain our estimates are, in practice. We have designed a model, implemented in software (named *Baymer*), that is able to use large windows of sequence context to accurately model mutation probabilities in a computationally efficient manner. We use Baymer to identify specific DNA sequences that have the biggest impacts on mutability and apply the model to find motifs that have potentially evolved mutability between different human populations. We also apply it to show that germline mutations observed as polymorphic sites in humans - those that have occurred in our recent evolutionary history - can model very young mutations (*de novo* mutations) as well as polymorphism observed in populations of closely related great ape species.

INTRODUCTION

Germline mutations are the primary source of genetic variation between and within species. Quantifying where, what type, and how frequently mutations arise is therefore of fundamental importance to population genetic inference and complex trait studies. Better estimates of mutation rates improve tools designed to quantify population divergence times[1], demographic history[2], and the effects of background selection[3]. Moreover, models for the underlying *de novo* mutation rate from which burden of mutations can be statistically assessed have enabled discovery of genes[4,5] and non-coding sequences[6,7] contributing to complex disease[4,5,8,9].

Our working hypothesis is that there exists an underlying structure to the context-dependent effects that shape the mutation rate. Here, we focus on polymorphism probabilities as a proxy for the mutation rate that we hypothesize share the same context-dependent architecture subject to genetic drift, demography, selection, biased gene conversion, or additional phenomenon that operate across population history. The frequency of polymorphisms varies widely across the genome[10] and correlates with several genomic features[11–13], with new mutations caused by both exogenous and endogenous sources[14]. There is considerable evidence to suggest that local nucleotide context directly relates to the probability that a nucleotide mutates. A classic example of this is the ~14-fold higher rate of C>T transitions at methylated CpG sites, owing to spontaneous deamination of 5-methylcytosine[15–17]. Long tracts of low-complexity DNA show elevated variability in mutation rates, which is hypothesized to be the result of slippage of DNA polymerase during replication[18]. This prior work suggests that local sequence context is integral to understanding variation in polymorphism rates across the genome, and that the most predictive models will be best positioned to guide elucidation of the underlying mutational mechanisms.

Our previous work demonstrated that a sequence context window of seven nucleotides (i.e., '7-mer') provided a superior model to explain patterns of genetic variation relative to smaller windows that are commonly used (e.g., 3-mers)[19]. While an advance, this model was fundamentally limited for three reasons: *scalability, regularization,* and *uncertainty*. First, the size of the model – which increases by a factor of four for each nucleotide included – presents intrinsic limits both computationally and in terms of statistical power. Second, while it is straightforward to assume that every sequence context is meaningful, a more parsimonious model – informed by biological intuition – might be that only a subset of contexts contributes meaningfully to the observed variation in data. This is particularly important for inference of somatic and *de novo* mutation rates or in other data-sparse situations (e.g., across species). Finally, while our previous model provided a point estimate of the mean polymorphism probability, it did not immediately emit uncertainty resulting from multinomial variance and heterogeneity in larger sequence contexts. As sequence context sizes are expanded, there is functionally less data and thus more uncertainty in estimates, making point estimates even more unreliable. Quantifying uncertainty is also required for detecting differences in probabilities across models, for example when comparing differences in rates across populations[20–22] or at functional genomic features[23]. Ideally, a method should scale the inferred context length proportional to the amount of data and the biological signal that may be present within that data while providing uncertainty in estimated parameters and underlying probabilities.

Previous work has sought to address these challenges, though methods introduced to date do not address all limitations simultaneously. Sparsity and scalability have been tackled through a deep-learning framework[24] as well as an IUPAC-motif-based clustering approach[25] which modeled polymorphism probabilities up through 9-mers. Another method explored polymorphism probabilities up through 7-mers using DNA shape covariates to reduce the parameter space[26]. All three methods are robust and effective at measuring point estimates of polymorphism probabilities in expanded sequence contexts, however none explicitly estimates the uncertainty of these parameters. Finally, the CIPI model[27] is a Bayesian method that addresses these issues, but focuses on applications with smaller context-window motifs (5-mer) in variant settings with fewer mutation events (e.g., somatic mutations in cancer or mutations in viral genomes) and is not obviously scalable computationally to larger size context windows and sizes of contemporary population genomics data sets in humans (e.g., hundreds of millions of polymorphic sites).

Here, we develop a method that addresses all three limitations embedded in a novel model. We construct a Bayesian tree-based method that integrates sequence context window size, handles sparse data, and captures uncertainty in estimates of mutation probability via the posterior distribution. We subsequently apply our approach in multiple ways. First, we quantify differences in polymorphism probabilities between continental populations and place bounds on the effect sizes of potential undescribed context-dependent differences in the 1000 Genomes dataset[28]. Second, we explore the use of polymorphism datasets to predict *de novo* mutations. We measure the effect of population history, variant age, and sequence context size on model performance with the aim of generating a meaningful proxy to estimate the germline mutation rate. Finally, we build models of different great ape species and assess the similarity to human polymorphism models.

DESCRIPTION OF THE METHOD

A tree-based sequence-context model captures variation in polymorphism probabilities

We began by developing a model to describe the hierarchical relationship of sequence context dependencies over increasing window sizes. We structured this as a rooted, tree-based graph, where each type of substitution class is represented distinctly (**Fig. 1A**). Each level of the tree represents an increasing window size of sequence considered, alternating between incorporating nucleotides to the window on the 3′ end for even-sized contexts and on the 5′ end for odd-sized contexts. We fold over reverse complementary contexts to reduce the parameter count (**S1 Text**). To ease readability, we denote each mutation with the sequence context, the nucleotide in scope bolded, and the polymorphism indicated with an arrow (e.g. T<u>C</u>C>T represents the polymorphism where the underlined cytosine has become a thymine). Each non-root edge represents the log-transformed, multiplicative shift in polymorphism probability captured by expanding sequence context. The root edge corresponds to an estimated base polymorphism probability for a given mutation type. For a given sequence context, each node in the tree represents the probability of

observing a polymorphic site in the central nucleotide (referred to hereafter as polymorphism probability), and is the product of all edges, starting from the root that leads to the node (**Fig. 1B**). As our previous work has shown for a specific level of sequence context, the distribution of observed counts for each sequence context can be modelled via independent multinomial distributions[19] facilitating likelihood calculation. The resulting multinomial probability vector corresponds to the combination of individual polymorphism probability estimates across each mutation type tree for each sequence context (**S1 Text**).

Within the model, we incorporate two features essential for downstream applications when comparing the outputs of competing models. First, we employ a Bayesian formulation which generates posterior distributions for polymorphism probabilities (**S1 Text**). This approach naturally provides uncertainty around parameter estimates which is essential for comparison of rates across different tabulated models. Second, we incorporate regularization in the parameter estimation procedure for tree edges. Previous sequence context models estimated parameters (ϕ) for all edges of the tree, meaning that all values of were effectively non-zero. However, our previous work suggested that perhaps only a fraction of edges meaningfully contribute information[19]. Hypothesizing that only a subset of edges is informative for estimating mutation probabilities, we regularize our tree model by incorporating a *spike-and-slab* prior on the ϕ parameters[29]. Our approach estimates the fraction of posterior samples in the slab, implying a non-zero effect on polymorphism probabilities, and in the spike, which implies no effect. Thus, the probability of an edge being included in the slab is the equivalent of the posterior inclusion probability (PIP) for our model. We tune the model such that the slab is favored when the evidence suggests a multinomial probability shift greater than 10% for a given context level (**Fig. 1C**). This value was choosen weighing the stability of model convergence with the goal of inferring the largest possible effects.

Because the posterior distribution is not analytically tractable, we implemented an adaptive Metropolis-within-Gibbs Markov Chain Monte Carlo (MCMC) sampling scheme[30] to sample from and thereby estimate the posterior distribution of this model. To further aid in convergence and enforce intermediate nodes to have identifiable mutation probabilities, we estimated parameters of the model level-by-level rather than all simultaneously, leveraging the conditional dependency structure of the hierarchical tree. Under this set-up, the unseen higher-order layers are assigned $\phi_{a,b} = 0$ edges until their level has been sampled. We embedded this model and sampling scheme into software (named Baymer) for further testing and applications.

VERIFICATION AND COMPARISON

Asymmetric context expansion improves parsimony and model inference

The hierarchical tree-based Baymer graphs are constructed such that the difference in length of flanking nucleotides on either side of the focal nucleotide is zero for odd-length contexts, and one for even-length contexts. It follows that these trees can be constructed in three ways:

expanding sequence contexts by including even-length contexts (termed “asymmetric models”) – alternating expansions starting at either the (1) left, or 5’, end, or (2) starting with the right, or 3’, end – or (3) expanding by exclusively using odd-length contexts (termed “symmetric models”; **Figure 2A**). When expanding an odd-length context by two nucleotides (e.g., 1-mer to 3-mer), symmetric models require 16 edges per context as compared to 20 edges per context (16 edges + 4 intermediate edges) in the asymmetric model. Despite more total edges in the asymmetric model tree architecture, we hypothesized the intermediate edges would more efficiently capture signal and provide greater resolution to detect inflection points in the tree where the local sequence context results in mutability changes. We implemented and tested all three models on the same dataset. First, we observed that for our most challenging model (9-mers), rate estimates are very strongly correlated amongst all models (**Figures 2B and 2C**). Furthermore, we observed that despite more total edges in the tree-graph, asymmetric models include approximately 38% fewer overall edges with high confidence (**Figure 2D**), suggesting greater parsimony. Finally, asymmetric models produced models that better fit holdout data than the symmetric models (**S1 Text, Figure 2E**). This improvement arises specifically in situations where there is sufficient data to estimate 8-mer edges, but insufficient data to confidently estimate 9-mer rates. Given our folding scheme, we opted for the biologically pragmatic choice of an even-length context tree architectural model that initiates alternating context expansions with the right end (3’), as this captures CpG effect(s) as early as possible in the tree without distributing the effect across more than one edge.

Evaluation of the model demonstrates robust inference of the underlying rates with uncertainty

A key feature of Baymer is that it estimates posterior distributions for each parameter, allowing for uncertainty in the probabilities of polymorphism at each sequence context. To evaluate the coverage of the estimated posterior probabilities, we used simulations to assess how often our posterior distribution captures simulated values. Using a pre-specified polymorphism probability table, we tested how frequently polymorphism probabilities estimated by Baymer captured the true value for each sequence context (**S1 Text**). We found that across all sequence context sizes, 89%, 93%, and 97% of context simulations contained the true polymorphism probability in the 90%, 95%, and 99% credible intervals, respectively (**S1 Text, S1 Table**).

A second important feature is that regularization is embedded into the method, allowing for the creation of parsimonious models that capture most of the information with the fewest non-zero parameters. This part is critical to address cases where the amount of data is not large and limits power, or when considering larger windows of sequence context that are rare and/or uninformative. If robustly calibrated, we would expect probabilities inferred in a holdout set to strongly correlate with those estimated during a test phase (i.e., minimal overfitting). To evaluate the robustness of the inferred rates, we partitioned the human genome reference into two sets – even and odd base-pairs – and used SNPs of allele count 2 or greater observed in the gnomAD[31] non-Finnish European (NFE) collection to independently train models (**S1 Text**). We compared the concordance of probabilities for models with sequence context windows up to 4 flanking nucleotides on either side (i.e., a 9-mer model) using the maximum likelihood estimate approach[19]

222 and Baymer (**S1 Fig.**). For each comparison, in addition to the Spearman correlation, we also
223 calculated the root mean squared perpendicular error (RMSPE) from each point to the x-y axis,
224 as a measure of the tightness of the distribution from the true, shared value (**S1 Text**). The
225 maximum likelihood estimates of polymorphism probabilities (**Fig. 3A,** Spearman correlation ρ =
226 0.915; RMSPE = 0.117) were less correlated and considerably less tightly distributed than those
227 for Baymer-derived models (**Fig. 3B,** ρ = 0.990; RMSPE = 0.035). This result occurred even after
228 omitting ~16,000 sequence contexts with zero mutations in either dataset (odd and even base
229 pairs) from the maximum likelihood model comparison, rendering practical use of large swaths of
230 the model useless due to substantial overfitting at the 9-mer level. If zero-mutation contexts
231 omitted from the maximum likelihood model were included, the correlations would perform
232 considerably worse (**S1 Text, S1D Fig.,** ρ = 0.876; RMSPE = 0.744), as these polymorphism
233 probabilities are exclusively determined by pseudo counts. Within the NFE dataset, Baymer
234 inferences were also robust across allele frequency bins (**S2 Fig.**).
235
236 We next sought to evaluate the transferability of inferred models between experimental
237 collections; while internally consistent, the above procedure could simply reflect data set specific
238 biases[32]. For this, we compared non-admixed, non-Finnish European (EUR) samples obtained
239 from the 1000 Genomes (1KG) Project (re-sequenced by the New York Genome Center)[33] with
240 the gnomAD NFE sample described above. As before, we split the data into even and odd base
241 pairs but also applied a variant down-sampling procedure to match total variant count and site-
242 frequency spectrum between both sets **(S1 Text)**. By comparing variants found in the even base-
243 pair genome of gnomAD with the odd base-pair genome of 1KG, this strategy ensures no variation
244 overlapped between data sets. We observed that the probabilities estimated from both sample
245 sets were strongly correlated (ρ = 0.984; RMSPE = 0.045; **Fig. 3C**) though were slightly weaker
246 than the correlations from each internal comparison and fit less tightly (gnomAD ρ = 0.990;
247 RMSPE = 0.035; **Fig. 3B**; 1KG ρ = 0.986; RMSPE = 0.042; **S3 Fig.**). This result demonstrates
248 that while some additional between-sample variation may exist, Baymer infers probabilities of
249 polymorphism that are broadly consistent with one another, supporting the notion of model
250 transferability across different data sets.
251
252 We next aimed to quantify how well the model selects meaningful context features. We expected
253 more proximal bases to the focal site to have a greater impact on polymorphism probabilities for
254 two reasons, (i) due to data richness, and (ii) that proximity to the polymorphic site would suggest
255 more direct impacts on mutability, e.g., the CpG context. Consistent with expectation, the fraction
256 of edges with a PIP > 0.95 monotonically decreases as the sequence context size is increased
257 (**Fig. 3D**). For any given 9-mer context, we find a median of 3 edges included with high confidence
258 in the model (**S4A Fig.**). The median window of context-dependence for each 9-mer was 5 bp
259 wide, although this inference is limited by the sparsity of the model (**S4B Fig.**).
260
261 **Larger contexts best explain patterns of variation genome-wide**
262
263 We note that over 61% of all edges with a PIP > 0.95 are found in the 8-mer and 9-mer levels of
264 our model of polymorphism observed in the gnomAD NFE data. While fewer than 2% of 9-mer
265 edges meaningfully impact the final estimates, they still account for the most total absolute edges

(7189 total edges > 0.95 PIP) and are enriched for larger effect sizes (**S4C Fig.**). This observation holds even after filters for data sparsity (**S1 Text, Fig. 2E**). This implies a considerable impact on polymorphism probabilities in extended sequence contexts, consistent with previous work[19,23–25]. This general trend is similarly consistent across mutation types (**Fig. 2F**), although with a variable degree of impact, most notably with less additional variability estimated in wider CpG>T edges (**S4D Fig.**). We thus evaluated the overall improvement in likelihood by expanding window sizes up to 9-mers. Compared to lower context models (e.g., 3-mer, 5-mer, or 7-mer) on holdout data, 9-mer Baymer models substantially improved the likelihood and best fit to the data (**S1 Text, S2 Table**).

APPLICATIONS

Sequence context motifs are correlated with changes in polymorphism probability

We next aimed to identify inflection points in the Baymer trees by examining the edges corresponding to the largest phis across each layer. Unsurprisingly, the CpG>T edge had the largest mean posterior phi magnitude (**S3 Table**). Consistent with our previous results[19], edges with the largest absolute mean posterior phi are largely localized at the intersection of poly-A repeat-rich sequences (lower rates of A>T substitutions) but particularly presented in 8-mer context by poly-A tract of length 4, where the mutation type extends one of the repeated patterns (e.g., CGCG<u>A</u>GAGA>C or CCC<u>A</u>AAA>C), and the CA<u>A</u>TN motif which increases A>G mutability (2.35 – 2.99x increase, **S3 Table**).

Next, we attempted to discover specific motifs that are enriched in the highest or lowest 1% of 9-mer polymorphism probabilities within each mutation type (**S1 Text, S4 Table**). We recapitulate almost all previously reported motifs[19,23,25]. Consistent with previous reports, we identify a preponderance of repeat-rich motifs, which is likely due to the impact of slippage in introducing mutations[18]. We discover numerous motifs with flanks extending 4 bp from the focal nucleotide that showed enrichment (21 motifs with $p < 0.0001$; **S4 Table**), emphasizing the utility of expanded sequence context windows for modeling mutability.

Frequency of polymorphism across populations do not differ substantially across levels of sequence context

Prior work has centered around evaluating whether mutation rates have changed over evolutionary time by evaluating differences in the proportions of sequence-context-dependent polymorphism between human populations[21,22,34–36]. To determine whether polymorphism probabilities differ across human populations, we analyzed individuals from the NYGC resequencing of 1000 Genomes Project (1KG) Phase III representing continental European, African, East Asian, and South Asian groups. We extracted variants private to these continental groups, down-sampling to match site-frequency spectra bins and overall sample sizes (**S1 Text**). We then applied Baymer to each individual dataset to model probabilities up to a 9-mer window

of sequence context. We compared estimates of polymorphism probabilities in each population by assessing the degree to which the posterior distribution of each population's model parameters overlapped. The fraction overlap of each distribution is a proxy for the probability that the underlying parameters are the same. Due to the implicit tree structure of sequence context models, polymorphism probability shifts in edges will affect all edges downstream of the context in question. Therefore, we identified contexts where both the estimated polymorphism probability and the immediate edge leading to that context were both considered very likely to be different between populations.

Specifically, we identified contexts whose posterior estimates of polymorphism probabilities and edges both overlapped less than 1% in pairwise comparisons between the four populations (**S5 Table**). This included all the most notable previously reported 3-mer shifts across continental groups, including the increase in T<u>C</u>C>C mutations found in European relative to Non-European ancestry populations[20–22,34,35].

We next focused on the remainder of 3-mer and wider extended sequence contexts (**Table 1**). While a handful of such sequence contexts have been implicated[34], these results are confounded by batch effects in the original 1KG sequencing data[37]. Since the data we use for our analysis is derived from the New York Genome Center resequencing project[33], we do not expect the same confounder. In our results, we observed the presence of nucleotide repeats, e.g., TA / CG dinucleotides; poly-C / poly-A in several of the divergent contexts, which could be explained by polymerase slippage[18].

While the population-specific polymorphism probabilities estimated and polymorphism counts are identical between each pairwise comparison and thus correlated, we still note that 15/28 pairwise differences are specific to a single continental group. Of these, only the two canonical European context mutation differences (T<u>C</u>C>T and T<u>C</u>T>T) are in 3-mer contexts, otherwise all are found in 5-mer and greater window sizes. In South Asian samples, we find that the mean CT<u>A</u>TA>T polymorphism probabilities are approximately 1.6 times higher than the remaining populations and in Africans TAT<u>A</u>TATC>G is approximately 1.9 times higher. The largest population-specific effect was discovered in East Asians where ATA<u>C</u>CTC>A polymorphism probabilities are roughly 2.7 times higher than in European, African, or South Asian models. None of these effects have been explicitly documented before.

Taken collectively, we observed relatively few instances of edges that were quantifiably different across continental groups, and those that were observed were largely confined to relatively small windows of context where we might have anticipated well-powered tests (e.g., 3- and 5-mers). To quantify the power of our specific analytic procedure for discovery and the sample size necessary to identify true differences in polymorphism probabilities, we performed simulations where true effect differences were 'spiked-in' between two populations over a range of weak to stronger effects and across a sampling of different sequence contexts (**S1 Text**). Differences in mutability between populations for this experiment are defined as the natural log of the polymorphism probabilities ratio (NLPPR) between each simulated population. This allowed us to construct credible sets of effects that we were reasonably well powered (>80%) to discover (**Table 2**).

352 Unsurprisingly, the power scaled proportional to the number of context instances, simulated
353 mutations in the dataset, and the size of the spiked-in differences (**S5 Fig.**). Notably, extremely
354 subtle shifts (NLPPR <= 0.01; 0.99 – 1.01-fold change) were not detectable at any sequence
355 context size. On the opposite side of the spectrum, we found that we were reasonably powered
356 to identify shift differences where NLPPR > 1.0 (fold decrease <= 0.37 or fold increase >= 2.72)
357 up through 5-mers and in 6-mers with large sample sizes. For reference, the TCC>T
358 polymorphism has an NLPPR = 0.291 (~1.34 fold increase) – the largest difference of any 3-mer
359 by our calculation.
360
361 In contrast, our experiment had essentially no power to discover 9-mer polymorphism probability
362 changes and extremely limited power for 8-mers, even for large differences. Thus, there may exist
363 large differences at these sizes that we could not reliably capture. These results are consistent
364 with our comparisons in the real data (**Table 1**), as only differences within the detectable range
365 at each mer-level were implicated. These power calculations suggest that, given the experiment
366 we performed grouping all mutations together (agnostic to allele frequency or age, see
367 **Discussion**), if any 3-mer differences greater than the TCC>T shift exist, we would have
368 discovered these effects for a broad range of modest to very strong effects across a range of
369 sequence contexts window sizes. This effectively sets bounds on the differences possible for this
370 analysis scheme in this data.
371
372 **A sequence context model that captures variability in *de novo* mutational rates**
373
374 Given its formulation in handling data sparsity, we next sought to apply Baymer to develop a
375 model that best captures rates of *de novo* mutations across the genome. We took advantage of
376 a recent collection of 2,976 WGS Icelandic trios that identified 200,435 *de novo* events[38] and,
377 analogous to the above, we partitioned *de novo* variants into even (for training) and odd (for
378 testing) base pairs. We observed substantial improvement in the overall likelihood in the testing
379 set for 5-mer size windows compared to 3-mers (3-mer vs 5-mer, delta-LL = 2,144), but only
380 minimal improvement for increasing windows sizes further (5-mer vs 9-mer, delta-LL = 265).
381 Indeed, Baymer did not select any sequence context feature beyond the 5-mer level with PIP >
382 0.95. This is not unexpected given our approach to regularization, as the number of events in
383 larger sequence contexts is increasingly sparse, it is desirable to only include informative contexts
384 to avoid overfitting.
385
386 We next used Baymer to improve upon this baseline model. Previous work has demonstrated that
387 inference of *de novo* mutational probabilities can be captured via rare variant polymorphism data
388 obtained from population sets as a proxy[23]. We hypothesized that a partitioned set of
389 polymorphism data based on: (i) larger sample sizes that (ii) closely matched the ancestry of the
390 *de novo* set and (iii) focused on rare variants as a proxy to capture the most recent mutation
391 events would generate the most transferrable model and robust rate estimates. To build variant
392 partitions, we used variant call set data from gnomAD, focused on either a population-matched
393 proxy (i.e., NFE, the non-Finnish European subset) or variant calls from all samples in gnomAD
394 regardless of ancestry (i.e., ALL). For each of these, we created three partitions focused (i)
395 exclusively on variants with one allele count (i.e., singletons; for NFE labeled NFE-1), (ii)

396 exclusively on variants with two allele counts (i.e., doubletons; for NFE labeled NFE-2), and (iii)
397 variants with allele count of two or greater (for NFE labeled NFE-2+). Beyond this, we also
398 identified a set of putatively derived substitutions in the human lineage by comparing the GRCh38
399 human reference genome with ancestral sequences obtained from primates[39].
400
401 We applied Baymer to each variant set independently, comparing the likelihoods of each model
402 to explain rates of *de novo* mutation in the test set after downscaling probabilities proportional to
403 the sample size. First, we observed that for 3-mer sequence context models, the set of variants
404 obtained from the *de novo* training set outperformed all other models despite 102 to 1,377 times
405 fewer variants contributing to them than the polymorphism datasets (**Fig. 4A, S6 Table**). In
406 contrast, for larger windows of context (i.e., 7-mer and 9-mer), several of the polymorphism
407 partitions explained the data better than one trained directly from *de novo* events. This result
408 indicates that increased sample size is required to detect meaningful shifts in polymorphism
409 probabilities in larger sequence context windows.
410
411 Despite evidence to suggest singleton datasets should best recapitulate *de novo* variation[4,23,31],
412 we observed that models that trained exclusively on singletons and ALL-2 performed considerably
413 worse than the rest across all windows of sequence context (**Fig. 4A, S6 Table**). While our prior
414 intuition that larger numbers of variants would have provided better rate estimates from increased
415 power deeper in the context tree, rate models exclusively estimated with singletons suffer the
416 most from the impact of recurrent mutations[20,40], especially at CpG sites, which include the highest
417 polymorphism probability mutation type (CpG>T) (**S6 Table**).
418
419 Alternatively, population concordance between training and test and/or the quality of variant calls
420 used in training the model could also impact performance. As such, we next sought to explore the
421 effect of noise in low allele count variants. Although we only used variants that passed gnomAD
422 quality control checks, this filter still included a large proportion of variants with a negative log-
423 odds ratio of being a true variant (AS_VQSLOD < 0; **S6 Fig.**). This pattern was also evident for
424 other variant allele counts but were most striking in singletons and the ALL-2 variant groups.
425 Stricter quality filters (AS_VQSLOD > 5-10) considerably improved model performance, but still
426 did not surpass the *de novo* training model at the 3-mer level (**S6 Table**). Our NFE singleton
427 Baymer model trained on the strictest quality filter tested (AS_VQSLOD > 10) nearly equaled our
428 best performing model, NFE-2+, with ~ 1/30th the number of variants, but came up just short. In
429 summary, we observed that training from a population matched sample which excluded
430 singletons, NFE-2+, best predicted rates of *de novo* mutations in 5-mer or larger contexts, better
431 than models trained on *de novo* events directly.
432
433 Next, we sought to determine which sample set best modelled the *de novo* test set adjusting for
434 the total number of variants within the partition. To control sample size differences, we
435 downsampled each partition to match the number of variants observed in the *de novo* training set
436 (n=70,364) five times. After down-sampling and when considering 9-mer context models, we
437 observed that the partitions which included NFE exclusively (noted in green, **Fig. 4B**) performed
438 on average better than using the entirety of gnomAD, "ALL" (noted in orange in **Fig. 4B**), which
439 included a more diverse panel of individuals within Europe (e.g., Finnish) but also beyond Europe

(e.g., East and South Asian, African and African American). This is consistent with prior belief that, after controlling for the total sample size, variants that derive from samples where ancestries more closely match are the most informative.

A grafted tree approach provides superior estimates of *de novo* mutational probabilities

Given the observations that *de novo* models only outperform polymorphism-based models when either small sequence contexts are used (**Fig. 4A**) or the sample size is controlled (**Fig. 4B**), we next sought to explore a transfer learning-inspired[41] strategy to improve upon our model performance. Transfer learning has previously been employed in a sequence context modelling setting[24]. We hypothesized that regularization means that *de novo* models have reduced performance with expanded sequence contexts due to low sample sizes. Indeed, our *de novo* model did not have the power necessary to confidently (PIP > 0.95) include any non-zero shifts in sequence contexts larger than 5-mers in the model (**Fig. 5A**). The larger polymorphism datasets, however, were well-powered to detect shifts in every level of the tree (**Fig. 5A**).

The nested tree structure of our polymorphism probability models provides a natural strategy where specific branches of the estimated trees can be interchanged, i.e., a “grafted” tree. We asked how similar estimates for edges in expanded sequence contexts are between our *de novo* model and the best-performing polymorphism model, NFE-2+. In edges in 2-mer and greater levels where the *de novo* training model is powered enough to detect shifts (PIP > 0.95), the mean posterior estimates of shifts are highly correlated (**Fig. 5B**). This suggests a grafted tree approach is feasible, leveraging the polymorphism datasets for those edges the *de novo* model is incapable of estimating properly due to sparsity (**Fig. 5C**). Therefore, we built a grafted tree model using 1- to 3-mer edges estimated in the *de novo* training data model, and 4- to 9-mer edges estimated using the NFE-2+ data model. The resulting combined model had a greater fit to the holdout *de novo* data than either the NFE-2+ model or *de novo* model alone (**Fig. 5D, S1 Text**).

Sequence context-dependent mutability is shared between closely related great ape species

Finally, we examined how well human polymorphism models could capture variability in polymorphism levels observed in populations of great ape. Using polymorphism data from the Great Ape Genome Project[42], we built Baymer models of *Pan troglodytes* and *Gorilla gorilla* (**S1 Text**). We note broad agreement in estimated polymorphism probabilities between humans and both chimpanzees (**S7A Fig.,** Spearman correlation $\rho = 0.950$; RMSPE = 0.089) and gorillas (**S7B Fig.,** Spearman correlation $\rho = 0.942$; RMSPE = 0.103). These results indicate that the rates of polymorphism at higher orders of sequences contexts are similar across closely related great ape species. As we were especially interested in how human polymorphism models compared with chimpanzee and gorilla models in predicting holdout data in each respective species, we then tested models on odd base pair data in each species, training models using even base pair data for the species in focus (**S1 Text**). For both chimpanzee (**S7C Fig.**) and gorilla (**S7D Fig.**) tests, species-matched 9-mer models outperformed all other models. While human-derived models are outperformed at the 9-mer level, it is notable that human 9-mer models are more likely than

chimpanzee 7-mer models against chimpanzee data and gorilla 5-mer models against gorilla data (**S7 Table**). Taken collectively, these results suggest the rates of polymorphism at higher orders of sequences contexts are similar across closely related great ape species, with within-species models best capturing variability in observed polymorphism levels.

DISCUSSION

Here, we present Baymer, a Bayesian method to model mutation rate variation that computationally scales to large windows of nucleotide sequence context, robustly manages sparse data through an efficient regularization strategy, and emits posterior probabilities that capture uncertainty in estimated probabilities. Consistent with previous studies[24–26], we show that expanded sequence context models in most current human datasets are overfit with classic empirical methods but considerably improve model performance when properly regularized. As a result, this method allows for renewed evaluation of experiments that originally were statistically limited to polymorphism probability models with small sequence context windows.

We examined differences in polymorphism probabilities between the continental populations in the 1KG project. While differences in 3-mer polymorphism probabilities have been well-documented[20–22] and expansions up to 7-mers have been tested[34], both methods rely on empirical models with frequentist measures of uncertainty. Here, we expanded the search space out to 9-mer windows and leverage the uncertainty estimated in the model to directly quantify differences in these populations. We note that many of the differences discovered contain poly-nucleotide repeats. There is some prior literature on the mechanism of slippage in polymerases during replication of such sequences[18], so differential efficiencies of these enzymes across populations could conceivably result in these patterns. However, it is also very possible that artifacts from sequencing errors with differential effects across populations could explain the differences.

Despite being well-powered to identify a large range of differences in 3-mer and smaller contexts, we identified very few contexts that differ with high probability between the populations tested. This implies that if large-scale population differences in the mutation spectrum do exist at these window context sizes, they may be comprised of numerous subtle shifts rather than a few large changes, in agreement with conclusions from prior work[22].

We also explicitly placed bounds on the magnitude of differences that could possibly exist in these data without being detected, quantifying what differences we can expect to be discovered given the way variants are grouped in this experiment. Even though the 1KG project is relatively small compared to current datasets, the number of sequence contexts available for modeling is dataset-independent and inherently limited by the sequence diversity of the human genome. Thus, while more polymorphism data could lead to the discovery of additional smaller shifts in the future, bigger datasets will not improve the power to detect larger shifts in this allele frequency agnostic setting. In fact, for very large samples, polymorphisms in some contexts can become saturated[43], reducing the information content in a similar manner as overly sparse data. Thus, both to increase power and to improve modeling resolution, it will become necessary to partition the data (e.g., by allele frequency or variant age[36], or other genomic features).

It remains a challenge to disentangle the contribution of demography[20,35,44] versus changes in the underlying mutation rate on the mutation spectrum. Here, we control for the site frequency spectrum of variants included, but the next stage of this model will need to incorporate more sophisticated demographic features. Integrating Baymer-derived trees with a joint mutation spectrum and demographic history method, such as mushi[35], is a promising future direction. While this work focuses on modeling mutability in regions minimally affected by background selection, constraint could also bias estimates. Given prior work[19] we do not expect the underlying sequence-context-mediated mutability to behave any differently in constrained regions, suggesting future Baymer-estimated codon-aware models to explicitly model expected variation in coding regions.

We also aimed to address the degree to which polymorphism datasets could be used to approximate the *de novo* mutation rate. Currently, true *de novo* mutation datasets are limited in size, which place bounds on the scope of inference for adequate sequence context modeling. We demonstrate that polymorphism datasets are accurate proxies for *de novo* mutation models and largely share the same context-dependent mutability shifts, though in contrast to reports in the literature[4,23,31], the focus exclusively on singleton variants (at least, using gnomAD calls) performed poorly relative to all other considered models. Indeed, our experiment indicates that it is preferable to use germline mutation models based on large polymorphism datasets that can estimate shifts through the 9-mer level than it is to use the largest 3-mer de novo dataset, as is commonly used in the literature[4,5,31]. Including exclusively variants from either polymorphism data or *de novo* data was also suboptimal, however, as the best possible model we built for estimating *de novo* mutation rates used *de novo* mutations in concert with polymorphism datasets. The success of this experiment implies a general context-dependent mutability architecture that underlies the human mutation spectrum. The similarity of the derived dataset, which in theory represents the oldest subset of variants tested, to the *de novo* variation further strengthens this argument. We note this dataset could in theory be biased towards European samples given the history of the Human Genome Project[45], and as such, refinements will need to be made as more diverse representations of the human genome are created. Overall, this work suggests that although there have been some well-documented small changes in context-dependent mutation rates, the general architecture remains largely conserved during modern human history.

Our experiments modeling great ape variation suggests this general architecture might be more pervasive across the tree of life. While some specific mutation spectra differences have been documented[22,46], we note broad agreement amongst closely related species as well as similar signals in extended sequence contexts. For those non-human species with WGS datasets, cohort-sizes are usually very small (< 100), however, Baymer is well-suited to handle these sparse data situations. Furthermore, for those species with very little data, this work opens the exciting possibility to leverage closely related species' models as priors for estimating variation in less well-characterized relatives. Further work is necessary to model species across the tree of life to determine the extent that sequence context-dependent mutability is shared and how transferable 9-mer models can be.

One limitation of the model is the treatment of multi-allelic sites. Currently, multi-allelic sites are treated as separate polymorphisms which violates assumptions of the multinomial model, where only one outcome is possible for each locus. When we excluded multi-allelic sites, we observed biases in the rates of CpG>A and CpG>G mutations, which are disproportionately filtered as a side-effect of sharing the same sequence contexts with CpG>T mutations. A more nuanced approach that models multiallelic and biallelic sites separately and then integrates jointly would deal with this issue, though multiple mutations at the same nucleotide position with the same allele change would require additional effort[40].

Finally, although we can identify regions of the tree where polymorphism probabilities diverge and thus infer critical points in the tree, this model is tailored towards polymorphism probability estimation rather than explicitly for motif discovery[27]. Our objective is to estimate polymorphism probabilities rather than finding those contexts with the largest effect sizes. Although including even-length contexts yields better-performing models, the current tree architecture only explicitly captures the effects of half of such contexts. While adding one nucleotide at a time pseudo-symmetrically for tree generation reduces the computational sampling load, it makes for more difficult interpretation of the resulting mono-nucleotide impacts. Baymer's formulation also does not model the mutability of target contexts independent of mutation type, which currently requires post-hoc analysis to identify motifs that have non-specific mutability signatures. Future work will therefore need to integrate all possible paths through the sequence context tree and share information across contexts between mutation type trees.

In all our experiments, we focused on the entirety of the non-coding genome that is accessible to sequencing. That said, Baymer can easily be applied to any genomic features of interest for both polymorphism probability estimates and comparisons of feature-dependent sequence context mutability changes. Our approach does not currently incorporate genomic features in the model, but given genomic area bounds, polymorphism probabilities can be tailored to a biological question of interest. Addressing questions regarding the impact of genomic features on observed polymorphisms will be enhanced with well-regularized models, as smaller genomic areas or specific variant conditions can induce considerable data sparsity by reducing the number of contexts and/or polymorphisms available. Therefore, Baymer paves the way for exciting possibilities to study the effects of genomic features, variant age, and smaller subpopulations on sequence context-dependent mutation rate variation.

Acknowledgements

We would like to thank Dr. Ziyue Gao for her very helpful feedback on the manuscript and in the development process.

Author Contributions

The experiments were conceived and designed by C.J.A., S.T.J., I.M., and B.F.V. C.J.A. and B.F.V. performed statistical analyses. C.J.A., M.C., B.J.A., and B.F.V. analyzed the data. C.J.A. and B.F.V. drafted the initial manuscript. All authors contributed and edited the final manuscript. The work was supervised by B.F.V.

Financial Disclosure

B.F.V. acknowledges support for this work from the NIH/NIDDK (DK101478 and DK126194). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

REFERENCES

1. Wang, Y. & Nielsen, R. Estimating population divergence time and phylogeny from single-nucleotide polymorphisms data with outgroup ascertainment bias. *Mol Ecol* (2012) doi:10.1111/j.1365-294X.2011.05413.x.
2. Gutenkunst, R. N., Hernandez, R. D., Williamson, S. H. & Bustamante, C. D. Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. *PLoS Genet* (2009) doi:10.1371/journal.pgen.1000695.
3. McVicker, G., Gordon, D., Davis, C. & Green, P. Widespread genomic signatures of natural selection in hominid evolution. *PLoS Genet* (2009) doi:10.1371/journal.pgen.1000471.
4. Lek, M. *et al.* Analysis of protein-coding genetic variation in 60,706 humans. *Nature* (2016) doi:10.1038/nature19057.
5. Havrilla, J. M., Pedersen, B. S., Layer, R. M. & Quinlan, A. R. A map of constrained coding regions in the human genome. *Nat Genet* (2019) doi:10.1038/s41588-018-0294-6.
6. Chen, S. *et al.* A genome-wide mutational constraint map quantified from variation in 76,156 human genomes. *bioRxiv* 2022.03.20.485034 (2022) doi:10.1101/2022.03.20.485034.
7. Petrovski, S. *et al.* The intolerance of regulatory sequence to genetic variation predicts gene dosage sensitivity. *PLoS Genet* **11**, e1005492 (2015).
8. He, X. *et al.* Integrated model of de novo and inherited genetic variants yields greater power to identify risk genes. *PLoS Genet* **9**, e1003671 (2013).
9. di Iulio, J. *et al.* The human noncoding genome defined by genetic diversity. *Nat Genet* (2018) doi:10.1038/s41588-018-0062-7.
10. Hodgkinson, A. & Eyre-Walker, A. Variation in the mutation rate across mammalian genomes. *Nature Reviews Genetics* Preprint at https://doi.org/10.1038/nrg3098 (2011).
11. Stamatoyannopoulos, J. A. *et al.* Human mutation rate associated with DNA replication timing. *Nat Genet* (2009) doi:10.1038/ng.363.
12. Fryxell, K. J. & Moon, W. J. CpG mutation rates in the human genome are highly dependent on local GC content. *Mol Biol Evol* (2005) doi:10.1093/molbev/msi043.
13. Schuster-Böckler, B. & Lehner, B. Chromatin organization is a major influence on regional mutation rates in human cancer cells. *Nature* (2012) doi:10.1038/nature11273.
14. Gonzalez-Perez, A., Sabarinathan, R. & Lopez-Bigas, N. Local Determinants of the Mutational Landscape of the Human Genome. *Cell* Preprint at https://doi.org/10.1016/j.cell.2019.02.051 (2019).
15. Holliday, R. & Grigg, G. W. DNA methylation and mutation. *Mutation Research/Fundamental and Molecular Mechanisms of Mutagenesis* **285**, 61–67 (1993).
16. Sung, W. *et al.* Asymmetric context-dependent mutation patterns revealed through mutation-accumulation experiments. *Mol Biol Evol* (2015) doi:10.1093/molbev/msv055.
17. Lujan, S. A. *et al.* Heterogeneous polymerase fidelity and mismatch repair bias genome variation and composition. *Genome Res* (2014) doi:10.1101/gr.178335.114.
18. Bzymek, M. & Lovett, S. T. Instability of repetitive DNA sequences: The role of replication in multiple mechanisms. *Proc Natl Acad Sci U S A* (2001) doi:10.1073/pnas.111008398.

19. Aggarwala, V. & Voight, B. F. An expanded sequence context model broadly explains variability in polymorphism levels across the human genome. *Nat Genet* (2016) doi:10.1038/ng.3511.
20. Mathieson, I. & Reich, D. Differences in the rare variant spectrum among human populations. *PLoS Genet* (2017) doi:10.1371/journal.pgen.1006581.
21. Harris, K. Evidence for recent, population-specific evolution of the human mutation rate. *Proceedings of the National Academy of Sciences* (2015) doi:10.1073/pnas.1418652112.
22. Harris, K. & Pritchard, J. K. Rapid evolution of the human mutation spectrum. *Elife* (2017) doi:10.7554/elife.24284.
23. Carlson, J. *et al.* Extremely rare variants reveal patterns of germline mutation rate heterogeneity in humans. *Nat Commun* (2018) doi:10.1038/s41467-018-05936-5.
24. Fang, Y., Deng, S. & Li, C. A deep learning-based framework for estimating fine-scale germline mutation rates. *bioRxiv* (2021).
25. Bethune, J., Kleppe, A. & Besenbacher, S. A method to build extended sequence context models of point mutations and indels. *Nat Commun* **13**, 7884 (2022).
26. Liu, Z. & Samee, M. A. H. Mutation rate variations in the human genome are encoded in DNA shape. *BioRxiv* (2021).
27. Ling, G., Miller, D., Nielsen, R. & Stern, A. A Bayesian Framework for Inferring the Influence of Sequence Context on Point Mutations. *Mol Biol Evol* (2020) doi:10.1093/molbev/msz248.
28. 1000 Genomes Project Consortium *et al.* A global reference for human genetic variation. *Nature* (2015) doi:10.1038/nature15393.
29. George, E. I. & McCulloch, R. E. Variable selection via Gibbs sampling. *J Am Stat Assoc* **88**, 881–889 (1993).
30. Roberts, G. O. & Rosenthal, J. S. Examples of adaptive MCMC. *Journal of computational and graphical statistics* **18**, 349–367 (2009).
31. Karczewski, K. J. *et al.* The mutational constraint spectrum quantified from variation in 141,456 humans. *Nature* (2020) doi:10.1038/s41586-020-2308-7.
32. Ross, M. G. *et al.* Characterizing and measuring bias in sequence data. *Genome Biol* **14**, 1–20 (2013).
33. Byrska-Bishop, M. *et al.* High coverage whole genome sequencing of the expanded 1000 Genomes Project cohort including 602 trios. *bioRxiv* 2021.02.06.430068 (2021) doi:10.1101/2021.02.06.430068.
34. Aikens, R. C., Johnson, K. E. & Voight, B. F. Signals of Variation in Human Mutation Rate at Multiple Levels of Sequence Context. *Mol Biol Evol* (2019) doi:10.1093/molbev/msz023.
35. DeWitt, W. S., Harris, K. D., Ragsdale, A. P. & Harris, K. Nonparametric coalescent inference of mutation spectrum history and demography. *Proceedings of the National Academy of Sciences* **118**, e2013798118 (2021).
36. Gao, Z., Zhang, Y., Cramer, N., Przeworski, M. & Moorjani, P. Limited role of generation time changes in driving the evolution of the mutation spectrum in humans. *Elife* **12**, e81188 (2023).
37. Anderson-Trocmé, L. *et al.* Legacy Data Confound Genomics Studies. *Mol Biol Evol* **37**, 2–10 (2020).

38. Halldorsson, B. v *et al.* Characterizing mutagenic effects of recombination through a sequence-level genetic map. *Science (1979)* **363**, eaau1043 (2019).
39. Ensembl. Ensembl, Data from “homo_sapiens_ancestor_GRCh38.” *http://ftp.ensembl.org/pub/release-97/fasta/ancestral_alleles/homo_sapiens_ancestor_GRCh38.tar.gz*.
40. Johnson, K. E., Adams, C. J. & Voight, B. F. Identifying rare variants inconsistent with identity-by-descent in population-scale whole-genome sequencing data. *Methods Ecol Evol* **13**, 2429–2442 (2022).
41. Weiss, K., Khoshgoftaar, T. M. & Wang, D. A survey of transfer learning. *J Big Data* **3**, 1–40 (2016).
42. Prado-Martinez, J. *et al.* Great ape genetic diversity and population history. *Nature* **499**, 471–475 (2013).
43. Agarwal, I. & Przeworski, M. Mutation saturation for fitness effects at human CpG sites. *Elife* **10**, e71513 (2021).
44. Gao, Z. *et al.* Overlooked roles of DNA damage and maternal age in generating human germline mutations. *Proceedings of the National Academy of Sciences* **116**, 9491–9500 (2019).
45. Lander, E. S. *et al.* Initial sequencing and analysis of the human genome. *Nature* **409**, 860–921 (2001).
46. Goldberg, M. E. & Harris, K. Mutational Signatures of Replication Timing and Epigenetic Modification Persist through the Global Divergence of Mutation Spectra across the Great Ape Phylogeny. *Genome Biol Evol* **14**, evab104 (2022).
47. Smit, A. F. A. Repeat-Masker Open-3.0. *http://www. repeatmasker. org* (2004).
48. Hernandez, R. D., Williamson, S. H. & Bustamante, C. D. Context Dependence, Ancestral Misidentification, and Spurious Signatures of Natural Selection. *Mol Biol Evol* **24**, 1792–1800 (2007).
49. George, E. I. & McCulloch, R. E. Approaches for Bayesian variable selection. *Stat Sin* 339–373 (1997).
50. Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. Equation of state calculations by fast computing machines. *J Chem Phys* **21**, 1087–1092 (1953).

Main Figures

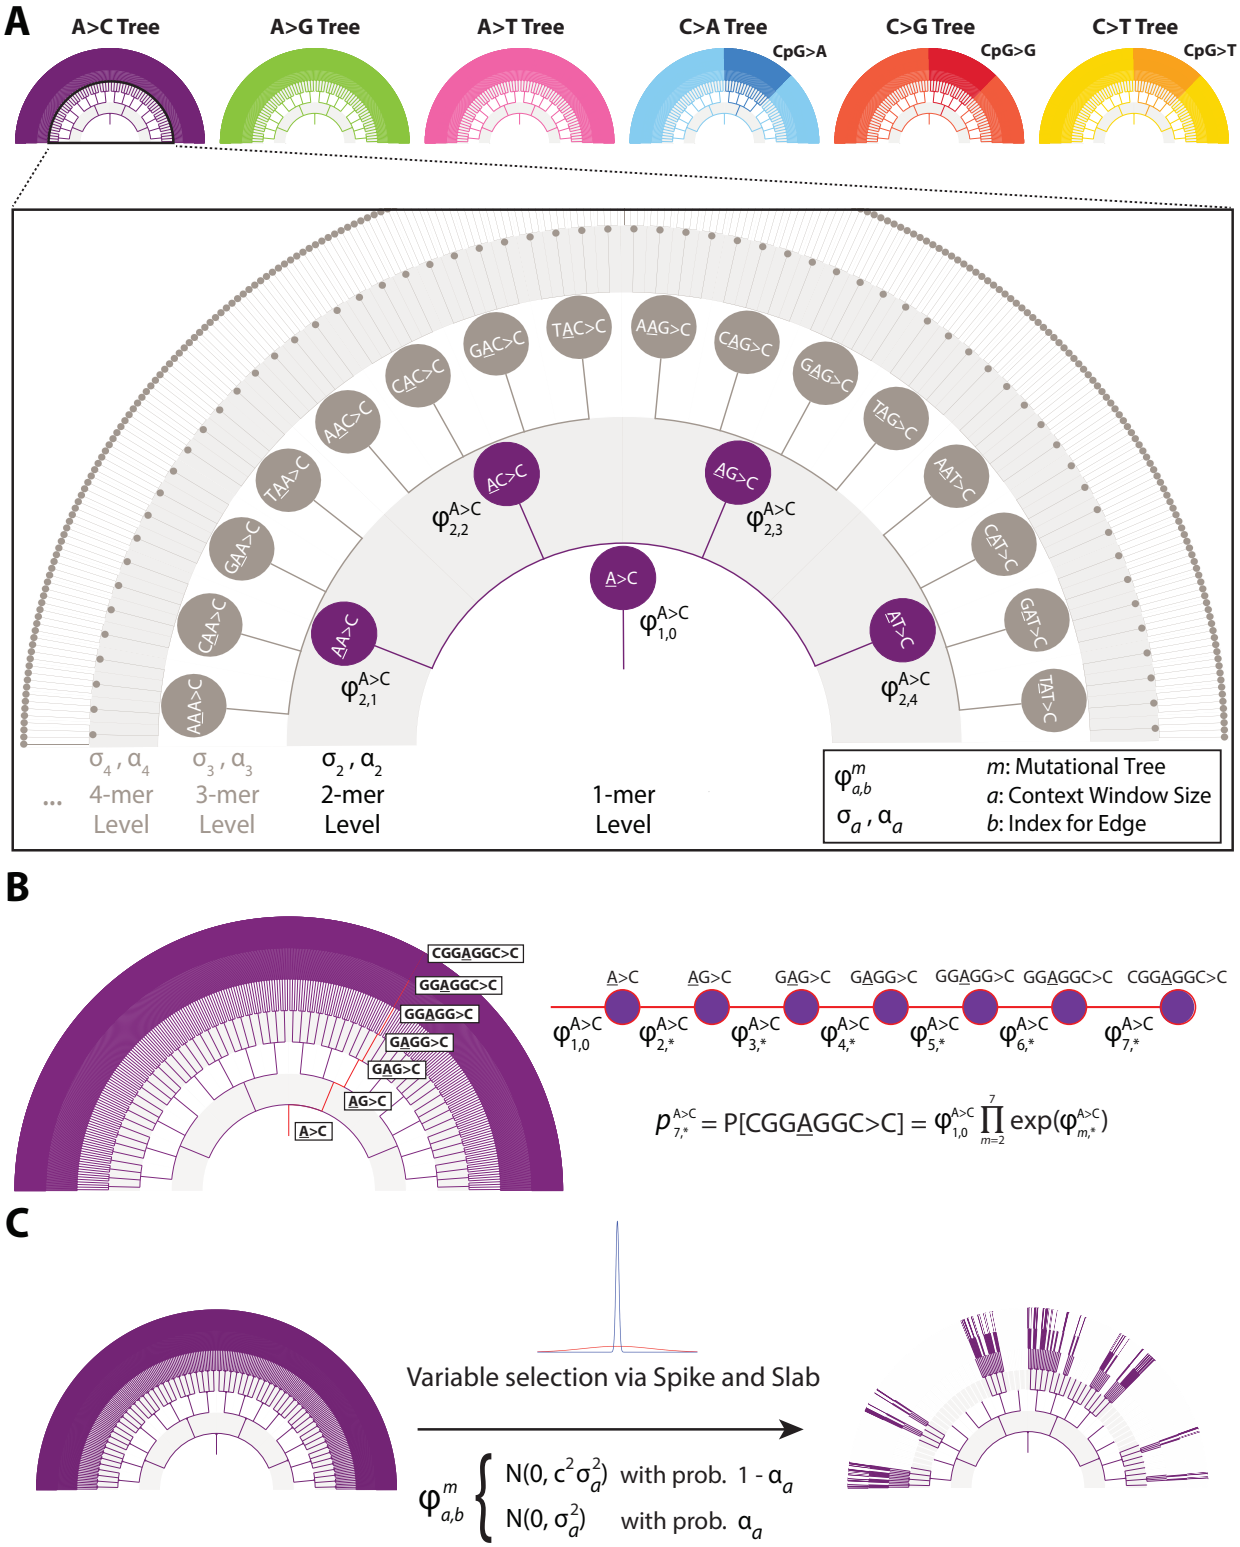

Fig. 1. Hierarchical relationship of sequence contexts and key algorithmic elements of Baymer. (A) Each mutation type is represented by a separate sequence context tree, related by

743 the shared mer level parameters and joint multinomial likelihood distribution. Each sequence
744 context tree has a nested structure where information is partially pooled across each shared
745 parent. (B) Polymorphism probabilities are parameterized as the product of the series of edges
746 that lead to the sequence context of interest. (C) Sequence context trees are regularized using
747 a spike-and-slab prior distribution.
748
749

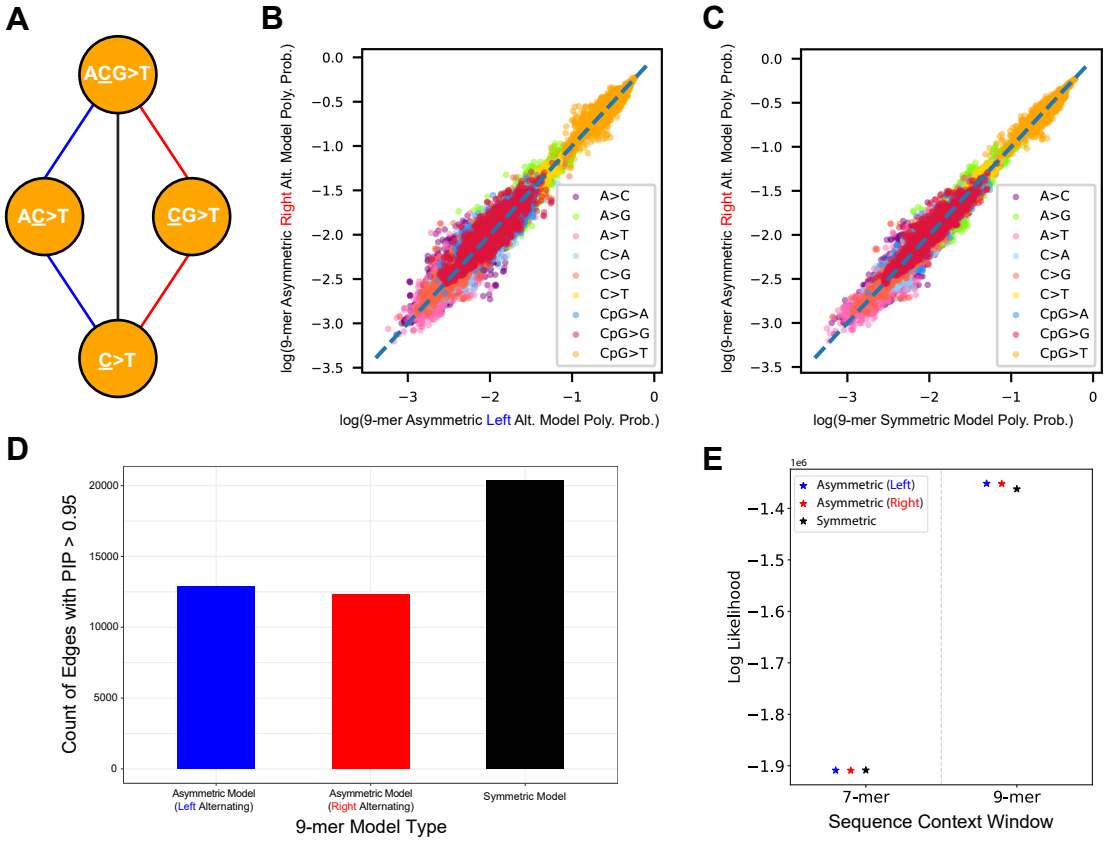

750
751 **Fig. 2. Exploration of different strategies to build out the Baymer sequence context**
752 **hierarchical trees using gnomAD non-Finnish European (NFE) polymorphisms with**
753 **derived allele count greater than or equal to 2 in non-coding accessible regions.** (A)
754 Sequence contexts can be built by starting to alternate adding nucleotides with the left, or 5',
755 end of each odd-length context (blue path), with the right, or 3', end of each context (red path),
756 or by only including odd-length contexts and adding a nucleotide to both sides of the growing
757 context (black path) for each expansion. (B) Baymer mean posterior estimates of 9-mer
758 polymorphism probabilities estimated using even base pair data with the right alternation pattern
759 and the left alternation pattern (Spearman correlation 0.997; $p < 10^{-100}$). (C) Baymer mean
760 posterior estimates of 9-mer polymorphism probabilities estimated using even base pair data
761 with the right alternation pattern and a model exclusively using odd-length sequence contexts
762 (Spearman correlation 0.998; $p < 10^{-100}$). (D) Absolute count of edges in each tree architecture
763 model with a PIP > 0.95 in a 9-mer model. (E) Multinomial likelihoods for each model are
764 calculated on odd bp NFE test data using 7-mer and 9-mer models. Polymorphism probability

estimates were linearly scaled to match the mean polymorphism probability of the holdout dataset.

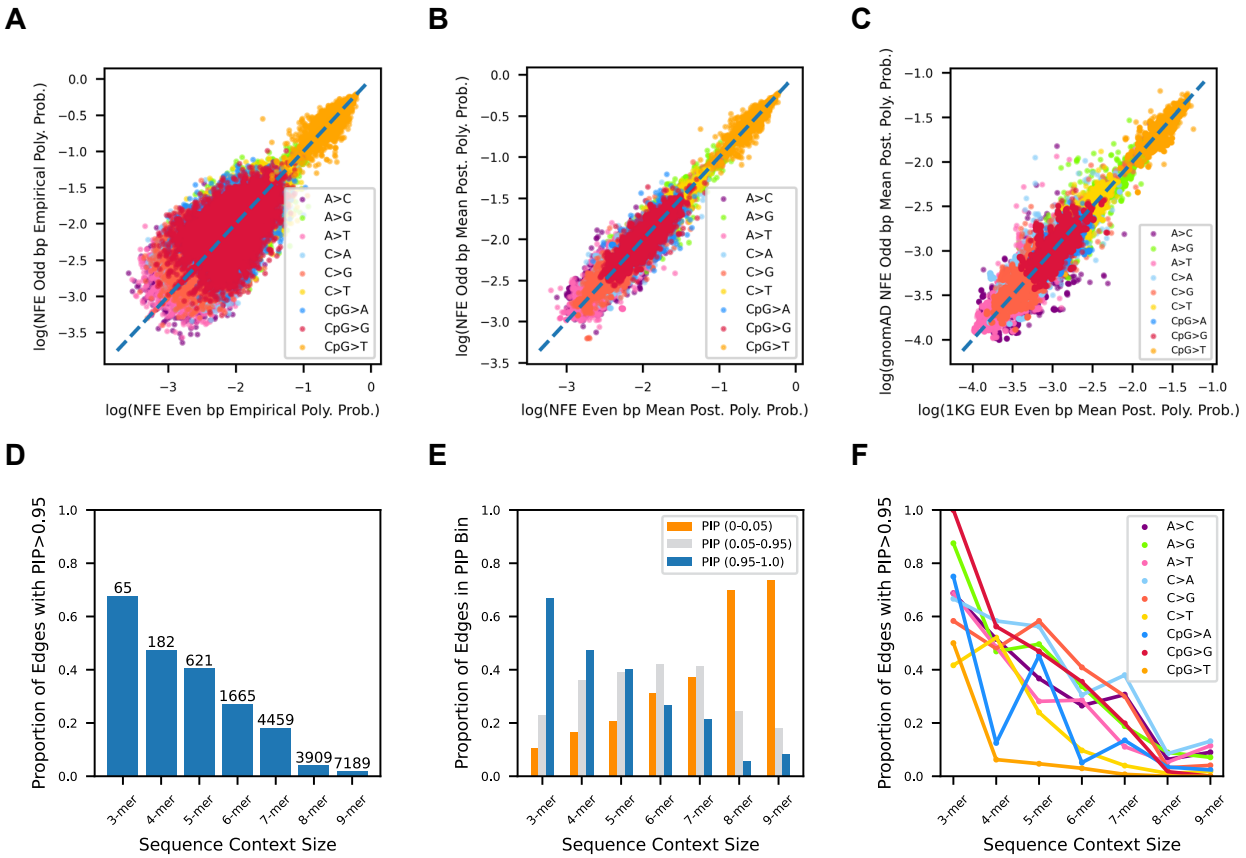

Fig. 3. Baymer model validation, transferability, and regularization in gnomAD non-Finnish European (NFE) polymorphisms with derived allele count greater than or equal to 2 in non-coding accessible regions. (A) Empirical 9mer polymorphism probabilities for context mutations with at least 1 occurrence in both datasets (15910 omitted context mutations) are plotted against one another (Spearman correlation 0.915; $p < 10^{-100}$; RMSPE = 0.12). (B) Baymer mean posterior estimates for 9mer polymorphism estimates in even and odd bp datasets (Spearman correlation 0.990; $p < 10^{-100}$; RMSPE = 0.035). (C) Baymer mean posterior estimates for 9mer polymorphism estimates in odd bp non-Finnish European gnomAD data and even bp NYGC 1KGPIII data, down-sampled to match total number of polymorphisms and site frequency spectrum (Spearman correlation 0.984; $p < 10^{-100}$; RMSPE = 0.045). (D) Fraction of edges in the NFE model with a PIP > 0.95 in each sequence context window layer. Absolute count of edges above bars. (E) For high-data contexts with at least 100,000 total instances in the non-coding genome and 50 total mutations, fraction of edges at each sequence context window size across PIP bins. (F) Proportion of high-data contexts within each mutation type at each sequence context window size with PIP>0.95.

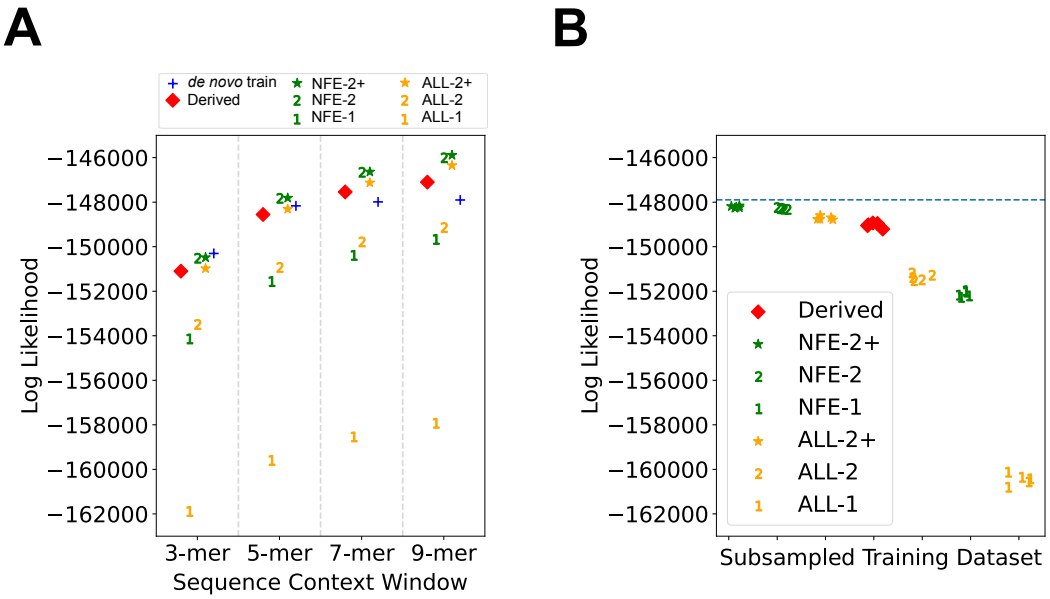

Fig. 4. Modeling de novo mutation probabilities using polymorphism datasets. Even bp Halldorsson et al. *de novo* training data modeled by Baymer is compared to Baymer-modelled polymorphism datasets partitioned by allele count. (A) Multinomial likelihoods for each model are calculated on odd *de novo* bp test data at various sequence context sizes. Polymorphism probability estimates were linearly scaled to match the mean polymorphism probability of the holdout dataset. (B) Polymorphism datasets were down-sampled to match the size of the even bp de novo data (70,364 variants) and multinomial likelihoods were calculated on odd de novo bp data. Each dataset was down-sampled using 5 different random seeds. The LL of the 9mer *de novo* training model is indicated with the blue dotted line.

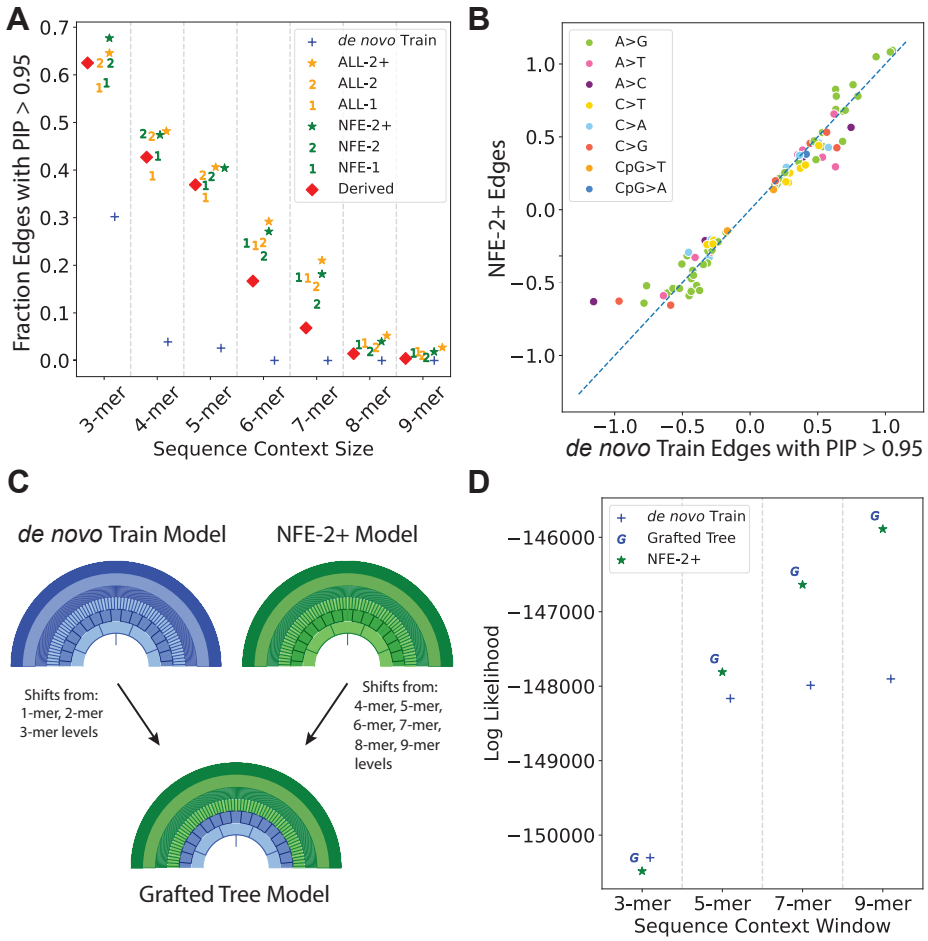

Fig. 5. Tree grafting strategy to share information between Baymer models. (A) For each de novo proxy model, we calculated the fraction of context polymorphism probability edges with a PIP > 0.95 in 2-mer – 9-mer mer-levels as a proxy for the degree of regularization in each model. (B) Edges in the de novo training model that are included with high-confidence (PIP>0.95) are very similar in magnitude and direction to their equivalents in the best-performing proxy model, NFE-2+, in 2-mer – 9-mer levels, implying a shared polymorphism probability shift structure. (C) Proposed tree-grafting schema for modeling de novo mutations that leverages mer-levels where de novo data is plentiful (1-mer – 3-mers) and uses polymorphism data to model the remainder of each model in larger mer-levels (4-mer – 9-mers) where the de novo model is underpowered. (D) The grafted tree method outperforms the previously best-performing model, NFE-2+.

819 **Main Tables**

820

Table 1. Baymer modeled 1KG private continental context mutations with extreme polymorphism probability differences

Population Comparison	Context Mutation	log(Poly. Prob. Fraction)	Poly. Prob. Fraction Overlap	Shift Difference	Shift Fraction Overlap	Population Specificity
European v African	TCC>T[abc]	0.291	0	-0.174	1.4E-157	European
	TCT>T[abc]	0.136	1.6E-18	-0.116	8.5E-16	European
	GCAATTA>G	0.569	4.7E-03	-0.668	2.4E-03	
	TATATATC>G	-0.660	7.2E-03	0.730	5.6E-03	African
European v South Asian	TCC>T[abc]	0.112	1.2E-09	-0.059	2.7E-03	European
	TCT>T[abc]	0.063	5.0E-03	-0.066	2.9E-03	European
	CTATA>T	-0.587	2.9E-03	0.493	7.3E-03	South Asian
	ATCTTC>G	-0.606	7.6E-03	0.668	5.4E-03	
European v East Asian	CCC>T[abc]	0.081	1.4E-03	0.075	6.6E-04	
	TCC>T[abc]	0.312	0	-0.156	2.4E-97	European
	GCT>T	-0.064	5.7E-03	0.095	6.1E-05	
	TCT>T[abc]	0.133	3.0E-19	-0.102	9.6E-06	European
	GCAACCA>G	1.056	5.3E-03	-1.104	5.0E-03	
	ATACCTC>A	-1.029	4.2E-03	0.830	5.0E-03	East Asian
African v South Asian	TCC>T[abc]	-0.179	1.7E-118	0.115	3.4E-12	
	CTATA>T	-0.507	6.1E-03	0.482	7.4E-03	South Asian
	CCCCCAG>G	-0.818	2.6E-03	0.767	2.7E-03	
	TATATATC>G	0.668	3.3E-03	-0.738	2.2E-03	African
African v East Asian	GCT>T	-0.063	9.1E-03	0.074	2.2E-03	
	CTCGCG>T	1.240	2.8E-03	-1.243	3.6E-03	
	TAAAATA>T	-1.160	3.9E-03	1.135	4.8E-03	
	ATACCTC>A	-1.061	4.6E-03	0.829	5.7E-03	East Asian
	TATATATC>G	0.712	3.9E-04	-0.748	1.3E-04	African
East Asian v South Asian	TCC>T[abc]	-0.200	2.4E-155	0.097	5.4E-05	
	CTATA>T	-0.519	5.3E-03	0.479	7.8E-03	South Asian
	CTCGCG>T	-1.244	2.0E-03	1.247	2.7E-03	
	ATACCTC>A	0.906	8.5E-03	-0.819	9.1E-03	East Asia
	CCCCCAG>G	-0.819	3.8E-03	0.764	4.4E-03	

[a] previously reported in Mathieson and Reich; [b] previously reported in Harris and Pritchard;

[c] previously reported in Aikens et al

821

822

823

824

825

Table 2. Power estimates for 1KG continental private polymorphism probabilities.

abs(log(adj. poly. prob / null poly. prob.))	# contexts sample size percentile	fraction of contexts with >80% power at each mer level						
		3mers	4mers	5mers	6mers	7mers	8mers	9mers
0.01	0-25%	0	0	0	0	0	0	0
	26-50%	0	0	0	0	0	0	0
	51-75%	0	0	0	0	0	0	0
	76-100%	0	0	0	0	0	0	0
0.1	0-25%	0.44	0.11	0	0	0	0	0
	26-50%	0.63	0.04	0	0	0	0	0
	51-75%	0.73	0.03	0	0	0	0	0
	76-100%	0.58	0.10	0	0	0	0	0
0.5	0-25%	1.00	0.92	0.30	0.21	0.01	0	0
	26-50%	1.00	1.00	0.68	0.15	0.01	0	0
	51-75%	1.00	1.00	0.76	0.27	0.02	0	0
	76-100%	1.00	1.00	0.87	0.20	0.03	0	0
1	0-25%	1.00	0.99	0.81	0.34	0.20	0.02	0
	26-50%	1.00	1.00	1.00	0.73	0.23	0.06	0
	51-75%	1.00	1.00	1.00	0.87	0.24	0.04	0
	76-100%	1.00	1.00	1.00	0.87	0.37	0.08	0
1.5	0-25%	1.00	1.00	0.96	0.61	0.25	0.08	0
	26-50%	1.00	1.00	1.00	0.91	0.39	0.18	0.02
	51-75%	1.00	1.00	1.00	0.99	0.59	0.20	0
	76-100%	1.00	1.00	1.00	0.99	0.69	0.26	0

826
827
828 Supplementary Figure Captions
829
830 **S1 Fig. Empirical even odd polymorphism probability scatter plots for the NFE dataset**
831 **including zero mut variants.** Baymer mean posterior estimates for (A) 3-mer models
832 (Spearman correlation 0.999; $p < 10^{-100}$; RMSPE = 0.0009), (B) 5-mer models (Spearman
833 correlation 0.999; $p < 10^{-100}$; RMSPE = 0.0063), (C) 7-mer models (Spearman correlation 0.992;
834 $p < 10^{-100}$; RMSPE = 0.0459), and (D) 9-mer models (Spearman correlation 0.876; $p < 10^{-100}$;
835 RMSPE = 0.7441) in even and odd bp datasets. Polymorphism probabilities in the bottom two,
836 and top left, quadrants correspond to those contexts where no mutations are present for the
837 given mutation type in the respective datasets. These polymorphism probabilities are
838 exclusively calculated using pseudocounts.
839
840 **S2 Fig. Comparison of Baymer mean posterior estimates for differing allele frequency**
841 **(AF) bins in the NFE dataset.** (A-D) AF 0.02-0.05 compared against 0.05-0.15 AF, 0.15-0.30

842 AF, 0.30-0.50 AF, and 0.50-0.85 AF, respectively. (E-G) AF 0.05-0.15 AF compared against
843 0.15-0.30 AF, 0.30-0.50 AF, and 0.50-0.85 AF, respectively. (H-I) 0.15-0.30 AF compared
844 against 0.30-0.50 AF and 0.50-0.85 AF, respectively. (J) 0.30-0.50 AF compared against 0.50-
845 0.85 AF.
846
847 **S3 Fig. Comparison of empirical and Baymer-derived 9-mer polymorphism probabilities**
848 **in NYGC-resequenced 1000 Genomes Phase III non-admixed non-Finnish European**
849 **(EUR) polymorphisms with derived AC >= 2 in non-coding accessible regions.** (A)
850 Empirical 9mer polymorphism probabilities for context mutations with at least 1 occurrence in
851 both datasets (102,875 omitted context mutations) are plotted against one another (Spearman
852 correlation 0.862; RMSPE = 0.175). (B) Baymer mean posterior estimates for 9mer
853 polymorphism estimates in even and odd bp datasets (Spearman correlation 0.986; RMSPE =
854 0.042).

855 **S4 Fig. Overview of the characteristics of edge mutability change dynamics in Baymer**
856 **models of the NFE dataset.** (A) Histogram of the number of edges per 9-mer that were inferred
857 to confidently change polymorphism probabilities (PIP > 0.95). (B) Histogram of the maximum
858 edge size per each 9-mer that was inferred to confidently change polymorphism probabilities
859 (PIP > 0.95). (C) Estimated distributions of phi for each mer size level. (D) The distribution of the
860 fractional differences of each 9-mer mean posterior polymorphism probability with their
861 respective nested 3-mer mean posterior polymorphism probability estimates, partitioned by
862 mutation type.
863
864 **S5 Fig. Fraction overlap of simulated datasets trained by Baymer at varying sequence**
865 **contexts and log changes to the null polymorphism probability.**
866
867 **S6 Fig. Distribution of gnomAD AS_VQSLOD quality scores in non-Finnish European**
868 **samples ("NFE"; A-C) and in all populations ("ALL"; D-F), separated into singletons**
869 **(A,D), doubletons (B,E), and variants with allele count greater than or equal to 3 (C,F).**
870
871 **S7 Fig. Comparison of *Homo sapiens* Baymer model (NFE-2+ model) estimates with *Pan***
872 ***troglodytes* and *Gorilla gorilla* great ape species.** (A) Mean polymorphism estimates of
873 *Homo sapiens* model plotted against mean polymorphism estimates of *Pan troglodytes* model
874 (Spearman correlation 0.957; RMSPE = 0.088). (B) Mean polymorphism estimates of *Homo*
875 *sapiens* model plotted against mean polymorphism estimates of *Gorilla gorilla* model (Spearman
876 correlation 0.950; RMSPE = 0.097). (C) Multinomial likelihoods for each model are calculated
877 on odd bp pan troglodytes test data at various sequence context sizes. *Pan troglodytes* model is
878 trained on even bp data only. (D) Multinomial likelihoods for each model are calculated on odd
879 bp gorilla gorilla test data at various sequence context sizes. *Gorilla gorilla* model is trained on
880 even bp data only. Polymorphism probability estimates were linearly scaled to match the mean
881 polymorphism probability of the holdout dataset.
882
883
884 Supplementary Table Names
885

S1 Table. Calibration of credible sets across mer levels by measuring number of simulations capturing the truth value.

S2 Table. Likelihood of even bp NFE Baymer models on 9-mer odd bp holdout NFE data.

S3 Table. Top 5 largest mean posterior phi estimates for each context size.

S4 Table. Motifs tested for enrichment in the top or bottom 1% of polymorphism probabilities for each mutation type.

S5 Table. Baymer modeled 1KG private continental context mutations with extreme polymorphism probability differences

S6 Table. Sample sizes and test likelihoods for each de novo comparison dataset.

S7 Table. Sample sizes and test likelihoods for each great ape comparison test.

S8 Table. Run-times for models with increasing sequence context windows.

Supplementary Text

S1 Text. Supplementary Methods

S2 Text. Computational Considerations

S3 Text. Data Availability