

Supplementary Information for “Reference-based phasing using the Haplotype Reference Consortium panel”

Po-Ru Loh, Petr Danecek, Pier Francesco Palamara, Christian Fuchsberger, Yakir A Reshef, Hilary K Finucane, Sebastian Schoenherr, Lukas Forer, Shane McCarthy, Goncalo R Abecasis, Richard Durbin, Alkes L Price

Contents

Supplementary Note: Eagle2 algorithm	3
1 Background	3
1.1 Existing phasing methods	3
1.2 New approach: Eagle2	4
2 Eagle2 core phasing algorithm	6
2.1 Diploptype probability model	6
2.1.1 Distribution of recombined haplotypes given reference haplotypes	6
2.1.2 Distribution of genotypes given a diploptype	8
2.2 Algorithm overview: Fast computation and focused search	8
2.2.1 Efficiently computing haploptype probabilities	9
2.2.2 Exploring only the likeliest diploptypes	10
2.2.3 Computational cost	12
2.3 Formal inference procedure	12
2.3.1 Approximating marginal posteriors using ensembles of diploptypes	12
2.3.2 Propagating diploptypes: Branching, merging, pruning, and calling	13
2.3.3 Refining phase via constrained search	14
2.4 Optimizations	14
2.4.1 Condensing blocks of markers homozygous in the target sample	14
2.4.2 Creating individual-specific condensed haploptype trees	15
3 Eagle2 missing genotype imputation algorithm	17
3.1 Recursion for posterior probability of a haploptype	17
3.2 Beam search for locally best haplotypes	18
3.3 Forward-backward algorithm for imputation	19
3.4 Optimizations	20

4	HapHedge data structure	21
4.1	Compact representation of haplotype prefix trees in memory	21
4.2	Constant-time haplotype prefix extension	22
4.3	Linear-time HapHedge construction	23
4.4	Extension to haplotypes with multiplicities: HapHedgeMulti	23
5	Appendix: Probability of recombining off a shared haplotype	24
	References	26
	Supplementary Figures	29
	Supplementary Tables	31

Supplementary Note: Eagle2 algorithm

1 Background

1.1 Existing phasing methods

To our knowledge, all general-purpose methods for genome-wide population-based phasing that have enjoyed widespread use in the past decade are built on the Li-Stephens haplotype-copying HMM [20]. Very simply, given a set of K reference haplotypes, the Li-Stephens model generates a new haplotype as an imperfect mosaic of reference haplotype segments. The states of the model (K states per SNP) designate reference haplotypes, and the Markov process moves left-to-right across a chromosome, choosing at each successive SNP whether to stay on the current reference haplotype or jump to a new haplotype (the transitions) and whether to emit the reference haplotype's allele or a mutant allele (the emissions). Thus, the Li-Stephens HMM has the virtues of (approximately) modeling haplotype frequencies, mutation, and recombination. Moreover, it naturally extends to a generative model for diplotypes (i.e., pairs of haplotypes) assuming independence of maternally- and paternally-derived haplotypes [5].

The diploid version of the Li-Stephens model is still Markovian and hence lends itself to phase inference (for a given sample conditional on either a phased reference panel or estimated haplotypes for other samples) via forward-backward, Viterbi, or Gibbs sampling approaches. The model was first used (very successfully) for phasing in the PHASE v2 algorithm [5]. However, as data sizes increased, the HMM computations for this model quickly became intractable; even with recursive computation, calculation of forward and backward probabilities requires $O(MK^2)$ time for M SNPs and K reference haplotypes [6]. This computational bottleneck led to a series of innovations over the past decade, all of which involve some type of approximation:

- fastPHASE [6] performs clustering on haplotypes and approximates each cluster with a single state, reducing K from the number of reference haplotypes to the number of clusters.
- Beagle [7, 9] applies a localized haplotype clustering approach [31, 32], reducing the state space by approximating the HMM as a haplotype graph and carefully estimating transition probabilities between clusters.
- HAPI-UR [11] operates on windows of dozens of SNPs at a time, collapsing haplotypes identical within a window into a single state (thus approximating the HMM), and discarding states inconsistent with the diploid individual being phased.
- SHAPEIT [10, 12, 14] conditions on only a subset of K locally best reference haplotypes (in version 2 and higher), building an individual-specific HMM on 3-het haplotype clusters and

performing Gibbs sampling in $O(MK)$ time per MCMC iteration.

- Eagle1 [13] (in its final HMM-like step) analyzes a rudimentary HMM, using a fast scoring scheme to search for an approximate Viterbi path through local subsets of K haplotypes and applying a branching-and-pruning beam search to limit computational cost to $O(MKP)$, where P is the beam width (i.e., the number of haplotype pairs retained after pruning).

Our recent work on Eagle1 [13] aimed to improve phasing speed at very large sample sizes by using a hybrid approach of long-range phasing [8] followed by approximate HMM-based phasing. (We note that long-range phasing alone suffices—and in fact should yield better accuracy than HMM-based inference—when a sizable fraction of a population has been genotyped; however, this threshold has only been reached in Iceland [8].) We demonstrated that when used to phase $N=150K$ UK Biobank samples, Eagle1 achieved 1–2 order of magnitude speedups over existing methods at equal or greater accuracy. However, for $N<50K$ (and for underrepresented ethnicities in larger data sets), we observed that Eagle1 achieved suboptimal accuracy, as expected given its makeshift HMM: our HMM approximation was good enough to quickly identify strong phase evidence from sharing of long haplotypes, but not good enough to properly evaluate regions with less-certain phase.

1.2 New approach: Eagle2

Given the above observation, we were motivated to develop a completely new statistical phasing algorithm, Eagle2, that can attain the speed of Eagle1 without losing accuracy at smaller sample sizes. Unlike Eagle1, Eagle2 analyzes a full probabilistic model similar to the diploid Li-Stephens model used by existing methods (but relaxing the Markovian assumption). The key difference between Eagle2 and previous methods is that whereas previous approaches *approximate the haplotype structure* (e.g., by merging haplotypes into local clusters) to produce a more tractable HMM, Eagle2 (1) *efficiently represents the full haplotype structure* in a way that losslessly condenses locally matching haplotypes; and (2) using this representation, *selectively explores the diplotype space* in a way that only expends computation on the most likely phase paths (i.e., diploypes with highest posterior probabilities). We achieve (1) by introducing a new data structure, the HapHedge, which requires only $O(MK)$ time to generate using the positional Burrows-Wheeler transform (PBWT) [19], and we achieve (2) by using a branching-and-pruning beam search (similar in spirit to the beam search used in Eagle1 [13] but very different in implementation).

Intuitively, the main advantage of performing approximate inference by limiting exploration of the diplotype space vs. simplifying the model itself is that the first approach (which we adopt in Eagle2) has the potential to lose little or no accuracy if most phase paths are extremely unlikely. The same can be said for the usual approach of reducing model complexity *only if* the model is

simplified in a manner that eliminates unlikely phase paths without accidentally losing (or damping) likely phase paths. SHAPEIT2 [12] seems to achieve this behavior by building its HMMs using small sets of locally-best references selected in an individual-specific manner, but in general, simplifying the HMM in a way that does not lose accuracy appears to be quite challenging (based on our benchmarks against existing methods). In particular, methods that create a single simplified HMM by locally clustering haplotypes and then modeling transition probabilities between consecutive clusters necessarily lose longer-range information.

A second advantage of the Eagle2 approach is that it no longer relies on standard forward-backward HMM recursions for computational tractability, and as such, allows us to introduce non-Markovian recombination probabilities that more accurately model the coalescent process. Specifically, as we move along the genome extending a shared haplotype, the probability of recombining off the shared haplotype (conditional on the length shared so far) should decrease as the sharing length grows. HMMs are unable to model this history-dependent behavior (unless the state space is augmented with an age parameter at the expense of increased model complexity), whereas Eagle2 has no trouble doing so. Empirically, we have observed that this modeling improvement reduces error rates by a few percent.

To be fair, the Eagle2 approach of exploring only a small portion of the diplotype space does have the theoretical downside of being heuristic: by choosing to consider only part of the space, we are forced to adopt less-rigorous approximate inference procedures and lose the theoretical guarantees of HMM posterior decoding or sampling. Having to make this trade-off is not ideal, but given our empirical evidence supporting the accuracy of Eagle2, its gains in speed over existing methods, and the fact that other methods make theoretically challenging approximations at the level of HMM-building, we believe the trade-off is worthwhile.

Thus far we have compared methods based on the core algorithm that they use to phase a single sample using a given set of reference haplotypes. In practice, input for phasing software typically contains many samples to be phased and may or may not include a phased reference panel. Procedures for handling these situations are fairly straightforward; e.g., algorithms typically iterate over samples, (re-)estimating haplotypes for each sample using the current haplotype estimates for all other samples along with any provided reference haplotypes. The precise top-level procedure that our software uses when analyzing a many-sample data set is described in Online Methods. In the remainder of this note, we stay one level lower and detail the core subroutines that Eagle2 applies to phase a single sample given a reference set and to impute missing genotypes in the same scenario. We then detail the core HapHedge data structure that makes these algorithms possible.

2 Eagle2 core phasing algorithm

Here we describe the mathematical model and core algorithm that Eagle2 uses to phase a single sample given a set of reference haplotypes.

2.1 Diplotype probability model

We consider a single “target” sample with diploid genotypes $g_1, \dots, g_M \in \{0, 1, 2\}$ at M biallelic markers. We wish to model the distribution over possible diplotypes—i.e., pairs of maternally- and paternally-derived haplotypes $h_1^{\text{mat}}, \dots, h_M^{\text{mat}} \in \{0, 1\}$ and $h_1^{\text{pat}}, \dots, h_M^{\text{pat}} \in \{0, 1\}$ —that could underlie the observed genotypes g_1, \dots, g_M , taking into account population haplotype structure, recombination, and the possibility of single-SNP discrepancies (due to genotyping error, mutation, or gene conversion, any of which could result in $g_m \neq h_m^{\text{mat}} + h_m^{\text{pat}}$). We assume we are given a set of K (phased) reference haplotypes $\{h_1^k, \dots, h_M^k\}_{k=1}^K$ with which to model population haplotype structure. For notational convenience, we will use colons to abbreviate lists, e.g., $g_{1:M}$ to abbreviate g_1, \dots, g_M , from here on.

To model each of the above features of haplotype distributions, we use the general haplotype-copying-with-error approach. Specifically,

- we put a prior $P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$ on diplotypes $(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$ based on the process of copying reference haplotypes with recombination,
- we put a distribution $P(g_{1:M} | h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$ on observed genotypes given putative maternally- and paternally-derived haplotypes (to model single-SNP discrepancies), and
- we infer (relative) posterior probabilities $P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}} | g_{1:M})$ using Bayes’ rule:

$$P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}} | g_{1:M}) \propto P(g_{1:M} | h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}}) \cdot P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}}). \quad (1)$$

We now define the probabilities $P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$ and $P(g_{1:M} | h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$.

2.1.1 Distribution of recombined haplotypes given reference haplotypes

To write down a prior $P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$ on diplotypes, we first make the independence approximation

$$P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}}) \approx P(h_{1:M}^{\text{mat}}) \cdot P(h_{1:M}^{\text{pat}}), \quad (2)$$

where the probabilities $P(h_{1:M}^{\text{mat}})$ and $P(h_{1:M}^{\text{pat}})$ for the maternally- and paternally-derived haplotypes come from the same distribution on haplotypes $h_{1:M}$. (We will thus be unable to distinguish the maternally- and paternally-derived haplotypes; however, we retain the notation $h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}}$ for ease of understanding even though all distributions will be symmetric with respect to $h_{1:M}^{\text{mat}}$ and $h_{1:M}^{\text{pat}}$.)

For a single haplotype $h_{1:M}$, we model $h_{1:M}$ as a mosaic of recombined haplotype segments $h_{1:A}, h_{A+1:B}, \dots, h_{L+1:M}$ copied from the reference haplotypes $\{h_{1:M}^k\}_{k=1}^K$. (We have denoted the recombination breakpoints partitioning the sites $1:M$ using the notation A, B, \dots, L —which should be taken to represent a list of arbitrary length—to avoid double subscripts.) We may thus decompose $P(h_{1:M})$ as a sum over possible sets of recombination breakpoints A, B, \dots, L :

$$P(h_{1:M}) = \sum_{A,B,\dots,L} P(h_{1:A}, \text{no rec}) \cdot P(\text{rec } A \mid \text{rec } 0) \cdot P(h_{A+1:B}, \text{no rec}) \cdot P(\text{rec } B \mid \text{rec } A) \cdots P(h_{L+1:M}, \text{no rec}) \cdot P(\text{rec } M \mid \text{rec } L), \quad (3)$$

where $P(h_{X+1:Y}, \text{no rec})$ denotes the probability of copying a given haplotype segment $h_{X+1:Y}$, and $P(\text{rec } Y \mid \text{rec } X)$ denotes the probability of a recombination between markers Y and $Y + 1$ given that the previous recombination occurred between markers X and $X + 1$.

The probability of copying the segment $h_{X+1:Y}$ is given by

$$P(h_{X+1:Y}, \text{no rec}) = f(h_{X+1:Y}) \cdot P(\text{no rec}_{X+1:Y}), \quad (4)$$

where

$$f(h_{X+1:Y}) = \frac{\#\{k : h_{X+1:Y}^k = h_{X+1:Y}\}}{K} \quad (5)$$

is the frequency of the haplotype segment $h_{X+1:Y}$ in the reference. Assuming the segment $X + 1 : Y$ is not extremely long, recombination is unlikely, so we can ignore the term $P(\text{no rec}_{X+1:Y})$ and approximate

$$P(h_{X+1:Y}, \text{no rec}) \approx f(h_{X+1:Y}). \quad (6)$$

As for the recombination probability $P(\text{rec } Y \mid \text{rec } X)$, we first note that the recombination model in the standard Li-Stephens HMM corresponds to the special (Markovian) case in which we assume $P(\text{rec } Y)$ is independent of the copied haplotype length and set

$$P(\text{rec } Y) = 1 - \exp\left(-\frac{4N_e d_Y}{K}\right), \quad (7)$$

where N_e is the effective population size and d_Y is the genetic distance between markers Y and $Y + 1$.

In Eagle2, we do not need to impose a Markovian requirement to achieve computational speed, allowing us to use a recombination probability that decreases with the copying length $d_{X+1:Y}$ (more accurately reflecting the coalescent; see derivation in Section 5):

$$P(\text{rec } Y \mid \text{rec } X) = 1 - \exp\left(-\frac{3d_Y}{\frac{K}{2N_e} + d_{X+1:Y}}\right). \quad (8)$$

2.1.2 Distribution of genotypes given a diplotype

To model single-SNP discrepancies between a diplotype $h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}}$ and the observed genotypes $g_{1:M}$, we use the following simple approximation:

$$P(g_{1:M} \mid h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}}) \propto \epsilon^{n_{\text{err}}}, \quad (9)$$

where ϵ is an “error probability” parameter that roughly represents the chance of an allele modification, and

$$n_{\text{err}} = \sum_{m=1}^M |g_m - (h_m^{\text{mat}} + h_m^{\text{pat}})| \quad (10)$$

is the total number of allele modifications required to transform the diplotype $(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$ to match the observed genotypes $g_{1:M}$. We note that this intuition does not precisely produce equation (9): strictly speaking, we should have an extra factor of 2 when $|g_m - (h_m^{\text{mat}} + h_m^{\text{pat}})| = 1$ (because the discrepancy could occur on either haplotype), and we should also have a term $(1 - \epsilon)^{2M - n_{\text{err}}}$ for non-error probabilities (which we have instead approximated to 1). However, in practice, we have observed that our inference is insensitive to the precise value of ϵ —and in reality, the distributions of genotyping errors, mutations, and gene conversions are much more complex anyway (see next paragraph)—so we adopt the approximation in equation (9) due to its simplicity.

Fleshing out the above point, we note that because our copying model for $h_{1:M}$ as a mosaic of segments $h_{X+1:Y}$ does not allow for mutation, the probability $P(g_{1:M} \mid h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$ must actually account for any mutations or gene conversions that occur along the lineages connecting copied haplotype segments to the target sample. (To be completely precise, in the model we have defined, $h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}}$ are actually best described as “mosaics of surrogate maternal and paternal haplotypes”—each separated from the target sample by some number of generations—rather than the maternally- and paternally-derived haplotypes of the target sample.) In theory, we could therefore improve the accuracy of our model by allowing the error parameter ϵ to be a function of haplotype copying length (which scales inversely with time to coalescence and thus with accumulated mutation rate along a lineage). However, for large reference panels, in which coalescences between target haplotypes and closest reference haplotypes are relatively recent, single-SNP discrepancies are most often produced by genotyping error, so we ignore this subtlety and use a fixed value of ϵ .

2.2 Algorithm overview: Fast computation and focused search

The Eagle2 phasing algorithm applies two main ideas to perform fast and accurate phase inference given the probabilistic model we have just described. The first idea is a new approach to efficient computation of haplotype probabilities under a copying model (in our case, equation (3)).

Naïvely, these computations require exponential time. The standard approach to overcoming this barrier is to observe that within a K_{HMM} -state HMM, recursion allows computation of *all* marginal probabilities (for all K_{HMM} states at each of M positions) in $O(MK_{\text{HMM}}^2)$ time. In Eagle2, we take a completely different recursive approach that computes the probability of a *single* haplotype $P(h_{1:M})$ in $O(M)$ time—independent of the number of reference haplotypes K —after creation of a new data structure, the HapHedge, in $O(MK)$ time. The HapHedge essentially consolidates all reference haplotypes sharing a common prefix (starting at any given position) into a single atom of data, thus eliminating future computation that scales with K .

Of course, being able to very rapidly compute the probability of a single haplotype is only useful if we can identify a small subset of haplotype probabilities that are worth computing; to this end, we need a second key idea. We perform a beam search from left to right across the genome, propagating a small set of P likely diplotypes that ideally represent most of the posterior probability mass in the full diplotype space. This approach essentially focuses computational effort on a small subset of the diplotype space (vs. expending computation evenly across the space as in HMM recursion), which is advantageous when most of the space is unlikely but difficult to discard *a priori*.

We now expand on each of these ideas; we fully describe the haplotype probability computation (aside from specification of the HapHedge, which we defer to Section 4), and we provide further details and intuition for the beam search, which we formally treat in Section 2.3.

2.2.1 Efficiently computing haplotype probabilities

Combining equations (3) and (6), we wish to compute

$$P(h_{1:M}) = \sum_{A,B,\dots,L} f(h_{1:A}) \cdot P(\text{rec } A \mid 0) \cdot f(h_{A+1:B}) \cdot P(\text{rec } B \mid A) \cdots f(h_{L+1:M}) \cdot P(\text{rec } M \mid L). \quad (11)$$

It turns out that each factor of each summand in the above formula can be computed in $O(1)$ time. This statement is obvious for the recombination probabilities $P(\text{rec } Y \mid \text{rec } X)$ (abbreviated as $P(\text{rec } Y \mid X)$ in the above), which are given by equation (8). As for the the haplotype segment frequencies $f(h_{X+1:Y})$, the HapHedge data structure precisely supports these lookups in constant time (assuming we have saved state information from previous lookups of the sub-segments $h_{X+1:Y-1}$, which is always the case for our algorithm).

The only remaining obstacle is the exponential sum over sets of recombination breakpoints A, B, \dots, L , but conveniently, equation (11) naturally admits the recursion

$$P(h_{1:Y}) = \sum_{X=1}^{Y-1} P(h_{1:X}) \cdot f(h_{X+1:Y}) \cdot P(\text{rec } Y \mid X). \quad (12)$$

This recursion yields an $O(M^2)$ dynamic programming algorithm for computing $P(h_{1:M})$.

We further reduce running time by restricting the sum over previous recombination points X to sites within a distance H of the current site Y :

$$P(h_{1:Y}) \approx \sum_{X=Y-H}^{Y-1} P(h_{1:X}) \cdot f(h_{X+1:Y}) \cdot P(\text{rec } Y \mid X). \quad (13)$$

That is, we allow the model to “forget” history older than H positions earlier in the genome, or equivalently, we limit the maximum length of copied haplotypes to H . This approximation reduces the cost of (approximately) computing $P(h_{1:M})$ to $O(MH)$. We set H to be the distance spanned by 100 heterozygous sites in the target sample; this choice of H corresponds to a few cM at the density of typical genotyping arrays. (In Section 2.4, we will describe optimizations that condense sites between consecutive target hets, hence the choice of this positional unit.)

2.2.2 Exploring only the likeliest diplotypes

The computational tricks above allow $O(MH)$ -time estimation of $P(h_{1:M})$ for any $h_{1:M}$ and thus also $P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$ and the desired posterior $P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}} \mid g_{1:M})$ for any given diplotype. We are left with the challenge of determining which diploypes to evaluate, which we treat as a search problem. That is, we build diploypes left-to-right across a chromosome, and at each successive marker, for each diplotype, we consider extending each of its haplotypes to either the 0 or 1 allele at the next marker m . This branching procedure increases the number of diploypes under consideration by a factor of 4, so to prevent an exponential explosion, we then prune the set of diploypes back down to at most P diploypes $\{(h_{1:m}^{\text{mat}(p)}, h_{1:m}^{\text{pat}(p)})\}_{p=1}^P$, keeping the most likely diploypes according to posterior probabilities $P(h_{1:m}^{\text{mat}}, h_{1:m}^{\text{pat}} \mid g_{1:m})$ based on information thus far (i.e., markers $1:m$). The total cost of the entire procedure is $O(MHP)$. The intuition behind the beam search is that it greedily prioritizes likely diploypes, and while future information may show that some of these diploypes have only low-probability extensions, the hope is that the beam will be wide enough (i.e., contain enough diploypes) to be robust to such local maxima.

The procedure that we have described is a good start, but we should be a bit skeptical: How can we expect a small number of greedily-selected diploypes to adequately represent the exponential space of phase paths, especially at the scale of entire chromosomes? Even if the posterior distribution on diploypes is locally well-concentrated, the space will inevitably grow exponentially, which causes two opposite types of problems for our beam search: parallelism and degeneracy.

- **Pitfall 1: Parallelism.** This problem amounts to a lack of beam diversity in the recent past (i.e., when restricted to the preceding local region of the genome): the beam can fragment into several paths with similar probabilities and then follow each path along parallel trajectories (i.e., identical phase paths after the split). This behavior is in fact highly likely if the

search space is locally well-constrained (as it tends to be), and it is problematic because fragmentation of the beam into parallel paths effectively reduces its width and thus its ability to overcome local maxima.

- **Pitfall 2: Degeneracy.** This problem is a lack of beam diversity in the more distant past. Our greedy search prioritizes best paths in an exponentially growing space, pruning away less-likely paths; as such, the set of P diplotypes that comprise the beam front at any given position will (hopefully) be a good representation of the diversity of paths through recent positions. However, if we trace these P diplotypes back toward the start of the chromosome, they will inevitably collapse into a single trajectory, providing a poor representation of the marginal posterior on diplotypes at positions far behind the current position.

To overcome these pitfalls, we need to exercise greater finesse as we conduct our beam search and perform phase inference. We make three main adjustments:

- **Solution 1: Merging.** To deal with parallelism, when we observe that two diplotypes under consideration agree at the last 64 target hets, we merge the diplotypes (by pruning one diplotype and adding its “weight” to the other diplotype; we defer a precise definition of this procedure to the formal exposition in Section 2.3).
- **Solution 2: Fixed-lag smoothing.** To deal with degeneracy, we do not attempt to use the final beam front of full diplotypes $(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}})$ to call phase throughout the chromosome. Instead, we call the relative phase of a pair of consecutive hets using only a limited amount of future (forward) information, a procedure similar to fixed-lag smoothing from signal processing. Specifically, at any position m , we use the P diplotypes $\{(h_{1:m}^{\text{mat}(p)}, h_{1:m}^{\text{pat}(p)})\}_{p=1}^P$ in the beam front to call phase at a distance $\Delta=20$ hets earlier (by examining the marginal posterior distribution at the 19th- and 20th-most recent hets). Effectively, this choice means that at any position m , we need the beam front to only be a good approximation of the marginal posterior over “recent history” in the interval $[m - \Delta, m]$.
- **Solution 3: Refinement via constrained search.** Finally, to improve the ability of the beam search to overcome local maxima, we apply a two-step search. The first step is to run the beam search described above; however, rather than simply making phase calls during this run, we also estimate marginal posteriors for each consecutive pair of hets, thus obtaining phase confidences. We then rank the phase confidences from most to least confident, identify high-confidence phase calls, and run a second beam search in which we constrain phase at high-confidence het pairs to match the first-iteration calls. The idea is that these phase constraints are “information passed back from the future” that make the second-iteration search better able to avoid obstacles while focusing attention on only the most difficult phase calls.

2.2.3 Computational cost

The total computational cost of the algorithm we have described for beam propagation and phase inference algorithm is $O(MHP)$, where M is the number of markers, H is the history parameter governing maximum copying length in the mosaic model, and P is the number of paths in the beam front. We also need to take into account the $O(MK)$ cost of building the HapHedge data structure on K reference haplotypes—but if we were to implement the exact algorithm described thus far, this data structure would only need to be built once (rather than once per target individual).

Because the constant factor on HapHedge construction is quite small, we have chosen to instead build individual-specific condensed HapHedge structures on subsets of $K = 10,000$ best reference haplotypes per individual; doing so has the advantage of allowing the beam search to proceed one het at a time rather than one marker at a time, thus reducing the constant factor on the $O(MHP)$ term. Given a panel of $2N$ reference haplotypes as input, our algorithm for selecting the best K reference haplotypes for a given target sample currently runs in $O(MN)$ time with a very small constant factor via bit-level parallelism; thus, the overall procedure takes $O(MN + MHP)$ time with small constant factors on both terms. Details of these optimizations are given in Section 2.4.

2.3 Formal inference procedure

We now present a formal exposition of our beam search-based inference procedure that implements the preceding intuition.

2.3.1 Approximating marginal posteriors using ensembles of diplotypes

For each $m = 1, \dots, M$, we wish to create an ensemble of $\leq P$ weighted diplotypes, denoted $\{(w^{(p)}, h_{1:m}^{\text{mat}(p)}, h_{1:m}^{\text{pat}(p)})\}_p$, with the property that the truncation of the ensemble to the “recent past” $[m - \Delta, m]$ provides a good approximation of the true marginal posterior on diplotypes given genotypes $g_{1:m}$ “seen so far.” That is, we want:

$$\begin{aligned}
 & P(h_{m-\Delta:m}^{\text{mat}}, h_{m-\Delta:m}^{\text{pat}} \mid g_{1:m}) \\
 & \propto \sum_p w^{(p)} \cdot \hat{P}(h_{1:m}^{\text{mat}(p)}, h_{1:m}^{\text{pat}(p)} \mid g_{1:m}) \cdot \mathbf{1}(h_{m-\Delta:m}^{\text{mat}(p)}, h_{m-\Delta:m}^{\text{pat}(p)} = h_{m-\Delta:m}^{\text{mat}}, h_{m-\Delta:m}^{\text{pat}}),
 \end{aligned} \tag{14}$$

where the \hat{P} term indicates the approximate posterior we described how to compute in equations (1), (2), (9), and (13), and $\mathbf{1}(\cdot)$ denotes the indicator function. (We could have absorbed these terms into the weights $w^{(p)}$, but we wish to view the weights as diplotype “multiplicities,” which will be intuitive when we later discuss merging diplotypes.)

Is there hope of achieving the approximation suggested in equation (14) using only a small ensemble of diplotypes? The feasibility of this task is essentially a question of whether or not the

posterior $P(h_{m-\Delta:m}^{\text{mat}}, h_{m-\Delta:m}^{\text{pat}} \mid g_{1:m})$ is concentrated on a small support—which should be the case if phase is relatively well-constrained over the interval $[m - \Delta, m]$ (i.e., only a few phase paths through this interval are likely).

2.3.2 Propagating diplotypes: Branching, merging, pruning, and calling

We construct the ensembles $\{(w^{(p)}, h_{1:m}^{\text{mat}(p)}, h_{1:m}^{\text{pat}(p)})\}_p$ for $m = 1, \dots, M$ by propagating a “beam front” across the chromosome. We initialize the beam front at $m = 0$ with a single diplotype on no markers (containing no haplotype information) and proceed left to right, creating beam front m (i.e., the m -th ensemble) from beam front $m - 1$ according to the following procedure.

Branching. For each diplotype $\{(w^{(p')}, h_{1:m-1}^{\text{mat}(p')}, h_{1:m-1}^{\text{pat}(p')})\}_{p'}$ in ensemble $m - 1$ (indexed by p'), we create its four possible extensions by appending $h_m^{\text{mat}} \in \{0, 1\}$ and $h_m^{\text{pat}} \in \{0, 1\}$ to $h_{1:m-1}^{\text{mat}(p')}$ and $h_{1:m-1}^{\text{pat}(p')}$. (We consider all four possible extensions to allow for single-SNP discrepancies due to genotyping error, etc.) We give its weight $w^{(p')}$ to all four possible extensions (for now), and we compute the posterior probabilities $\hat{P}(h_{1:m-1}^{\text{mat}(p')}, h_m^{\text{mat}}, h_{1:m-1}^{\text{pat}(p')}, h_m^{\text{pat}} \mid g_{1:m})$ in $O(H)$ time by making use of the recursion in equation (13), looking up previously computed values (i.e., applying dynamic programming).

Merging and pruning. We now have up to $4P$ diploypes on $1:m$ and need to reduce to an ensemble of size $\leq P$. To do so, we employ two techniques: merging and pruning. First, if two weighted diploypes $(w^{(p)}, h_{1:m}^{\text{mat}(p)}, h_{1:m}^{\text{pat}(p)})$ and $(w^{(q)}, h_{1:m}^{\text{mat}(q)}, h_{1:m}^{\text{pat}(q)})$ are identical on the interval $[m - \Delta, m]$, i.e.,

$$h_{m-\Delta:m}^{\text{mat}(p)}, h_{m-\Delta:m}^{\text{pat}(p)} = h_{m-\Delta:m}^{\text{mat}(q)}, h_{m-\Delta:m}^{\text{pat}(q)}, \quad (15)$$

then we *merge* them by eliminating diplotype q and giving its weight to diplotype p (after adjusting for the relative posterior probabilities of diploypes p and q). Explicitly, we augment the weight, or “multiplicity,” of diplotype p as follows:

$$w^{(p)} \leftarrow w^{(p)} + w^{(q)} \cdot \frac{\hat{P}(h_{1:m}^{\text{mat}(q)}, h_{1:m}^{\text{pat}(q)} \mid g_{1:m})}{\hat{P}(h_{1:m}^{\text{mat}(p)}, h_{1:m}^{\text{pat}(p)} \mid g_{1:m})}. \quad (16)$$

Second, we sort all remaining diploypes by weighted posterior probability $w^{(p)} \cdot \hat{P}(h_{1:m}^{\text{mat}(p)}, h_{1:m}^{\text{pat}(p)} \mid g_{1:m})$ and *prune* the ensemble to the top P weighted diploypes. Additionally, we prune any diploypes with weighted probability less than ϵ^2 times that of the leading diplotype, where ϵ is the error probability parameter for single-SNP discrepancies (so ϵ^2 corresponds to the penalty of two errors).

Calling phase. For any pairs of consecutive heterozygous sites u, v in the interval $[m - \Delta, m]$, we estimate the relative phase of those sites (overwriting any previous estimates) by using the 2-dimensional marginal of the ensemble at the pair of positions (u, v) :

$$P(h_{u:v}^{\text{mat}}, h_{u:v}^{\text{pat}} \mid g_{1:m}) \propto \sum_p w^{(p)} \cdot \hat{P}(h_{1:m}^{\text{mat}(p)}, h_{1:m}^{\text{pat}(p)} \mid g_{1:m}) \cdot \mathbf{1}(h_{u:v}^{\text{mat}(p)}, h_{u:v}^{\text{pat}(p)} = h_{u:v}^{\text{mat}}, h_{u:v}^{\text{pat}}). \quad (17)$$

2.3.3 Refining phase via constrained search

The procedure that we have described above works quite well, but we can do better via the following two-pass scheme. We first run the propagation algorithm described above (with reduced computational parameters $H=30$, $P=30$, and $\Delta=10$), obtaining phase calls and confidences at all pairs of consecutive hets and also obtaining error probabilities at homozygous genotypes (by marginalizing to single homozygous sites). We then identify phase calls and homozygous genotypes of confidence $>99\%$, after which we perform a second, more accurate run of the propagation algorithm with increased computational parameters ($H=100$, $P=50$, and $\Delta=20$) but much more limited branching: instead of considering all four haplotype extension combinations at each step, we consider only extensions consistent with the high-confidence inferences from the first step.

2.4 Optimizations

Having described the core theoretical components of the Eagle2 phasing algorithm, we now discuss a few key optimizations that achieve important constant-factor speedups in practice.

2.4.1 Condensing blocks of markers homozygous in the target sample

Thus far, we have framed the phase inference problem in terms of computing posterior probabilities $P(h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}} \mid g_{1:M})$, but we actually only wish to output inferred phase at markers m for which the target sample is heterozygous, i.e., $g_m = 1$. We note that the parental haplotypes $h_{1:M}^{\text{mat}}, h_{1:M}^{\text{pat}}$ are not completely determined when $g_m \in \{0, 2\}$ because of the possibility of single-SNP errors. However, we have found that in practice, such errors are sufficiently strongly probabilistically disfavored that simply forcing copied haplotypes to agree with the target allele at homozygous sites results in almost no loss of accuracy (which seems reasonable given that alternative reference haplotypes without errors are usually available).

We make use of this observation to “horizontally” condense the haplotype space conditional on the genotypes of a given target sample. Specifically,

- we retain full haplotype information $h_{s(1:T)}$ at a subset of markers $\{s(1), \dots, s(T)\} \subseteq \{1, \dots, M\}$ containing all target-het sites (plus “spacer” sites that maintain a maximum dis-

tance of 0.5cM between consecutive sites $s(t), s(t + 1)$ to avoid excessive coarsening of the model in runs of homozygosity), and

- within each interval $(s(t), s(t + 1))$, we record only 1 bit e_t that indicates whether or not discrepancies exist between the haplotype segment $h_{s(t)+1:s(t+1)-1}$ and the (homozygous) target genotypes $g_{s(t)+1:s(t+1)-1}$.

This procedure produces a condensed representation $\{e_0 h_{s(1)} e_1 \cdots h_{s(T)} e_T\}$ of the reference haplotypes (Supplementary Fig. 1), as we describe in Section 2.4.2, but more importantly, it also enables a more efficient inference algorithm that skips forward one het at a time. In target-homozygous intervals $(s(t), s(t + 1))$, we simply need to check the bit e_t : if $e_t = 1$, then we must end the haplotype segment currently being copied (and pay the transition cost of a recombination).

2.4.2 Creating individual-specific condensed haplotype trees

Performing inference on condensed haplotypes as described above substantially reduces the number of propagation iterations (from the number of markers M to roughly the number of hets in the target sample). However, because the condensed representation of the reference haplotypes is target sample-dependent, we need to rapidly generate a new HapHedge (on the condensed reference) for each target sample analyzed, and we also need each such HapHedge to be compact (because during multithreaded computation, each thread will need to build its own HapHedge for the target sample it is analyzing). Both of these goals are reasonable given the $O(MK)$ -time and memory cost of building a HapHedge on K reference haplotypes, but a few optimizations help reduce the constant factors.

Representing condensed haplotypes in a HapHedge. As described above, a condensed haplotype $e_0 h_{s(1)} e_1 \cdots h_{s(T)} e_T$ directly corresponds to a bit string in which we alternately represent each actual allele $h_{s(t)}$ with 1 bit and each error bit e_t with 1 bit (Supplementary Fig. 1). Thus, the binary representation of a condensed haplotype requires ≈ 2 bits per target het.

Converting reference haplotypes into condensed bit string representations gives a condensed bit array; we next wish to build a “hedge” of haplotype prefix trees on this array rooted at every other bit (specifically, at the bits directly encoding alleles $h_{s(t)}$). We could directly build a HapHedge on the bit array, but one additional optimization saves time and memory. Given that we prohibit extension of a haplotype through any target-homozygous interval containing one or more discrepancies, we may truncate a prefix upon reaching an error bit $e_t = 1$. Upon truncating prefixes in this manner, many reference haplotypes collapse to the same prefix (for a given root location). Consequently, we use a slightly different data structure (HapHedgeMulti, described in Section 4.4) that efficiently represents trees of haplotype prefixes with multiplicities.

Masking double-IBD regions. Previous authors have observed that some care is needed in scenarios where (i) the $2N$ reference haplotypes consist of imperfectly phased haplotype pairs from N individuals *and* (ii) the target sample contains “double-IBD” regions in which both chromosomal segments are shared identical-by-descent with a reference sample, such that the diploid genotypes of the reference and target samples match exactly in the region [12]. The first situation is guaranteed to arise when performing phasing without an external reference panel (because in this case, the “reference haplotypes” are simply current haplotype estimates for other samples). The second situation typically arises due to relatedness. Whatever the cause of this phenomenon, the result is that a model built on all reference haplotypes as-is will likely copy both double-IBD reference haplotypes when phasing the target sample, bringing along any phasing errors.

We apply the following procedure to guard against this pitfall. For a given target sample, we pre-process each pair of reference haplotypes (derived from a single individual) from left to right, identifying long regions of exact diploid genotype agreement. When such a region spans >20 split points $s(t)$, we mask both reference haplotypes throughout the double-IBD region (by setting error bits $e_t = 1$, which forces a switch upon entering the region). When such a region spans 10–20 split points $s(t)$, we mask the reference haplotype with the more recent error $e_t = 1$.

Selecting K best reference haplotypes. All of the computation that we have described in the above paragraphs (to create a HapHedgeMulti data structure on a set of masked condensed reference haplotypes) scales linearly in the number of reference haplotypes and the number of markers. This computation is quite efficient, but as described in Section 2.2.3, when the number of reference haplotypes ($2N$) is very large, restricting to a subset of K individual-specific conditioning haplotypes saves time. To identify K “best” conditioning haplotypes, we rank reference haplotypes by computing the number of discrepancies between each reference haplotype and the homozygous genotypes of the target sample. As in our previous work [13], we perform computation on blocks of up to 64 SNPs at once using bit arithmetic; thus, the total computational cost of subset selection is $O(MN)$ with a very small constant factor (plus $O(N \log N)$ for sorting, which is negligible in practice).

We note that our discrepancy metric does not make use of inferred phase of the target genotypes and produces a single set of conditioning haplotypes to use for the entire region being phased; as such, our method for selecting conditioning haplotypes is much less sophisticated than that of SHAPEIT2 [12]. However, Eagle2 is able to condition on thousands of haplotypes (vs. $K=100$ by default for SHAPEIT2), which we suspect makes selection of conditioning haplotypes much less important. We also note that our $O(MN)$ cost for selection of conditioning haplotypes is asymptotically slower than the method used by SHAPEIT3 (based on divisive clustering) [14]. Improving the scaling is a direction for future work (geared at million-sample data sets); for now, the constant factor is so small that this cost negligible for $N < 100K$.

3 Eagle2 missing genotype imputation algorithm

The phasing algorithm we described in the previous section ignores genotype data at markers for which the target genotype is missing. This behavior is ideal for the task of phasing a single target sample using a reference panel, but if we wish to perform phasing without a reference (or if we wish to use phased target samples to aid the phasing of other target samples), we also need to impute missing genotypes.

To this end, we have implemented a separate algorithm that imputes missing genotypes in a phased target haplotype (with some missing data)—e.g., output from the Eagle2 core phasing algorithm above—given a panel of phased reference haplotypes (with no missing data). We note that abstractly, this task has the same structure as the standard problem of GWAS imputation using a reference panel after pre-phasing [26]. However, the engineering considerations involved in imputing a very small fraction of missing sites among genotyped SNPs vs. imputing a very large fraction of sequenced but untyped SNPs are very different; as such, the algorithm that we describe below should not be viewed as a GWAS imputation method.

3.1 Recursion for posterior probability of a haplotype

We consider a haplotype $h_{1:M}^{\text{miss}}$ containing missing data at some markers, i.e., for each marker m , $h_m^{\text{miss}} \in \{0, 1, *\}$, where $*$ indicates an unknown allele. (The Eagle2 phasing algorithm described in Section 2 produces a pair of such haplotypes when given a target sample with missing genotypes; the non-missing genotypes are phased, and the missing genotypes are left unknown.) We wish to infer an imputed haplotype $h_{1:M}$ in which each missing allele is filled in with its most likely value (i.e., its marginal posterior rounded to 0 or 1). To perform this inference, we again use Bayes' rule:

$$P(h_{1:M} | h_{1:M}^{\text{miss}}) \propto P(h_{1:M}^{\text{miss}} | h_{1:M}) \cdot P(h_{1:M}). \quad (18)$$

We denote the unnormalized posterior on the right side by $q(h_{1:M})$, i.e.,

$$q(h_{1:M}) := P(h_{1:M}^{\text{miss}} | h_{1:M}) \cdot P(h_{1:M}). \quad (19)$$

Following equations (3)–(6), we may decompose $q(h_{1:M})$ as a sum over possible sets of recombination breakpoints A, B, \dots, L , the only difference being that we now also decompose the term

$P(h_{1:M}^{\text{miss}} | h_{1:M})$:

$$q(h_{1:M}) \approx \sum_{A,B,\dots,L} \underbrace{P(h_{1:A}^{\text{miss}} | h_{1:A})f(h_{1:A})P(\text{rec } A | 0)}_{1:A} \cdot \underbrace{P(h_{A+1:B}^{\text{miss}} | h_{A+1:B})f(h_{A+1:B})P(\text{rec } B | A)}_{A+1:B} \cdots \underbrace{P(h_{L+1:M}^{\text{miss}} | h_{L+1:M})f(h_{L+1:M})P(\text{rec } M | L)}_{L+1:M}. \quad (20)$$

Following equations (9)–(10), we approximate the first term in each triplet with

$$P(h_{X+1:Y}^{\text{miss}} | h_{X+1:Y}) \approx \epsilon^{n_{\text{err}}}, \quad (21)$$

where n_{err} is the number of markers at which $h_{X+1:Y}^{\text{miss}}$ does not match $h_{X+1:Y}$ (ignoring sites at which $h_{X+1:Y}^{\text{miss}}$ contains missing data). The second term, $f(h_{X+1:Y})$, is the frequency of the haplotype $h_{X+1:Y}$ in the reference panel (given in equation (5)), and the third term is the recombination probability from equation (8).

As before, equation (20) admits a natural recursion, allowing us to write

$$q(h_{1:Y}) \approx \sum_{X=Y-H}^{Y-1} q(h_{1:X}) \cdot \underbrace{P(h_{X+1:Y}^{\text{miss}} | h_{X+1:Y})f(h_{X+1:Y})P(\text{rec } Y | X)}_{X+1:Y}, \quad (22)$$

where we have required the previous recombination point X to fall within distance H of the current site Y as in equation (13).

The above recursion allows us to rapidly compute the (unnormalized) posterior probability $q(h_{1:M})$ of any given haplotype $h_{1:M}$, leaving us once again with the problem of deciding which haplotypes to evaluate. *Locally*—i.e., when determining which haplotype *segments* $h_{X+1:Y}$ to consider copying—we apply a branching-and-pruning beam search. However, *globally*, we are able to integrate over all possible combinations of selected haplotype segments via a forward-backward algorithm. We detail each of these algorithmic components below.

3.2 Beam search for locally best haplotypes

Restricting our attention for now to local segments $X + 1 : Y$, we wish to find for each X and Y a set of $\leq P$ candidate haplotypes $h_{X+1:Y}$ with maximal values of $P(h_{X+1:Y}^{\text{miss}} | h_{X+1:Y})f(h_{X+1:Y})$. Exact computation is intractable given that there are 2^{Y-X} possible haplotypes $h_{X+1:Y}$, so instead, we use a heuristic beam search to find approximate top haplotype sets. Explicitly, for each X , we construct sets $\{h_{X+1:Y}^{(p)}\}_p$ for $Y = X + 1, X + 2, \dots, X + H$ in turn using the following iterative procedure. Given the set $\{h_{X+1:Y-1}^{(p')}\}_{p'}$, we create the two possible extensions of each

haplotype $h_{X+1:Y-1}^{(p')}$ by appending $h_Y \in \{0, 1\}$. For each extension $h_{X+1:Y-1}^{(p')}h_Y$, we then compute the objective function $P(h_{X+1:Y}^{\text{miss}} | h_{X+1:Y-1}^{(p')}h_Y)f(h_{X+1:Y-1}^{(p')}h_Y)$ in $O(1)$ time (making use of the previously computed value $P(h_{X+1:Y-1}^{\text{miss}} | h_{X+1:Y-1}^{(p')})$ and the saved HapHedge state information for $h_{X+1:Y-1}^{(p')}$). Finally, we sort the extended haplotypes accordingly and prune to $\leq P$ top haplotypes $h_{X+1:Y}^{(p)}$, also pruning any haplotypes with posterior probability less than ϵ^2 times less that of the leading haplotype.

For a beam width of $\leq P$ haplotypes, the computational cost of each branching-and-pruning iteration is $O(P)$. This iteration must be performed H times for each of M starting locations X , giving a total cost of $O(MHP)$ to perform the complete set of beam searches needed to generate all candidate haplotype sets $\{h_{X+1:Y}^{(p)}\}_p$ with $X + 1 \leq Y \leq X + H$.

3.3 Forward-backward algorithm for imputation

Having identified a library of candidate haplotype segments $\{h_{X+1:Y}^{(p)}\}_p$, we wish to perform posterior inference on the set of all possible stitchings of these segments into haplotypes $h_{1:M}$. More precisely, for every set of breakpoints A, B, \dots, L , we consider all composite haplotypes $h_{1:M}$ formed by concatenating haplotype segments $h_{1:A}, h_{A+1:B}, \dots, h_{L+1:M}$ chosen from the candidate sets for these intervals. Each such composite haplotype has a posterior probability given $h_{1:M}^{\text{miss}}$, and we wish to compute the weighted mean over all haplotype-breakpoint combinations $(h_{1:M}; A, B, \dots, L)$.

To do so, we make use of equation (20), which gives the approximate (unnormalized) posterior probability of the haplotype $h_{1:M}$. To simplify notation, we define

$$r(h_{X+1:Y}) := P(h_{X+1:Y}^{\text{miss}} | h_{X+1:Y})f(h_{X+1:Y})P(\text{rec } Y | X), \quad (23)$$

so equation (20) becomes

$$q(h_{1:M}) = \sum_{A,B,\dots,L} r(h_{1:A})r(h_{A+1:B}) \cdots r(h_{L+1:M}). \quad (24)$$

(We have dropped the ‘‘approximately equals’’ sign from equation (20) for clarity in the derivations that follow.) For each marker m , we are interested in the relative values of

$$q(h_m = 0) := \sum_{h_{1:M}|h_m=0} q(h_{1:M}) \quad \text{vs.} \quad q(h_m = 1) := \sum_{h_{1:M}|h_m=1} q(h_{1:M}), \quad (25)$$

which give the posterior odds of the m -th allele being 0 vs. 1. For each $i \in \{0, 1\}$, we write

$$\begin{aligned}
& q(h_m = i) \\
&= \sum_{h_{1:M}|h_m=i} q(h_{1:M}) \\
&= \sum_{h_{1:M}|h_m=i} \sum_{A,B,\dots,L} r(h_{1:A})r(h_{A+1:B}) \cdots r(h_{L+1:M}) \\
&= \sum_{h_{1:M}|h_m=i} q(h_{1:X}) \left(\sum_{X,Y|X+1 \leq m \leq Y} r(h_{X+1,Y}) \right) q(h_{Y+1:M}) \\
&= \sum_{X,Y|X+1 \leq m \leq Y} \left(\sum_{h_{1:X}} q(h_{1:X}) \right) \left(\sum_{h_{X+1,Y}|h_m=i} r(h_{X+1,Y}) \right) \left(\sum_{h_{Y+1:M}} q(h_{Y+1:M}) \right). \quad (26)
\end{aligned}$$

This final form is now amenable to forward-backward computation: the left sum $\sum_{h_{1:X}} q(h_{1:X})$ can be computed via the forward recursion in equation (22), and the right sum $\sum_{h_{Y+1:M}} q(h_{Y+1:M})$ can be computed via an analogous backward recursion. Each of these recursions has total cost $O(MHP)$, after which each of the terms in equation (26) takes $O(1)$ time to compute. The total number of unique terms over all sites m is just the number of candidate haplotype segments $h_{X+1,Y}$, which is again $O(MHP)$. (Technically, for each of the $O(MHP)$ haplotype segments $h_{X+1,Y}$, after computing its term in equation (22), we still need to augment $q(h_m = i)$ for each marker $m \in \{X + 1, \dots, Y\}$ for which h_m^{miss} is missing, and for a given missing rate, the number of such markers grows with H . However, in practice, with missing rates on the order of 1% and $H=500$ (for imputation, in which we do not condense sites), this final cost is negligible. We therefore view the total cost of the imputation algorithm is $O(MHP)$.)

We remark that this $O(MHP)$ -time forward-backward algorithm is specific to the haploid imputation problem; no equally fast analog exists for the diploid phasing problem, hence our use of beam search and backward smoothing heuristics in the Eagle2 phasing algorithm.

3.4 Optimizations

While the $O(MHP)$ computational cost of the algorithm we have described is quite fast, one optimization allows a useful constant-factor speedup. Rather than considering the possibility of a recombination between each pair of consecutive SNPs, we allow recombinations (i.e., transitions between copied reference haplotypes) only after every 25th SNP. At the density of typical genotyping chips, this approximation has little effect on accuracy, and it saves both computation time and memory (because it allows us to store only $M/25$ trees in the HapHedge). With this optimization, we are able to perform missing imputation (with non-condensed history length $H=500$ and beam width $P=16$) for a small fraction of the computational cost of phasing.

4 HapHedge data structure

The HapHedge data structure represents a sequence of haplotype prefix trees (i.e., binary trees on haplotype prefixes) rooted at a given set of starting positions along a chromosome (Figure 1). The key features of the HapHedge are linear-time construction, linear-memory representation, and constant-time prefix extension (all with small constant factors). Specifically, the computational costs associated with a HapHedge on K haplotypes and M markers are as follows:

- $O(MK)$ -time construction using the positional Burrows-Wheeler transform (PBWT) [19].
- $O(MK)$ -memory representation of the full “hedge” using radix trees. More precisely, to represent M_t trees, $12M_tK + MK/8$ bytes are required. (If we begin a new tree every s positions, $M_t = M/s$.)
- $O(1)$ -time extension of a haplotype prefix (with haplotype frequency lookup).

We first describe the in-memory representation of a HapHedge and show how it supports constant-time haplotype extension; we then describe HapHedge construction.

4.1 Compact representation of haplotype prefix trees in memory

The HapHedge data structure represents a sequence of haplotype prefix trees using *radix trees* backed by a *shared bit array* containing all input haplotypes. In general, a radix tree compactifies a prefix tree by merging any non-branching internal node with its parent. Each edge of a radix tree is thus labeled with a string (representing merged edges) rather than a single character.

In our case, we wish to represent many trees representing different suffixes of the same underlying set of K haplotype bit strings. As such, we can represent an edge label in constant space using only a pointer to the appropriate reference haplotype and the length of the label. We employ this strategy as follows.

- We store the reference haplotypes in a `bitArray` using MK bits.
- For each prefix start location, we store a radix tree using the following data layout:
 - 1 integer `hap0`: the index of the lexicographically first haplotype prefix
 - $K - 1$ nodes, each containing:
 - * 1 integer `mSplit`: the location of the next branch (i.e., the next location polymorphic among haplotype prefixes in the current subtree)
 - * 1 integer `count0`: the number of haplotype prefixes in the 0-subtree at `mSplit`
 - * 1 integer `hap1`: the index of the lexicographically first haplotype prefix in the 1-subtree at `mSplit`.

We note that the specification above does not contain any explicit pointers to indices of child nodes. We are able to avoid storing such pointers by storing the $K - 1$ nodes of each radix tree in the order that they are encountered during a pre-order traversal of the tree. Assuming this ordering is chosen, the left child of a node is simply the node that follows it in memory, and the right child of a node is the node `count` nodes after it. We present an example in Supplementary Fig. 2. (In the orientation of Supplementary Fig. 2, “left” and “right” children are depicted as upper and lower child nodes.)

4.2 Constant-time haplotype prefix extension

To follow a haplotype prefix along a radix tree represented in the above form, we need to keep track of some information beyond the current marker and the index of the current node. Specifically, the complete state information that we maintain during prefix lookup is:

- 1 integer `m`: the current marker
- 1 integer `node`: the index of the node at the end of the current edge
- 1 integer `hap`: the index of the lexicographically first haplotype prefix in the current subtree
- 1 integer `count`: the number of haplotype prefixes in the current subtree.

In particular, we note that the state for the root node of the radix tree starting at position `m0` is $\{m=m0, node=0, hap=hap0, count=K\}$. (We could of course have included the `hap` and `count` data fields in the nodes of the HapHedge, but doing so would increase its memory footprint.)

The logic for extending a haplotype prefix (with state information) to the `nextBit` (either 0 or 1) is probably easiest understood in pseudocode.

Pseudocode for extending a haplotype prefix.

```

INPUT:
- state.{m,node,hap,count}: current state information
- nextBit: bit (0 or 1) with which to attempt extension of prefix

OUTPUT:
- if successful extension, update state and return SUCCESS
- else, return FAILURE

function extend(state, nextBit) {
  if (state.count == 1 || nodes[state.node].mSplit > state.m) {
    if (bitArray(state.hap, state.m) != nextBit)
      return FAILURE
  }
  else {

```

```

    if (nextBit == 0) {
        state.count = nodes[state.node].count0
        state.node++
    }
    else {
        state.hap = nodes[state.node].hap1
        c0 = nodes[state.node].count0
        state.node += c0
        state.count -= c0
    }
}
state.m++
return SUCCESS
}

```

4.3 Linear-time HapHedge construction

The key insight for efficient HapHedge construction is that the positional Burrows-Wheeler transform [19] already encodes all the information we need to construct the haplotype radix tree at each desired root marker in linear time. Specifically, Algorithm 2 of ref. [19] shows how to iteratively construct a *positional suffix array* $a[]$ and *divergence array* $d[]$ at each marker. (Ref. [19] works from left to right and hence constructs positional prefix arrays rather than suffix arrays, but we need only reverse the direction of processing.) At each marker m , the positional suffix array $a[]$ precisely gives the lexicographic order on haplotype prefixes starting from m —i.e., the in-order traversal of leaves in the radix tree we wish to construct—and the divergence array $d[]$ encodes the branching depths m_{Split} of internal nodes. Thus, if we wish to construct a radix tree on prefixes starting at m , we can perform an in-order traversal of the tree (before explicitly representing it) by walking through the $a[]$ and $d[]$ arrays, recovering the tree topology as we go based on relative divergence depths d .

The observations above give the following two-step algorithm for constructing a radix tree rooted at marker m (on reaching m while running PBWT Algorithm 2 from right to left). Given the matrices $a[]$ and $d[]$ for marker m , we first construct a temporary (less memory-efficient) explicit representation of the radix tree using the in-order traversal described above. Second, we perform a pre-order traversal of our temporary representation to generate the memory-efficient representation we desire. For details, see the implementation in `HapHedge.cpp`.

4.4 Extension to haplotypes with multiplicities: HapHedgeMulti

As we noted in Section 2.4.2, it is advantageous to modify our radix tree representation when representing condensed haplotypes because the number of unique condensed haplotype prefixes (after

truncating at error bits $e_t = 1$) is typically much smaller than the number of input haplotypes. We call the modified data structure the HapHedgeMulti. Conceptually, the ideas behind the representation of haplotypes and the construction and prefix extension algorithms remain unchanged, so we only highlight a few main implementational differences.

- A HapHedgeMulti radix tree contains only $K_{\text{uniq}} - 1$ internal nodes (vs. $K - 1$ nodes in a HapHedge radix tree), where K_{uniq} is the number of unique prefixes.
- Each node of a HapHedgeMulti radix tree contains 5 integer data fields: the 3 fields `mSplit`, `count0`, and `hap1` from before, plus 2 new fields `node0` and `node1` giving the indices of child nodes. (These additional fields are needed because node indices can no longer be inferred from counts due to haplotype multiplicities.)
- Truncation of prefixes at bits indicating errors at target homs (Section 2.4.2) is easily accomplished while running PBWT Algorithm 2 by changing values at the end of the divergence array `d[]` to M . Specifically, if bit m encodes an error bit e_t , then after the first prefix with 1 at bit m , we change `d[*]` to M for all subsequent prefixes (which also must have 1 at bit m). This modification effectively equates all prefixes with 1 at bit m , thus accomplishing the truncation, and the equivalence is naturally propagated to radix tree construction and to subsequent PBWT iterations (both of which only see information about already-processed markers as represented by the matrices `a[]` and `d[]`).

Again, implementation details can be found in the file `HapHedge.cpp`.

5 Appendix: Probability of recombining off a shared haplotype

Here we derive equation (8) for the probability of ending a copied haplotype as a function of the copying length ℓ . Recombinations occur along the lineage connecting two shared haplotypes at a rate proportional to the number of generations t to the most recent common ancestor, so equivalently, we want to compute the expected age of an IBD segment given its length $E[t|\ell]$.

We begin with some background on the length distribution of IBD segments as a function of age. Assume we are given a randomly chosen genomic site. At this site, two individuals share a common ancestor t generations in the past. Because recombination is well-modeled by a Poisson process, we know that the distance (in Morgans) to the first recombination event to either the left or the right of the chosen genomic location is an exponential random variable with mean $2t$, i.e. the total number of generations separating the two individuals. We assume that IBD segments are delimited by recombination events. In the full coalescent model, this may not be true, leading to more complex modeling [33], but for long segments this represents a very good approximation [34] and corresponds to adopting the SMC model [35]. In this scenario, the total length of the IBD

segment is given by the sum of two exponential random variables with mean $2t$, leading to an Erlang-2 distribution with mean $2t$.

If we are given an IBD segment of length ℓ Morgans and we know the segment was sampled from a population of constant size N , we can use the above considerations to derive a distribution for the age of the segment. If the segment comes from generation t , the probability that its length is ℓ is given by the Erlang-2 distribution $p(\ell|t) = (2t)^2 \ell e^{-2t\ell}$. Note, however, that this distribution describes the probability that a randomly drawn site is spanned by a segment of length ℓ . Here, we are interested in the age of a randomly drawn segment, which spans exactly ℓ Morgans. To get a distribution for the age of a segment, rather than a site within it, we divide the above distribution by ℓ , obtaining the (unnormalized) distribution $(2t)^2 e^{-2t\ell}$, and switch the unit from that of a site to that of a segment. The distribution of the segment age is then obtained by considering the probability of the two individuals sharing a common ancestor at generation t for the considered site, which for a constant population is given by $p(t|N) = N^{-1} e^{t/N}$. Putting things together and renormalizing,

$$\begin{aligned} p(t|\ell) &= \frac{(2t)^2 e^{-2t\ell} N^{-1} e^{t/N}}{\int_0^\infty (2t)^2 e^{-2t\ell} N^{-1} e^{t/N} dt} \\ &= \frac{t^2}{2} (N^{-1} + 2\ell)^3 e^{-t(N^{-1} + 2\ell)}, \end{aligned} \quad (27)$$

which is an Erlang-3 distribution with rate $(N^{-1} + 2\ell)$ and mean


$$E[t|\ell] = \frac{3}{N^{-1} + 2\ell}. \quad (28)$$

References



1. Tewhey, R., Bansal, V., Torkamani, A., Topol, E. J. & Schork, N. J. The importance of phase information for human genomics. *Nature Reviews Genetics* **12**, 215–223 (2011).
2. Browning, S. R. & Browning, B. L. Haplotype phasing: existing methods and new developments. *Nature Reviews Genetics* **12**, 703–714 (2011).
3. Stephens, M., Smith, N. J. & Donnelly, P. A new statistical method for haplotype reconstruction from population data. *American Journal of Human Genetics* **68**, 978–989 (2001).
4. Halperin, E. & Eskin, E. Haplotype reconstruction from genotype data using imperfect phylogeny. *Bioinformatics* **20**, 1842–1849 (2004).
5. Stephens, M. & Scheet, P. Accounting for decay of linkage disequilibrium in haplotype inference and missing-data imputation. *American Journal of Human Genetics* **76**, 449–462 (2005).
6. Scheet, P. & Stephens, M. A fast and flexible statistical model for large-scale population genotype data: applications to inferring missing genotypes and haplotypic phase. *American Journal of Human Genetics* **78**, 629–644 (2006).
7. Browning, S. R. & Browning, B. L. Rapid and accurate haplotype phasing and missing-data inference for whole-genome association studies by use of localized haplotype clustering. *American Journal of Human Genetics* **81**, 1084–1097 (2007).
8. Kong, A. *et al.* Detection of sharing by descent, long-range phasing and haplotype imputation. *Nature Genetics* **40**, 1068–1075 (2008).
9. Browning, B. L. & Browning, S. R. A unified approach to genotype imputation and haplotype-phase inference for large data sets of trios and unrelated individuals. *American Journal of Human Genetics* **84**, 210–223 (2009).
10. Delaneau, O., Marchini, J. & Zagury, J.-F. A linear complexity phasing method for thousands of genomes. *Nature Methods* **9**, 179–181 (2012).
11. Williams, A. L., Patterson, N., Glessner, J., Hakonarson, H. & Reich, D. Phasing of many thousands of genotyped samples. *American Journal of Human Genetics* **91**, 238–251 (2012).
12. Delaneau, O., Zagury, J.-F. & Marchini, J. Improved whole-chromosome phasing for disease and population genetic studies. *Nature Methods* **10**, 5–6 (2013).
13. Loh, P.-R., Palamara, P. & Price, A. Fast and accurate long-range phasing and imputation in a uk biobank cohort. *Nature Genetics* (2016). URL <http://biorxiv.org/content/early/2015/12/18/028282>.
14. O’Connell, J. *et al.* Haplotype estimation for biobank scale datasets. *Nature Genetics* (in press).

15. Snyder, M. W., Adey, A., Kitzman, J. O. & Shendure, J. Haplotype-resolved genome sequencing: experimental methods and applications. *Nature Reviews Genetics* **16**, 344–358 (2015).
16. van de Geijn, B., McVicker, G., Gilad, Y. & Pritchard, J. K. WASP: allele-specific software for robust molecular quantitative trait locus discovery. *Nature Methods* **12**, 1061–1063 (2015).
17. Kumasaka, N., Knights, A. J. & Gaffney, D. J. Fine-mapping cellular QTLs with RASQUAL and ATAC-seq. *Nature Genetics* (2015).
18. McCarthy, S. *et al.* A reference panel of 64,976 haplotypes for genotype imputation. *bioRxiv* 035170 (2016).
19. Durbin, R. Efficient haplotype matching and storage using the positional Burrows–Wheeler transform (PBWT). *Bioinformatics* **30**, 1266–1272 (2014).
20. Li, N. & Stephens, M. Modeling linkage disequilibrium and identifying recombination hotspots using single-nucleotide polymorphism data. *Genetics* **165**, 2213–2233 (2003).
21. Sudlow, C. *et al.* UK Biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLOS Medicine* **12**, 1–10 (2015).
22. Kvale, M. N. *et al.* Genotyping informatics and quality control for 100,000 Subjects in the Genetic Epidemiology Research on Adult Health and Aging (GERA) Cohort. *Genetics* **200**, 1051–1060 (2015).
23. Banda, Y. *et al.* Characterizing race/ethnicity and genetic ancestry for 100,000 subjects in the Genetic Epidemiology Research on Adult Health and Aging (GERA) cohort. *Genetics* **200**, 1285–1295 (2015).
24. Browning, B. L. & Browning, S. R. Genotype imputation with millions of reference samples. *The American Journal of Human Genetics* **98**, 116–126 (2016).
25. 1000 Genomes Project Consortium *et al.* A global reference for human genetic variation. *Nature* **526**, 68–74 (2015).
26. Howie, B., Fuchsberger, C., Stephens, M., Marchini, J. & Abecasis, G. R. Fast and accurate genotype imputation in genome-wide association studies through pre-phasing. *Nature Genetics* **44**, 955–959 (2012).
27. He, D., Han, B. & Eskin, E. Hap-seq: an optimal algorithm for haplotype phasing with imputation using sequencing data. *Journal of Computational Biology* **20**, 80–92 (2013).
28. Delaneau, O., Howie, B., Cox, A. J., Zagury, J.-F. & Marchini, J. Haplotype estimation using sequencing reads. *American Journal of Human Genetics* **93**, 687–696 (2013).
29. Sharp, K., Kretzschmar, W., Delaneau, O. & Marchini, J. Phasing for medical sequencing using rare variants and large haplotype reference panels. *Bioinformatics* (2016).

30. Chang, C. C. *et al.* Second-generation PLINK: rising to the challenge of larger and richer datasets. *GigaScience* **4**, 1–16 (2015).
31. Browning, S. R. Multilocus association mapping using variable-length Markov chains. *American Journal of Human Genetics* **78**, 903–913 (2006).
32. Browning, B. L. & Browning, S. R. Efficient multilocus association testing for whole genome association studies using localized haplotype clustering. *Genetic Epidemiology* **31**, 365–375 (2007).
33. Harris, K. & Nielsen, R. Inferring demographic history from a spectrum of shared haplotype lengths. *PLOS Genetics* **9**, e1003521 (2013).
34. Palamara, P. F., Lencz, T., Darvasi, A. & Pe'er, I. Length distributions of identity by descent reveal fine-scale demographic history. *American Journal of Human Genetics* **91**, 809–822 (2012).
35. McVean, G. A. & Cardin, N. J. Approximating the coalescent with recombination. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* **360**, 1387–1393 (2005).



Supplementary Figure 1. Conversion of reference haplotypes to condensed form. This figure illustrates the procedure used to create condensed reference haplotypes given the genotypes of a specific target sample (Supplementary Note Section 2.4.2). Alleles are directly coded in 1 bit each at a subset of sites $s(1), \dots, s(T)$ consisting of target het sites plus spacer sites as needed to break up runs of homozygosity (ROH). In each region between directly-coded alleles, all target genotypes are either homozygous or missing; we encode the presence or absence of discrepancies between the reference haplotype and the target in 1 bit. (Finally, we pad each condensed haplotype with $h_{s(0)} = 0$ and $h_{s(T+1)} = 0, e_{T+1} = 0$ for implementational convenience.)



Supplementary Figure 2. A haplotype prefix tree in radix format. This example shows the HapHedge representation of a single tree on $K = 5$ haplotypes rooted at the first of $M = 6$ markers. A HapHedge will typically contain many trees rooted at different starting positions, all sharing the same bitArray, but each with its own hap0 and nodes.

Supplementary Table 1. Data sets analyzed.

Data set	Samples (<i>N</i>)	Indep. trios	Markers (<i>M</i>)	Heterozygosity	Missingness
UK Biobank	152,248	72	707,524	0.173	0.9%
GERA European chip	62,318	200	657,184	0.257	0.6%
GERA African chip	3,826	3	851,970	0.216	1.1%
GERA East Asian chip	5,188	7	694,877	0.240	0.6%
GERA Latino chip	7,154	3	776,817	0.225	1.4%

We report basic sample and SNP statistics for the UK Biobank and GERA data sets analyzed. The UK Biobank statistics reflect the light QC we performed (described in Online Methods); the GERA statistics correspond exactly to the data we obtained from dbGaP. “Independent trios” indicate numbers of non-overlapping mother-father-child trios. “Markers” indicate autosomal SNP counts. Heterozygosity and missingness are means over all samples and SNPs.

Supplementary Table 2. Running time and accuracy of reference-based phasing in UK Biobank simulations.

(a) CPU time per target genome

Method	$N_{\text{ref}}=15\text{K}$	$N_{\text{ref}}=30\text{K}$	$N_{\text{ref}}=50\text{K}$
Eagle1	1.2 min	1.4 min	1.6 min
Eagle2 $K=5\text{K}$	1.1 min	1.1 min	1.1 min
Eagle2 $K=10\text{K}$	1.5 min	1.5 min	1.5 min
Eagle2 $K=20\text{K}$	2.4 min	2.4 min	2.4 min
Eagle2 $K=40\text{K}$	3.3 min	4.0 min	4.1 min
SHAPEIT2 $K=50$ –no-mcmc	2.1 min	2.6 min	3.6 min
SHAPEIT2 $K=100$ –no-mcmc	2.6 min	3.1 min	3.7 min
SHAPEIT2 $K=200$ –no-mcmc	3.7 min	4.1 min	4.9 min
SHAPEIT2 $K=400$ –no-mcmc	5.6 min	6.1 min	6.9 min
SHAPEIT2 $K=50$	13.7 min	20.8 min	30.3 min
SHAPEIT2 $K=100$	18.0 min	26.0 min	36.0 min
SHAPEIT2 $K=200$	26.2 min	34.2 min	44.9 min
SHAPEIT2 $K=400$	41.6 min	49.7 min	60.9 min

(b) Mean switch error rate (s.e.m.)

Method	$N_{\text{ref}}=15\text{K}$	$N_{\text{ref}}=30\text{K}$	$N_{\text{ref}}=50\text{K}$
Eagle1	1.35% (0.04%)	0.88% (0.03%)	0.65% (0.03%)
Eagle2 $K=5\text{K}$	0.98% (0.03%)	0.75% (0.03%)	0.63% (0.03%)
Eagle2 $K=10\text{K}$	0.86% (0.03%)	0.65% (0.02%)	0.53% (0.02%)
Eagle2 $K=20\text{K}$	0.78% (0.02%)	0.58% (0.02%)	0.47% (0.02%)
Eagle2 $K=40\text{K}$	0.76% (0.02%)	0.53% (0.02%)	0.43% (0.02%)
SHAPEIT2 $K=50$ –no-mcmc	3.82% (0.06%)	3.64% (0.07%)	3.58% (0.08%)
SHAPEIT2 $K=100$ –no-mcmc	2.35% (0.05%)	2.14% (0.05%)	2.09% (0.05%)
SHAPEIT2 $K=200$ –no-mcmc	1.65% (0.04%)	1.39% (0.04%)	1.30% (0.04%)
SHAPEIT2 $K=400$ –no-mcmc	1.28% (0.03%)	1.01% (0.03%)	0.89% (0.03%)
SHAPEIT2 $K=50$	1.23% (0.03%)	0.86% (0.03%)	0.69% (0.03%)
SHAPEIT2 $K=100$	1.04% (0.03%)	0.74% (0.03%)	0.58% (0.02%)
SHAPEIT2 $K=200$	0.93% (0.03%)	0.65% (0.02%)	0.52% (0.02%)
SHAPEIT2 $K=400$	0.85% (0.02%)	0.60% (0.02%)	0.47% (0.02%)

(This table provides numeric data plotted in Figure 2 along with additional benchmarks.) We performed reference-based phasing using reference panels generated from $N_{\text{ref}} = 15,000, 30,000,$ or $50,000$ UK Biobank samples. We analyzed chromosomes 1, 5, 10, 15, and 20 ($\approx 25\%$ of the genome) in 4-threaded computation on 2.27 GHz Intel Xeon L5640 processors. To obtain CPU time per target genome, we then scaled up computation time by a factor of 4; see Supplementary Table 3 for details. We assessed accuracy on the 70 UK Biobank European-ancestry trio children, aggregating switch errors over the five chromosomes analyzed, and computing means and s.e.m. over the 70 samples. The default values of K for SHAPEIT2 and Eagle2 are 100 and 10,000, respectively.

Supplementary Table 3. Computational cost of reference-based phasing in UK Biobank simulations at different target sample sizes.

(a) CPU time and memory for phasing chr1,5,10,15,20 of $N_{\text{target}} = 72$ samples

Method	$N_{\text{ref}} = 15\text{K}$	$N_{\text{ref}} = 30\text{K}$	$N_{\text{ref}} = 50\text{K}$
Eagle1	0.6 hr / 2.8 GB	0.9 hr / 4.8 GB	1.2 hr / 6.7 GB
Eagle2 $K=5\text{K}$	0.4 hr / 1.2 GB	0.4 hr / 1.8 GB	0.4 hr / 2.9 GB
Eagle2 $K=10\text{K}$	0.5 hr / 1.4 GB	0.5 hr / 1.9 GB	0.6 hr / 3.1 GB
Eagle2 $K=20\text{K}$	0.8 hr / 1.6 GB	0.8 hr / 2.2 GB	0.8 hr / 3.2 GB
Eagle2 $K=40\text{K}$	1.0 hr / 1.8 GB	1.3 hr / 2.7 GB	1.3 hr / 3.8 GB
SHAPEIT2 $K=50$ –no-mcmc	1.1 hr / 0.6 GB	2.1 hr / 1.0 GB	3.0 hr / 1.7 GB
SHAPEIT2 $K=100$ –no-mcmc	1.4 hr / 0.7 GB	2.3 hr / 1.1 GB	3.4 hr / 1.8 GB
SHAPEIT2 $K=200$ –no-mcmc	1.7 hr / 0.8 GB	2.7 hr / 1.2 GB	3.7 hr / 2.0 GB
SHAPEIT2 $K=400$ –no-mcmc	2.3 hr / 1.2 GB	3.3 hr / 1.6 GB	4.4 hr / 2.3 GB
SHAPEIT2 $K=50$	4.3 hr / 0.6 GB	6.9 hr / 1.0 GB	10.7 hr / 1.7 GB
SHAPEIT2 $K=100$	5.4 hr / 0.6 GB	8.4 hr / 1.0 GB	12.2 hr / 1.7 GB
SHAPEIT2 $K=200$	7.8 hr / 0.6 GB	10.7 hr / 1.0 GB	14.5 hr / 1.7 GB
SHAPEIT2 $K=400$	12.4 hr / 0.6 GB	15.4 hr / 1.0 GB	19.4 hr / 1.7 GB

(b) CPU time and memory for phasing chr1,5,10,15,20 of $N_{\text{target}} = 1,072$ samples

Method	$N_{\text{ref}} = 15\text{K}$	$N_{\text{ref}} = 30\text{K}$	$N_{\text{ref}} = 50\text{K}$
Eagle1	5.7 hr / 3.0 GB	6.7 hr / 4.7 GB	7.8 hr / 6.8 GB
Eagle2 $K=5\text{K}$	5.0 hr / 1.3 GB	4.9 hr / 2.0 GB	5.0 hr / 3.0 GB
Eagle2 $K=10\text{K}$	6.9 hr / 1.4 GB	6.9 hr / 2.1 GB	6.8 hr / 3.1 GB
Eagle2 $K=20\text{K}$	10.7 hr / 1.6 GB	10.7 hr / 2.3 GB	10.7 hr / 3.3 GB
Eagle2 $K=40\text{K}$	14.6 hr / 1.9 GB	18.2 hr / 2.9 GB	18.5 hr / 3.8 GB
SHAPEIT2 $K=50$ –no-mcmc	9.8 hr / 1.6 GB	12.9 hr / 2.0 GB	18.1 hr / 2.7 GB
SHAPEIT2 $K=100$ –no-mcmc	12.2 hr / 1.7 GB	15.3 hr / 2.1 GB	18.7 hr / 2.8 GB
SHAPEIT2 $K=200$ –no-mcmc	16.9 hr / 1.9 GB	19.8 hr / 2.3 GB	23.9 hr / 3.0 GB
SHAPEIT2 $K=400$ –no-mcmc	25.6 hr / 2.2 GB	28.8 hr / 2.6 GB	33.0 hr / 3.3 GB
SHAPEIT2 $K=50$	61.6 hr / 1.6 GB	93.5 hr / 2.0 GB	137.1 hr / 2.7 GB
SHAPEIT2 $K=100$	80.3 hr / 1.6 GB	116.7 hr / 2.0 GB	162.3 hr / 2.7 GB
SHAPEIT2 $K=200$	117.1 hr / 1.6 GB	153.1 hr / 2.0 GB	201.7 hr / 2.7 GB
SHAPEIT2 $K=400$	185.5 hr / 1.6 GB	222.4 hr / 2.0 GB	273.0 hr / 2.7 GB

To benchmark per-sample computational cost of reference-based phasing, we performed two sets of analyses: one in which we phased only $N_{\text{target}} = 72$ UK Biobank trio children, and another in which we phased an additional 1,000 randomly selected samples. We then subtracted the $N_{\text{target}} = 72$ CPU times from the $N_{\text{target}} = 1,072$ CPU times to obtain the incremental cost of phasing 1,000 samples. This procedure was necessary to adjust for initialization costs (e.g., reading the reference data and synchronizing it with the target data), which account for a non-negligible fraction of total computational cost when N_{target} is small. Finally, we divided by 1,000 to obtain per-sample costs and multiplied by 4 to scale up from 25% of the genome (chr20,15,10,15,1) to obtain the numbers plotted in Fig. 2a and reported in Supplementary Table 2.

We performed runs using 4 cores of a 2.27 GHz Intel Xeon L5640 processor. We chose to report CPU times rather than wall clock run times because we observed less run-to-run variability in CPU times on the shared compute cluster we used for these benchmarks. (Both Eagle and SHAPEIT perform efficient multithreading, so the choice of CPU time vs. wall time has little effect on relative performance.)

Supplementary Table 4. Accuracy of reference-based phasing in GERA simulations using varying numbers of conditioning haplotypes.

(a) Mean switch error rate (s.e.m.), trio parents

Method	European chip	African chip	East Asian chip	Latino chip
	$N_{\text{ref}} = 61,684$ $N_{\text{target}} = 400$	$N_{\text{ref}} = 3,817$ $N_{\text{target}} = 6$	$N_{\text{ref}} = 5,164$ $N_{\text{target}} = 14$	$N_{\text{ref}} = 7,144$ $N_{\text{target}} = 6$
SHAPEIT2 $K=100$	1.31% (0.04%)	3.21% (0.08%)	2.24% (0.12%)	2.42% (0.05%)
SHAPEIT2 $K=200$	1.23% (0.04%)	2.93% (0.07%)	2.11% (0.11%)	2.29% (0.05%)
SHAPEIT2 $K=400$	1.17% (0.04%)	2.75% (0.06%)	2.02% (0.11%)	2.18% (0.04%)
SHAPEIT2 $K=100$ –no-mcmc	2.83% (0.05%)	5.29% (0.15%)	3.34% (0.14%)	3.91% (0.08%)
SHAPEIT2 $K=200$ –no-mcmc	2.11% (0.05%)	4.16% (0.11%)	2.80% (0.13%)	3.14% (0.06%)
SHAPEIT2 $K=400$ –no-mcmc	1.72% (0.05%)	3.51% (0.09%)	2.48% (0.12%)	2.72% (0.06%)
Eagle1	1.36% (0.04%)	3.87% (0.09%)	2.71% (0.14%)	2.65% (0.05%)
Eagle2 $K=5K$	1.32% (0.04%)	2.64% (0.06%)	2.00% (0.10%)	2.17% (0.04%)
Eagle2 $K=10K$	1.24% (0.04%)	2.48% (0.05%)	1.94% (0.10%)	2.08% (0.04%)
Eagle2 $K=20K$	1.18% (0.04%)	2.48% (0.05%)	1.96% (0.10%)	2.06% (0.04%)

(b) Mean switch error rate (s.e.m.), trio children

Method	European
	$N_{\text{ref}} = 61,684$ $N_{\text{target}} = 200$
SHAPEIT2 $K=100$	0.90% (0.04%)
SHAPEIT2 $K=200$	0.82% (0.04%)
SHAPEIT2 $K=400$	0.75% (0.04%)
SHAPEIT2 $K=100$ –no-mcmc	2.46% (0.06%)
SHAPEIT2 $K=200$ –no-mcmc	1.73% (0.05%)
SHAPEIT2 $K=400$ –no-mcmc	1.31% (0.05%)
Eagle1	0.95% (0.04%)
Eagle2 $K=5K$	0.93% (0.04%)
Eagle2 $K=10K$	0.84% (0.04%)
Eagle2 $K=20K$	0.78% (0.04%)

(This table provides numeric data plotted in Figure 3 along with additional benchmarks.)

(a) We phased trio parents in each GERA sub-cohort using a reference panel generated from all other non-familial samples in the same sub-cohort. We ran each method on all 22 autosomes and computed aggregate mean switch error rates (s.e.m.). Standard errors for the European-ancestry sub-cohort are over 400 parent samples. Standard errors for the other three sub-cohorts are over 25 SNP blocks.

(b) We phased GERA trio children on the European chip sub-cohort using a reference panel generated from all other non-familial European chip samples. (We did not perform this analysis on the other sub-cohorts due to the low numbers of trios.) We note that accuracy is higher across all methods for trio children vs. trio parents, perhaps due to greater levels of (within-Europe) admixture in the children.

Supplementary Table 5. Accuracy of reference-based phasing using the 1000 Genomes and HRC panels.

(a) Mean switch error rates (s.e.m.) using the 1000 Genomes Project Phase 3 panel

Method	CEU $N_{\text{target}} = 32$	CHS $N_{\text{target}} = 31$	PEL $N_{\text{target}} = 30$	PJL $N_{\text{target}} = 15$	YRI $N_{\text{target}} = 19$
SHAPEIT2	3.52% (0.06%)	6.51% (0.11%)	4.45% (0.16%)	4.75% (0.24%)	3.71% (0.11%)
SHAPEIT2 –no-mcmc	4.83% (0.05%)	7.79% (0.11%)	5.71% (0.17%)	6.41% (0.26%)	4.97% (0.11%)
Eagle1	5.40% (0.08%)	8.98% (0.14%)	6.57% (0.21%)	6.89% (0.35%)	6.72% (0.11%)
Eagle2	3.29% (0.06%)	6.41% (0.11%)	4.49% (0.19%)	4.38% (0.21%)	3.67% (0.11%)

(b) Mean switch error rates (s.e.m.) using the HRC panel

Method	CEU $N_{\text{target}} = 32$	CHS $N_{\text{target}} = 31$	PEL $N_{\text{target}} = 30$	PJL $N_{\text{target}} = 15$	YRI $N_{\text{target}} = 19$
SHAPEIT2	1.60% (0.05%)	7.78% (0.08%)	4.65% (0.15%)	5.97% (0.19%)	4.68% (0.10%)
SHAPEIT2 –no-mcmc	3.39% (0.08%)	9.56% (0.09%)	6.92% (0.10%)	9.00% (0.13%)	6.51% (0.09%)
Eagle1	1.89% (0.06%)	10.04% (0.10%)	6.30% (0.21%)	7.35% (0.17%)	7.76% (0.15%)
Eagle2	1.35% (0.04%)	7.54% (0.08%)	4.80% (0.17%)	5.25% (0.15%)	4.84% (0.10%)

We phased 1000 Genomes Phase 3 trio children on chromosome 1 using the 1000 Genomes Project Phase 3 panel or the Haplotype Reference Consortium panel (excluding these trios). To simulate a typical use case, we restricted the data to 31,853 markers (genotyped by 23andMe). We report mean switch error rates (s.e.m.) over trio samples in each population with >1 trio.

Supplementary Table 6. Accuracy of phasing non-European UK Biobank samples using increasing numbers of European reference samples.

N_{matched}	Method	$N_{\text{mismatched}}$				
		0	5,000	15,000	50,000	147,000
1,000	Eagle2	5.25% (0.21%)	5.25% (0.25%)	5.25% (0.25%)	4.97% (0.22%)	4.72% (0.22%)
	SHAPEIT2	5.89% (0.21%)	5.67% (0.23%)	5.87% (0.24%)	6.01% (0.24%)	5.83% (0.23%)
2,000	Eagle2	4.12% (0.18%)	4.18% (0.19%)	4.15% (0.20%)	4.10% (0.18%)	4.01% (0.19%)
	SHAPEIT2	4.82% (0.17%)	4.60% (0.22%)	4.80% (0.22%)	4.90% (0.22%)	4.88% (0.21%)
4,000	Eagle2	3.37% (0.15%)	3.32% (0.16%)	3.29% (0.16%)	3.28% (0.15%)	3.23% (0.15%)
	SHAPEIT2	3.96% (0.16%)	3.75% (0.16%)	3.83% (0.17%)	3.88% (0.18%)	4.00% (0.17%)

To investigate the effect of including ancestry-mismatched reference samples on reference-based phasing accuracy, we phased trio parents in the two non-European (specifically, Indian and Caribbean) UK Biobank trios using subsets of UK Biobank haplotypes (from a run of Eagle1 on all samples together) containing increasing numbers of European haplotypes. Specifically, we partitioned the non-trio samples into an “ancestry-matched” set of 4,708 samples that self-reported Mixed, Asian, or Black ancestry and an “ancestry-mismatched” set of the remaining 147,324 predominantly European samples. We then created reference panels containing combinations of 1,000, 2,000, or 4,000 ancestry-matched samples and 0, 5,000, 15,000, 50,000, or 147,000 ancestry-mismatched samples. We used these reference panels to phase the 4 non-European parents using either Eagle2 or SHAPEIT2. We report mean switch error rates over chromosomes 15–22 (s.e.m.).

Supplementary Table 7. Computational cost and accuracy of non-reference-based phasing in the UK Biobank cohort.

(a) Wall clock time and memory cost per genome

Method	$N = 5K$	$N = 15K$	$N = 50K$	$N = 150K$
SHAPEIT2	55.7 hr / 5.3 GB	230.9 hr / 15.2 GB	–	–
Eagle1	8.1 hr / 4.4 GB	30.4 hr / 8.0 GB	138.6 hr / 21.4 GB	715.7 hr / 56.7 GB
Eagle2 $K=5K$	10.4 hr / 4.6 GB	32.8 hr / 7.4 GB	115.9 hr / 17.1 GB	485.4 hr / 45.8 GB
Eagle2 $K=10K$	12.2 hr / 5.3 GB	39.7 hr / 7.8 GB	134.7 hr / 17.1 GB	540.5 hr / 45.8 GB
Eagle2 $K=20K$	13.9 hr / 5.1 GB	50.7 hr / 8.4 GB	172.5 hr / 17.1 GB	656.2 hr / 45.8 GB
Eagle2 $K=40K$	14.3 hr / 4.6 GB	64.7 hr / 9.1 GB	251.8 hr / 17.2 GB	890.6 hr / 45.8 GB

(b) Mean switch error rate (s.e.m.)

Method	$N = 5K$	$N = 15K$	$N = 50K$	$N = 150K$
SHAPEIT2	1.38% (0.03%)	0.86% (0.03%)	–	–
Eagle1	2.38% (0.05%)	1.30% (0.04%)	0.59% (0.03%)	0.31% (0.02%)
Eagle2 $K=5K$	1.35% (0.03%)	0.99% (0.03%)	0.64% (0.03%)	0.40% (0.02%)
Eagle2 $K=10K$	1.23% (0.03%)	0.83% (0.02%)	0.52% (0.02%)	0.33% (0.02%)
Eagle2 $K=20K$	1.23% (0.03%)	0.75% (0.02%)	0.46% (0.02%)	0.29% (0.02%)
Eagle2 $K=40K$	1.23% (0.03%)	0.73% (0.02%)	0.41% (0.02%)	0.26% (0.02%)

(This table provides numeric data plotted in Figure 5 along with benchmarks for additional parameter settings of Eagle2.) We benchmarked Eagle2 and other available phasing methods on $N=5,000$, $15,000$, $50,000$, and $150,000$ UK Biobank samples (including trio children and excluding trio parents). (a) Total wall clock time and maximum memory used for genome-wide phasing on a 16-core 2.60 GHz Intel Xeon E5-2650 v2 processor. (We analyzed a total of 174,595 markers on chromosomes 1, 5, 10, 15, and 20, representing $\approx 25\%$ of the genome, and scaled up running times by a factor of 4; see Supplementary Table 8 for per-chromosome data.) SHAPEIT2 was unable to complete the $N=50,000$ chr1 and chr5 analyses and was unable to complete any of the $N=150,000$ analyses in 5 days, the run time limit for single compute jobs. (b) Mean switch error rate (s.e.m.) over 70 European-ancestry trios.

Supplementary Table 8. Per-chromosome computational cost and accuracy of non-reference-based phasing in the UK Biobank cohort.

(a) Running time and memory cost

<i>N</i>	chrom	SHAPEIT2	Eagle1	Eagle2			
				<i>K</i> =5K	<i>K</i> =10K	<i>K</i> =20K	<i>K</i> =40K
5K	chr20	1.5 hr / 1.8 GB	0.2 hr / 1.3 GB	0.3 hr / 1.6 GB	0.3 hr / 1.9 GB	0.4 hr / 1.8 GB	0.4 hr / 1.8 GB
5K	chr15	1.8 hr / 2.1 GB	0.3 hr / 1.7 GB	0.4 hr / 2.0 GB	0.4 hr / 2.2 GB	0.5 hr / 2.2 GB	0.4 hr / 2.2 GB
5K	chr10	2.8 hr / 3.2 GB	0.4 hr / 2.4 GB	0.5 hr / 2.8 GB	0.6 hr / 3.1 GB	0.7 hr / 3.2 GB	0.8 hr / 2.7 GB
5K	chr5	3.4 hr / 3.9 GB	0.4 hr / 3.0 GB	0.6 hr / 3.4 GB	0.7 hr / 3.9 GB	0.8 hr / 3.8 GB	0.8 hr / 3.9 GB
5K	chr1	4.5 hr / 5.3 GB	0.7 hr / 4.4 GB	0.8 hr / 4.6 GB	1.0 hr / 5.3 GB	1.2 hr / 5.1 GB	1.2 hr / 4.6 GB
15K	chr20	5.9 hr / 5.1 GB	0.8 hr / 2.9 GB	0.9 hr / 2.7 GB	1.0 hr / 2.8 GB	1.4 hr / 3.1 GB	1.8 hr / 3.6 GB
15K	chr15	7.4 hr / 6.1 GB	1.0 hr / 3.4 GB	1.1 hr / 3.2 GB	1.3 hr / 3.3 GB	1.6 hr / 3.8 GB	2.1 hr / 4.1 GB
15K	chr10	10.8 hr / 9.5 GB	1.5 hr / 5.1 GB	1.7 hr / 4.5 GB	1.9 hr / 4.7 GB	2.5 hr / 5.3 GB	3.2 hr / 5.7 GB
15K	chr5	13.8 hr / 11.3 GB	1.7 hr / 6.1 GB	1.9 hr / 5.4 GB	2.2 hr / 5.7 GB	3.0 hr / 6.0 GB	3.9 hr / 6.8 GB
15K	chr1	19.8 hr / 15.2 GB	2.6 hr / 8.0 GB	2.6 hr / 7.4 GB	3.5 hr / 7.8 GB	4.3 hr / 8.4 GB	5.2 hr / 9.1 GB
50K	chr20	46.5 hr / 16.8 GB	3.9 hr / 7.7 GB	3.1 hr / 5.7 GB	3.6 hr / 5.7 GB	4.6 hr / 5.7 GB	6.8 hr / 5.8 GB
50K	chr15	59.0 hr / 20.2 GB	4.6 hr / 9.0 GB	3.8 hr / 7.0 GB	4.4 hr / 7.0 GB	5.7 hr / 7.0 GB	8.3 hr / 7.0 GB
50K	chr10	93.8 hr / 31.4 GB	6.6 hr / 12.9 GB	5.7 hr / 10.3 GB	6.6 hr / 10.3 GB	8.5 hr / 10.3 GB	12.4 hr / 10.3 GB
50K	chr5	–	8.2 hr / 15.4 GB	6.8 hr / 12.3 GB	8.0 hr / 12.3 GB	10.3 hr / 12.3 GB	15.0 hr / 12.3 GB
50K	chr1	–	11.4 hr / 21.4 GB	9.5 hr / 17.1 GB	11.0 hr / 17.1 GB	14.0 hr / 17.1 GB	20.5 hr / 17.2 GB
150K	chr20	–	19.9 hr / 20.1 GB	12.8 hr / 15.0 GB	14.4 hr / 15.0 GB	17.4 hr / 15.0 GB	23.7 hr / 15.0 GB
150K	chr15	–	24.4 hr / 24.1 GB	15.8 hr / 18.3 GB	17.7 hr / 18.3 GB	21.6 hr / 18.3 GB	29.2 hr / 18.3 GB
150K	chr10	–	35.2 hr / 34.2 GB	23.8 hr / 27.7 GB	26.7 hr / 27.6 GB	32.4 hr / 27.6 GB	44.0 hr / 27.7 GB
150K	chr5	–	41.3 hr / 40.5 GB	28.6 hr / 33.0 GB	32.1 hr / 32.9 GB	39.0 hr / 32.9 GB	53.2 hr / 33.0 GB
150K	chr1	–	58.1 hr / 56.7 GB	40.3 hr / 45.8 GB	44.3 hr / 45.8 GB	53.7 hr / 45.8 GB	72.4 hr / 45.8 GB

(b) Mean switch error rate (s.e.m.)

<i>N</i>	chrom	SHAPEIT2	Eagle1	Eagle2			
				<i>K</i> =5K	<i>K</i> =10K	<i>K</i> =20K	<i>K</i> =40K
5K	chr20	1.73% (0.06%)	2.86% (0.09%)	1.62% (0.06%)	1.46% (0.05%)	1.46% (0.05%)	1.46% (0.05%)
5K	chr15	1.66% (0.06%)	2.81% (0.09%)	1.60% (0.05%)	1.48% (0.05%)	1.46% (0.05%)	1.46% (0.05%)
5K	chr10	1.33% (0.04%)	2.29% (0.07%)	1.29% (0.04%)	1.18% (0.04%)	1.20% (0.04%)	1.20% (0.04%)
5K	chr5	1.21% (0.04%)	2.18% (0.07%)	1.22% (0.04%)	1.11% (0.04%)	1.11% (0.04%)	1.11% (0.04%)
5K	chr1	1.31% (0.04%)	2.26% (0.06%)	1.29% (0.04%)	1.16% (0.04%)	1.16% (0.03%)	1.16% (0.03%)
15K	chr20	1.05% (0.05%)	1.54% (0.06%)	1.11% (0.05%)	0.95% (0.04%)	0.87% (0.04%)	0.85% (0.04%)
15K	chr15	1.04% (0.05%)	1.52% (0.07%)	1.12% (0.05%)	0.99% (0.04%)	0.92% (0.04%)	0.88% (0.04%)
15K	chr10	0.82% (0.03%)	1.28% (0.05%)	0.97% (0.04%)	0.82% (0.03%)	0.74% (0.03%)	0.72% (0.03%)
15K	chr5	0.79% (0.03%)	1.18% (0.05%)	0.93% (0.04%)	0.76% (0.03%)	0.67% (0.03%)	0.65% (0.03%)
15K	chr1	0.81% (0.03%)	1.24% (0.04%)	0.95% (0.03%)	0.80% (0.03%)	0.71% (0.03%)	0.69% (0.03%)
50K	chr20	0.53% (0.04%)	0.67% (0.05%)	0.68% (0.04%)	0.56% (0.04%)	0.50% (0.04%)	0.45% (0.03%)
50K	chr15	0.57% (0.03%)	0.70% (0.04%)	0.72% (0.04%)	0.59% (0.04%)	0.53% (0.03%)	0.48% (0.03%)
50K	chr10	0.47% (0.02%)	0.59% (0.03%)	0.60% (0.03%)	0.51% (0.03%)	0.46% (0.03%)	0.41% (0.02%)
50K	chr5	–	0.53% (0.03%)	0.58% (0.03%)	0.47% (0.03%)	0.41% (0.02%)	0.36% (0.02%)
50K	chr1	–	0.55% (0.03%)	0.67% (0.03%)	0.54% (0.02%)	0.46% (0.03%)	0.40% (0.02%)
150K	chr20	–	0.32% (0.03%)	0.36% (0.03%)	0.31% (0.03%)	0.28% (0.02%)	0.26% (0.02%)
150K	chr15	–	0.36% (0.03%)	0.44% (0.03%)	0.39% (0.03%)	0.33% (0.02%)	0.30% (0.02%)
150K	chr10	–	0.32% (0.02%)	0.38% (0.03%)	0.31% (0.02%)	0.28% (0.02%)	0.25% (0.02%)
150K	chr5	–	0.29% (0.02%)	0.38% (0.02%)	0.31% (0.02%)	0.26% (0.02%)	0.24% (0.02%)
150K	chr1	–	0.29% (0.02%)	0.43% (0.03%)	0.34% (0.03%)	0.30% (0.02%)	0.27% (0.02%)

This table provides per-chromosome breakdowns of computational cost and phasing accuracy plotted in Figure 5 along with benchmarks for additional parameter settings of Eagle2. See the captions of Figure 5 and Supplementary Table 7 for benchmarking details.

Supplementary Table 9. Accuracy of non-reference-based phasing in the GERA sub-cohorts.

Method	European chip <i>N</i> =62,084	African chip <i>N</i> =3,823	East Asian chip <i>N</i> =5,178	Latino chip <i>N</i> =7,150
SHAPEIT2 <i>K</i> =100	–	2.85% (0.07%)	2.02% (0.11%)	2.23% (0.05%)
Eagle1	1.34% (0.05%)	4.07% (0.10%)	3.09% (0.16%)	2.76% (0.06%)
Eagle2 <i>K</i> =5K	1.36% (0.04%)	2.81% (0.06%)	2.07% (0.11%)	2.21% (0.04%)
Eagle2 <i>K</i> =10K	1.26% (0.04%)	2.52% (0.05%)	1.97% (0.10%)	2.08% (0.04%)
Eagle2 <i>K</i> =20K	1.18% (0.04%)	2.50% (0.05%)	1.99% (0.11%)	2.06% (0.04%)

We benchmarked Eagle2 and other available phasing methods on the four GERA sub-cohorts (excluding all non-trio-parent samples from each trio family). For the European chip sub-cohort, we analyzed a total of 161,232 markers on chromosomes 1, 5, 10, 15, and 20; for the other sub-cohorts, we analyzed all 22 autosomes. Standard errors for the European-ancestry sub-cohort are over 400 parent samples. Standard errors for the other three sub-cohorts are over 25 SNP blocks. We did not benchmark SHAPEIT2 on the European-ancestry sub-cohort, as it would have required more than two weeks of computation.