

DEPARTMENT OF COMPUTER SCIENCE
SERIES OF PUBLICATIONS C
REPORT C-2007-57

Combining haplotypers

Matti Kääriäinen

Niels Landwehr

Sampsa Lappalainen

Taneli Mielikäinen

UNIVERSITY OF HELSINKI
FINLAND

Combining haplotypers

Matti Kääriäinen
HIIT Basic Research Unit
Department of Computer Science
University of Helsinki
matti.kaariainen@cs.helsinki.fi

Niels Landwehr
Machine Learning Lab
Department of Computer Science
Katholieke Universiteit Leuven
niels.landwehr@cs.kuleuven.be

Sampsa Lappalainen
HIIT Basic Research Unit
Department of Computer Science
University of Helsinki
sampsa.lappalainen@cs.helsinki.fi

Taneli Mielikäinen
Department of Computer Science
University of Helsinki and
Nokia Research Center Palo Alto
taneli.mielikainen@iki.fi

Department of Computer Science, University of Helsinki
Technical report, Series of Publications C, Report C-2007-57
Helsinki, September 2007, ii + 24 pages

Abstract

Statistically resolving the underlying haplotype pair for a genotype measurement is an important intermediate step in gene mapping studies, and has received much attention recently. Consequently, a variety of methods for this problem have been developed. Different methods employ different statistical models, and thus implicitly encode different assumptions about the nature of the underlying haplotype structure. Depending on the population sample in question, their relative performance can vary greatly, and it is unclear which method to choose for a particular sample. Instead of choosing a single method, we explore combining predictions returned by different methods in a principled way, and thereby circumvent the problem of method selection.

We propose several techniques for combining haplotype reconstructions and analyze their computational properties. In an experimental study on real-world haplotype data we show that such techniques can provide more accurate and robust reconstructions, and are useful for outlier detection. Typically, the combined prediction is at least as accurate as or even more accurate than the best individual method, effectively circumventing the method selection problem.

Computing Reviews (1998) Categories and Subject Descriptors:

F.2.2 Analysis of Algorithms and Problem Complexity: Nonnumerical Algorithms and Problems

I.2.6 Artificial Intelligence: Learning

J.3 Life and Medical Sciences: Biology and Genetics

General Terms:

Algorithms, Experimentation, Theory

Additional Key Words and Phrases:

Haplotyping, Ensemble Methods, Sequence Prediction

1 Introduction

Complex diseases such as Diabetes or Alzheimer’s disease are often linked to individual genetic variations. The analysis of genetic variation in human populations is therefore critical for understanding individual risk factors for such diseases. Most of the human genome is invariant among individuals, and it is sufficient to concentrate on small parts of the whole genome sequence to analyze genetic variation. Frequently studied differences are *single nucleotide polymorphisms* (SNPs), which are single-nucleotide variations at a particular location in the genome. The positions in the sequence are called *markers* and the different possible values *alleles*. A *haplotype* is a sequence of SNP alleles along a chromosome, and concisely represent the variable genetic information in that region. In the search for DNA sequence variants which are related to common diseases, haplotype-based approaches have become a central theme [The05].

Diploid human cells have two *homologous* (i.e., almost identical) copies of each chromosome. Current practical laboratory measurement techniques produce a *genotype*—for m markers, a sequence of m unordered pairs of alleles. A genotype reveals the two alleles that are present at each marker, but not their respective chromosome origin. To obtain haplotypes from genotype data, this hidden phase information has to be reconstructed. Two alternative approaches exist: If family trios are available, most of the ambiguity in the haplotype pair can be resolved analytically. Otherwise, population-based statistical methods have to be used to estimate the haplotype pair. Because trios are more difficult to recruit and more expensive to genotype, the population-based haplotyping approach is often the only cost-effective method for large-scale studies.

The haplotyping problem has received a lot of attention recently, and many different haplotyping methods have been proposed [SS05, SS06, KS05a, EGT06, RKMU05]. All of these methods employ different statistical models, which reflect different assumptions about the underlying distribution over haplotypes in a population sample. Furthermore, the methods offer different trade-offs in terms of reconstruction accuracy and scaling behavior in the number of markers and individuals in the sample. On the other hand, the statistical properties of haplotype “datasets” (a particular set of markers genotyped for a particular set of individuals) vary depending on marker spacing, sample size and population characteristics. In fact, some haplotyping methods have been specifically tailored to particular dataset characteristics. For example, the HIT system [RKMU05] is especially effective for population isolates, and the HaploRec system [EGT06] for reconstruction of large, possibly genome-wide marker maps.

It is therefore unlikely that there is one haplotyping method which is generally superior. Instead, the relative performance of different methods will vary depending on the characteristics of the dataset to be haplotyped. In contrast to

other statistical modeling tasks, in haplotyping there is typically no “training data” available for which the ground truth is known. This precludes the use of model selection techniques such as cross-validation (although it is possible to use cross-validation estimates of performance on related tasks such as missing genotype imputation for model selection, see e.g. [SS06]). Nevertheless, one often has to commit to just one haplotype reconstruction in the end. Hence, it is natural to ask whether the predictions of the different methods could be combined in a simple way to give more accurate and robust haplotype reconstructions without having to know in advance which of the baseline methods performs well on the dataset at hand.

In this paper we study how to combine haplotype reconstructions produced by various methods. We formulate several approaches for combining haplotypes, study the algorithmics of the problem, and experimentally validate that combining haplotypes is beneficial.

2 Population-based haplotyping

A haplotype h can be represented as a sequence of alleles $h[i]$ in markers $i = 1, \dots, m$. For most SNP markers, only two alternative nucleotides (alleles) occur in a population, so we can assume $h \in \{0, 1\}^m$. A genotype g for an individual can be represented as a sequence of unordered pairs $g[i] = \{h_g^1[i], h_g^2[i]\}$ of alleles in markers $i = 1, \dots, m$. Hence, $g \in \{\{0, 0\}, \{1, 1\}, \{0, 1\}\}^m$. A marker with alleles $\{0, 0\}$ or $\{1, 1\}$ is *homozygous* whereas a marker with alleles $\{0, 1\}$ is *heterozygous*. We denote the number of heterozygous markers by m' , and their positions in the haplotype sequence by $i_1, \dots, i_{m'}$.

The haplotyping problem arises from the fact that while each haplotype pair corresponds to a unique genotype, a genotype may correspond to a large number of different haplotype pairs. Population-based haplotyping is the task of statistically resolving this ambiguity:

The haplotype reconstruction problem: Given a multiset \mathcal{G} of genotypes, find for each genotype $g \in \mathcal{G}$ the haplotypes h_g^1 and h_g^2 that have generated g .

For the rest of the paper we will denote the two individual haplotypes in a haplotype pair as h^1 and h^2 , and use h as a shorthand to denote the pair $\{h^1, h^2\}$ when there is no ambiguity. Furthermore, we denote a substring $s[i]s[i+1] \dots s[i+k]$ of a string s by $s[i, k]$.

For each genotype $g \in \{0, 1\}^m$, there are $2^{m'-1}$ different haplotype reconstructions. Only one of these reconstructions is correct, so inferring the haplotypes is clearly impossible without additional information or assumptions. These assumptions are typically inspired by population genetics, and can take either

a combinatorial or a probabilistic form. The models borrowed from population genetics are often rather simplistic abstractions of the complicated reality. Furthermore, additional simplifications and heuristics may be needed to make haplotype inference computationally tractable. The number of ways to combine these choices—which of the imperfect population genetics models to build on and which computational strategies to use—has led to the development of a large and diverse set of different haplotyping methods, each with their own advantages. The following lists just a few prominent examples.

The currently most widely used method PHASE [SSD01, SD03, SS05] is based on quite sophisticated probabilistic models and is computationally expensive; fastPHASE [SS06], a more efficient but still almost as accurate method has been published recently. Several other methods have recently been developed. GERBIL [KS05a, KS05b] is based on reconstructing block partitioning and resolving the haplotypes simultaneously. HAP [HE04] implements a method based on imperfect phylogeny. HIT [RKMU05] and HINT [KS05c] use HMM founder models for haplotyping. HAPLOREC [EGT04, EGT06] is based on variable-length Markov chains. SPAMM [LME⁺06, LME⁺07] is an approach based on levelwise construction of constrained Hidden Markov Models.

3 Combining haplotypers

In practice, genetics researchers often face the problem that different haplotype reconstruction methods give different results and there is no straightforward way to decide which method to choose. Due to the varying characteristics of haplotyping datasets, it is unlikely that one haplotyping method is generally superior. Instead, different methods have different relative strengths and weaknesses, and will fail in different parts of the reconstruction.

The promise of ensemble methods lies in “averaging out” those errors, as far as they are specific to a small subset of methods (rather than a systematic error affecting all methods). This intuition can be made precise by making probabilistic assumptions about how the reconstruction methods err: If the errors in the reconstructions were small random perturbations of the true haplotype pair, taking a majority vote (in an appropriate sense depending on the type of perturbations) of sufficiently many reconstructions would with high probability correct all the errors. While such probabilistic assumptions are not true in practice, they serve as a guideline and motivation for the combination methods we derive next.

The idea of using ensemble methods in haplotyping is not entirely new. It is used in existing systems for combining results from several random restarts of a method [SS06, Gus02], or to obtain a point estimate from an inferred posterior distribution on haplotypes [SSD01]. However, to the best of our knowledge, our approach of combining unrelated haplotypers—and thus gaining the benefits

of their potentially orthogonal strengths—has not been studied before.

The haplotyper combination problem can be viewed as an instance of the general problem of finding a consensus object for a given collection of objects. In the simplest case the objects are individual predictions as in ensemble methods in machine learning [Bre96, CSS02, SG02]. The objects can also be more complicated structures such as sequences [JXL04, LP02, SP03], rankings [DKNS01, FKM⁺04, FISS03], clusterings [ACN05, GMT05, LB05], or segmentations [MTT06]. Although sequential prediction has been studied a lot, there exists little work on ensemble methods for sequence prediction. Our approach to haplotyper combination resembles closely the work on combining part-of-speech taggers [Sjö03] to improve tagging accuracy. However, due to the nature of haplotype data, we need more refined strategies than simple position-wise voting.

3.1 Problem definitions

To combine the haplotypings suggested by l given baseline haplotype reconstruction methods, we formulate two computational problems. We limit ourselves to combination methods that process each individual separately, thus enabling immediate parallelization of the combination strategies for large populations.

Problem 1 (Haplotyper combination). Given the haplotype reconstructions $\{h_1^1, h_1^2\}, \dots, \{h_l^1, h_l^2\} \subseteq \{0, 1\}^m$, and a distance function $d : \{0, 1\}^m \times \{0, 1\}^m \rightarrow \mathbb{R}_{\geq 0}$, find:

- HVP: a reconstruction $\{h^1, h^2\} \subseteq \{0, 1\}^m$ minimizing the sum of distances, i.e., find

$$\{h^1, h^2\} = \underset{h_i^1, h_i^2 \in \{0, 1\}}{\operatorname{argmin}} \sum_{i=1}^l d(\{h_i^1, h_i^2\}, \{h^1, h^2\}).$$

- HSP: a reconstruction $\{h_i^1, h_i^2\}, i \in \{1, \dots, l\}$ minimizing the sum of distances, i.e., find

$$i = \underset{j \in \{1, \dots, l\}}{\operatorname{argmin}} \sum_{i=1}^l d(\{h_i^1, h_i^2\}, \{h_j^1, h_j^2\}).$$

In both cases, ties are broken arbitrarily.

The difference between the Haplotyper Voting Problem (HVP) and the Haplotyper Selection Problem (HSP) is that in the latter, the solution is required to be one of the input haplotype reconstructions. Using clustering terminology, the Haplotyper Voting Problem (HVP) seeks for the average haplotype

reconstruction based on the input haplotypings, whereas the Haplotype Selection Problem (HSP) selects the median haplotype reconstruction as the most plausible haplotyping. The exact meaning of average and median depends, of course, on the properties of d . Ideally, d should be such that the solutions to HVP and HSP can be found efficiently and are close to the unknown true haplotypes. We will discuss viable candidates for such d later in Section 3.2. While HSP can be solved efficiently by brute force provided that d can be computed efficiently, the computational aspects of HVP depend heavily on d . Thus, their discussion will be postponed to Section 3.3.

HVP and HSP are closely related for all distance functions d . A solution to HSP is a 2-approximation of a solution to HVP and the solution to HVP can be transformed into a 2-approximation of a solution to HSP.

Proposition 1. *Let d be a distance function between haplotype pairs satisfying the triangle inequality. Let $h_1 = \{h_1^1, h_1^2\}, \dots, h_l = \{h_l^1, h_l^2\}$ be the haplotype pairs to combine and $h_{\text{HVP}} = \{h_{\text{HVP}}^1, h_{\text{HVP}}^2\}$, $h_{\text{HSP}} = \{h_{\text{HSP}}^1, h_{\text{HSP}}^2\}$ the optimal HVP and HSP solutions. Then*

1. h_{HVP} is a feasible solution of HVP and

$$\sum_{i=1}^l d(h_i, h_{\text{HVP}}) \leq \sum_{i=1}^l d(h_i, h_{\text{HSP}}) \leq 2 \sum_{i=1}^l d(h_i, h_{\text{HVP}}).$$

2. $h_j = \operatorname{argmin}_{i=1, \dots, l} d(h_i, h_{\text{HVP}})$ is a feasible solution of HSP and

$$\sum_{i=1}^l d(h_i, h_{\text{HSP}}) \leq \sum_{i=1}^l d(h_i, h_j) \leq 2 \sum_{i=1}^l d(h_i, h_{\text{HSP}}).$$

Proof. We have $\sum_{i=1}^l d(h_i, h_{\text{HVP}}) \leq \sum_{i=1}^l d(h_i, h_{\text{HSP}})$ because

$$\begin{aligned} \sum_{i=1}^l d(h_i, h_{\text{HVP}}) &= \min_{h^1, h^2 \in \{0,1\}^m} \sum_{i=1}^l d(h_i, h) \\ &\leq \min_{j \in \{1, \dots, l\}} \sum_{i=1}^l d(h_i, h_j) \\ &= \sum_{i=1}^l d(h_i, h_{\text{HSP}}). \end{aligned}$$

To see that $\sum_{i=1}^l d(h_i, h_{\text{HSP}}) \leq 2 \sum_{i=1}^l d(h_i, h_{\text{HVP}})$, note that there must be a haplotype pair $h_j = \{h_j^1, h_j^2\}$, $j \in \{1, \dots, l\}$ such that

$$d(h_j, h_{\text{HVP}}) \leq \frac{1}{l} \sum_{i=1}^l d(h_i, h_{\text{HVP}}).$$

Hence,

$$\sum_{i=1}^l d(h_i, h_j) \leq \sum_{i=1}^l d(h_i, h_{\text{HVP}}) + ld(h_j, h_{\text{HVP}}) \leq 2 \sum_{i=1}^l d(h_i, h_{\text{HVP}}).$$

Similarly $\sum_{i=1}^l d(h_i, h_{\text{HSP}}) \leq \sum_{i=1}^l d(h_i, h_j)$ because

$$\sum_{i=1}^l d(h_i, h_{\text{HSP}}) = \min_{j \in \{1, \dots, l\}} \sum_{i=1}^l d(h_i, h_j) \leq \sum_{i=1}^l d(h_i, h_{j'})$$

for any $j' \in \{1, \dots, l\}$.

To see that $\sum_{i=1}^l d(h_i, h_j) \leq 2 \sum_{i=1}^l d(h_i, h_{\text{HSP}})$, note that

$$\min_{j \in \{1, \dots, l\}} d(h_j, h_{\text{HVP}}) \leq \frac{1}{l} \sum_{i=1}^l d(h_i, h_{\text{HVP}}).$$

Hence,

$$\begin{aligned} \sum_{i=1}^l d(h_i, h_j) &\leq \sum_{i=1}^l d(h_i, h_{\text{HSP}}) + \sum_{i=1}^l d(h_j, h_{\text{HVP}}) \\ &= \sum_{i=1}^l d(h_i, h_{\text{HSP}}) + ld(h_j, h_{\text{HVP}}) \\ &\leq \sum_{i=1}^l d(h_i, h_{\text{HSP}}) + \sum_{i=1}^l d(h_i, h_{\text{HVP}}) \\ &\leq 2 \sum_{i=1}^l d(h_i, h_{\text{HSP}}). \end{aligned}$$

□

3.2 Distance functions

In order to define average and median haplotypings, we need to choose a distance function d for measuring the similarity between haplotype sequences. To satisfy the intuition that the solutions to HSP and HVP should be on average close to the baseline haplotype reconstructions, we will focus only on a small set of distance measures d that are reasonable candidates for measuring genetic distance between haplotype pairs.

Hamming distance and other distances induced by distances on sequences. The most common distance measure between sequences $s, t \in \Sigma^m$

is the Hamming distance that counts the number of disagreements between s and t :

$$d_H(s, t) = |\{i \in \{1, \dots, m\} : s[i] \neq t[i]\}|.$$

The Hamming distance is not directly applicable as a measure of genetic distance between individuals, because the haplotypes corresponding to an individual's genotype form an unordered pair. To define a Hamming distance between unordered pairs of haplotypes, let us consider haplotype pairs $\{h_1^1, h_1^2\}$ and $\{h_2^1, h_2^2\}$. The distance between the pairs should be zero if the sets $\{h_1^1, h_1^2\}$ and $\{h_2^1, h_2^2\}$ are the same. Hence, we should try both ways to pair the haplotypes and take the one with the smaller distance, i.e.,

$$d_H(\{h_1^1, h_1^2\}, \{h_2^1, h_2^2\}) = \min\{d_H(h_1^1, h_2^1) + d_H(h_1^2, h_2^2), d_H(h_1^1, h_2^2) + d_H(h_1^2, h_2^1)\}.$$

Note that a similar construction can be used to map any distance function between haplotype sequences to a distance function between pairs of haplotypings. Furthermore, the next proposition shows that if the distance function between the sequences satisfies the triangle inequality, so does the corresponding distance function for haplotype reconstructions.

Proposition 2. *Let $d: \Sigma^m \times \Sigma^m \rightarrow \mathbb{R}_{\geq 0}$ be a distance function between sequences of length Σ^m and let*

$$d(\{h_1^1, h_1^2\}, \{h_2^1, h_2^2\}) = \min\{d(h_1^1, h_2^1) + d(h_1^2, h_2^2), d(h_1^1, h_2^2) + d(h_1^2, h_2^1)\}$$

for all $h_1^1, h_1^2, h_2^1, h_2^2 \in \Sigma^m$. If d satisfies the triangle inequality for comparing sequences, i.e., $d(s, t) \leq d(s, u) + d(t, u)$ for all $s, t, u \in \Sigma^m$, then d satisfies the triangle inequality for comparing unordered pairs of sequences $d(h_1, h_2) \leq d(h_1, h_3) + d(h_2, h_3)$ for all $h_1^1, h_1^2, h_2^1, h_2^2, h_3^1, h_3^2 \in \Sigma^m$.

Proof. Choose arbitrary sequences $h_1^1, h_1^2, h_2^1, h_2^2, h_3^1, h_3^2 \in \Sigma^m$. We show that the claim holds for them and hence for all sequences of length m over the alphabet Σ .

Assume, without loss of generality, that $d(\{h_1^1, h_1^2\}, \{h_2^1, h_2^2\}) = d(h_1^1, h_2^1) + d(h_1^2, h_2^2)$ and $d(\{h_1^1, h_1^2\}, \{h_3^1, h_3^2\}) = d(h_1^1, h_3^1) + d(h_1^2, h_3^2)$.

For $d(\{h_2^1, h_2^2\}, \{h_3^1, h_3^2\})$ there are two cases as it is the minimum of $d(h_2^1, h_3^1) + d(h_2^2, h_3^2)$ and $d(h_2^2, h_3^1) + d(h_2^1, h_3^2)$.

If $d(\{h_2^1, h_2^2\}, \{h_3^1, h_3^2\}) = d(h_2^1, h_3^1) + d(h_2^2, h_3^2)$, then

$$\begin{aligned} & d(\{h_1^1, h_1^2\}, \{h_3^1, h_3^2\}) + d(\{h_2^1, h_2^2\}, \{h_3^1, h_3^2\}) \\ &= d(h_1^1, h_3^1) + d(h_1^2, h_3^2) + d(h_2^1, h_3^1) + d(h_2^2, h_3^2) \\ &= [d(h_1^1, h_3^1) + d(h_2^1, h_3^1)] + [d(h_1^2, h_3^2) + d(h_2^2, h_3^2)] \\ &\geq d(h_1^1, h_2^1) + d(h_1^2, h_2^2). \end{aligned}$$

If $d(\{h_2^1, h_2^2\}, \{h_3^1, h_3^2\}) = d(h_2^2, h_3^1) + d(h_2^1, h_3^2)$, then

$$\begin{aligned} & d(\{h_1^1, h_1^2\}, \{h_3^1, h_3^2\}) + d(\{h_2^1, h_2^2\}, \{h_3^1, h_3^2\}) \\ &= d(h_1^1, h_3^1) + d(h_1^2, h_3^2) + d(h_2^2, h_3^1) + d(h_2^1, h_3^2) \\ &= [d(h_1^1, h_3^1) + d(h_2^2, h_3^1)] + [d(h_1^2, h_3^2) + d(h_2^1, h_3^2)] \\ &\geq d(h_1^1, h_2^2) + d(h_1^2, h_2^1) \geq d(h_1^1, h_2^1) + d(h_1^2, h_2^2). \end{aligned}$$

Thus, the claim holds. \square

Switch distance. The approach of defining distance functions between haplotype pairs based on distance functions between haplotypes has some limitations, independently of the distance function used. This is because much of the variance in haplotypes originates from *chromosomal crossover* during meiosis, which breaks up chromosomes and reconnects the resulting segments to form new chromosomes for the offspring. The chromosome pair resulting from a crossover could be seen as genetically close to the original pair even if the individual sequences do not match very well. *Switch distance* is a distance measure for haplotype pairs that takes such similarities into account. It is defined as the number of *switches* that are needed to transform a haplotype pair to another haplotype pair with the same homozygous and heterozygous markers. A switch between markers i and $i + 1$ for a haplotype pair $\{h^1, h^2\}$ transforms the pair $\{h^1, h^2\} = \{h^1[1, i]h^1[i + 1, m], h^2[1, i]h^2[i + 1, m]\}$ into the pair $\{h^1[1, i]h^2[i + 1, m], h^2[1, i]h^1[i + 1, m]\}$. It is easy to see that for any pair of haplotype reconstructions corresponding to the same genotype, there is a sequence of switches transforming one into the other. Thus, switch distance is well defined for the cases we are interested in.

The switch distance has the advantage over the Hamming distance that the order of the haplotypes in the haplotype pair does not matter in the distance computation: the haplotype pair can be encoded uniquely as a bit sequence consisting of just the switches between the consecutive heterozygous markers, i.e., as a *switch sequence*:

Definition 1 (Switch sequence). Let $h^1, h^2 \in \{0, 1\}^m$ and let $i_1 < \dots < i_{m'}$ be the heterozygous markers in $\{h^1, h^2\}$. The switch sequence of a haplotype pair $\{h^1, h^2\}$ is a sequence $s(h^1, h^2) = s(h^2, h^1) = s \in \{0, 1\}^{m'-1}$ such that

$$s[j] = \begin{cases} 0 & \text{if } h^1[i_j] = h^1[i_{j+1}] \text{ and } h^2[i_j] = h^2[i_{j+1}] \\ 1 & \text{if } h^1[i_j] \neq h^1[i_{j+1}] \text{ and } h^2[i_j] \neq h^2[i_{j+1}] \end{cases}$$

The switch distance between haplotype reconstructions can be defined in terms of the Hamming distance between switch sequences as follows.

Definition 2 (Switch distance). Let $h_1 = \{h_1^1, h_1^2\}$ and $h_2 = \{h_2^1, h_2^2\}$ be haplotype pairs corresponding to the same genotype. The switch distance between the pairs is $d_s(h_1, h_2) = d_H(s(h_1^1, h_1^2), s(h_2^1, h_2^2))$.

As switch distance is the Hamming distance between the switch sequences, the following proposition is immediate:

Proposition 3. *The switch distance satisfies the triangle inequality.*

k -Hamming distance. Switch distance considers only a very small neighborhood of each marker, namely only the previous and the next heterozygous marker in the haplotype. On the other extreme, the Hamming distance uses the complete neighborhood (via the min operation), i.e., the whole haplotypes for each marker. The intermediate cases are covered by the following k -Hamming distance in which all windows of a chosen length $k \in \{2, \dots, m\}$ are considered. The intuition behind the definition is that each window of length k is a potential location for a gene, and we want to measure how close the haplotype reconstruction $\{h^1, h^2\}$ gets to the true haplotype $\{h_1^1, h_2^1\}$ in predicting each of these potential genes.

Definition 3 (k -Hamming distance). Let $\{h_1^1, h_1^2\}$ and $\{h_2^1, h_2^2\}$ be pairs of haplotype sequences corresponding to the same genotype with m' heterozygous markers in positions i_1, \dots, i_m . The k -Hamming distance d_{k-H} between $\{h_1^1, h_1^2\}$ and $\{h_2^1, h_2^2\}$ is defined by

$$d_{k-H}(h_1, h_2) = \sum_{j=1}^{m'-k+1} d_H(h_1[i_j, \dots, i_{j+k-1}], h_2[i_j, \dots, i_{j+k-1}])$$

unless $m' < k$, in which case $d_{k-H}(h_1, h_2) = d_H(h_1, h_2)$.

It is easy to see that $d_{2-H} = 2d_S$, and that for haplotyping pairs with m' heterozygous markers, we have $d_{m'-H} = d_{m-H} = d_H$. Thus, the switch distance and the Hamming distance are the two extreme cases between which d_{k-H} interpolates for $k = 2, \dots, m$.

3.3 Algorithms and complexity

The HSP problem is easily solved by trying out each of the l reconstructions as a candidate solution, and choosing the best one. The complexity of this straightforward strategy is roughly l^2 times the time needed to evaluate the distance function d . Thus, there seems to be no need for more efficient algorithms for HSP in practice.

The complexity of HVP depends on d in a more involved way. As we will show next, for $d = d_S$ a simple voting scheme gives the solution. The rest of the distances considered in Section 3.2 are more challenging. If $d = d_{k-H}$ and k is small, the solution can be found by dynamic programming. For $d = d_{k-H}$ with large k and $d = d_H$, we are aware of no efficient general solutions. However, we will outline methods that can solve most of the problem instances that one may encounter in practice.

Switch distance: $d = d_S$. For the switch distance, the solution to HVP can be found by the following voting scheme:

- Transform the haplotype reconstructions $\{h_i^1, h_i^2\} \subseteq \{0, 1\}^m$, $i = 1, \dots, l$ into switch sequences $s_1, \dots, s_l \in \{0, 1\}^{m'-1}$.
- Return the pair $\{h^1, h^2\}$ that shares the homozygous markers with the reconstructions $\{h_i^1, h_i^2\}$ and whose switch sequence $s \in \{0, 1\}^{m'-1}$ is defined by $s[j] = \underset{b \in \{0,1\}}{\operatorname{argmax}} |\{j \in \{1, \dots, m' - 1\} : s_i[j] = b\}|$.

The time complexity of this method is $O(lm)$.

k -Hamming distance: $d = d_{k-H}$. The optimal solution $h_{\text{HVP}} = \{h_{\text{HVP}}^1, h_{\text{HVP}}^2\}$ of HVP is given by

$$h_{\text{HVP}} = \underset{\{h^1, h^2\} \subseteq \{0,1\}^m}{\operatorname{argmin}} \sum_{i=1}^l d_{k-H}(h_i, h).$$

The number of potentially optimal solutions is $2^{m'}$, but the solution can be constructed incrementally based on the following observation:

$$\begin{aligned} h_{\text{HVP}} &= \underset{\{h^1, h^2\}}{\operatorname{argmin}} \sum_{i=1}^l d_{k-H}(h_i, h) \\ &= \underset{\{h^1, h^2\}}{\operatorname{argmin}} \sum_{i=1}^l \sum_{j=1}^{m'-k+1} d_H(h_i[i_j, \dots, i_{j+k-1}], h[i_j, \dots, i_{j+k-1}]) \end{aligned}$$

Hence, the cost of any solution is a sum of terms

$$D_j(\{x, \bar{x}\}) = \sum_{i=1}^l d_H(h_i[i_j, \dots, i_{j+k-1}], \{x, \bar{x}\}),$$

$j = 1, \dots, m' - k + 1$, $x \in \{0, 1\}^k$ and \bar{x} denotes the complement of x . There are $(m' - k + 1)2^{k-1}$ such terms. Furthermore, the cost of the optimal solution can be computed by dynamic programming using the recurrence relation

$$T_j(\{x, \bar{x}\}) = \begin{cases} 0 & \text{if } j = 0 \\ D_j(\{x, \bar{x}\}) + \min_{b \in \{0,1\}} T_{j-1}(\{bx, \overline{bx}\}) & \text{if } j > 0 \end{cases}$$

Namely, the cost of the optimal solution is $\min_{x \in \{0,1\}^k} T_{m'}(\{x, \bar{x}\})$ and the optimal solution itself can be reconstructed by backtracking the path that leads to this position. The total time complexity for finding the optimal solution using dynamic programming is $\mathcal{O}(lm + 2^k kl(m' - k))$: the heterozygous markers

can be detected and the data can be projected onto them in time $\mathcal{O}(lm)$, and the optimal haplotype reconstruction for the projected data can be computed in time $\mathcal{O}(2^k kl(m' - k))$. So the problem is fixed-parameter tractable¹ in k .

Hamming distance: $d = d_H$. An ordering (h^1, h^2) of an optimal solution $\{h^1, h^2\}$ to HVP with Hamming distance determines an ordering of the unordered input haplotype pairs $\{h_1^1, h_1^2\}, \dots, \{h_l^1, h_l^2\}$. This ordering can be represented by a binary vector $o = (o_1, \dots, o_l) \in \{0, 1\}^l$ that states for each $i = 1, \dots, l$ that the ordering of $\{h_i^1, h_i^2\}$ is $(h_i^{1+o_i}, h_i^{2-o_i})$. Thus, $o_i = \operatorname{argmin}_{b \in \{0,1\}} d_H(h^1, h_i^{1+b})$, where ties are broken arbitrarily.

If the ordering o is known and l is odd, the optimal haplotype reconstruction can be determined in time $\mathcal{O}(lm)$ using the formulae

$$h^1[i] = \operatorname{argmax}_{b \in \{0,1\}} \left| \left\{ j \in \{1, \dots, l\} : h_j^{1+o_j}[i] = b \right\} \right| \quad (1)$$

and

$$h^2[i] = \operatorname{argmax}_{b \in \{0,1\}} \left| \left\{ j \in \{1, \dots, l\} : h_j^{2-o_j}[i] = b \right\} \right|. \quad (2)$$

Hence, solving HVP is polynomial-time equivalent to the task of determining the ordering vector o corresponding to the best haplotype reconstruction $\{h^1, h^2\}$.

The straightforward way to find the optimal ordering is to evaluate the quality of each of the 2^{l-1} non-equivalent orderings. The quality of a single ordering can be evaluated in time $\mathcal{O}(lm)$. Hence, the HVP problem can be solved in total time $\mathcal{O}(lm + 2^l lm')$. The runtime can be reduced to $\mathcal{O}(lm + 2^l m')$ by using Gray codes [Sav97] to enumerate all bit vectors o in such order that consecutive bit vectors differ only by one bit. Hence, the problem is fixed-parameter tractable in l (i.e., in the number of methods). If l is large, however, a more clever strategy is needed. We are unaware of a tractable efficient general solution and suspect that HVP for $d = d_H$ is NP-complete in general. However, we have efficient solutions to two special cases of practical relevance:

Small number of heterozygous markers. If the number of heterozygous positions m' is small, we can simply enumerate all the $2^{m'-1}$ non-equivalent possible solutions to the problem, and pick the optimal one from among them. The time complexity of this approach is $\mathcal{O}(2^{m'} lm')$. Thus, the problem is fixed-parameter tractable also in m' (the number of heterozygous markers).

All reconstructions close to the optimal solution for HVP. Fixing an ordering to any one of the input haplotype reconstructions $\{h_i^1, h_i^2\}$ induces an ordering

¹A problem is called fixed-parameter tractable in a parameter k , if the running time of the algorithm is $f(k)\mathcal{O}(n^c)$ where k is some parameter of the input and c is a constant (and hence not depending on k .) For a good introduction to fixed-parameter tractability and parameterized complexity, see [FG06].

to the remaining input haplotypes. This ordering can be used to compute a solution to HVP through equations 1 and 2. The next proposition shows that the solution obtained in this way is provably optimal if all input haplotype reconstructions are within $m'/2$ of the optimal solution $\{h_{\text{HVP}}^1, h_{\text{HVP}}^2\}$ of HVP.

Proposition 4. *If $d_H(\{h_i^1, h_i^2\}, \{h_{\text{HVP}}^1, h_{\text{HVP}}^2\}) < m'/2$ for each $i \in \{1, \dots, l\}$, then the ordering induced by any of the input haplotype pairs is equivalent to the ordering corresponding to the optimal solution to HVP.*

Proof. By assumption, $d_H(h_i^{1+o_i}, h_{\text{HVP}}^1) < m'/4$ for each $i = 1, \dots, l$ for one of the choices $o_i \in \{0, 1\}$. Then

$$d_H(h_i^{1+o_i}, h_j^{1+o_j}) \leq d_H(h_i^{1+o_i}, h_{\text{HVP}}^1) + d_H(h_j^{1+o_j}, h_{\text{HVP}}^1) < m'/4 + m'/4 = m'/2.$$

Thus, if we use $(h_i^{1+o_i}, h_i^{2-o_i})$ as a reference point, the induced ordering for the haplotypes will be the same o that is induced by using $(h_{\text{HVP}}^1, h_{\text{HVP}}^2)$ as a reference point. Switching the ordering in the reference point to $(h_i^{2-o_i}, h_i^{1+o_i})$ will induce the equivalent ordering $1 - o$. \square

4 Experiments

To investigate the haplotype reconstruction combination problem empirically, real-world genotype data was phased with different haplotyping systems and their reconstructions evaluated. The data was obtained from three sources: a collection of datasets from the **Yoruba** population in Ibadan, Nigeria [The05], the well-known dataset derived from a European population of **Daly** et al. [DRS⁺01], and samples from the recently published **D-HaploDB** haplotype database [HMK⁺07] derived from a Japanese population. For the Yoruba and Daly data, true haplotype pairs were inferred from family trios. Furthermore, the nontransmitted parental chromosomes of each trio were combined to form additional artificial haplotype pairs. For the HaploDB dataset, definite haplotypes were determined from complete hydatidiform moles (CHMs). The 74 available CHMs haplotypes were paired to form 37 diploid individuals.

For all datasets markers with minor allele frequency of less than 5% and genotypes with more than 15% missing values were removed. For the Yoruba population, information on 3.8 million SNPs spread over the whole genome is available. We sampled 100 sets of 100 markers each from distinct regions on chromosome 1. There are 60 individuals in these datasets after preprocessing as described above, with an average fraction of missing values of 3.6% and 32.2% heterozygous markers. For the Daly dataset, there is information on 103 markers and 174 individuals available after data preprocessing, the average fraction of missing values is 7.9% and the average fraction of heterozygous markers is 30.6%. In HaploDB, a genome-wide set of 281 439 SNP markers

Table 1: Switch (top-right triangle) and Hamming (bottom-left triangle) distances between the truth and the baseline methods for the Daly dataset.

	TRIO	fP	HIT	S	HR	G	P
TRIO	-	105	121	127	131	132	145
fP	480	-	82	82	104	85	118
HIT	514	414	-	88	116	103	146
S	510	434	508	-	118	117	122
HR	736	676	784	716	-	119	150
G	568	478	522	546	810	-	143
P	654	590	728	650	850	718	-

is available, from which we sampled 100 sets of 100 markers each from distinct regions on chromosome 1. The average fraction of missing values is 3.1% and the fraction of heterozygous markers is 39.9%. All datasets were phased with each of the following 6 publicly available haplotyping systems, yielding 6 different reconstructed haplotype pairs for every genotype: PHASE version 2.1.1. [SS05], fastPHASE version 1.1. [SS06], GERBIL as included in GEVALT version 1.0. [KS05a], HIT [RKMU05], HAPLOREC version 2.0. [EGT06] and SPAMM version 1.0. [LME⁺06]. All methods were run using their default parameters.

Let us first consider how the reconstructions produced by the baseline methods differ on the Daly dataset. Table 1 shows the switch and Hamming distances between the different haplotype reconstructions, including the reconstructions inferred from the family trios (TRIO) as the ground truth, and the methods fastPHASE (fP), HIT (HIT), SPAMM (S), HAPLOREC (HR), GERBIL (G), and PHASE (P). The fastPHASE system clearly has the smallest reconstruction error with respect to switch and Hamming distances on the Daly dataset. While the accuracy performance of the other methods is worse, the distances between all the methods are of the same order of magnitude. This indicates that it makes sense to try to combine the haplotypers.

We tested the haplotyper selection and voting techniques using the Daly, Yoruba and HaploDB datasets. As it is not clear which combination of selection/voting and internal distance measure (switch distance, Hamming distance, k -Hamming distance) yields best results, systematic experiments using all different combinations were performed. The quality of the resulting reconstructions is measured by switch distance only, as this is the standard way of measuring the quality of reconstructions in haplotyping experiments.

The main goals of the experimental study are as follows. First, the goal is to evaluate whether the simple combination approaches like selection and voting can be used to find a more robust solution when the best-performing method is not known. Second, the goal is to see whether the combination methods

Table 2: The total switch error between true haplotypes and the haplotype reconstructions over all individuals for the baseline methods. For Yoruba and HaploDB, the reported numbers are averages over the 100 datasets.

Method	Daly	Yoruba	HaploDB
PHASE	145	37.61	108.36
fastPHASE	105	45.87	110.45
SPAMM	127	54.69	120.29
HAPLOREC	131	56.62	130.28
HIT	121	73.23	123.95
GERBIL	132	75.05	134.22

improve over the baseline methods when using different subsets of the baseline methods. For this purpose, we consider leaving out one of the baseline methods PHASE (the most accurate on the Yoruba and HaploDB datasets on average and on the HaploDB dataset), fastPHASE (most accurate on the Daly dataset), and GERBIL (slow and least accurate on all datasets), and also leaving out all three of them simultaneously. The results using these subsets are representative of results for other subsets we experimented with but do not report on here. Third, the goal is to find out how the haplotyper selection results compare to haplotyper voting results and how the different distance functions affect the quality of the solutions.

The results for the baseline methods are summarized in Table 2 and results for the combination methods in Table 3 and Table 4. Let us first consider the Daly dataset. The best baseline method is fastPHASE, resulting in 105 switch errors. The selection and voting methods applied to the set of all baseline methods produce results comparable to fastPHASE, and are consistently better than the haplotype reconstructions produced by any other baseline method. Thus, by employing the haplotyper combination approach, we can achieve performance comparable to the best baseline method without having to know which of the baseline methods is best in advance. Leaving out one of the methods PHASE, fastPHASE, or GERBIL has no significant effects on the results of the combination methods. Thus, the combination methods seem to be quite robust against small perturbations of the set of baseline methods, even if they lead to the exclusion of the best performing method. If PHASE, fastPHASE, and GERBIL are left out simultaneously, the results degrade below the level of fastPHASE, but are still better than those of any other method.

The results on the Yoruba datasets follow a similar pattern, except that now PHASE—the baseline method with worst performance on the Daly dataset—is the best on average. The combination methods provide solutions comparable to those of the best method (PHASE) and better than those of any other baseline method. When only subsets of the baseline methods are used, the performance of the combination methods drops, but not significantly unless

Table 3: The total switch error between true haplotypes and the haplotype reconstructions over all individuals for the haplotyper selection methods for different combinations of baseline haplotypers. For Yoruba and HaploDB, the reported numbers are averages over the 100 datasets.

Haplotyper Selection				
Methods	Distance	Daly	Yoruba	HaploDB
all methods	d_s	103	37.67	103.43
	d_{3-H}	103	38.29	104.10
	d_{4-H}	103	38.41	104.52
	d_{5-H}	105	38.35	104.76
	d_H	107	40.14	110.84
w/o PHASE	d_s	106	43.58	107.16
	d_{3-H}	107	43.42	107.40
	d_{4-H}	109	43.99	108.55
	d_{5-H}	107	44.16	108.55
	d_H	102	48.36	117.00
w/o fP	d_s	108	40.00	105.06
	d_{3-H}	105	40.13	105.74
	d_{4-H}	110	40.82	106.91
	d_{5-H}	114	41.59	107.32
	d_H	115	45.30	116.27
w/o GERBIL	d_s	103	38.47	103.91
	d_{3-H}	105	38.53	104.69
	d_{4-H}	104	38.98	105.52
	d_{5-H}	118	39.07	106.05
	d_H	113	42.46	111.57
w/o P, fP, G	d_s	116	47.94	113.95
	d_{3-H}	108	47.87	114.39
	d_{4-H}	104	48.48	115.57
	d_{5-H}	116	48.66	116.57
	d_H	117	53.47	122.61

Table 4: The total switch error between true haplotypes and the haplotype reconstructions over all individuals for the haplotyper voting methods for different combinations of baseline haplotypes. For Yoruba and HaploDB, the reported numbers are averages over the 100 datasets.

Haplotyper Voting				
Methods	Distance	Daly	Yoruba	HaploDB
all methods	d_s	104	39.86	103.06
	d_{3-H}	107	39.15	102.24
	d_{4-H}	107	40.08	104.00
	d_{5-H}	107	39.56	104.29
	d_H	106	51.07	134.16
w/o PHASE	d_s	107	43.18	105.68
	d_{3-H}	107	43.15	106.41
	d_{4-H}	114	43.67	107.14
	d_{5-H}	107	44.14	107.67
	d_H	105	50.29	119.99
w/o fP	d_s	109	39.71	103.77
	d_{3-H}	107	39.92	104.26
	d_{4-H}	106	40.79	105.42
	d_{5-H}	112	41.34	105.78
	d_H	117	47.59	119.19
w/o GERBIL	d_s	105	38.27	102.76
	d_{3-H}	104	38.19	103.38
	d_{4-H}	104	38.70	104.39
	d_{5-H}	112	38.93	104.62
	d_H	110	43.91	114.93
w/o P, fP, G	d_s	112	46.28	110.58
	d_{3-H}	109	46.58	110.99
	d_{4-H}	107	48.09	113.25
	d_{5-H}	111	48.60	113.91
	d_H	114	53.92	122.99

PHASE, fastPHASE, and GERBIL are left out simultaneously.

On the HaploDB dataset the advantage of using combination methods is even more evident. The best baseline method PHASE is clearly outperformed by all combination methods except voting and selection with Hamming distance. The results are only slightly degraded when PHASE, fastPHASE, or GERBIL is left out of the ensemble, and when they are all left out simultaneously, the performance of the combination strategies is still significantly better than that of the best remaining baseline method SPAMM.

In summary, our results indicate that using the haplotyper combination approach sometimes significantly increases the haplotyping accuracy, and never significantly decreases the accuracy in comparison to the best baseline method. Of course, in practice the identity of the best baseline method is not known and changes from dataset to dataset. In a more realistic comparison to the baseline method that does best on average on all the datasets (fastPHASE), all the proposed haplotyper combination methods are clearly more accurate on average. Hence, our experiments suggest that it is indeed better to combine the predictions of all the baseline methods than to (blindly) choose and use any one of them.

In general, combination methods using switch distance as the distance function tend to produce most accurate results. This suggests that the errors of the baseline methods resemble random switches rather than random single nucleotide mutations. The performance of different distance functions also depends on the density of the used marker map. For dense marker maps, larger windows are beneficial, whereas in sparse maps considering dependencies between consecutive markers probably suffices. Furthermore, the selection methods seem to perform slightly better than voting methods. A potential explanation is that the median haplotype reconstruction is more tolerant to random errors in the baseline methods than the mean haplotype reconstruction. Further analysis is needed in order to fully understand the differences between the combination methods, but it seems safe to conclude that haplotyper selection with switch distance is the best choice (among combination methods and baseline methods) at least when no additional information about the problem at hand is available.

The computational price for the potential improvements in accuracy is the added effort of first running all the baseline methods and then solving the HVP or HSP problem. This may be a problem if some of the baseline methods are very slow. In such cases, we suggest the strategy of computing the predictions of as many baseline methods as time constraints permit, and combining the resulting reconstructions using one of the combination methods. The running times of the baseline methods vary greatly, so running, e.g., all but the slowest baseline method may well be much more efficient than running the slowest method alone.




Figure 1: The relative Hamming distance vs the number of heterozygous markers in the Daly dataset. The plot shows that the relative error being higher than $1/2$ is mainly the problem of individuals with small number of heterozygous markers. The fraction of individuals with the relative hamming error at least $1/2$ is $20/174 \approx 11.5\%$.

A major computational difficulty with Hamming voting is that the basic method for computing it scales exponentially in the number of haplotype reconstructions per individual. In Section 3.3 we showed that if the number of heterozygous markers is small or the relative Hamming score of the solution is at most $1/2$ then the ordering of the haplotypes in the pairs can be determined efficiently. Figure 1 illustrates that in the Daly dataset most of the individuals have either very small number of heterozygous markers or small relative Hamming error. This supports the hypothesis that the Hamming voting problem can be solved sufficiently efficiently in practice even with a larger number of baseline methods.

Depending on the combination method there can be multiple solutions that have the same score but different distance to the ground truth. In Table 3 and Table 4 ties are broken by selecting the solution with optimal score that was found first. The rule for breaking ties may have a significant effect on the final accuracy: For example, when applying switch voting to all the 6 methods on the Daly dataset, there are a total of 35 ties. Thus, depending on how




Figure 2: The switch error of PHASE vs the sum of the switch distances between the baseline methods for Yoruba datasets. Each point corresponds to one of the Yoruba datasets, x-coordinate being the sum of distances between the reconstructions obtained by the baseline methods, and y-coordinate corresponding to the switch errors of the reconstructions by PHASE.

ties are broken, the error may be anywhere between 89 and 124. Breaking the ties randomly results on average in 106.5 errors, which is quite close to the result 105 obtained by our arbitrary tie breaking. The situation with other combination methods and datasets is similar. Better results might be obtained by more sophisticated tie breaking rules, e.g., by following the overall leader fastPHASE in case of ties. We leave the exploration of such advanced tie breaking rules for future work.

When the experimental results are analyzed in detail at the level of individuals, it turns out that the combination methods tend to fail only when many of the baseline methods perform rather badly. Even though it is hard to recover the true haplotypes in such cases, the fact that the baseline methods are in wide disagreement can be used to identify such problematic individuals. We have observed that the sum of distances between the baseline haplotype reconstructions has very high correlation (between 0.95 and 0.99) with the error in the final reconstructions for the individual. Figure 2 illustrates this for PHASE using the Yoruba datasets. This indicates that combining haplotypes can

also be a strong method for outlier detection, which is helpful for removing probably incorrectly haplotyped individuals from further consideration.

5 Conclusions

Haplotype reconstruction is an important intermediate task in the study of genetic variations in human populations. Various techniques to reconstruct haplotypes from measurable genotype data have been proposed. Different methods typically return substantially different reconstructions, and there seems to be no method that is generally superior on all datasets.

To overcome these difficulties, we have studied the problem of combining haplotypers in order to improve the haplotype reconstructions. More specifically, we have considered two variants of the problem: haplotyper voting, where the goal is to find a consensus haplotype reconstruction given multiple haplotype reconstructions, and haplotyper selection, where the goal is to find the best haplotype reconstruction for each individual. We have developed algorithms for using various internal distance functions. The experiments show that combining haplotypers provides improvements over the average performance of the haplotype reconstruction methods, and the reconstruction quality is even comparable to or better than the best method for each dataset. Hence, using the combination methods virtually never degrades the performance, and sometimes gives clear advantages on accuracy in comparison to the best baseline method.

According to the experiments, haplotyper selection with switch distance is consistently close to the best combination method. Thus, the original problem of having to choose a baseline methods has not been lifted to an analogous problem of choosing a combination method as haplotyper selection with switch distance seems to be a good choice always.

Combining haplotypers opens many avenues to improved techniques for haplotype reconstruction. First, an obvious direction of refinements would be to use more complex combinator functions. For example, there could be a-priori knowledge about the performance of the methods on some part of the data or on other, similar datasets. This knowledge could be used to reweight methods in the combination algorithms, or pursue more complex prediction approaches such as decision trees or support vector machines. Such approaches would be especially useful for combining a large number of reconstructions of varying quality. Second, the different methods could be used to guide haplotype reconstruction techniques, e.g., to detect potentially problematic regions of the data where the reconstruction model should be refined. Third, assessing the quality of haplotype reconstructions by combining haplotypers is a promising direction. Often it will be acceptable to discard part of the reconstructed haplotypes or markers to avoid errors. Our preliminary results suggest that

multiple haplotype reconstructions could be used to detect individuals which are likely to be haplotyped erroneously and even problematic regions of the marker map. Fourth, more refined measures of the reconstruction quality are also of interest, for example modeling the dependence structure of the markers in more detail or taking genetic background information into account. Finally, we intend to evaluate the approach with further genotype datasets with known haplotypes as they become available.

References

- [ACN05] Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: ranking and clustering. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 684–693. ACM, 2005.
- [Bre96] Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [CSS02] Michael Collins, Robert E. Schapire, and Yoram Singer. Logistic regression, adaboost and bregman distances. *Machine Learning*, 48(1-3):253–285, 2002.
- [DKNS01] Cynthia Dwork, Ravi Kumar, Moni Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the Tenth International World Wide Web Conference, WWW 10, Hong Kong, China, May 1-5, 2001*, pages 613–622. ACM, 2001.
- [DRS⁺01] Mark J. Daly, John D. Rioux, Stephen F. Schaffner, Thomas J. Hudson, and Eric S. Lander. High-Resolution Haplotype Structure in the Human Genome. *Nature Genetics*, 29:229–232, 2001.
- [EGT04] Lauri Eronen, Floris Geerts, and Hannu Toivonen. A Markov Chain Approach to Reconstruction of Long Haplotypes. In Russ B. Altman, A. Keith Dunker, Lawrence Hunter, Tiffany A. Jung, and Teri E. Klein, editors, *Biocomputing 2004, Proceedings of the Pacific Symposium, Hawaii, USA, 6-10 January 2004*, pages 104–115. World Scientific, 2004.
- [EGT06] Lauri Eronen, Floris Geerts, and Hannu Toivonen. Haplorec: Efficient and Accurate Reconstruction of Long Haplotypes. *BMC Bioinformatics*, 7, 2006.
- [FG06] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. EATCS Texts in Theoretical Computer Science. Springer, 2006.

- [FISS03] Yav Freund, Raj Iyer, Rober E. Schapire, and Yoram Singer. An efficient boosting algorithm for combining preferences. *Journal of Machine Learning Research*, 4:933–969, 2003.
- [FKM⁺04] Ronald Fagin, Ravi Kumar, Mohammad Mahdian, D. Sivakumar, and Erik Vee. Comparing and aggregating rankings with ties. In Alin Deutsch, editor, *Proceedings of the Twenty-third ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 14-16, 2004, Paris, France*, pages 47–58. ACM, 2004.
- [GMT05] Aristides Gionis, Heikki Mannila, and Panayiotis Tsaparas. Clustering aggregation. In *Proceedings of the 21st International Conference on Research in Data Engineering*, pages 51–60, Tokyo, Japan, 2005.
- [Gus02] Dan Gusfield. An overview of combinatorial methods for haplotype inference. In Sorin Istrail, Michael S. Waterman, and Andrew G. Clark, editors, *Computational Methods for SNPs and Haplotype Inference*, volume 2983 of *Lecture Notes in Computer Science*, pages 9–25. Springer, 2002.
- [HE04] Eran Halperin and Eleazar Eskin. Haplotype reconstruction from genotype data using imperfect phylogeny. *Bioinformatics*, 20(12):1842–1849, 2004.
- [HMK⁺07] Koichiro Higasa, Katsuyuki Miyatake, Yoji Kukita, Tomoko Tahira, and Kenshi Hayashi. D-HaploDB: a database of definitive haplotypes determined by genotyping complete hydatidiform mole samples. *Nucleic Acids Research*, 35:D685–D689, 2007.
- [JXL04] Yishan Jiao, Jingyi Xu, and Ming Li. On the k-closest substring and k-consensus pattern problems. In Süleyman Cenk Sahinalp, S. Muthukrishnan, and Ugur Dogrusöz, editors, *Combinatorial Pattern Matching, 15th Annual Symposium, CPM 2004, Istanbul, Turkey, July 5-7, 2004, Proceedings*, volume 3109 of *Lecture Notes in Computer Science*, pages 130–144. Springer, 2004.
- [KS05a] Gad Kimmel and Ron Shamir. A Block-Free Hidden Markov Model for Genotypes and Its Applications to Disease Association. *Journal of Computational Biology*, 12(10):1243–1259, 2005.
- [KS05b] Gad Kimmel and Ron Shamir. GERBIL: Genotype Resolution and Block Identification Using Likelihood. *Proceedings of The National Academy of Sciences*, 102(1):158–162, 2005.

- [KS05c] Gad Kimmel and Ron Shamir. The incomplete perfect phylogeny haplotype problem. *Journal of Bioinformatics and Computational Biology*, 3(2):359–384, 2005.
- [LB05] Tilman Lange and Joachim M. Buhman. Combining partitions by probabilistic label aggregation. In Robert Grossman, Roberto Bayardo, and Kristin P. Bennett, editors, *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Chicago, Illinois, USA, August 21-24, 2005*, pages 147–156. ACM, 2005.
- [LME⁺06] Niels Landwehr, Taneli Mielikäinen, Lauri Eronen, Hannu Toivonen, and Heikki Mannila. Constrained hidden markov models for population-based haplotyping. In *Probabilistic Modeling and Machine Learning in Structural and Systems Biology, Workshop Proceedings, Tuusula, Finland, June 17-18, 2006*, 2006.
- [LME⁺07] Niels Landwehr, Taneli Mielikäinen, Lauri Eronen, Hannu Toivonen, and Heikki Mannila. Constrained hidden markov models for population-based haplotyping. *BMC Bioinformatics*, 2007. Accepted.
- [LP02] Rune B. Lyngsø and Christian N. S. Pedersen. The consensus string problem and the complexity of comparing hidden Markov models. *Journal of Computer and System Sciences*, 65(3):545–569, 2002.
- [MTT06] Taneli Mielikäinen, Evimaria Terzi, and Panayiotis Tsaparas. Aggregating time partitions. In Mark Craven and Dimitrios Gunopulos, editors, *Proceedings of The Twelfth Annual SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2006), Philadelphia, USA, August 20*. ACM, 2006.
- [RKMU05] Pasi Rastas, Mikko Koivisto, Heikki Mannila, and Esko Ukkonen. A Hidden Markov Technique for Haplotype Reconstruction. In Rita Casadio and Gene Myers, editors, *Algorithms in Bioinformatics, 5th International Workshop, WABI 2005, Mallorca, Spain, October 3-6, 2005, Proceedings*, volume 3692 of *Lecture Notes in Computer Science*, pages 140–151. Springer, 2005.
- [Sav97] Carla Savage. A survey of combinatorial gray codes. *SIAM Review*, 39(4):605–629, 1997.
- [SD03] Matthew Stephens and Peter Donnelly. A comparison of bayesian methods for haplotype reconstruction from population genotype data. *The American Journal of Human Genetics*, 73:1162–1169, 2003.

- [SG02] Alexander Strehl and Joydeep Ghosh. Cluster ensembles – A knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3, 2002.
- [Sjö03] Jonas Sjöbergh. Combining POS-taggers for improved accuracy on swedish text. In *Proceedings of the 14th Nordic Conference on Computational Linguistics (NoDaLiDa 2003), Reykjavik, Iceland, May 30-31, 2003*, 2003.
- [SP03] Jeong Seop Sim and Kunsoo Park. The consensus string problem for a metric is NP-complete. *Journal of Discrete Algorithms*, 1(1):111–117, 2003.
- [SS05] Matthew Stephens and Paul Scheet. Accounting for Decay of Linkage Disequilibrium in Haplotype Inference and Missing-Data Imputation. *The American Journal of Human Genetics*, 76:449–462, 2005.
- [SS06] Paul Scheet and Matthew Stephens. A Fast and Flexible Statistical Model for Large-Scale Population Genotype Data: Applications to Inferring Missing Genotypes and Haplotypic Phase. *The American Journal of Human Genetics*, 78:629–644, 2006.
- [SSD01] Matthew Stephens, Nicholas J. Smith, and Peter Donnelly. A New Statistical Method for Haplotype Reconstruction from Population Data. *The American Journal of Human Genetics*, 68:978–989, 2001.
- [The05] The International HapMap Consortium. A Haplotype Map of the Human Genome. *Nature*, 437:1299–1320, 2005.