



(86) **Date de dépôt PCT/PCT Filing Date:** 2015/10/14
 (87) **Date publication PCT/PCT Publication Date:** 2016/04/21
 (85) **Entrée phase nationale/National Entry:** 2017/04/13
 (86) **N° demande PCT/PCT Application No.:** US 2015/055579
 (87) **N° publication PCT/PCT Publication No.:** 2016/061260
 (30) **Priorité/Priority:** 2014/10/14 (US62/063,849)

(51) **Cl.Int./Int.Cl. G06F 19/24** (2011.01),
G06F 19/10 (2011.01), **G06F 19/18** (2011.01),
C12Q 1/68 (2006.01)
 (71) **Demandeur/Applicant:**
 ANCESTRY.COM DNA, LLC, US
 (72) **Inventeurs/Inventors:**
 BARBER, MATHEW J., US;
 WANG, YONG, US;
 NOTO, KEITH D., US;
 CHAHINE, KENNETH G., US;
 BALL, CATHERINE ANN, US
 (74) **Agent:** SMART & BIGGAR

(54) **Titre : REDUCTION D'ERREUR DANS DES RELATIONS GENETIQUES PREDITES**
 (54) **Title: REDUCING ERROR IN PREDICTED GENETIC RELATIONSHIPS**

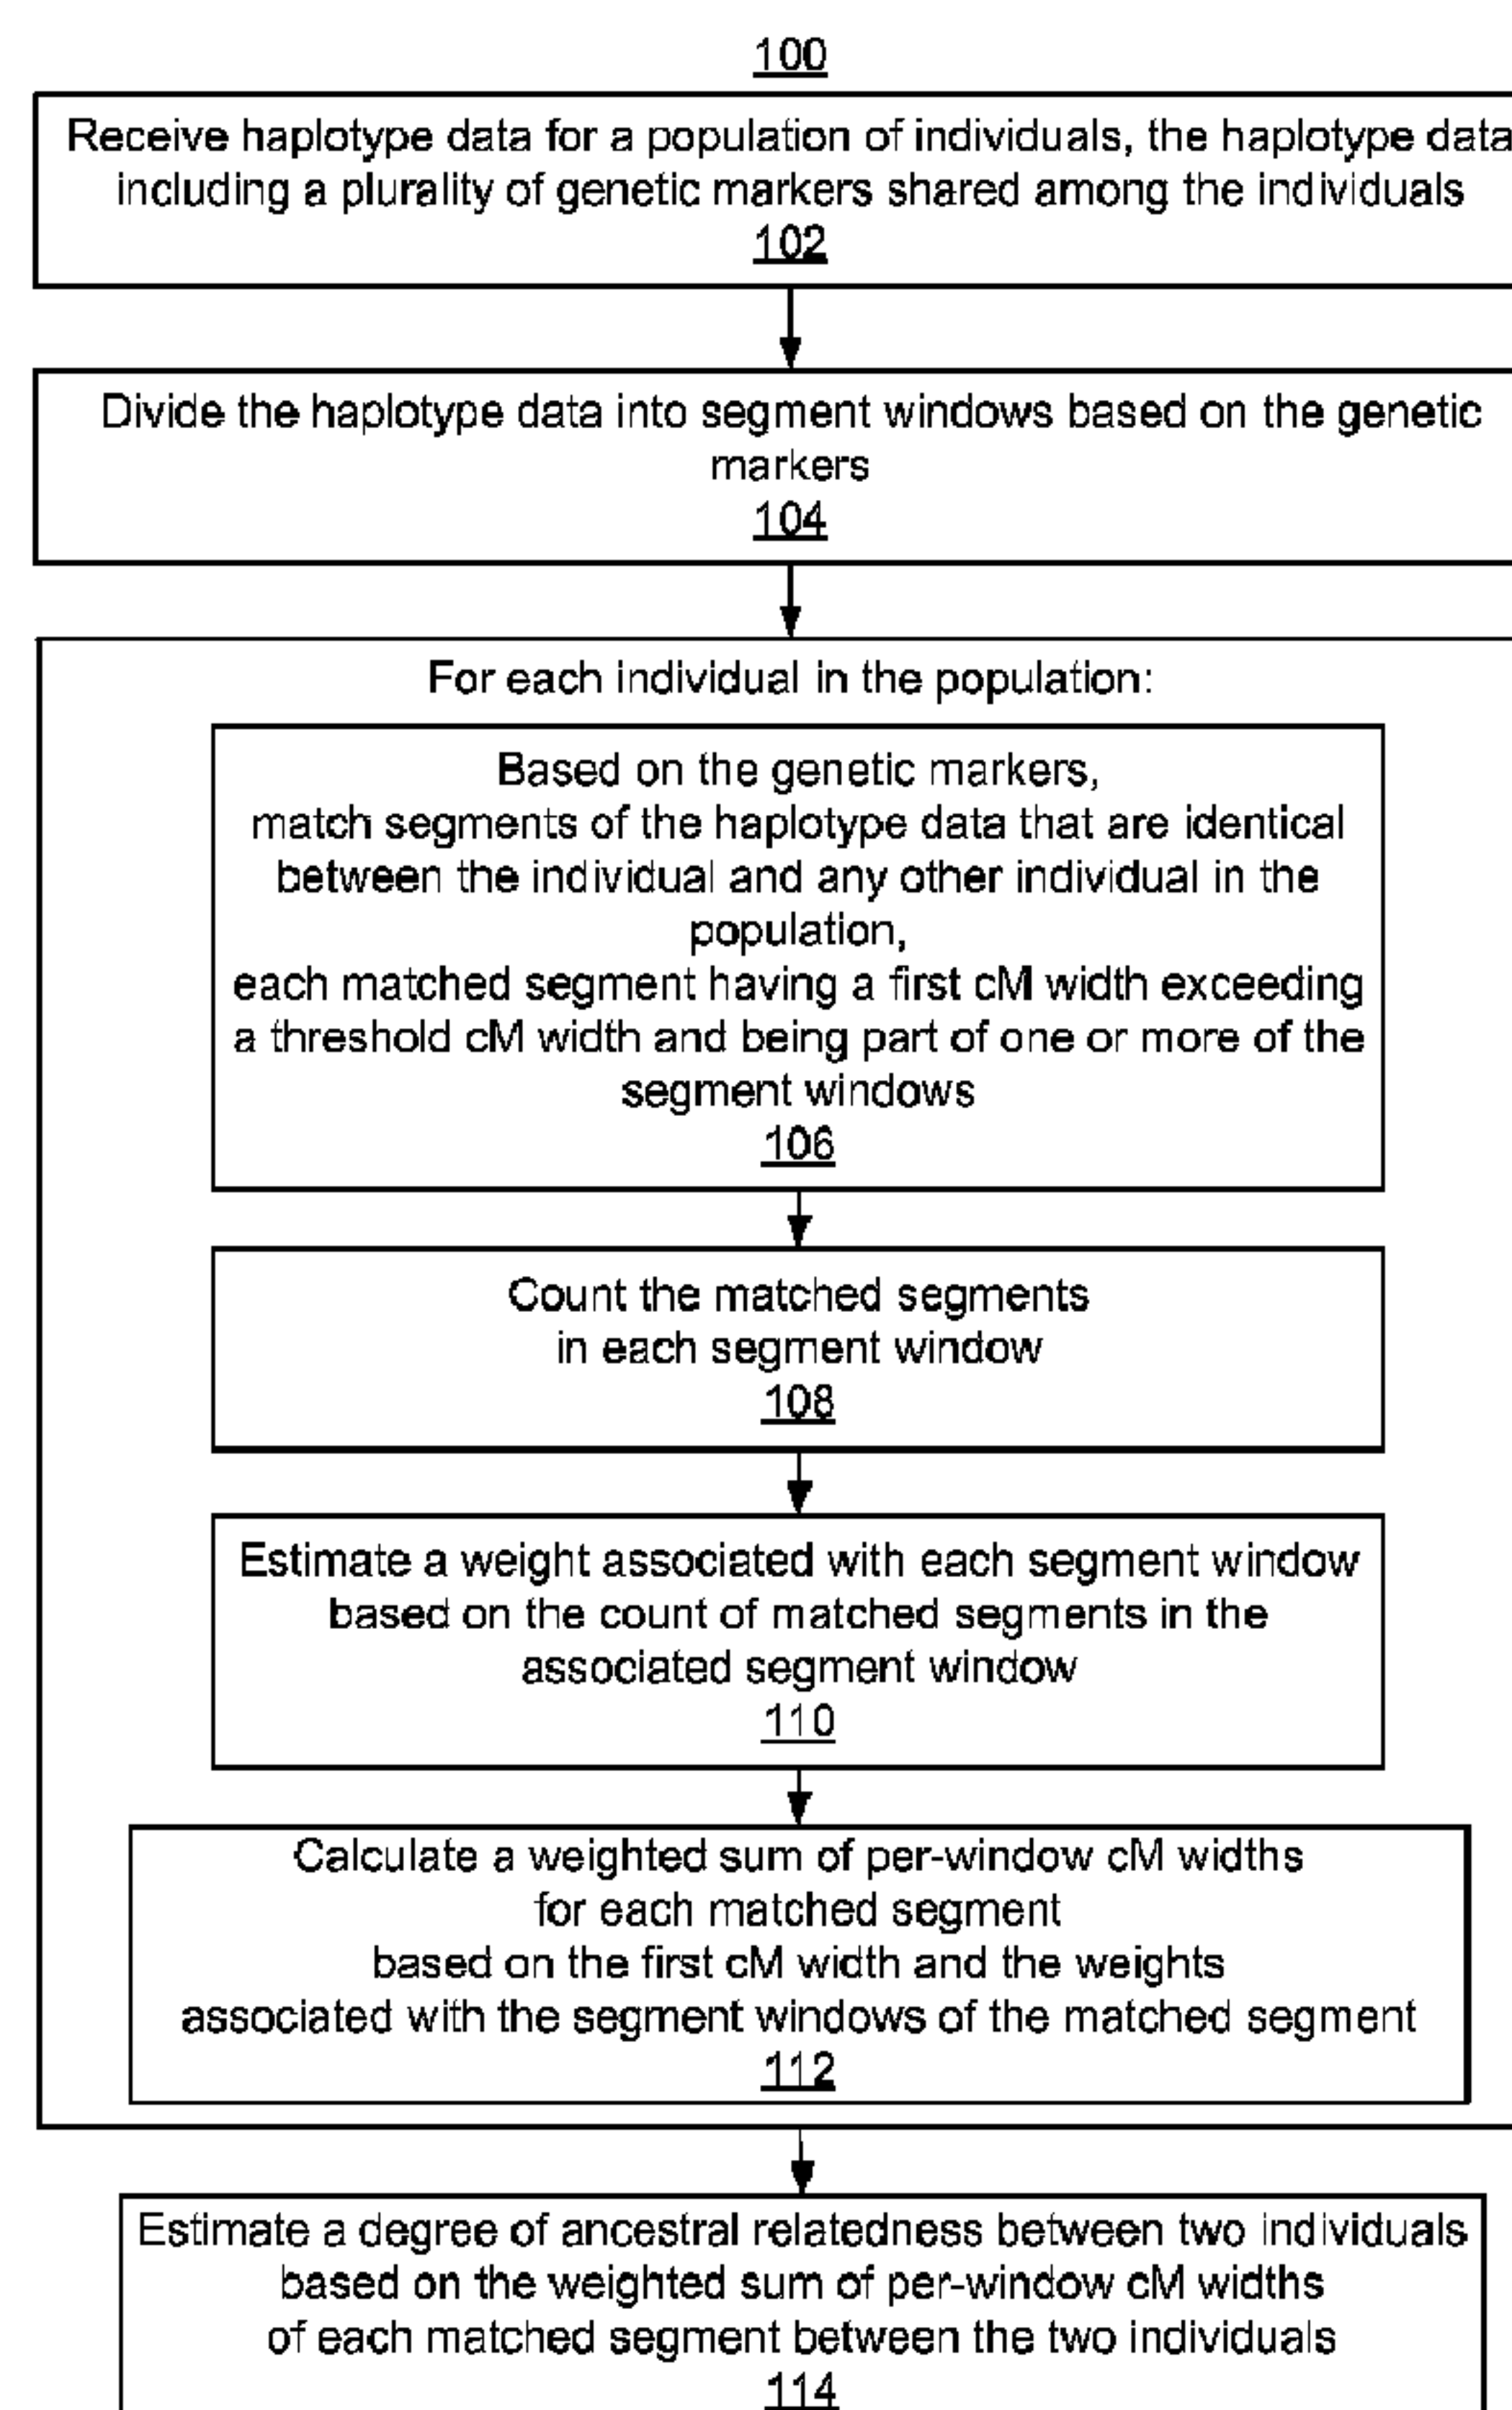


FIG. 1A

(57) **Abrégé/Abstract:**

System, computer program products, and methods are disclosed for estimating a degree of ancestral relatedness between two individuals. The haplotype data for a population of individuals is divided into segment windows based on genetic markers, and

(57) Abrégé(suite)/Abstract(continued):

matched segments for the haplotype data are generated. Each matched segment having a first cM width that exceeds a threshold cM width is included in counting the matched segments in each segment window. A weight associated with each segment window is estimated based on the count of matched segments in the associated segment window. A weighted sum of per-window cM widths for each matched segment is calculated based on the first cM width and the weights associated with the segment windows of the matched segment. The weighted sum of per-window cM widths are used to estimate a degree of ancestral relatedness between two individuals.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau(10) International Publication Number
WO 2016/061260 A1(43) International Publication Date
21 April 2016 (21.04.2016)

(51) International Patent Classification:

G06F 19/12 (2011.01) *C12Q 1/68* (2006.01)
G06F 19/28 (2011.01)

(21) International Application Number:

PCT/US2015/055579

(22) International Filing Date:

14 October 2015 (14.10.2015)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/063,849 14 October 2014 (14.10.2014) US

(71) Applicant: ANCESTRY.COM DNA, LLC [US/US]; 360 W 4800 N, Provo, UT 84604 (US).

(72) Inventors: BARBER, Mathew, J.; c/o Ancestry.Com DNA, LLC, 360 W 4800 N, Provo, UT 84604 (US). WANG, Yong; c/o Ancestry.Com DNA, LLC, 360 W 4800 N, Provo, UT 84604 (US). NOTO, Keith, D.; c/o Ancestry.Com DNA, LLC, 360 W 4800 N, Provo, UT

84604 (US). CHAHINE, Kenneth, G.; c/o Ancestry.Com DNA, LLC, 360 W 4800 N, Provo, UT 84604 (US). BALL, Catherine, Ann; c/o Ancestry.Com DNA, LLC, 360 W 4800 N, Provo, UT 84604 (US).

(74) Agents: BROWNSTONE, Daniel, R. et al.; Fenwick & West LLP, Silicon Valley Center, 801 California Street, Mountain View, CA 94041 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

[Continued on next page]

(54) Title: REDUCING ERROR IN PREDICTED GENETIC RELATIONSHIPS

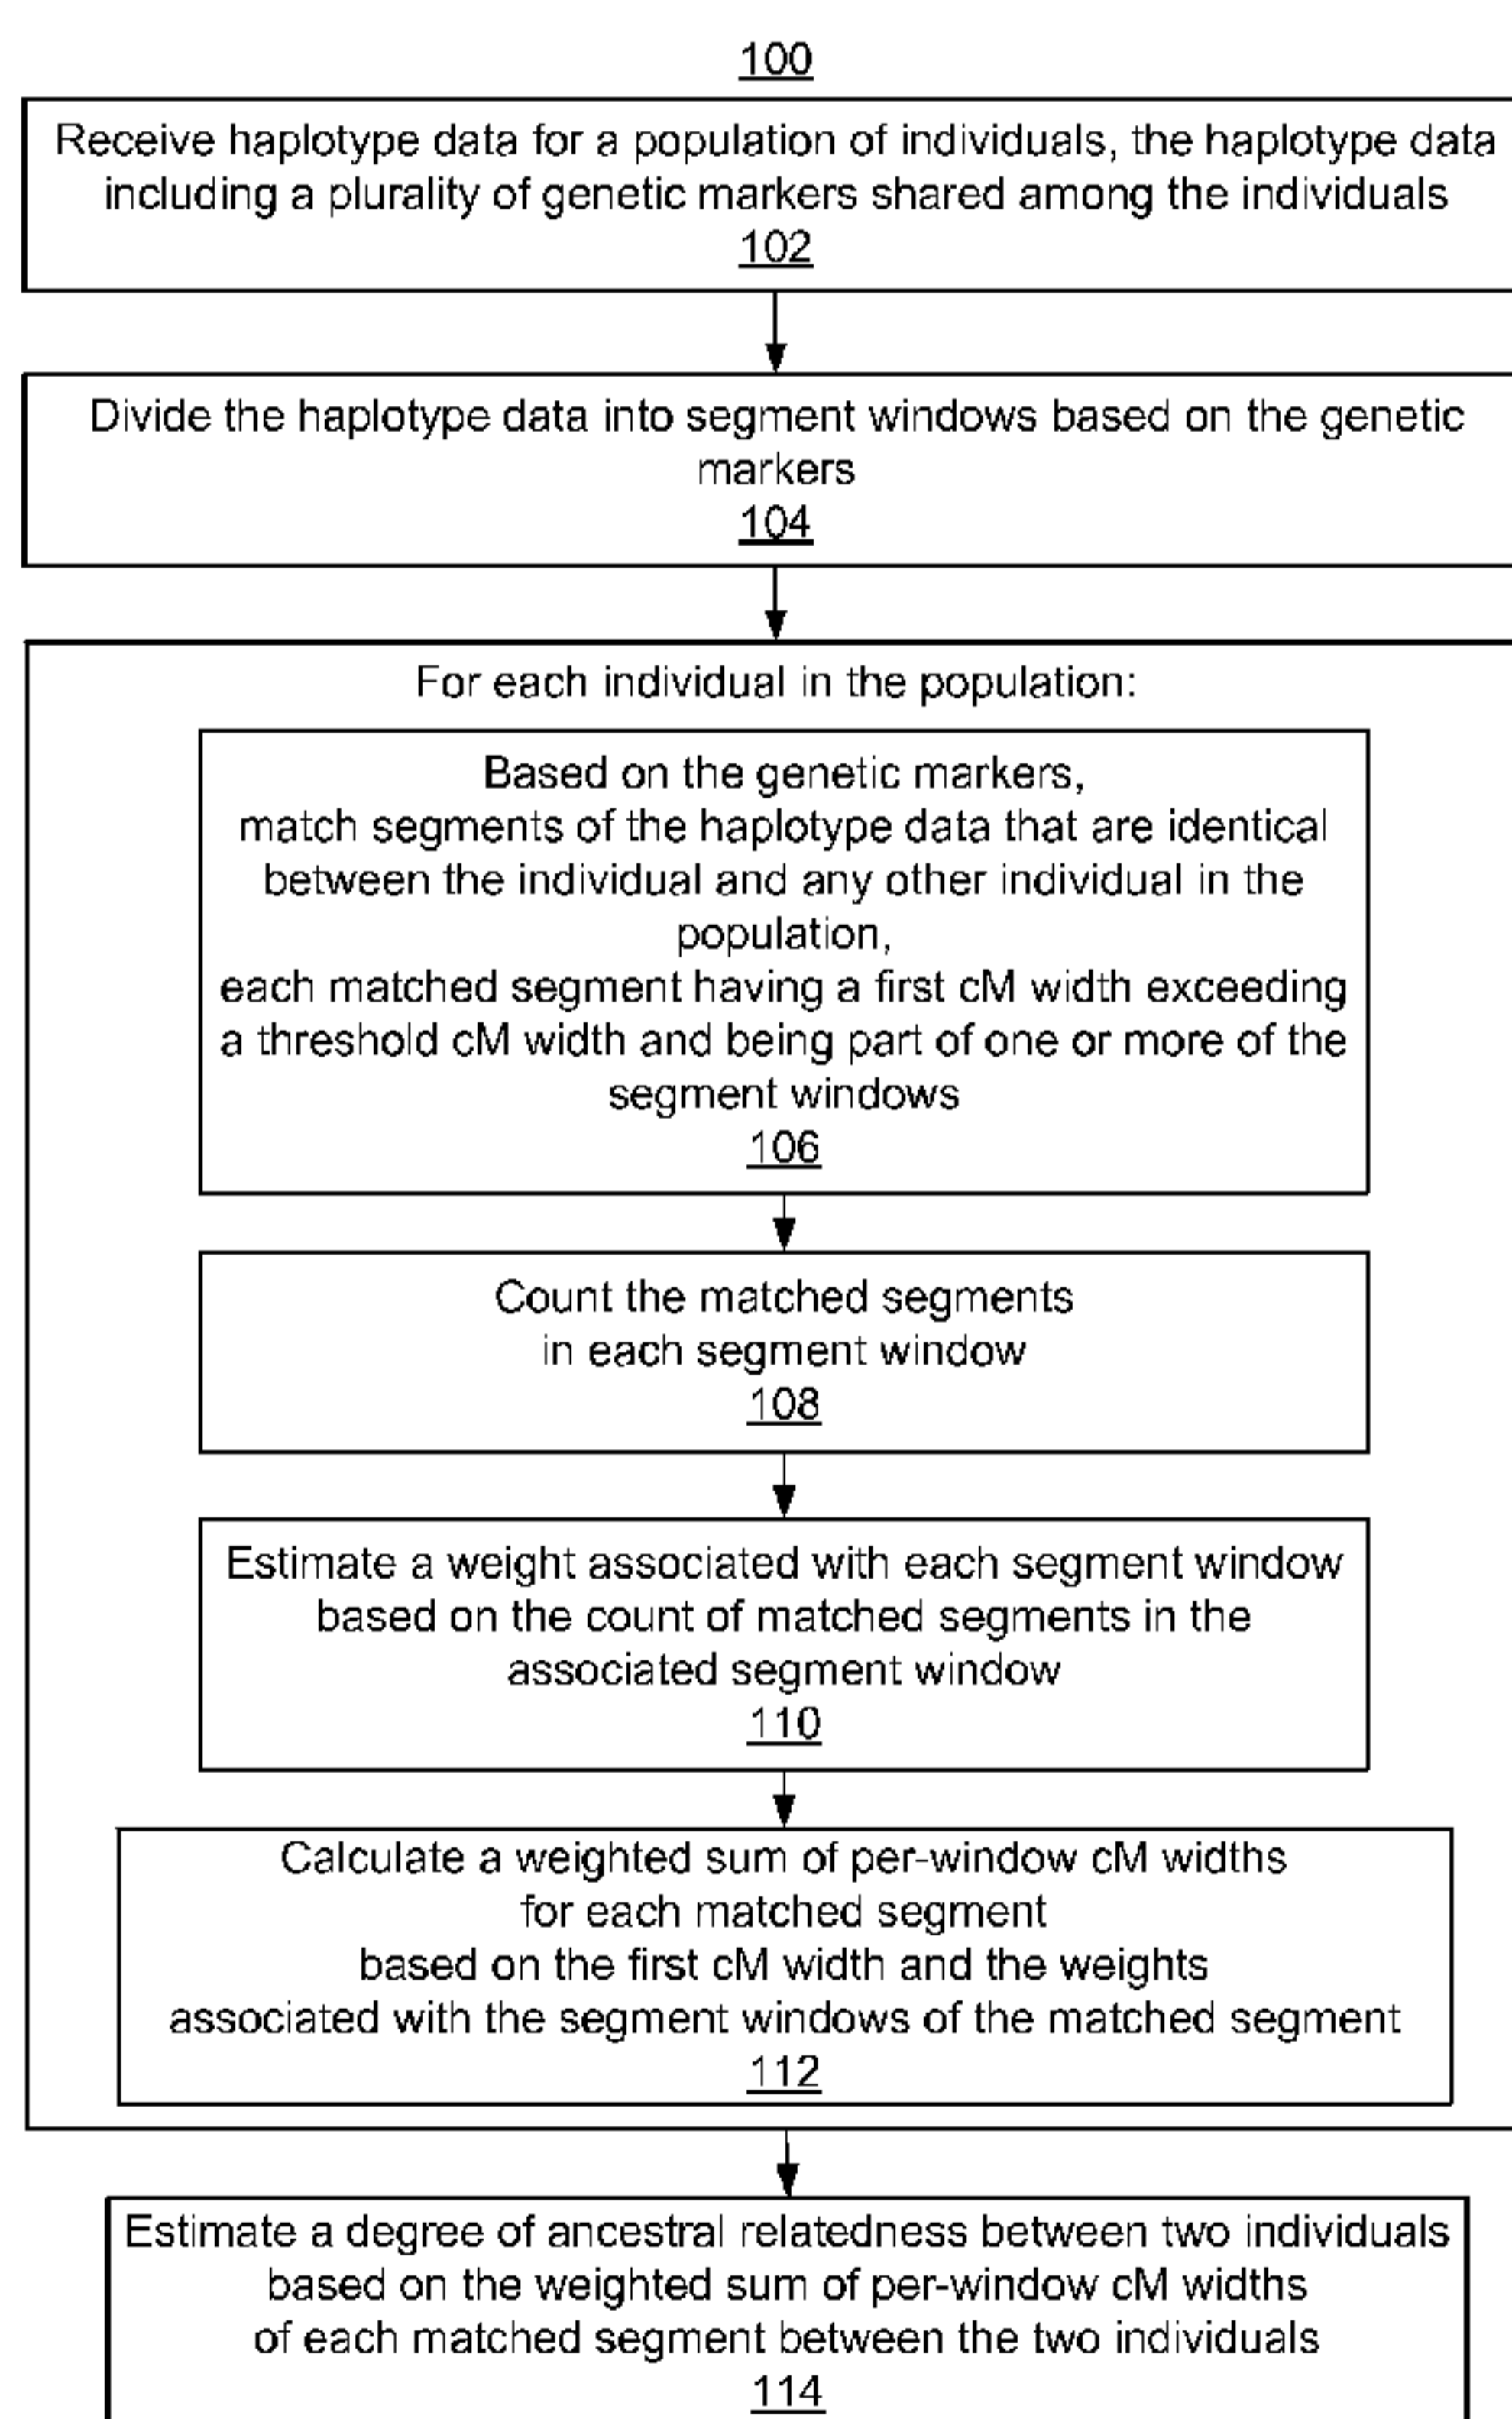


FIG. 1A

(57) Abstract: System, computer program products, and methods are disclosed for estimating a degree of ancestral relatedness between two individuals. The haplotype data for a population of individuals is divided into segment windows based on genetic markers, and matched segments for the haplotype data are generated. Each matched segment having a first cM width that exceeds a threshold cM width is included in counting the matched segments in each segment window. A weight associated with each segment window is estimated based on the count of matched segments in the associated segment window. A weighted sum of per-window cM widths for each matched segment is calculated based on the first cM width and the weights associated with the segment windows of the matched segment. The weighted sum of per-window cM widths are used to estimate a degree of ancestral relatedness between two individuals.

WO 2016/061260 A1



GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,
TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,
TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE,
DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT,
LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE,

SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

REDUCING ERROR IN PREDICTED GENETIC RELATIONSHIPS

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] The application claims the benefit of U.S. Provisional Application No. 62/063,849, filed on October 14, 2014, the contents of which are incorporated herein by reference.

BACKGROUND

1. FIELD

[0002] The disclosed embodiments relate to computer program products, systems, and methods used to identify individuals in a population who are ancestrally related based on the individuals' genetic data.

2. DESCRIPTION OF THE RELATED ART

[0003] Although humans are, genetically speaking, almost entirely identical, small differences in our DNA are responsible for much of the variation between individuals. Stretches of DNA that are determined to be relevant for some purpose are referred to as haplotypes. Haplotypes are identified based on consecutive single nucleotide polymorphisms (SNPs) of varying length. Certain haplotypes shared by individuals suggests a familial relationship between those individuals based on a principal known as identity-by-descent (IBD).

[0004] Because identifying segments of IBD DNA between pairs of genotyped individuals is useful in many applications, numerous methods have been developed to perform IBD analysis (Purcell et al. 2007, Gusev A. et al., The Architecture of Long-Range Haplotypes Shared within and across Populations, *Mol. Biol. Evol.*, 29(2):473–86, 2012; Browning S.R. and 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*, 91:1084-96, 2007; Browning S.R. and Browning B.L., Identity by descent between distant relatives: detection and applications, *Annu. Rev. Genet.*, 46:617-33, 2012). However, these approaches do not scale for continuously growing very large datasets. For example, the existing GERMLINE implementation is designed to take a single input file containing all individuals to be compared against one another. While appropriate for the case in which all samples are genotyped and analyzed simultaneously, this approach is not practical when samples are collected incrementally. The program suite

GERMLINE (Gusev A. et al., Whole population, genome-wide mapping of hidden relatedness, *Genome Res.*, 19:318–26, 2009) offers an “ibs filter”, which removes highly frequent matches (defined by chromosome, as well as the start and end position on the chromosome). Like GERMLINE’s matched segment discovery approach, the “ibs filter” approach is designed to be fast, and is relatively simplistic as a consequence. The more accurate of these methods, such as Refined IBD, are much more accurate than the GERMLINE “ibs filter”, but they do not scale computationally and would be difficult to integrate into an analytical pipeline even if they did. There are many existing methods that assess evidence for a matched segment not just by centimorgan width, as is done within GERMLINE. Examples include Refined IBD, fast IBD, SLRP, and PARENTE. They emphasize that differences between these approaches are a tradeoff between model complexity and computational speed (and feasibility).

SUMMARY

[0005] Methods, systems, and computer program products are disclosed for estimating a degree of ancestral relatedness between two individuals. The computer program products include TIMBER. To estimate the ancestral relatedness of two individuals methods include receiving haplotype data from a population of individuals. The haplotype data include a plurality of genetic markers that are shared among the individuals in the population. The haplotype data is then divided into segment windows based on the genetic markers. In some embodiments, the genetic markers include single nucleotide polymorphisms (SNPs) and the haplotype data is divided into K segment windows including an equal number d of SNPs. In some embodiments, the haplotype data is divided into 4105 segment windows of 96 SNPs.

[0006] For each individual, the method includes matching segments of the haplotype data that are identical between the individual and any other individual in the population, wherein the matching is based on the genetic markers. Each matched segment has a first centimorgan (cM) width that exceeds a threshold cM width. In some embodiments, the threshold cM width is 5 cM. Each matched segment is part of one or more of the segment windows. The matched segments in each segment window are then counted. The count of matched segments in a segment window is also referred to as a per-window match count.

[0007] For each individual, the method includes estimating a weight associated with each segment window based on the count of matched segments in the associated segment window. In some embodiments, the weight associated with a segment window is decreased as the

count of matched segments increases. The benefit of decreasing weights for increasing per-window match counts includes reducing the effect of matched segments that are likely not from the recent genealogical history (RGH) of the individuals, but rather from a more distant common ancestry at the human, ethnicity, or sub-ethnicity level.

[0008] For each individual, the method includes calculating a weighted sum of per-window cM widths for each matched segment based on the first cM width and the weights associated with the segment windows of the matched segment. A degree of ancestral relatedness between two individuals is estimated based on the weighted sum of per-window cM widths of each matched segment between the two individuals. In some embodiments, the degree equals the weighted sum of of per-window cM widths. In some embodiments, the weighted sum of per-window cM widths is the sum of the first cM widths for each segment window of a matched segment between the two individuals multiplied by the two weights for each individual associated with these segment windows.

[0009] In some embodiments, TIMBER, an ancestry prediction machine matching genetic markers, is a procedure operating on a computer for refining each individual's list of matched segments and prioritizing the matched segments that are most likely to be from the individuals' recent genealogical history. TIMBER uses the matched segments to remove the effect of "noisy" segment windows within the haplotype data that display an "excessive" count of matched segments between numerous individuals. In some embodiments, a count is excessive if the count is larger than 10 or 20. It is less likely that a matched segment is from recent genealogical history if a matched segment is mainly part of "noisy" segment windows. TIMBER estimates weights from the matched segment data and estimates a weighted sum of per-window cM widths of a matched segment based on discounting "noisy" segment windows. TIMBER is computationally efficient and scalable, allowing reevaluating an entire population of individuals each time new individuals are added to the population.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] Figure (Fig.) 1A illustrates a flowchart of a method for estimating a degree of ancestral relatedness between two individuals, according to some embodiments.

[0011] Fig. 1B is a block diagram of a computing environment for estimating a degree of ancestral relatedness between two individuals, according to one embodiment.

[0012] Fig. 2 illustrates an example of per-window match counts on a chunk of the genome for one individual, according to some embodiments.

[0013] Fig. 3 illustrates an example of the histogram of per-window match counts for all of the windows in the genome for one individual, according to some embodiments.

[0014] Fig. 4 is an example of the histogram of per-window match count for all the non-zero count windows where the maximum viewable per-window count is 40, according to some embodiments.

[0015] Fig. 5 is an example of the histogram of per-window match counts for all the non-zero and also low match count windows, according to some embodiments.

[0016] Fig. 6 is an example of the estimated per-window weight (in the y-axis) as a function of the possible per-window match count (in the x-axis), according to some embodiments.

[0017] Fig. 7 is an example of the estimated per-window weight (in the y-axis) as a function of the possible per-window match count (in the x-axis), according to some embodiments.

[0018] Fig. 8 illustrates an example for the weights (solid line) given the per-window counts (dashed line) from the original example in FIG. 1, according to some embodiments.

[0019] Fig. 9 is an example of per-window match counts on a chunk of the genome for one individual both pre-TIMBER (dashed line) and post-TIMBER (solid line), according to some embodiments.

[0020] Fig. 10 illustrates results from TIMBER using different unweighted cM width scores, i.e., first cM width filters, and weighted cM sum scores, including the matched percentages of segments kept for the known and unknown meiosis, according to some embodiments.

[0021] The figures depict an embodiment for purposes of illustration only. One skilled in the art will readily recognize from the following description that alternative embodiments of the structures and methods illustrated herein may be employed without departing from the principles described herein.

DETAILED DESCRIPTION

I. OVERVIEW

[0022] Methods, systems, and computer program products are disclosed for estimating a degree of ancestral relatedness between two individuals. Estimating the ancestral relatedness of individuals includes identifying and scoring identical-by-descent (IBD) matched segments among the haplotype data of these individuals. To identify IBD segments, the method compares genetic markers among the individuals' haplotypes. In some embodiments, genetic

markers include single-nucleotide polymorphisms (SNPs). Segments from two individuals are considered identical by state (IBS) if the genetic markers along the individuals' haplotype sequences in these segments are identical at the same loci along the haplotypes. Throughout the disclosure unless otherwise stated, "matched haplotype segments" or "matched segments" refer to identical haplotype segments shared between two or more individuals. Generally, an IBS segment shared between two individuals is identical by descent (IBD) if the individuals inherited the IBS segment from a common ancestor, sharing the same ancestral origin. Thus, any IBD segment by definition also represents an IBS segment, while the reverse is typically not true, i.e., an IBS segment might not represent an IBD segment. Moreover, many IBD segments are not from the recent genealogical history (RGH) of individuals, but rather from a more distant common ancestry at the human, ethnicity, or sub-ethnicity level. The disclosed method allows for prioritizing matched segments that are more likely to be from the individuals' RGH over those segments that are from a more distant common ancestry, thus belonging to their distant ancestral past, i.e., non-recent genealogical history (non-RGH).

[0023] Fig. 1A is a flowchart illustrating a method 100 for estimating a degree of ancestral relatedness between two individuals, according to some embodiments. The method allows the user to input segments that are categorized as matched or discovered, i.e., having a first centimorgan (cM) width of over 5 cM. The method in form of the TIMBER program then uses those matched segments to calculate a weighted sum of per-window cM widths. The weighted sum of per-window cM widths takes into account the count of matched segments to other individuals of the population in the segment windows, down weighting segment windows that display a high degree of matched segments to many individuals.

[0024] In some embodiments, the inputs to TIMBER program are the pairwise matched segments between all individuals in a population stored in a database. The pairwise matched segments are translated into weights for each individual's matchable segment windows of the genome by the TIMBER program. The TIMBER program then uses the weights to recalibrate or re-score the original pairwise matched segments.

[0025] In some embodiments, TIMBER programs, wherein the TIMBER programs are stored in memory and configured to be executed by one or more processors of a computing device, the TIMBER programs including instructions when executed by the computing device cause the device to:

1. calculate the counts of matched segments in each segment window of the haplotype data, for each individual of the population, where the matched segments are between the individual and every other person in the database,

2. calculate weights for each individual and segment window, and
3. calculate a weighted sum of per-window cM widths for each matched segment between two individuals based on the weights.

[0026] The method 100 is performed at a computing device, such as the computing device, as may be controlled by specially programmed code (computer programming instructions) contained, for example, in the TIMBER program, wherein such specially programmed code is or is not natively present in the computing device. Embodiments of the computing device include, but are not limited, general-purpose computers, e.g., a desktop computer, a laptop computer, computing servers, tablets, mobile devices, or any similar computing devices. Once programmed to execute the methods described here, such a computing device becomes a special-purpose computer. Some embodiments of the method 100 may include fewer, additional, or different steps than those shown in Fig. 1A, and the steps may be performed in different orders. The steps of the method 100 are described with respect to example haplotype data illustrated in Figures (Figs.) 2 through 9.

[0027] Fig. 1B is a block diagram of an environment for using a computer system 120 to estimate a degree of ancestral relatedness between two individuals, according to some embodiments. Depicted in Fig. 1B are individuals 122 (i.e. a human or other organism), a DNA extraction service 124, and a DNA quality control (QC) and matching preparation service 126.

[0028] Individuals 122 provide DNA samples for analysis of their genetic data. In some embodiment, an individual uses a sample collection kit to provide a sample, e.g., saliva, from which genetic data can be reliably extracted according to conventional methods. DNA extraction service 124 receives the sample and genotypes the genetic data, for example by extracting the DNA from the sample and identifying values of SNPs present within the DNA. The result is a diploid genotype. DNA QC and matching preparation service 126 assesses data quality of the diploid genotype by checking various attributes such as genotyping call rate, genotyping heterozygosity rate, and agreement between genetic and self-reported gender. System 120 receives 102 the haplotype data from DNA extraction service 124 and optionally stores the haplotype data in a database 128 containing unphased DNA diploid genotypes, phased haplotypes, and other genomic data. Unless otherwise stated, haplotype data refers to any genetic or genome data obtained from the individuals 122, which is optionally stored in database 128.

[0029] In some embodiments, the partitioning module 130 divides 104 the haplotype data into segment windows based on the genetic markers. In some embodiments, the matching

module 132 matches 106 segments of the haplotype data that are identical between the individual and any other individual in the population, where each matched segment has a first cM width that exceeds a threshold cM width and is part of one or more of the segment window.

[0030] In some embodiments, the count/weight estimation module 134 counts 108 the matched segments in each segment window and estimates 110 a weight associated with each segment window based on the count of matched segments in the associated segment window.

[0031] The scoring module 136 then calculates 112 a weighted sum of per-window cM widths for each matched segment based on the first cM width and the weights associated with the segment windows of the matched segment. In some embodiments, the scoring module 136 estimates 114 a degree of ancestral relatedness between two individuals based on the weighted sum of per-window cM widths of each matched segment between the two individuals.

II. EMBODIMENTS OF THE ANCESTRAL RELATIONSHIP PREDICTION PROGRAMS

[0032] In some embodiments, matched segments among a population of individuals are generated based on the individuals' haplotype data. In some embodiments, the matched segments are stored in a database for later retrieval by the TIMBER program. The TIMBER program is configured to receive all the matched segments among a population of individuals and prioritize those matched segments that are from the individuals' recent genealogical history (RGH).

II.A. Match Haplotype Segments

[0033] The method 100 includes receiving 102 haplotype data for a population of individuals, the haplotype data including a plurality of genetic markers shared among the individuals, according to some embodiments. In some embodiments, to identify (match) and score IBD segments among the haplotype data, the method 100 uses a HADOOP® reimplementation of a matching algorithm. The method benefits from being computationally fast and scalable for larger-sized populations. In some embodiments, the matched IBS segments based on the haplotype data of individuals include the individuals' RGH and non-RGH segments. The matched IBS segments are generally referred to as matched segments. In some embodiments, matched segments are generated using methods that are well known in the art. Using, for example, the TIMBER program, the method 100 then prioritizes the matched segments into segments that are more likely from the RGH of two individuals compared to from their non-RGH by calculating a TIMBER score that quantifies the likelihood of two individuals sharing a common recent ancestral relationship.

[0034] The method for determining the TIMBER score is based on the assumption that the locations (loci) of matched segments from an individual's RGH are evenly distributed across an individual's genome. For example, by dividing the SNPs of the chromosomes across the individuals' genome into discrete windows, the counts of the segments in a specific window on chromosome 1 are independent of the counts of the segments in a window on chromosome 14. Consequently, matches between segments of individuals in a window on chromosome 1 are therefore independent from matches to segments in the window on chromosome 14, resulting in an even distribution across all windows.

[0035] Furthermore, matched segments that do not originate from an individual's RGH exhibit spikes at certain windows of the genome while being evenly distributed across the remaining windows. In some instances, these spikes can be attributed to particular reasons for the genetic variation in these windows, e.g., at a person level or at the database level. At the person level, the spikes may result from the segments of a window displaying a high level of sequence similarity across a particular ethnicity group, while at a database level, the individuals of population in the database may not possess any SNP variations across a particular window based on how the population was selected. In particular, it is unclear whether the distribution of the counts of matched segments in a window originates from an individual's RGH, and whether factors that confound local spikes of matching are due to unknown non-RGH reasons, which are difficult to model. To overcome the problems associated with these spikes, the method introduces weights for each window for rescoring all matched segments, in particular, the ones that contribute to the spikes in windows. In some instances, matched segments commonly occur at short and specific location of an individual's genome while matching a very large number of other individuals, e.g., larger than 1,000 individuals.

[0036] The method 100 includes dividing 104 the haplotype data into segment windows based on the genetic markers, according to some embodiments. The haplotype data includes, but is not limited to, SNPs observed across the individual's genome. In some embodiments, the haplotype data includes the haplotype data of an individual's entire or partial genome. In some embodiments, the method 100 divides the observed SNPs into K windows of equal size d , with each window, for example, including 96 SNPs. Other examples for window sizes include 50, 60, 70, 80, 90, 100, 110, 120, 130, 140, 150 or any number that falls within the range of 50 to 500. In some embodiments, the size of each window, i.e., the number of SNPs per window varies for across the windows.

[0037] For example, some windows may include only 50 SNPs depending on the

sequence length of these SNPs, while other windows include 96 SNPs. In some embodiments, the windows include other genetic markers besides SNPs that are used to identify matched segments. These markers include, but are not limited to, restriction fragment length polymorphisms, simple sequence length polymorphisms, amplified fragment length polymorphism, random amplification of polymorphic DNA, variable number tandem repeat, simple sequence repeat microsatellite polymorphism, short tandem repeats, single feature polymorphisms, restriction site associated DNA markers, and the like.

[0038] In some embodiments, the method 100 includes using phased haplotype data, i.e. data for which the phase has been estimated, as input to identify matched segments. For this, the method uses the haplotype data for a population of n individuals. In some embodiments, the input of the haplotype data is represented as a $2n \times s$ matrix \mathbf{H} with rows corresponding to $2n$ haplotypes and columns to s SNPs. By vertically slicing \mathbf{H} into non-overlapping, equal width submatrices \mathbf{H}_i of d columns, each submatrix \mathbf{H}_i then represents a different segment window i , where $i = 0 \dots K$ and $s = d \cdot K$. In some embodiments, the haplotype data includes n implicit non-phased haplotypes of the population using the population's genotypes to determine possible haplotype matches without explicit haplotype matching that would require the phase of the haplotypes to be known. Haplotype matching therefore refers to implied as well as explicit haplotype matching, where the former is based on non-phased genomic data of a population.

[0039] The method 100 includes for each individual in the population, based on the genetic markers, matching 106 segments of the haplotype data that are identical between the individual and any other individual in the population, according to some embodiments. In some embodiments, the Each matched segment has a first cM width that exceeds a threshold cM width and is part of one or more of the segment windows. The matching 104 includes identifying segment windows of exact haplotype matching between two individuals in the population. Windows of exact haplotype matching are used to anchor the identifying the entire matched segment that in some instances extends beyond the initial exact window match. In some embodiments, the method 100 includes extending the exact window match until two homozygous mismatching SNPs are observed on either side of the original exact window match. As a result, the method 100 determines the segment width by determining the minimum and maximum of the start and end locations of the windows with no homozygous mismatches and extending the exact window match.

[0040] In case that the determined segment width exceeds a threshold cM width threshold width, the method 100 identifies the corresponding segment as a matched segment. In some

embodiments, the threshold cM width is 5 cM. In some embodiments, the threshold cM width threshold is 6, 7, 8, 9, 10, 11, 12, 13, 14, 15 cM or any value larger than 5 cM. In some embodiments, the method 100 matches 104 segments of individuals included in a database of individuals that includes the haplotype data of each individual's genome.

II.B. Per-window Match Count of Matched Segments

[0041] The method 100 includes for each individual in the population counting 108 the matched segments in each segment window, according to some embodiments. For each individual in the population and every window, the method 100 determines a per-window match count of matched segments. A particular per-window match count k_i refers to the number of matched segments identified within the population with i indicating the window of the count. In some embodiments, since every matched segment between two individuals spans a number of windows, the method 100 translates matched segments of all individuals into the window identifiers and counts how many times each window is part of a matched segment. Matched segments of close relatives are not included in the per-window match count, which increases the likelihood that all matched segment share similar levels of uncertainty of whether or not the matched segment is part of the individuals' RGH. In some embodiments, a close relative is defined as an individual with whom the related individual has a raw TIMBER score that equals or exceeds a pre-defined close relative threshold. In some embodiments, the pre-defined close relative threshold is 50 cM. In some embodiments, the pre-defined close relative threshold is 30, 40, 50, 60, 70, 80, 90, 100 cM or any value from the range of 30 to 200 cM. A raw TIMBER score between two individuals is defined as the sum of the first cM widths of matched segments over all matched segments between the two individuals. The score is referred to as a "raw" score, since the cM widths used are the first cM widths that are not down-weighted.

[0042] In some embodiments, the method of determining the per-window match count for each individual includes the following steps:

1. initializing a per-window match count vector $\{k_i\}_{i=0\dots K}$ to zero counts (one value k_i for every matching window in the genome, e.g., with K equal to 4105 windows, and each window including 96 SNP markers; and
2. for every matched segment:
 - (a) skipping matched segment for close relatives and move to the next matched segment,
 - (b) translating the matched segment into a vector of window indices $\{i\}$ (that the matched segment spans), and

- (c) incrementing the per-window match count vector entries $\{k_i\}$ by one for respective values of the vector of window indices $\{i\}$.
3. removing any entries from the per-window match count vector $\{k_i: k_i > 0\}_{i=0\dots K}$ having a matched segment count of zero.

[0043] Fig. 2 illustrates an example of per-window match counts on a portion of the genome of a particular individual in the population. The x-axis shows the number of a particular window i along the haplotype sequence, and the y-axis represents the total number of matched segments k_i for a particular window i . In this example, the displayed windows range from number 1 to number 145. Fig. 2 displays windows with zero matched segments, while these windows are excluded from the per-window match count vector $\{k_i: k_i > 0\}_{i=0\dots K}$. The largest match counts are observed for window number 1 and number 2 in this example, reaching close to 60 matched segments.

[0044] Fig. 3 illustrates an example of the histogram of the per-window match counts $\{c: c = 0 \dots C_{max}, C_{max} \leq n\}$ for all the windows in the entire genome for an individual, which only shows a portion of the individual's genome. The histogram indicates the frequency of observing matched segments against the entire population n in each genomic window, wherein n therefore limits the maximal count per window C_{max} . Only windows including less than 305 matched segments are displayed with the count of windows rapidly decreasing for increasing per-window match count values. In this example, windows including zero matched segments are not shown, since these windows are not further analyzed in the method 100 of calculating the TIMBER score.

[0045] As shown in Fig. 2, certain segment windows or regions within the individual's genome display a very large count of matched segment, whether due to a high level of non-RGH matched segments for unknown reason or due to truly high distribution of matched RGH segments in these windows. The method 100 attempts to differentiate between these two possibilities by evaluating the per-window match count for every individual in the population. Furthermore, the method 100 determines a TIMBER score for each matched segment by down weighting matching windows of a matched segment that have a decreased likelihood to originate from the RGH between the two matched individuals. This down weighting, in some embodiments, includes determining a weight for each individual for a matched segment in a window.

II.C. Estimating Weights $\{w_i^A\}$

[0046] In some embodiments, the method 100 includes estimating 110 a weight

associated with each segment window based on the count of matched segments in the associated segment window, according to some embodiments. In some embodiments, the estimating 110 includes determining for each individual in the population a weight for each window that counts at least one matched segment. In some embodiments, the weight is approximated by the probability that matching in that window for that individual provides evidence for RGH. This probability is related to the count of matched segments in a window. A window with an extremely high matched segment count for an individual is very unlikely due to the RGH that the individual shares with other individuals in the population. Unknown factors other than RGH may account for a very high matched segment count as described above.

[0047] To estimate the weights, the method 100 determines the probability of RGH, $\text{Prob}(\text{RGH}|C = c)$, given the measured count c of matched segments in a window. The random variable C represents all possible counts of matched segments in a window and is assumed to be the identical across all windows. By measuring the actual counts c in particular windows, the method 100 determines the probability of RGH on the condition that C equals c for this window. To determine $\text{Prob}(\text{RGH}|C = c)$, the method 100 uses Bayes theorem that provides:

$$\text{Prob}(\text{RGH}|C = c) = \frac{\text{Prob}(C = c|\text{RGH})\text{Prob}(\text{RGH})}{\text{Prob}(C = c)} \quad (1),$$

where $\text{Prob}(C = c|\text{RGH})$ is the probability of having c matched segments in the window and all matched segments are due to the individuals' RGH, $\text{Prob}(C = c)$ is the probability of having c matched segments in the window regardless of the matched segment being from RGH or non-RGH, and $\text{Prob}(\text{RGH})$ is the probability that the matched segment is from the individuals' RGH. Based on the Bayes theorem, the method 100 determines estimates of $\text{Prob}(C = c|\text{RGH})$, $\text{Prob}(\text{RGH})$, and $\text{Prob}(C = c)$, wherein $\text{Prob}(\widehat{C} = c|\text{RGH})$, $\text{Prob}(\widehat{\text{RGH}})$, $\text{Prob}(\widehat{C} = c)$ are the estimates of $\text{Prob}(C = c|\text{RGH})$, $\text{Prob}(\text{RGH})$ and $\text{Prob}(C = c)$, respectively. The method 100 then determines the weight w_i^A for an individual A in a specific window i according to:

$$w_i^A = \frac{\text{Prob}(\widehat{C} = k_i|\text{RGH})\text{Prob}(\text{RGH})}{\text{Prob}(\widehat{C} = k_i)} \quad (2).$$

[0048] In some embodiments, since it is difficult to estimate $\text{Prob}(C = c|\text{RGH})$, the method 100 generates at least two slightly different estimates of $\text{Prob}(C = c|\text{RGH})$, and then selecting the estimate of the at least two estimates that results in the greatest down weighting

of C for determining the weight w_i^A .

II.C.1. Determine $\text{Prob}(\widehat{C} = c)$ Estimate

[0049] In some embodiments, the method 100 includes determining an estimate of the probability distribution of C , $\text{Prob}(\widehat{C} = c)$, which provides the likelihood that a window has a given matched segment count c for all the possible counts. These embodiments provide a more accurate prediction of the probability in the number of people that an individual matches in one window based on the probability of matching each individual in the population. For example, the value of $\text{Prob}(\widehat{C} = 20)$ is the probability of counting 20 matched segment in a given window for the entire haplotype data, including all matched segment of the population. $\text{Prob}(\widehat{C} = c)$ includes contributions from matched segments that are from RGH and non-RGH of the population's individuals. Fig. 4 illustrates $\text{Prob}(\widehat{C} = c)$ of one individual of the population based on the individual's entire genome. Given $\text{Prob}(\widehat{C} = c)$ and $\text{Prob}(C = c|\text{RGH})$, the method 100 is able to quantify the likelihood that a window convey information about RGH for a given count c . In some embodiments, both distributions estimate distribution of counts for counts that are greater than zero, i.e. windows that have at least one matched segment within them.

[0050] Figs. 3 and 4 illustrate $\text{Prob}(\widehat{C} = c)$ in form of a histogram of per-window match counts for all non-zero count windows without counting any windows that include zero segments. In some embodiments, only windows with zero segments are counted for at least the following two reasons: 1) different biological or observational process are the likely cause for discovering a matched segment within a window as compared to the number of discovered matched segments within that window for a given population; and 2) the weighing of windows should be based on matched segments that are actually present in and not missing from a population. Thus, these embodiments avoid effectively assigning a weight of one to windows with zero discovered matched segments. In comparison, Fig. 2 illustrates values of the counts for a specific window (and not the frequency with which that count is observed across all the non-zero windows).

[0051] To determine a distribution to the likelihood that a particular count is observed in a window, i.e., $\text{Prob}(C = c)$, the method 100 fits a distribution of observed counts in all the non-zero windows illustrated as a histogram in Fig. 4.

[0052] In some embodiments, the method 100 uses a first beta-binomial distribution to approximate the distribution of the per-window counts, $\text{Prob}(\widehat{C} = c)$, which represents the probability that each individual matches another individual in the population. The advantage

of using a beta-binomial distribution is that it is able to account for the underlying heterogeneity in the probability of matching individuals without identifying the reason for the heterogeneity. As illustrated in Fig. 4, the beta-binomial distribution provides a good fit to per-window counts in an individual. In comparison, if all individuals in the population are assumed to match with an equal chance to any other individuals in the population, the binomial distribution provides a model of the observed counts in all the non-zero windows. A binomial distribution would typically be used to model the probability of a number of events being successful, if the known probability of success is shared across all the independent events.

[0053] Fig. 4 further illustrates an example of fitting a binomial distribution to the per-window counts in an individual as compared to a beta-binomial distribution. Generally, the beta-binomial distribution is preferred for modeling $\text{Prob}(\widehat{C} = c)$. In particular, Fig. 4 illustrates an example of the histogram of per-window match count for all the non-zero count windows where the maximum displayed per-window match count is 40. The binomial distribution modeling $\text{Prob}(\widehat{C} = c)$ is shown as a dashed line with the solid line indicating the fit of the beta-binomial distribution to the histogram data.

[0054] In some embodiments, the method 100 determines two parameters α and β of the beta-binomial distribution to determine the optimal fit between the beta-binomial distribution and the per-window match count of matched segments for one individual based on the entire population of individuals. Since beta-binomial distribution is defined for counts from zero to a maximal count n that equals the population size, the method 100 uses a modified per-window match count vector that is obtained by subtracting one from each element of $\{k_i: k_i > 0\}_{i=0\dots K}$. The number of observations used to determine the beta-binomial distribution equals the number of windows K across the haplotype data. For example, the haplotype data is divided into 4105 windows, each window including 96 markers. The joint beta-binomial distribution $f(\{k_i\}|n, \alpha, \beta, k_i > 0)$ is given by:

$$f(\{k_i\}|n, \alpha, \beta, k_i > 0) = \prod_{i=1}^K \binom{n}{k_i - 1} \frac{B(k_i - 1 + \alpha, n - k_i + 1 + \beta)}{B(\alpha, \beta)} \quad (3),$$

where $\{k_i: k_i > 0\}_{i=0\dots K}$ is the vector of per-window counts (for the K windows with at least one matched segment), n is the population size, α and β are parameters of the distribution, and B is the beta function. $\text{Prob}(C = c)$ for a population size n is then given by:

$$\text{Prob}(C = c | c > 0) = \binom{n}{c} \frac{B(c + \alpha, n - c + \beta)}{B(\alpha, \beta)} \quad (4).$$

[0055] The method 100 uses the haplotype data of all individuals who match an individual to determine the parameters α and β of the distribution associated with this individual. This determination is based in part on the assumption that count of each window is independent from the count of any other window. For typical examples, this assumption provides a good estimation of actual data. In some instances, this assumption might not hold true, and other distribution functions can provide better approximations to the data.

[0056] In some embodiments, the method 100 uses maximum-likelihood estimation to determine the two parameters α and β of the joint beta-binomial distribution $f(\{k_i\}|n, \alpha, \beta)$. The observed per-window match count vector $\{k_i: k_i > 0\}_{i=0\dots K}$ represents at most K fixed parameters of the likelihood function, $\mathcal{L}(\alpha, \beta|\{k_i\}, n) = f(\{k_i\}|n, \alpha, \beta)$. The parameters α and β are real numbers larger than zero that maximize the average logarithm of the likelihood function.

[0057] In some embodiments, the method 100 uses an optimization method 100 for estimating the maximum-likelihood $\mathcal{L}(\alpha, \beta|\{k_i\}, n)$ of the joint beta-binomial distribution for given $\{k_i\}$ and n . Various optimization methods are well-known in the art, each of which can be used for the maximum-likelihood estimation. In some embodiments, the method 100 applies the “Nelder-Mead” simplex optimization (Nelder J.A. and Mead R., A simplex algorithm for Function Minimization, *Computer J.*, 7:308-13, 1965) algorithm as the optimization method 100. The parameters α and β are initially set to starting values denoted by $\tilde{\alpha}$ and $\tilde{\beta}$ provided by:

$$\tilde{\alpha} = \frac{nE1 - E2}{n\left(\frac{E2}{E1} - E1 - 1\right) + E1} \quad (5),$$

$$\tilde{\beta} = \frac{(n - E1) \cdot \left(n - \frac{E2}{E1}\right)}{n\left(\frac{E2}{E1} - E1 - 1\right) + E1} \quad (6),$$

$$\text{where } E1 = \frac{\sum_{i=1}^K (k_i - 1)}{n} \quad \text{and} \quad E2 = \frac{\sum_{i=1}^K (k_i - 1)^2}{n}.$$

[0058] The parameter α and β obtained by the optimization method 100 are denoted by $\hat{\alpha}$ and $\hat{\beta}$. All individuals thus have an estimated beta-binomial distribution to describe $\text{Prob}(C = c | c > 0) = \binom{n}{c} \frac{B(c+\hat{\alpha}, n-c+\hat{\beta})}{B(\hat{\alpha}, \hat{\beta})}$, where $\hat{\alpha}$ and $\hat{\beta}$ are specific to each individual in a population of size n , illustrated in the example of Fig. 4.

II.C.2. Determine $\text{Prob}(\widehat{C} = c | \text{RGH})$ Estimates

[0059] In some embodiments, the method 100 determines an estimate of the probability distribution $\text{Prob}(\widehat{C} = c | \text{RGH})$ by fitting a second beta-binomial distribution to the per-window match counts that only include windows with a low match count while excluding windows with a higher match count. The first beta-binomial distribution is based on the fit to $\text{Prob}(\widehat{C} = c)$, as described above.

[0060] In some embodiments, the method 100 excludes any per-window match counts that exceed a threshold value V . Excluding windows with higher match counts is based on the rationale that windows with low match count are mainly due to matched segments from the individuals' RGHs, while high match count windows likely include matched segment that are a mixture of RGH and non-RGH. Generally, $\text{Prob}(\widehat{C} = c | \text{RGH})$ can be estimated from $\text{Prob}(C = c)$, as the later distribution represents the sum of two conditional distributions $\text{Prob}(C = c | \text{RGH})$ and $\text{Prob}(C = c | \text{non-RGH})$. However, neither of these conditional distributions are known with any confidence or can be easily determined from $\text{Prob}(C = c)$, in part because one cannot distinguish between matched segments that are from the individuals' RGHs and matched segments that are not. Attempts in estimating $\text{Prob}(\widehat{C} = c | \text{RGH})$ directly from $\text{Prob}(C = c)$ therefore result often in poor and sometimes very poor estimates, while adding various computational problems.

[0061] Furthermore, $\text{Prob}(C = c)$ is well estimated by a beta-binomial distribution as described above. To estimate $\text{Prob}(C = c | \text{RGH})$, the method 100 therefore fits a second beta-binomial distribution to the per-window match counts based on windows having a match count below or equal to the threshold value V .

[0062] As illustrated in Fig. 5, both fitted distributions of $\text{Prob}(C = c)$ and $\text{Prob}(C = c | \text{RGH})$ are similar to each other, if the distribution of counts for the high value windows that are not included in the fit of $\text{Prob}(\widehat{C} = c | \text{RGH})$ is consistent with the distribution estimated from only the low match count windows. Fig. 5 shows the first beta-binomial distribution based on all the non-zero count windows (shown by the solid line), the second beta-binomial distribution based on only the low match count windows (shown by the dashed line), and the histogram based only on the data from the low match count windows. The maximum displayed per-window match count is 40.

[0063] In some embodiments, since the estimate $\text{Prob}(\widehat{C} = c | \text{RGH})$ is likely sensitive to the user-specified threshold value V , the method 100 determines at least two estimates of

Prob ($C = c|RGH, V$) using at least two different values of V . The method 100 then selects the estimate that results in smaller weights for the windows, more effectively down-weighting the individuals' matched segment counts as described in more detail below. In some embodiments, the method 100 determines only one estimate of Prob($\widehat{C} = c|RGH$). In some embodiments, V is specified as the maximum value of a minimum threshold value V_{min} and a quantile of all the counts in windows with at least one matched segment. In some embodiments, the minimum threshold value V_{min} is set to 11. In some embodiments, the minimum threshold value V_{min} is 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 13, 14, 15, 20, 25, 30, or larger than 30. In some embodiments, at least two quantiles are specified to determine the at least two estimates of Prob($\widehat{C} = c|RGH, V$). In some embodiments, two specified quantiles are 75% and 90%, respectively. In some embodiments, the first specified quantile equals to 50%, 55%, 60%, 65%, or 70%, and the second specified quantile equals to 65%, 70%, 75%, 80%, or 85%. In some embodiments, the two quantiles are specified so that the difference between the two quantiles is in the range of 10%-20%, and the smaller quantile equals to 50%, 55%, 60%, 65%, 70%, 75% or any value in the range of 40%-80%.

[0064] In some embodiments, to estimate Prob($C = c|RGH, V$) the method 100 uses a joint beta-binomial distribution $f(\{m_i\}|n, \alpha, \beta, V, m_i > 0)$ that depends on V , α and β and is given by:

$$f(\{m_i\}|n, \alpha, \beta, V, m_i > 0) = \frac{\prod_{i=1}^M \binom{n}{m_i-1} \frac{B(m_i-1+\alpha, n-m_i+1+\beta)}{B(\alpha, \beta)}}{\text{Prob}(0 < C \leq V|\alpha, \beta)} \quad (7).$$

[0065] where $\{m_i: 0 < m_i \leq V\}_{i=0\dots M}$ is the vector of per-window counts (for the M windows with at least one matched segment and less than or equal to V counts), n is the population size, α and β are the parameters for the beta function B , and Prob ($0 < C \leq V|\alpha, \beta$) is the probability that a per-window count of matched segment is greater than zero and less than or equal to V conditional to α and β .

[0066] The probability distribution of C conditional to RGH is then given by:

$$\text{Prob}(C = c|RGH, V, c > 0) = \frac{\binom{n}{c} \frac{B(c+\alpha, n-c+\beta)}{B(\alpha, \beta)}}{\text{Prob}(0 < C \leq V|\alpha, \beta)} \quad (8),$$

wherein Prob($0 < C \leq V|\alpha, \beta$) represents a normalization factor that is a function of α and β and given by:

$$\text{Prob}(0 < C \leq V | \alpha, \beta) = \sum_{c=1}^V \binom{n}{c} \frac{B(c + \alpha, n - c + \beta)}{B(\alpha, \beta)} \quad (9).$$

[0067] In some embodiments, the method 100 uses maximum-likelihood estimation to determine the two parameters α and β of the joint beta-binomial distribution $f(\{m_i\} | n, \alpha, \beta, V)$ using the same or similar optimization algorithms that are used in estimating $\text{Prob}(C = c)$ as described above. In some embodiments, the method 100 uses the same or similar starting values for α and β as are used to estimate $\text{Prob}(C = c)$. Using the same starting values for the first and the second joint beta-binomial distribution minimizes the effect of the starting values on the differences between these two distributions.

II.C.3. Determine $\widehat{\text{Prob}}(\text{RGH})$ Estimate

[0068] In some embodiments, the method 100 then determines an estimate of $\widehat{\text{Prob}}(\text{RGH})$ based on $\text{Prob}(\widehat{C} = c)$ and $\text{Prob}(\widehat{C} = c | \text{RGH})$. In some embodiments, $\widehat{\text{Prob}}(\text{RGH})$ is set to

be maximum of $\frac{1}{\text{Prob}(\widehat{C} = c | \text{RGH})}$, which is equal to $\frac{\text{Prob}(\widehat{C} = m)}{\text{Prob}(\widehat{C} = m | \text{RGH})}$, where m is the point

at which $\frac{\text{Prob}(\widehat{C} = c | \text{RGH})}{\text{Prob}(\widehat{C} = c)}$ is maximized. Thus, the method 100 only implicitly estimates $\widehat{\text{Prob}}(\text{RGH})$ by evaluating the ratio $\text{Prob}(\widehat{C} = c | \text{RGH}) / \text{Prob}(\widehat{C} = c)$. In some embodiments, the method 100 determines at least two estimates of $\widehat{\text{Prob}}(\text{RGH}, V)$ based on the at least two $\text{Prob}(C = c | \text{RGH}, V)$. In some embodiments, the method 100 determines two $\widehat{\text{Prob}}(\text{RGH}, V)$, estimates based on the two user-specified threshold values V , where each V are determined on two specified quantiles of all the counts in windows with at least one matched segment as described above. Determining the estimate of $\widehat{\text{Prob}}(\text{RGH})$ as a ratio of two probabilities ensures that the estimate as well as the corresponding weights fall with the range of zero to one, since the weight is also determined as a ratio of three probabilities that fall within the range of zero to one. The weight would be undefined, if $\text{Prob}(\widehat{C} = c)$ is zero. The above described estimation using a beta-binomial distribution ensures that any values of $\text{Prob}(\widehat{C} = c)$ are larger than zero.

II.C.4. Estimating Temporary Weights $\{w_c\}$ For Each Estimate of $\text{Prob}(C = c | \text{RGH})$

[0069] The method 100 then determines temporary weights based on the estimates of $\text{Prob}(\widehat{C} = c)$, $\text{Prob}(\widehat{C} = c | \text{RGH})$, and $\widehat{\text{Prob}}(\text{RGH})$. These weights are temporary, since the method 100 uses these weights to determine the final weight for each window. The temporary weights can be represented by a vector $\{w_c\}_{c=0 \dots n}$ that is a series of values for different match

counts and w_c is the raw weight for a match count of c . Given the at least two estimates of $\text{Prob}(C = c|RGH)$ in some embodiments, the method 100 determines temporary weights for each estimate of $\text{Prob}(C = c|RGH)$ and then selects one series of temporary weights to determine the final weight.

[0070] If any optimization fails for one quantile it is ignored in the decision step, hence if temporary weights can only be estimated for one choice of quantile, then those temporary weights are the final weights. Considering multiple options for the quantile ensures that we are not missing a good weight vector to down weight purely because of the choice of a fixed quantile for all individuals. The choice of quantile for the estimation of $\text{Prob}(C = c|RGH)$ is made for each individual by the observation which quantile most down weights the counts.

[0071] To determine the temporary weights, the method 100 initially determines raw weights $\{r_c\}_{c=0\dots n}$ based on the estimates of $\text{Prob}(\widehat{C} = c)$, $\text{Prob}(\widehat{C} = c|RGH)$, and $\text{Prob}(\widehat{RGH})$, wherein r_c is the raw weight for a match count of c . In some embodiments, the initial values of raw weights $\{r_c\}_{c=0\dots n}$ are determined by using equation 2. Subsequently, the method 100 determines the temporary weights from the raw weights so that the temporary weights satisfy the following three conditions:

1. the temporary weight of a window with one match in it has a weight of one, i.e. $w_{c=1} = 1$;
2. the values of the temporary weights monotonically decreases with increasing match count c in a specific window, i.e. $w_c > w_{c+1}$ for all match counts c ; and
3. the temporary weights default to one for all windows, if the estimates of $\text{Prob}(C = c)$, $\text{Prob}(C = c|RGH)$ or $\text{Prob}(RGH)$ are poor.

[0072] In some embodiments, estimates of these probabilities are considered poor if the optimization algorithm used to determine the beta-binomial distributions fails, the number of per-window match counts falls below threshold value, e.g., 20, or the number of points used for fitting the beta-binomial distribution is below a minimum number.

[0073] The first two conditions can be met by enforcing that the weight monotonically decreases as the match count in a specific window increases. The rationale for a default weight of one is to avoid introducing more rather than removing noise in the TIMBER score calculation, since estimating probabilities of the weights based on underrepresented matched segment counts likely introduce noise in the calculation. Furthermore, since only low per-window match counts are measured in this case, the method 100 would take all matched segment counts, which are mainly low per-window match counts, into consideration without

weighting down any particular matched segment counts. More specifically, the estimation the probabilities using the beta-binomial distributions is limited when interpreting low match count levels.

[0074] In some embodiments, the temporary weights $\{w_c\}_{c=0\dots n}$ are determined by the following two steps that are consistent with the first two above conditions, i.e. the first temporary weight equaling one and the temporary weights monotonically decreasing with increasing match count. In the first step, the method 100 sets the temporary weights to one for all windows with a match count less than or equal to M , wherein M is the count for which the ratio between the two beta-binomial estimated distributions is highest. In the second step, after the count for which $\text{Prob}(\text{RGH})$ was “estimated,” any increase in the weight with respect to increased match count is changed to be a zero increase. In some embodiments, the method 100 performs the two steps for all c by applying the following algorithm (except when applying the default weight because of poor probability estimates):

$$1. \quad \tilde{w}_c = 1, \quad \text{if } c \leq M \quad (10),$$

$$2. \quad d_c = \begin{cases} 0, & \text{if } r_c > r_{c-1} \\ r_c - r_{c-1}, & \text{else} \end{cases} \quad (11),$$

$$\tilde{w}_c = 1 + \sum_{i=M+1}^c d_i, \quad \text{if } c > M \quad (12).$$

[0075] Figs. 6 and 7 illustrate an example of the per-window temporary weights $\{w_c\}$ (y-axis) as a function of the possible per-window match count c (x-axis). In particular, Fig. 7 shows more detail than Fig. 6, since the maximum value of the per-window match count c displayed along the x-axis is set to 40. In this example, M is 4.

II.C.5. Estimating the final weights $\{w_i\}$

[0076] The method 100 includes determining final weights $\{w_i\}_{i=1\dots K}$ for weighting all matched segment windows based on the temporary weights $\{\tilde{w}_{c,V}\}_{c=0\dots C_{max,V}}$. In some embodiments, the final weight w_i is the temporary weights for a given estimate of $\text{Prob}(C = c | \text{RGH}, V)$, i.e. a given V , which minimizes the sum of the weighted per-window segment count:

$$\{\bar{w}_c\}_{c=0\dots C_{max}} = \arg \min_{\{w_{c,V}\}_V} \sum_{i=1}^K k_i \cdot \tilde{w}_{c=k_i,V} \quad (13),$$

$$\{w_i\}_{i=1\dots K} = \{w_i: w_i = \bar{w}_{k_i}, i = 1 \dots K\} \quad (14).$$

[0077] In summary, given the estimates for specific parameters for $\text{Prob}(C = c)$, multiple estimates of $\text{Prob}(C = c | \text{RGH})$, $\text{Prob}(\text{RGH})$ and some post-processing, we can find

the weight for a given window in a given individual i . Fig. 8 illustrates an example for the weights (solid line) given the per-window counts from the original example in Fig. 2 (dashed line). For comparative purposes only, the weight is re-scaled so that a weight of 1 has a value of 59 according to the y-axis. The weight calculation generates for every individual a weight value (within the interval 0 and 1) for the K segment windows.

II.D. Calculate weighted sum of per-window cM widths $cM_2^{A,B}$

[0078] The method 100 includes calculating 112 a weighted sum of per-window cM widths for each matched segment based on the first cM width and the weights associated with the segment windows of the matched segment, according to some embodiments. In particular, the method calculates the weighted sum of per-window cM widths for a matched segment, between person A and person B, given the individual-specific window estimated weights for person A and person B. The weighted sum of per-window cM widths $cM_2^{A,B}$ is the sum of the first cM widths $cM_{1,i}$ for each window i that the matched segment spans weighted by the product of the weights for both individuals, A and B, in those segment windows:

$$cM_2^{A,B} = \sum_{i=WIN_{start}}^{WIN_{end} \leq K} w_i^A \cdot w_i^B \cdot cM_{1,i} \quad (15),$$

for a segment between individuals A and B , starting at window WIN_{start} and ending at window WIN_{end} with w_{win}^A and w_{win}^B being the weights for both individuals, respectively, in window i . If the windows are either the start or the end window of the matched segment, the window widths are updated to be from either the first genetic marker in the window or to the last genetic marker in the window, respectively.

[0079] The weights are in the interval between 0 and 1 and can intuitively be thought of as a probability that this window should contribute to the new “width”. Taking the product of the weights for individuals A and B ensures that the window is valid in both individuals to be able to contribute to the weighted sum of per-window cM widths. The first and weighted sum of per-window cM widths would be identical, if all weights are equal to one. The new “width” or weighted sum of per-window cM widths typically is smaller than the raw or first cM width, and, therefore, down weights those matches in windows where there are a high number of matches either in individual A or individual B for a large population. Thus, the down weighting results in down-weight matches, which are less likely to be from recent genealogical history of the individuals. Fig. 9 illustrates an example of per-window match counts on a chunk of the genome for one individual both pre-TIMBER (dashed line) and post-TIMBER (solid line).

III. ANCESTRAL RELATIONSHIP PREDICTION

[0080] The method 100 includes estimating a degree of ancestral relatedness between two individuals based on the weighted sum of per-window cM widths of each matched segment between the two individuals, according to some embodiments. If a pair of individuals have a total sum of first cM widths of less than 60 cM, which effectively removes close relatives from the relatedness evaluation, the relationship distance is predicted based on the sum of all the weighted sum of per-window cM widths of the shared segments.

Otherwise, the summation of the first cM widths is used. This results in a relationship prediction that is more accurate for more distant relationships, while being as accurate for close relationships. This is especially true for certain ethnic groups (for example Jewish people) and whether they are assigned to be distantly related or not.

[0081] The method 100 by using, for example, the degree of ancestral relatedness, allows people to find their recent relatives and provide them with new information about their genealogy within a network of relatives (with known genealogy). In some embodiments, the degree of relatedness between two individuals represents a probability that the two individuals are ancestrally related and is equal to the weighted sum of per-window cM widths of the individuals' matched segments. In some embodiments, the degree of relatedness between two individuals is a binary yes or no answer whether the two individuals are ancestrally related based on the weighted sum of per-window cM widths. For example, if the weighted sum of per-window cM widths exceeds a relatedness threshold value, the two individuals are said to be ancestrally related. In some embodiments, the relatedness threshold value is 20, 25, 30, 35, 40, 45, 50 cM, or any value larger than 20 cM. In some embodiments, the relatedness threshold value is 30, 40 or 50 cM.

[0082] Previously, the prediction of the distance of relationship between two individuals (e.g. cousins, third cousins, etc.) was based purely on the total width of all IBD segments, where the width of the IBD segment was determined by its width in recombination distance (in cM). In this method 100, a total score is calculated based on all IBD segments between two individuals to provide relationship distance. However, method 100 uses a sum of reweighted weighted sum of per-window cM widths in this calculation to predict relatedness, if the sum of first cM widths of less than 60 cM.

[0083] In one examples, weight profiles are generated from matches to a static reference set of just over 300K samples. Those weights are then used to re-score, i.e., re-weight, all matches between any pair of individuals in the database. The weights used by method 100 are stored in their own database. In one example embodiment, based on the analysis of test sets

and real data, a threshold cM width of 5 cM was the minimum width, at which a matched segment is included in the weight estimation 110 and weighted sum of per-window cM widths calculation 112.

IV. EXAMPLE

[0084] In an example, TIMBER behavior was analyzed with a simulated test set. The test set consisted of exactly 3703 pairs of genotypes (7406 individual samples) representing relationships from parent/child (1 meiosis) to 5th cousins (12 meioses) and all in-between. Each relationship was created independently of all others by simulating meioses to create genotypes that represent individuals in the relevant part of the pedigree. The “founders” from which non-simulated genotypes come were a set of approximately 24,000 genotypes from the database that have no close estimated relationships among them based on the analysis of the non-weighted first cM widths. “Founders” used to create a given relationship were discarded and not used at all for simulating any of the other relationships. The fact that the initial founders had very few genuine relationships and were not reused helped to minimize the likelihood of relationships among synthetic genotypes that we were unable to document, but that were still possible. The test set might not be ideal for a real world scenario, since individuals were randomly paired to be parents. However, since in the simulation the real RGH segments were known, the simulation provided a way to verify how well TIMBER helped refining the analysis of matched segments. In the test, just over 300 pairs of each meiosis level were represented between from parent-child relationships (1 meiosis) to 5th cousins (12 meioses).

Table 1: TIMBER results for different minimum cutoffs (= threshold cM width)

	Percentage of matched segments kept		
Segment	Min. of 5 cM	Min. of 6 cM	Min. of 7 cM
True	92	88	84
False	3	1	1

[0085] Table 1 illustrates that TIMBER kept the vast majority (around 90%) of initially discovered IBD segments that overlap with real IBD segments. The results did not largely vary with respect to different first cM widths of 5, 6 and 7 and were used to determine if a discovered IBD segment was retained or not. TIMBER only kept at most 3% of the initially discovered IBD segments that are false positives. Thus, TIMBER presents a very useful filter

for keeping the real signals while removing the false positive signals of IBD segments between pairs of individuals due to the individuals' recent genealogical history. No "ibs" filtering was used to filter the raw and non-weighted output of matched IBD segments.

[0086] Table 2 shows that TIMBER was slightly more accurate for closer relatives, but generally worked well across the spectrum of recent genealogical history from parent-child relationships (1 meiosis) to 5th cousins (12 meioses). Fig. 10 illustrates results from TIMBER using different raw cM width, i.e., first cM width filters, including the matched percentages of segments kept for the known and unknown meioses.

Table 2: TIMBER results for true segments by meiosis and different minimum cutoffs

Meiosis	Percentage of matched segments kept		
	Min. of 5 cM	Min. of 6 cM	Min. of 7 cM
1	98	97	95
2	95	91	89
3	92	87	83
4	88	82	76
5	87	80	73
6	85	76	69
7	82	72	64
8	81	74	65
9	84	73	66
10	80	69	58
11	75	68	57
12	88	71	59

V. ADDITIONAL CONSIDERATIONS

[0087] Computing system 120 is implemented using one or more computers having one or more processors executing application code to perform the steps described herein, and data may be stored on any conventional non-transitory storage medium and, where appropriate, include a conventional database server implementation. For purposes of clarity and because they are well known to those of skill in the art, various components of a computer system, for example, processors, memory, input devices, network devices and the like are not shown in Fig. 1B. In some embodiments, a distributed computing architecture is used to implement the described features. One example of such a distributed computing platform is the Apache HADOOP® project available from the Apache Software Foundation.

[0088] In addition to the embodiments specifically described above, those of skill in the art will appreciate that the invention may additionally be practiced in other embodiments. Within this written description, the particular naming of the components, capitalization of terms, the attributes, data structures, or any other programming or structural aspect is not mandatory or significant unless otherwise noted, and the mechanisms that implement the described invention or its features may have different names, formats, or protocols. Further, the system may be implemented via a combination of hardware and software, as described, or entirely in hardware elements. Also, the particular division of functionality between the various system components described here is not mandatory; functions performed by a single module or system component may instead be performed by multiple components, and functions performed by multiple components may instead be performed by a single component. Likewise, the order in which method steps are performed is not mandatory unless otherwise noted or logically required. It should be noted that the process steps and instructions of the present invention could be embodied in software, firmware or hardware, and when embodied in software, could be downloaded to reside on and be operated from different platforms used by real time network operating systems.

[0089] Algorithmic descriptions and representations included in this description are understood to be implemented by computer programs. Furthermore, it has also proven convenient at times, to refer to these arrangements of operations as modules or code devices, without loss of generality.

[0090] Unless otherwise indicated, discussions utilizing terms such as “selecting” or “computing” or “determining” or the like refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0091] The present invention also relates to an apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general-purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a non-transitory computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, DVDs, CD-ROMs, magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, application specific integrated circuits (ASICs), or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus. Furthermore, the

computers referred to in the specification may include a single processor or may be architectures employing multiple processor designs for increased computing capability.

[0092] The algorithms and displays presented are not inherently related to any particular computer or other apparatus. Various general-purpose systems may also be used with programs in accordance with the teachings above, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description above. In addition, a variety of programming languages may be used to implement the teachings above.

[0093] Finally, it should be noted that the language used in the specification has been principally selected for readability and instructional purposes, and may not have been selected to delineate or circumscribe the inventive subject matter. Accordingly, the disclosure of the present invention is intended to be illustrative, but not limiting, of the scope of the invention.

What is claimed is:

1. A computer program product for estimating a degree of ancestral relatedness between two individuals, the computer program product stored on a non-transitory computer readable medium and including instructions configured to cause a processor to execute steps comprising:

receiving haplotype data for a population of individuals, the haplotype data including a plurality of genetic markers shared among the individuals;

dividing the haplotype data into segment windows based on the genetic markers; for each individual in the population:

based on the genetic markers, matching segments of the haplotype data that are identical between the individual and any other individual in the population, each matched segment having a first cM width exceeding a threshold cM width and being part of one or more of the segment windows;

counting the matched segments in each segment window;

estimating a weight associated with each segment window based on the count of matched segments in the associated segment window;

calculating a weighted sum of per-window cM widths for each matched segment based on the first cM width and the weights associated with the segment windows of the matched segment; and

estimating a degree of ancestral relatedness between two individuals based on the weighted sum of per-window cM widths of each matched segment between the two individuals.

2. The computer program product of claim 1, wherein threshold cM width is 5 cM, 6 cM, 7 cM, 8 cM, 9 cM, 10 cM, or any real number within the range of 5 cM to 10 cM.

3. The computer program product of claim 1, wherein the weight associated with a segment window for individual A is approximated as:

$$w_i^A = \frac{\text{Prob}(\widehat{C} = k_i | \text{RGH}) \text{Prob}(\text{RGH})}{\text{Prob}(\widehat{C} = k_i)},$$

wherein $\text{Prob}(\widehat{C} = c | \text{RGH})$, $\text{Prob}(\text{RGH})$, $\text{Prob}(\widehat{C} = c)$ are the estimates of a probability of an RGH segment given the measured count c of matched segments in a window, a probability of an RGH segment in a window, and the probability of measuring a

count c of matched segments in a window, respectively.

4. The computer program product of claim 3, wherein $\text{Prob}(\widehat{\text{RGH}})$ is approximated by the maximum of:

$$\frac{1}{\frac{\text{Prob}(\widehat{\text{C}} = c | \text{RGH})}{\text{Prob}(\widehat{\text{C}} = c)}}$$

5. The computer program product of claim 1, wherein weighted sum of per-window cM widths for a segment between two individual A and B is approximated as:

$$cM_2^{A,B} = \sum_{i=\text{WIN}_{\text{start}}}^{\text{WIN}_{\text{end}} \leq K} w_i^A \cdot w_i^B \cdot cM_{1,i},$$

the segment between individuals A and B starts at window $\text{WIN}_{\text{start}}$ and ends at window WIN_{end} with w_{win}^A and w_{win}^B being the weights associated with segment window i for individual A and B , respectively.

6. The computer program product of claim 1, wherein estimating a weight comprises calculating temporary weights \tilde{w}_c for the count c of matched segments in the associated segment window is approximated as:

$$\tilde{w}_c = 1, \quad \text{if } c \leq M,$$

$$d_c = \begin{cases} 0, & \text{if } r_c > r_{c-1}, \\ r_c - r_{c-1}, & \text{else} \end{cases},$$

$$\tilde{w}_c = 1 + \sum_{i=M+1}^c d_i, \quad \text{if } c > M,$$

wherein r_c is the weight based on the count of matched segments c and approximated as:

$$r_c = \frac{\text{Prob}(\widehat{\text{C}} = c | \text{RGH}) \text{Prob}(\text{RGH})}{\text{Prob}(\widehat{\text{C}} = c)}.$$

7. The computer program product of claim 1, wherein the weights associated with each segment window decrease if count of matched segments in the associated segment window increases.

8. The computer program product of claim 3 or 6, wherein $\text{Prob}(\widehat{\text{C}} = c)$ is approximated as:

$$\text{Prob}(\widehat{\text{C}} = c | c > 0) = \binom{n}{c} \frac{B(c + \alpha, n - c + \beta)}{B(\alpha, \beta)},$$

wherein n is the size of the population and α, β are parameters of the Beta function B .

9. The computer program product of claim 3 or 6, wherein $\text{Prob}(C = c | \text{RGH}, V)$ is approximated as:

$$\text{Prob}(C = c | \text{RGH}, V, c > 0) = \frac{\binom{n}{c} \frac{B(c + \alpha, n - c + \beta)}{B(\alpha, \beta)}}{\text{Prob}(0 < C \leq V | \alpha, \beta)},$$

wherein n is the size of the population, α, β are parameters of the Beta function B .

10. The computer program product of claim 8, wherein α and β is estimated by using a maximum likelihood estimation of a joint distribution approximated as:

$$f(\{k_i\} | n, \alpha, \beta, k_i > 0) = \prod_{i=1}^K \binom{n}{k_i - 1} \frac{B(k_i - 1 + \alpha, n - k_i + 1 + \beta)}{B(\alpha, \beta)}.$$

11. The computer program product of claim 9, wherein α and β is estimated by using a maximum likelihood estimation of a joint distribution approximated as:

$$f(\{m_i\} | n, \alpha, \beta, V, m_i > 0) = \frac{\prod_{i=1}^M \binom{n}{m_i - 1} \frac{B(m_i - 1 + \alpha, n - m_i + 1 + \beta)}{B(\alpha, \beta)}}{\text{Prob}(0 < C \leq V | \alpha, \beta)}.$$

12. The computer program product of claim 1, wherein a pre-defined close relative threshold is 50 cM.

13. The computer program product of claim 1, wherein a pre-defined close relative threshold is 30, 40, 50, 60, 70, 80, 90, 100 cM or any value from the range of 30 to 200 cM.

14. The computer program product of claim 1, wherein a size of the segment windows comprises 50, 60, 70, 80, 90, 100, 110, 120, 130, 140, 150 genetic markers or any number that falls within the range of 50 to 500 genetic marker.

15. The computer program product of claim 1, wherein the two specified quantiles are 75% and 90%, respectively, or the first specified quantile equals to 50%, 55%, 60%, 65%, or 70%, and the second specified quantile equals to 65%, 70%, 75%, 80%, or 85%.

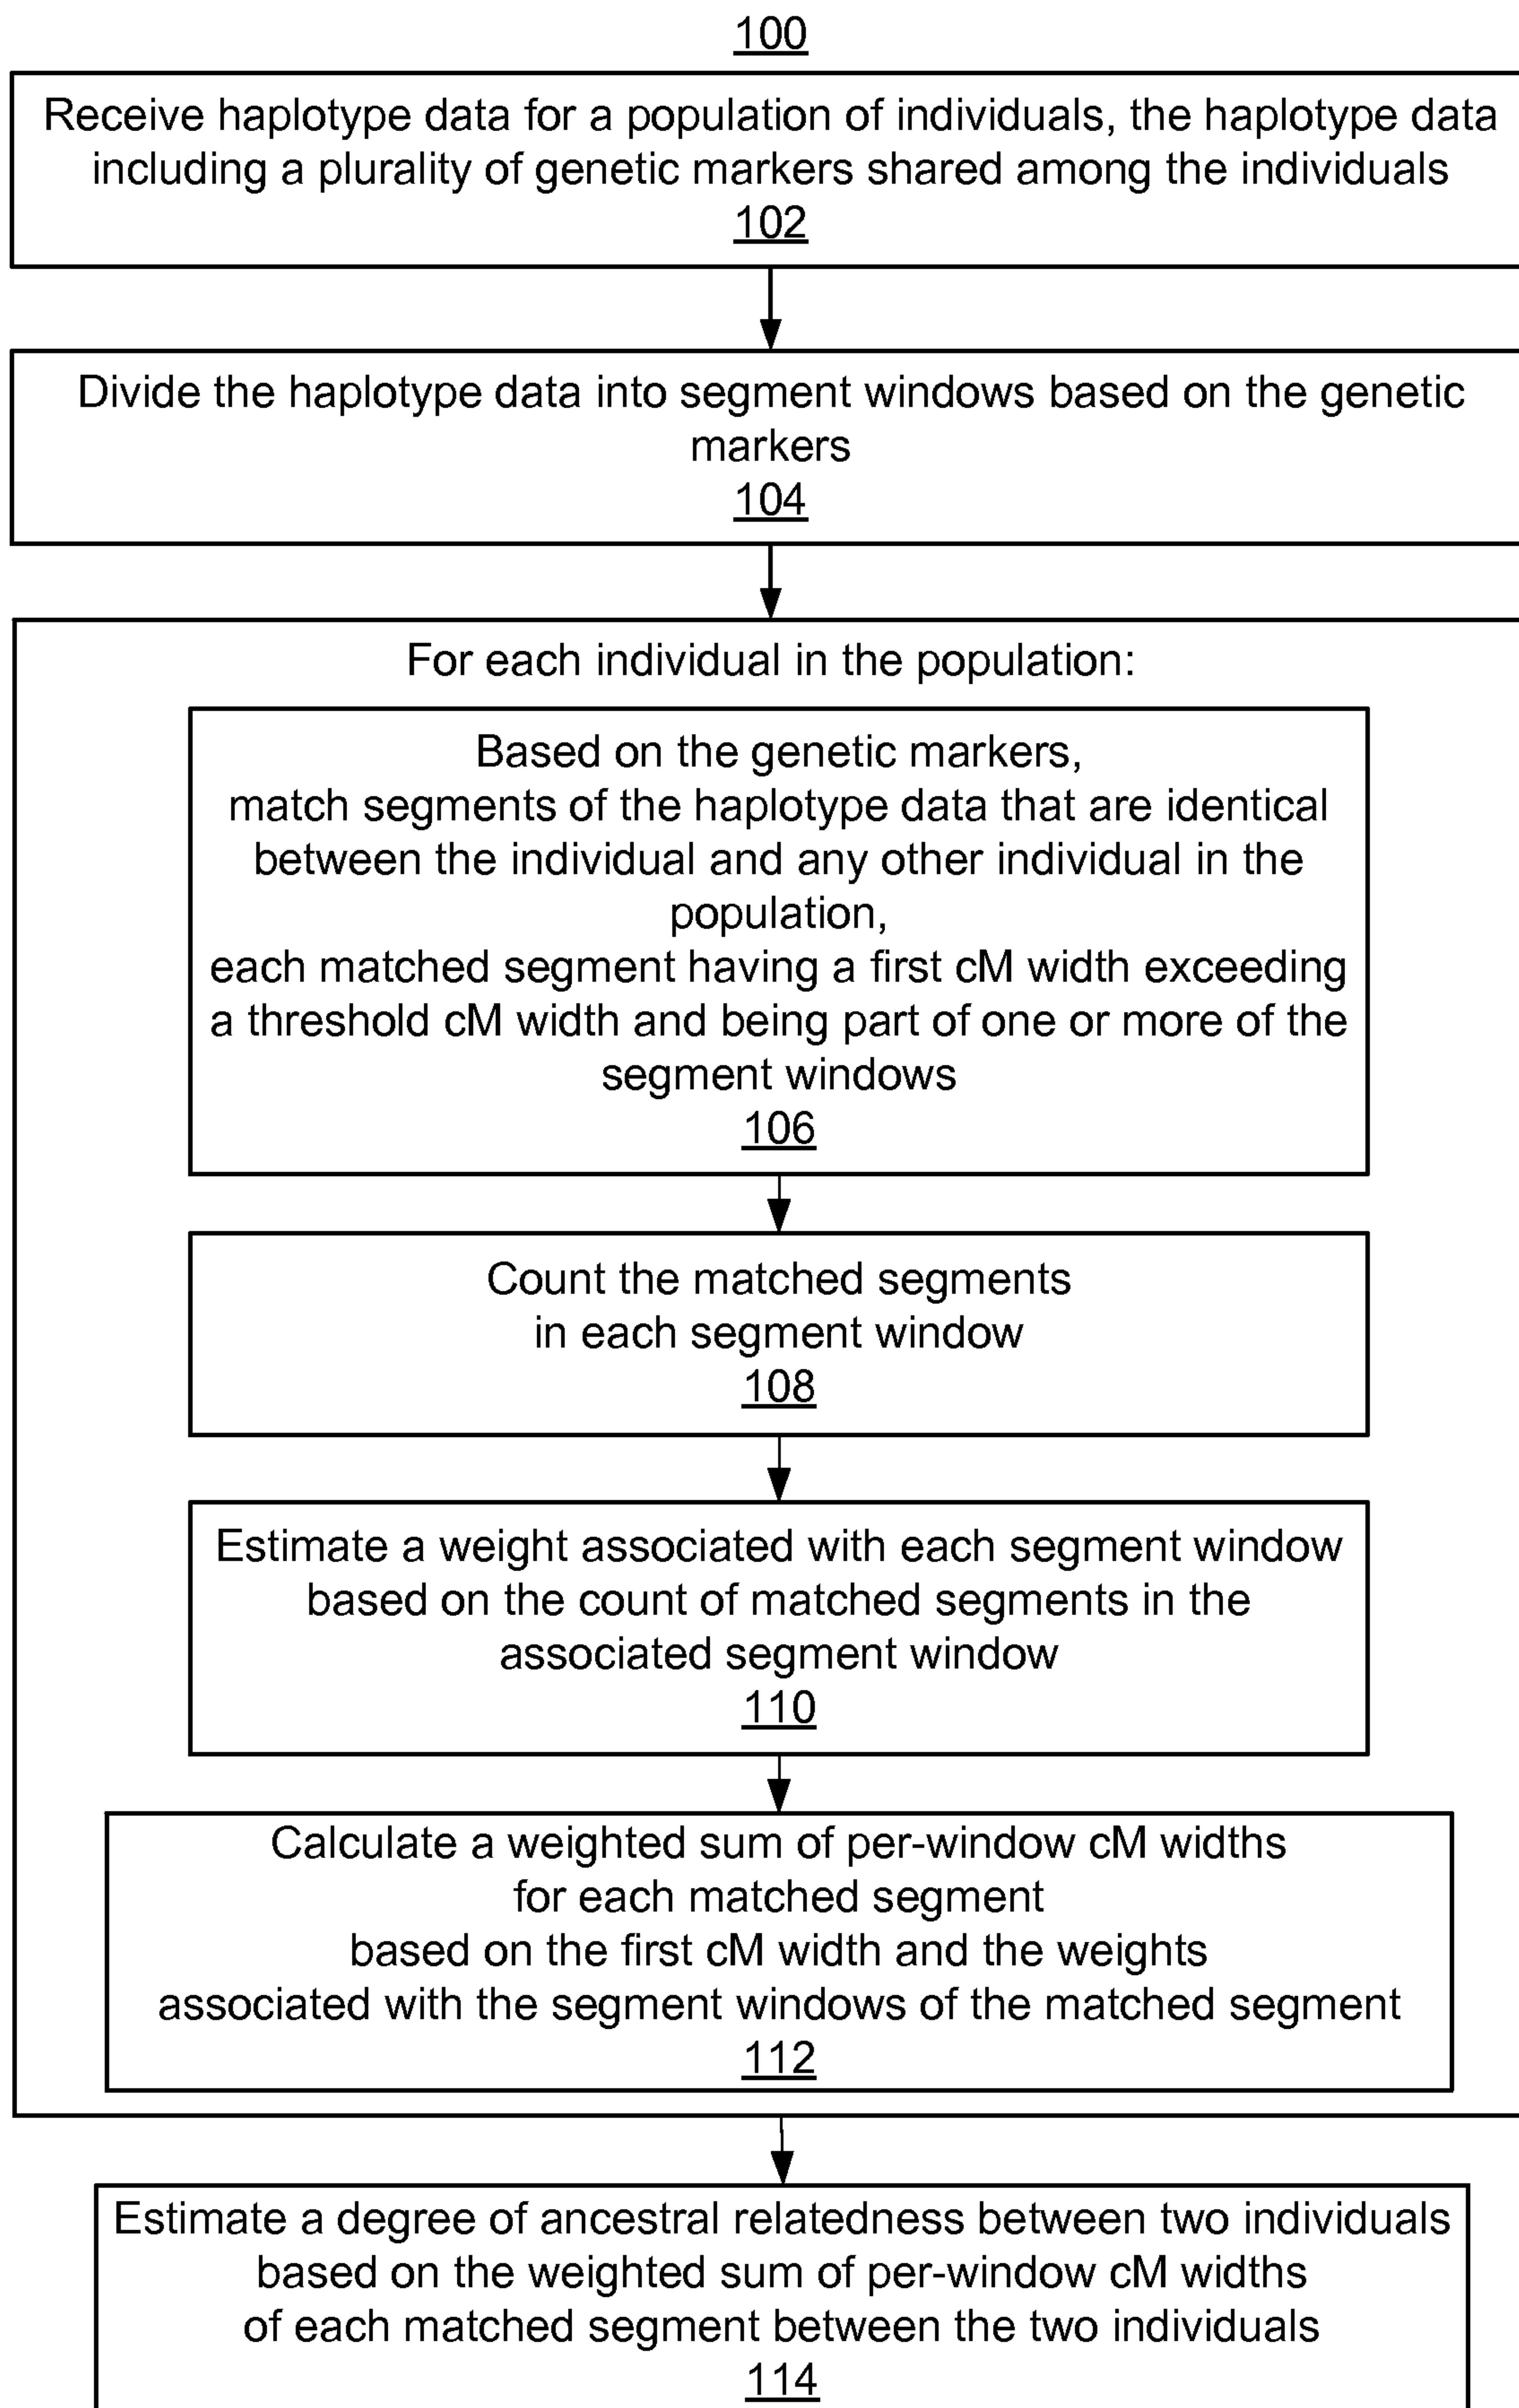
16. The computer program product of claim 1, wherein a size of the population is larger than 300,000.

17. The computer program product of claim 1, wherein the population size n , number of genetic markers per segment window d and number of segment windows K is approximated as:

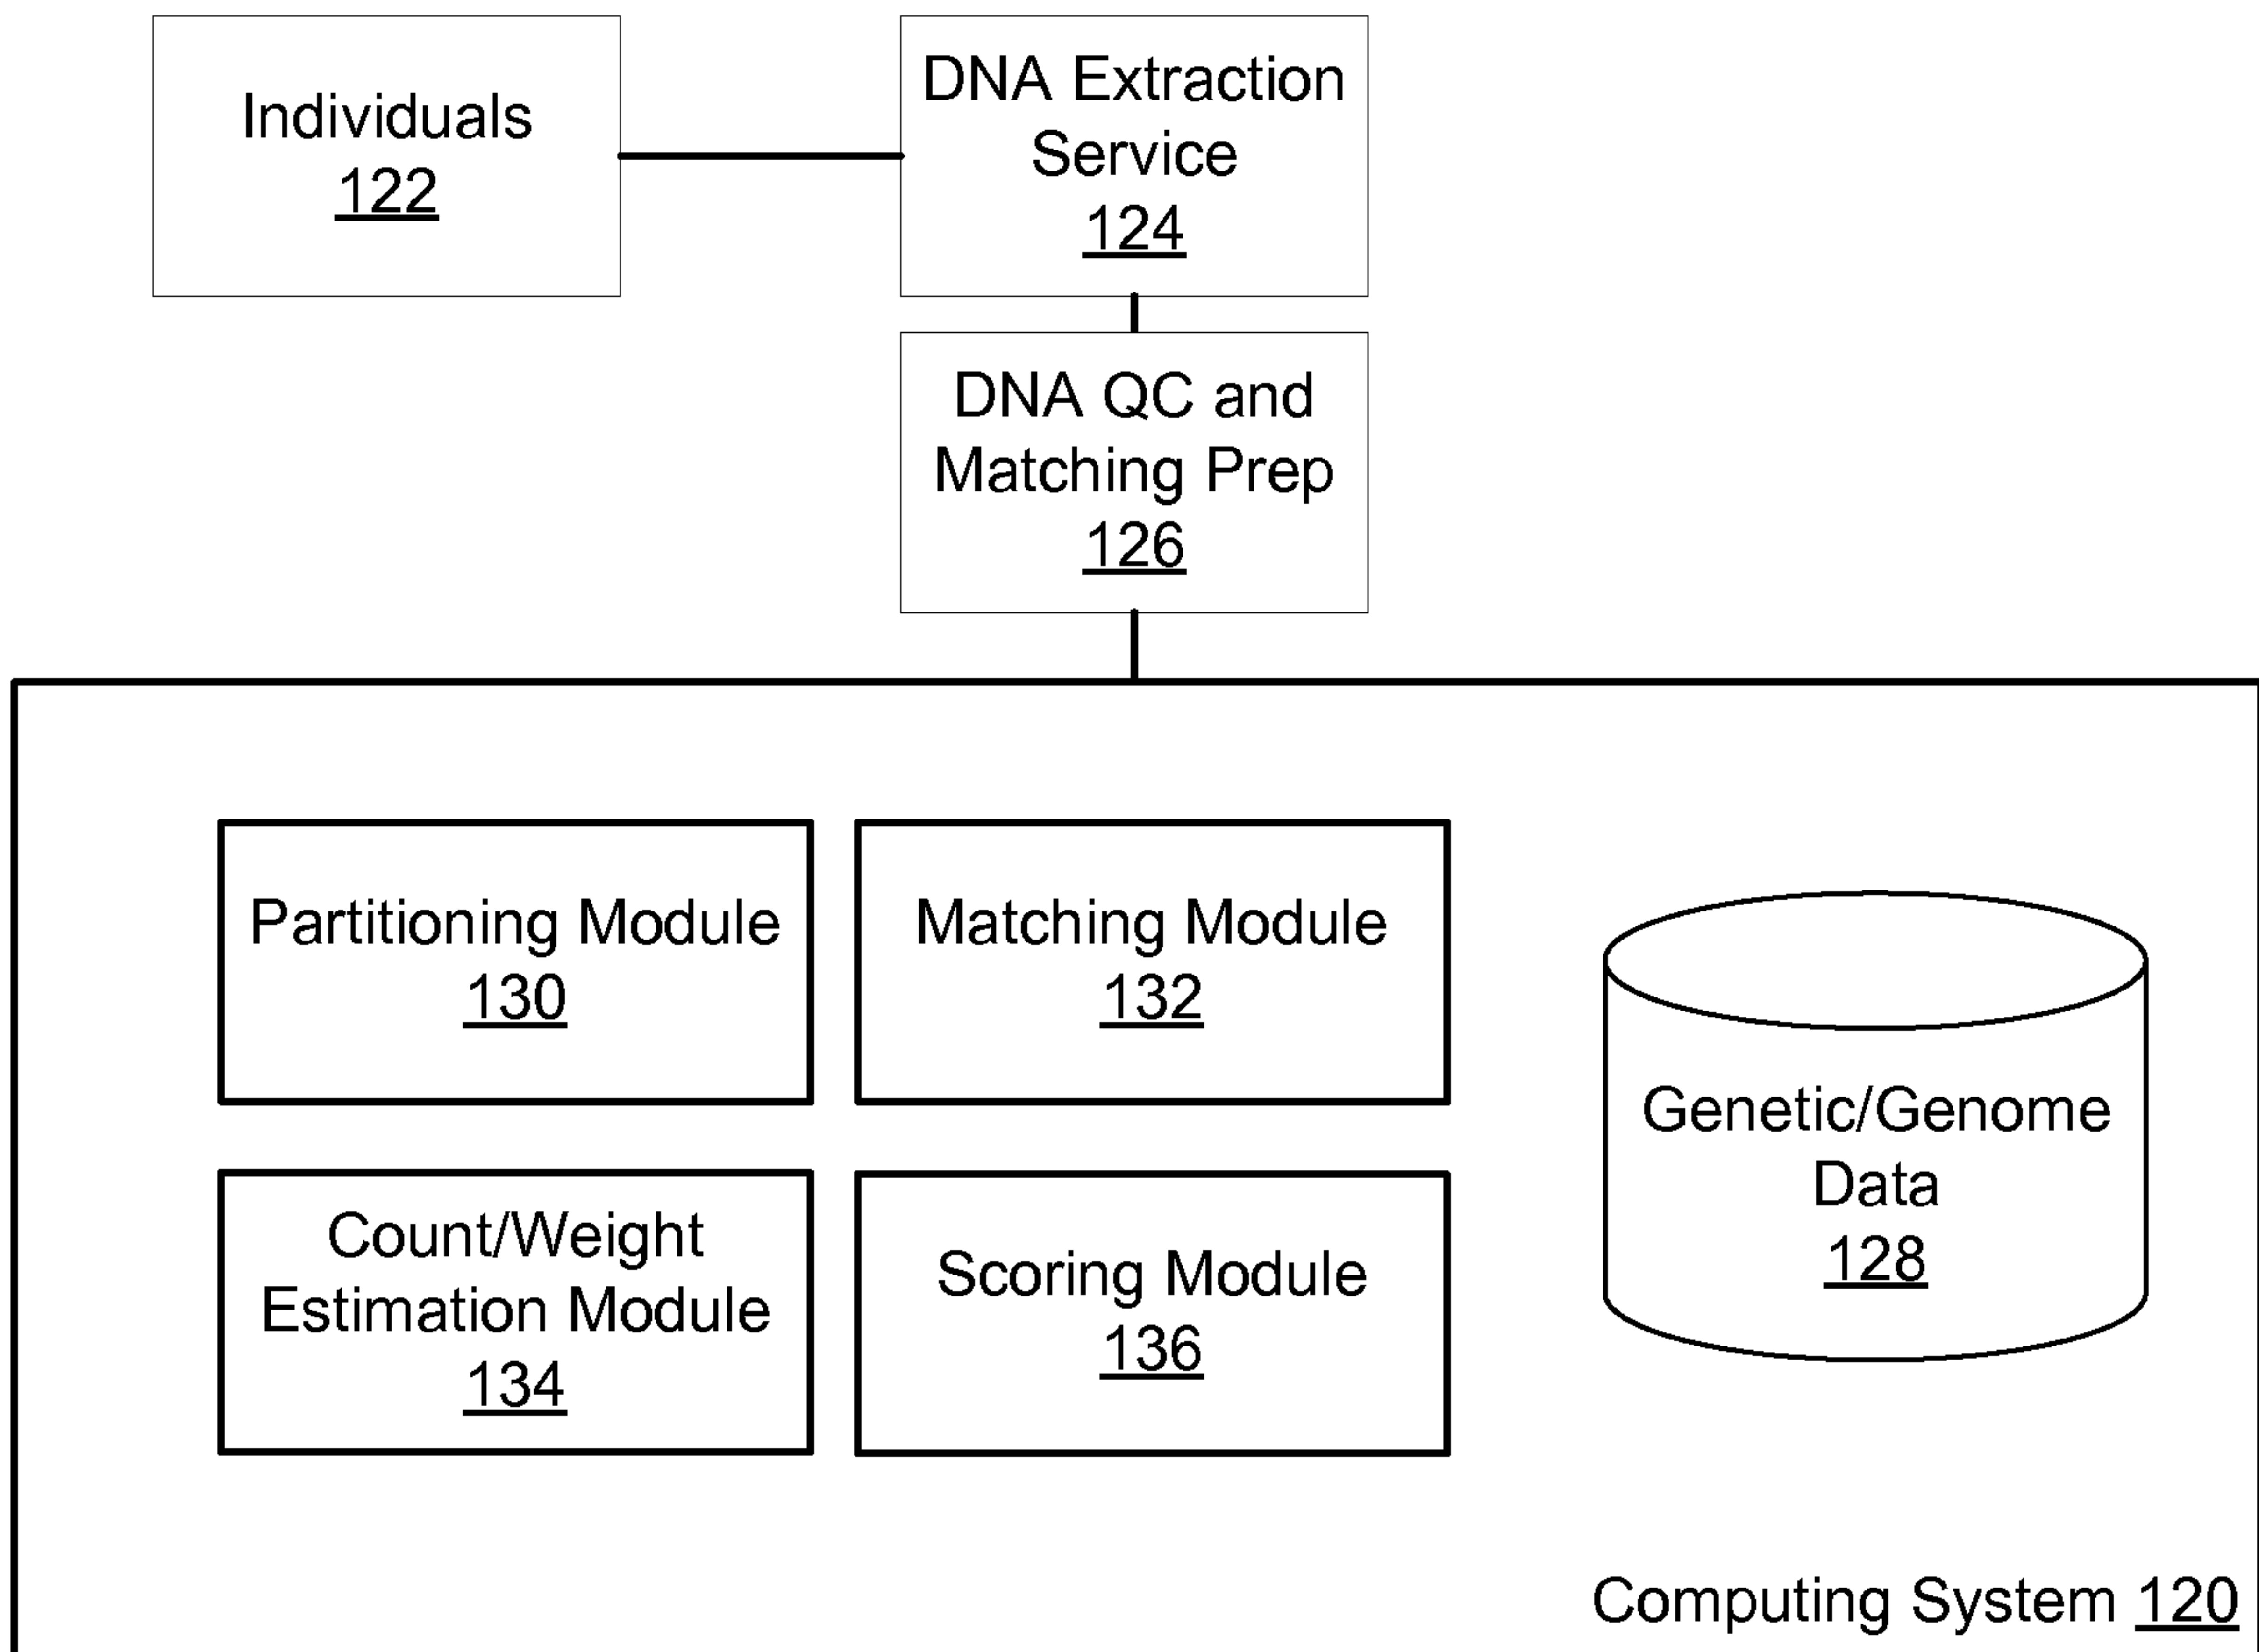
$$K = n/d.$$

18. The computer program product of claim 1, wherein close relatives are removed from the population.
19. The computer program product of claim 16, wherein close relatives comprises two individuals having a total first cM width larger than 60 cM.
20. The computer program product of claim 1, wherein the degree of relatedness between two individuals comprises a probability that the two individuals are ancestrally related.
21. The computer program product of claim 1, wherein the degree of relatedness between two individuals comprises a binary yes or no answer whether the two individuals are ancestrally related.

1/11

**FIG. 1A**

2/11

**FIG. 1B**




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

9/11




FIG. 8

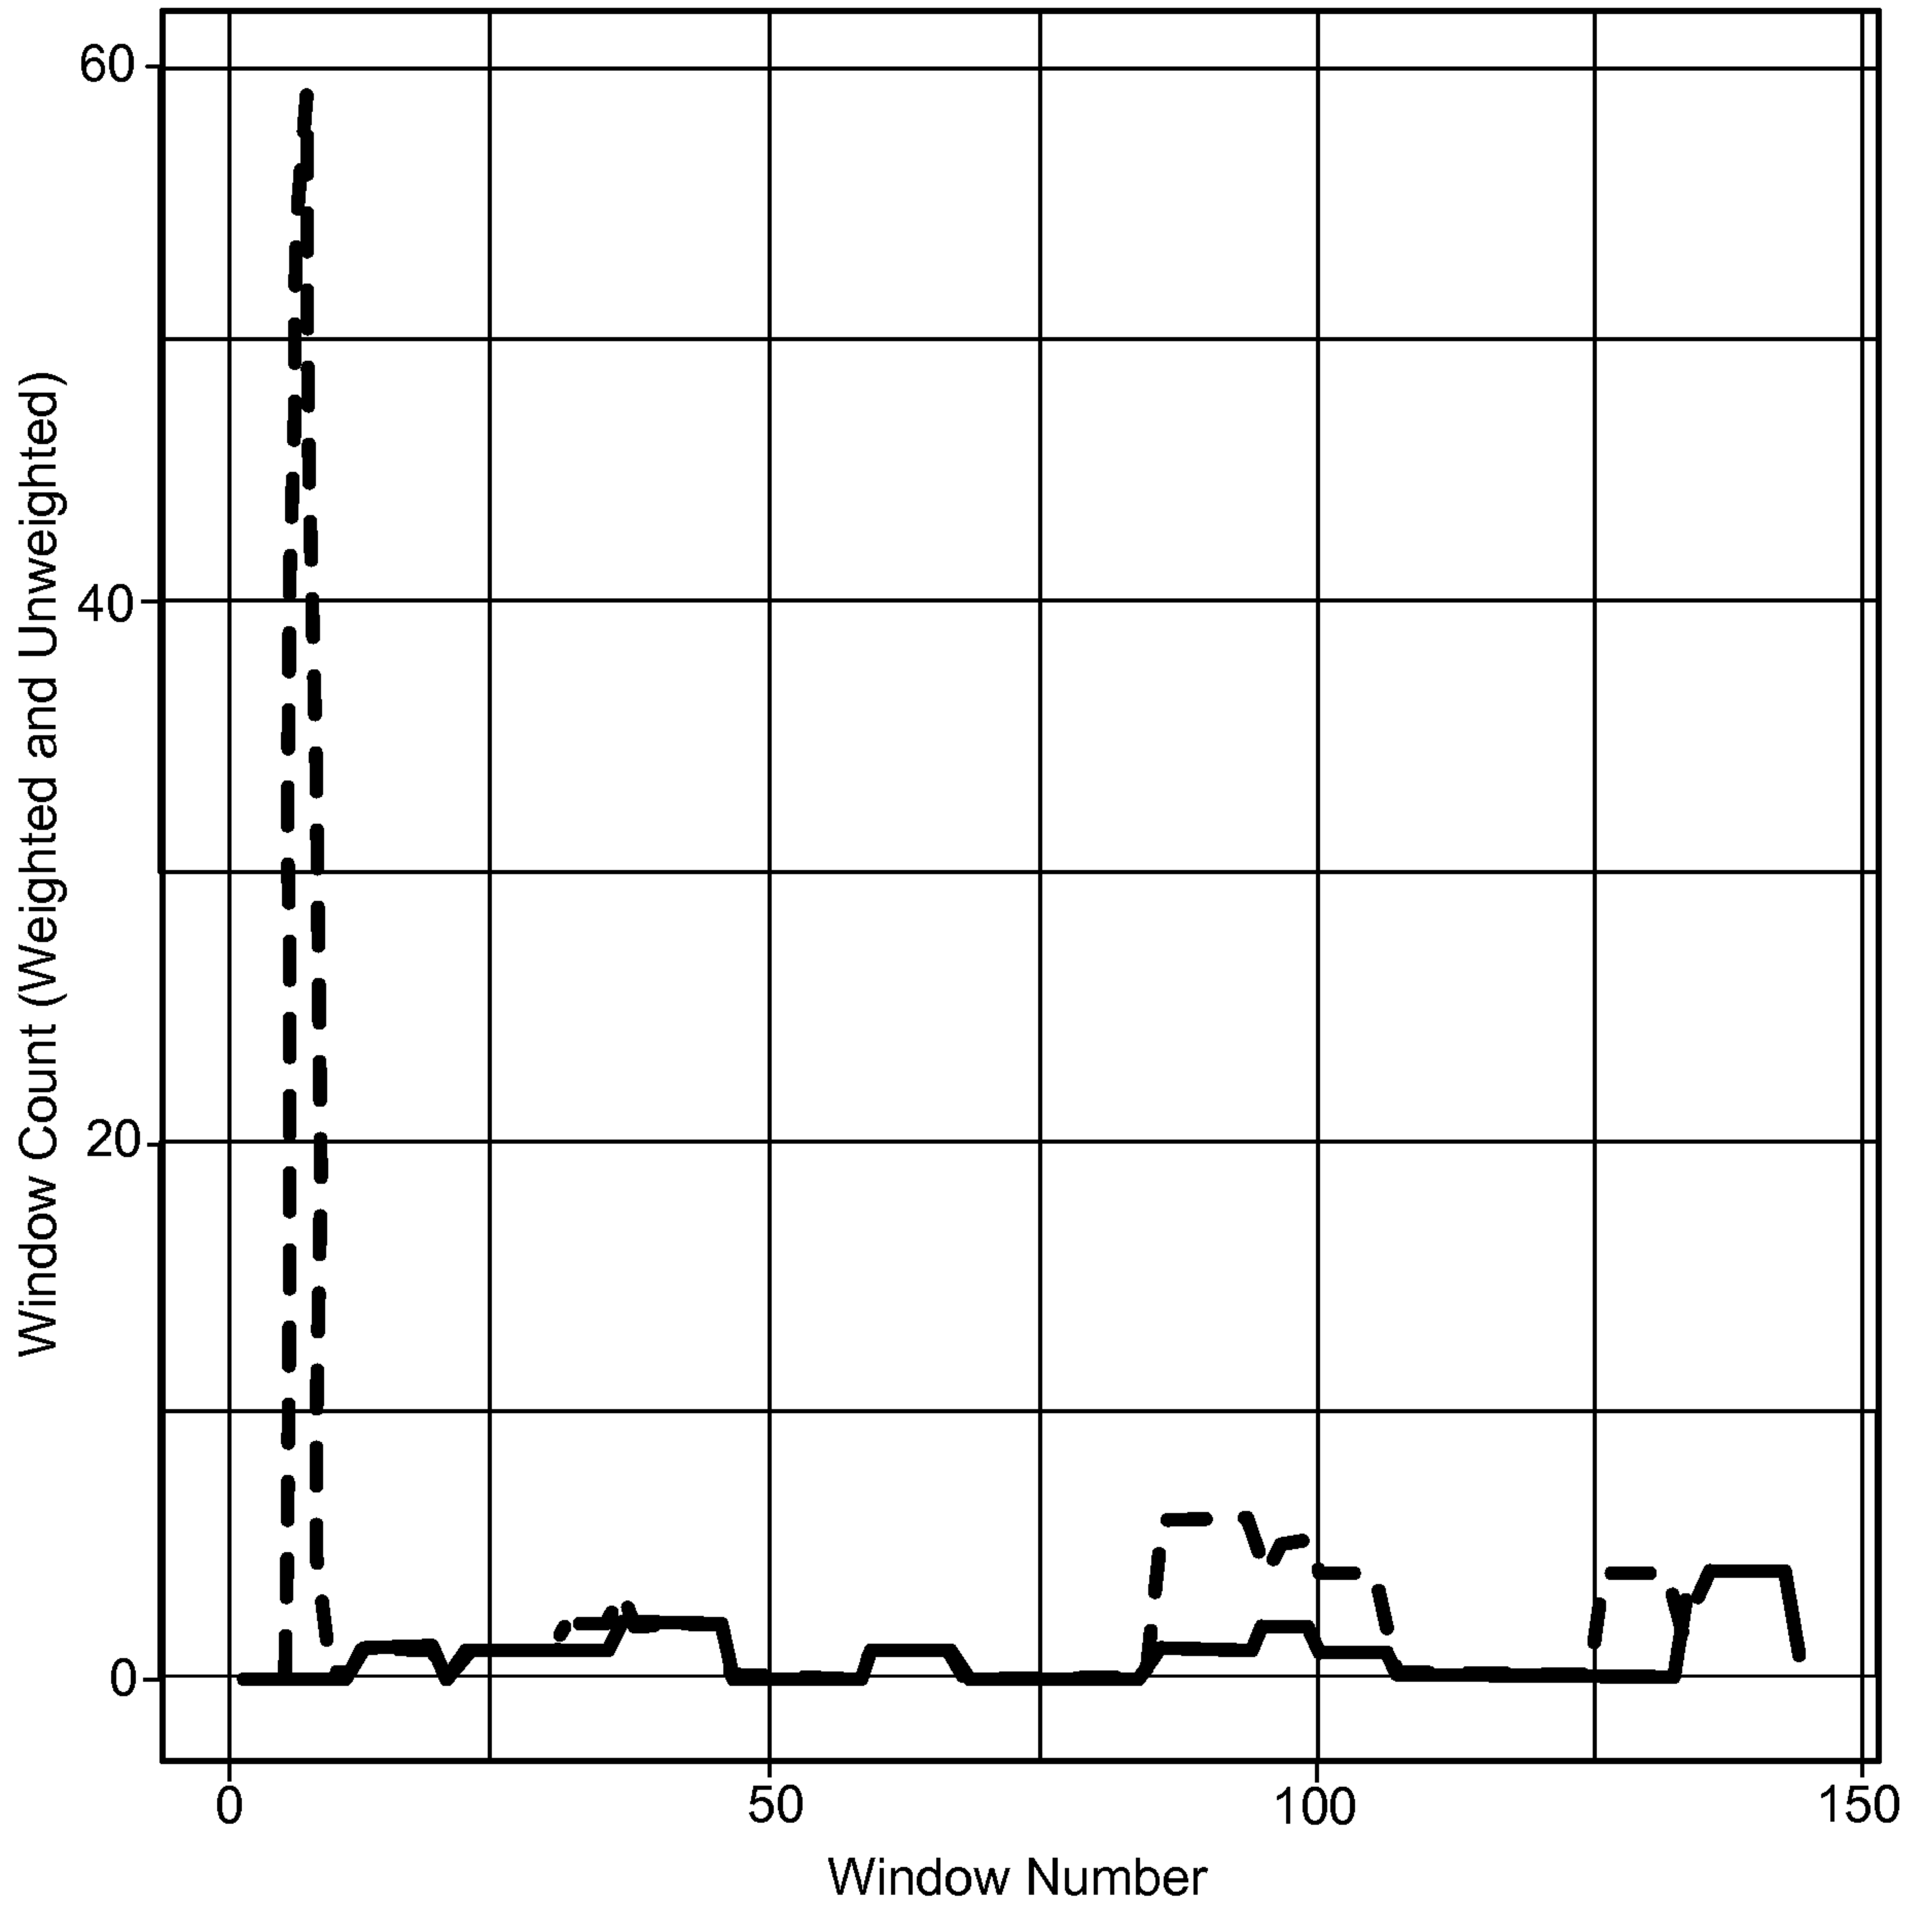


FIG. 9

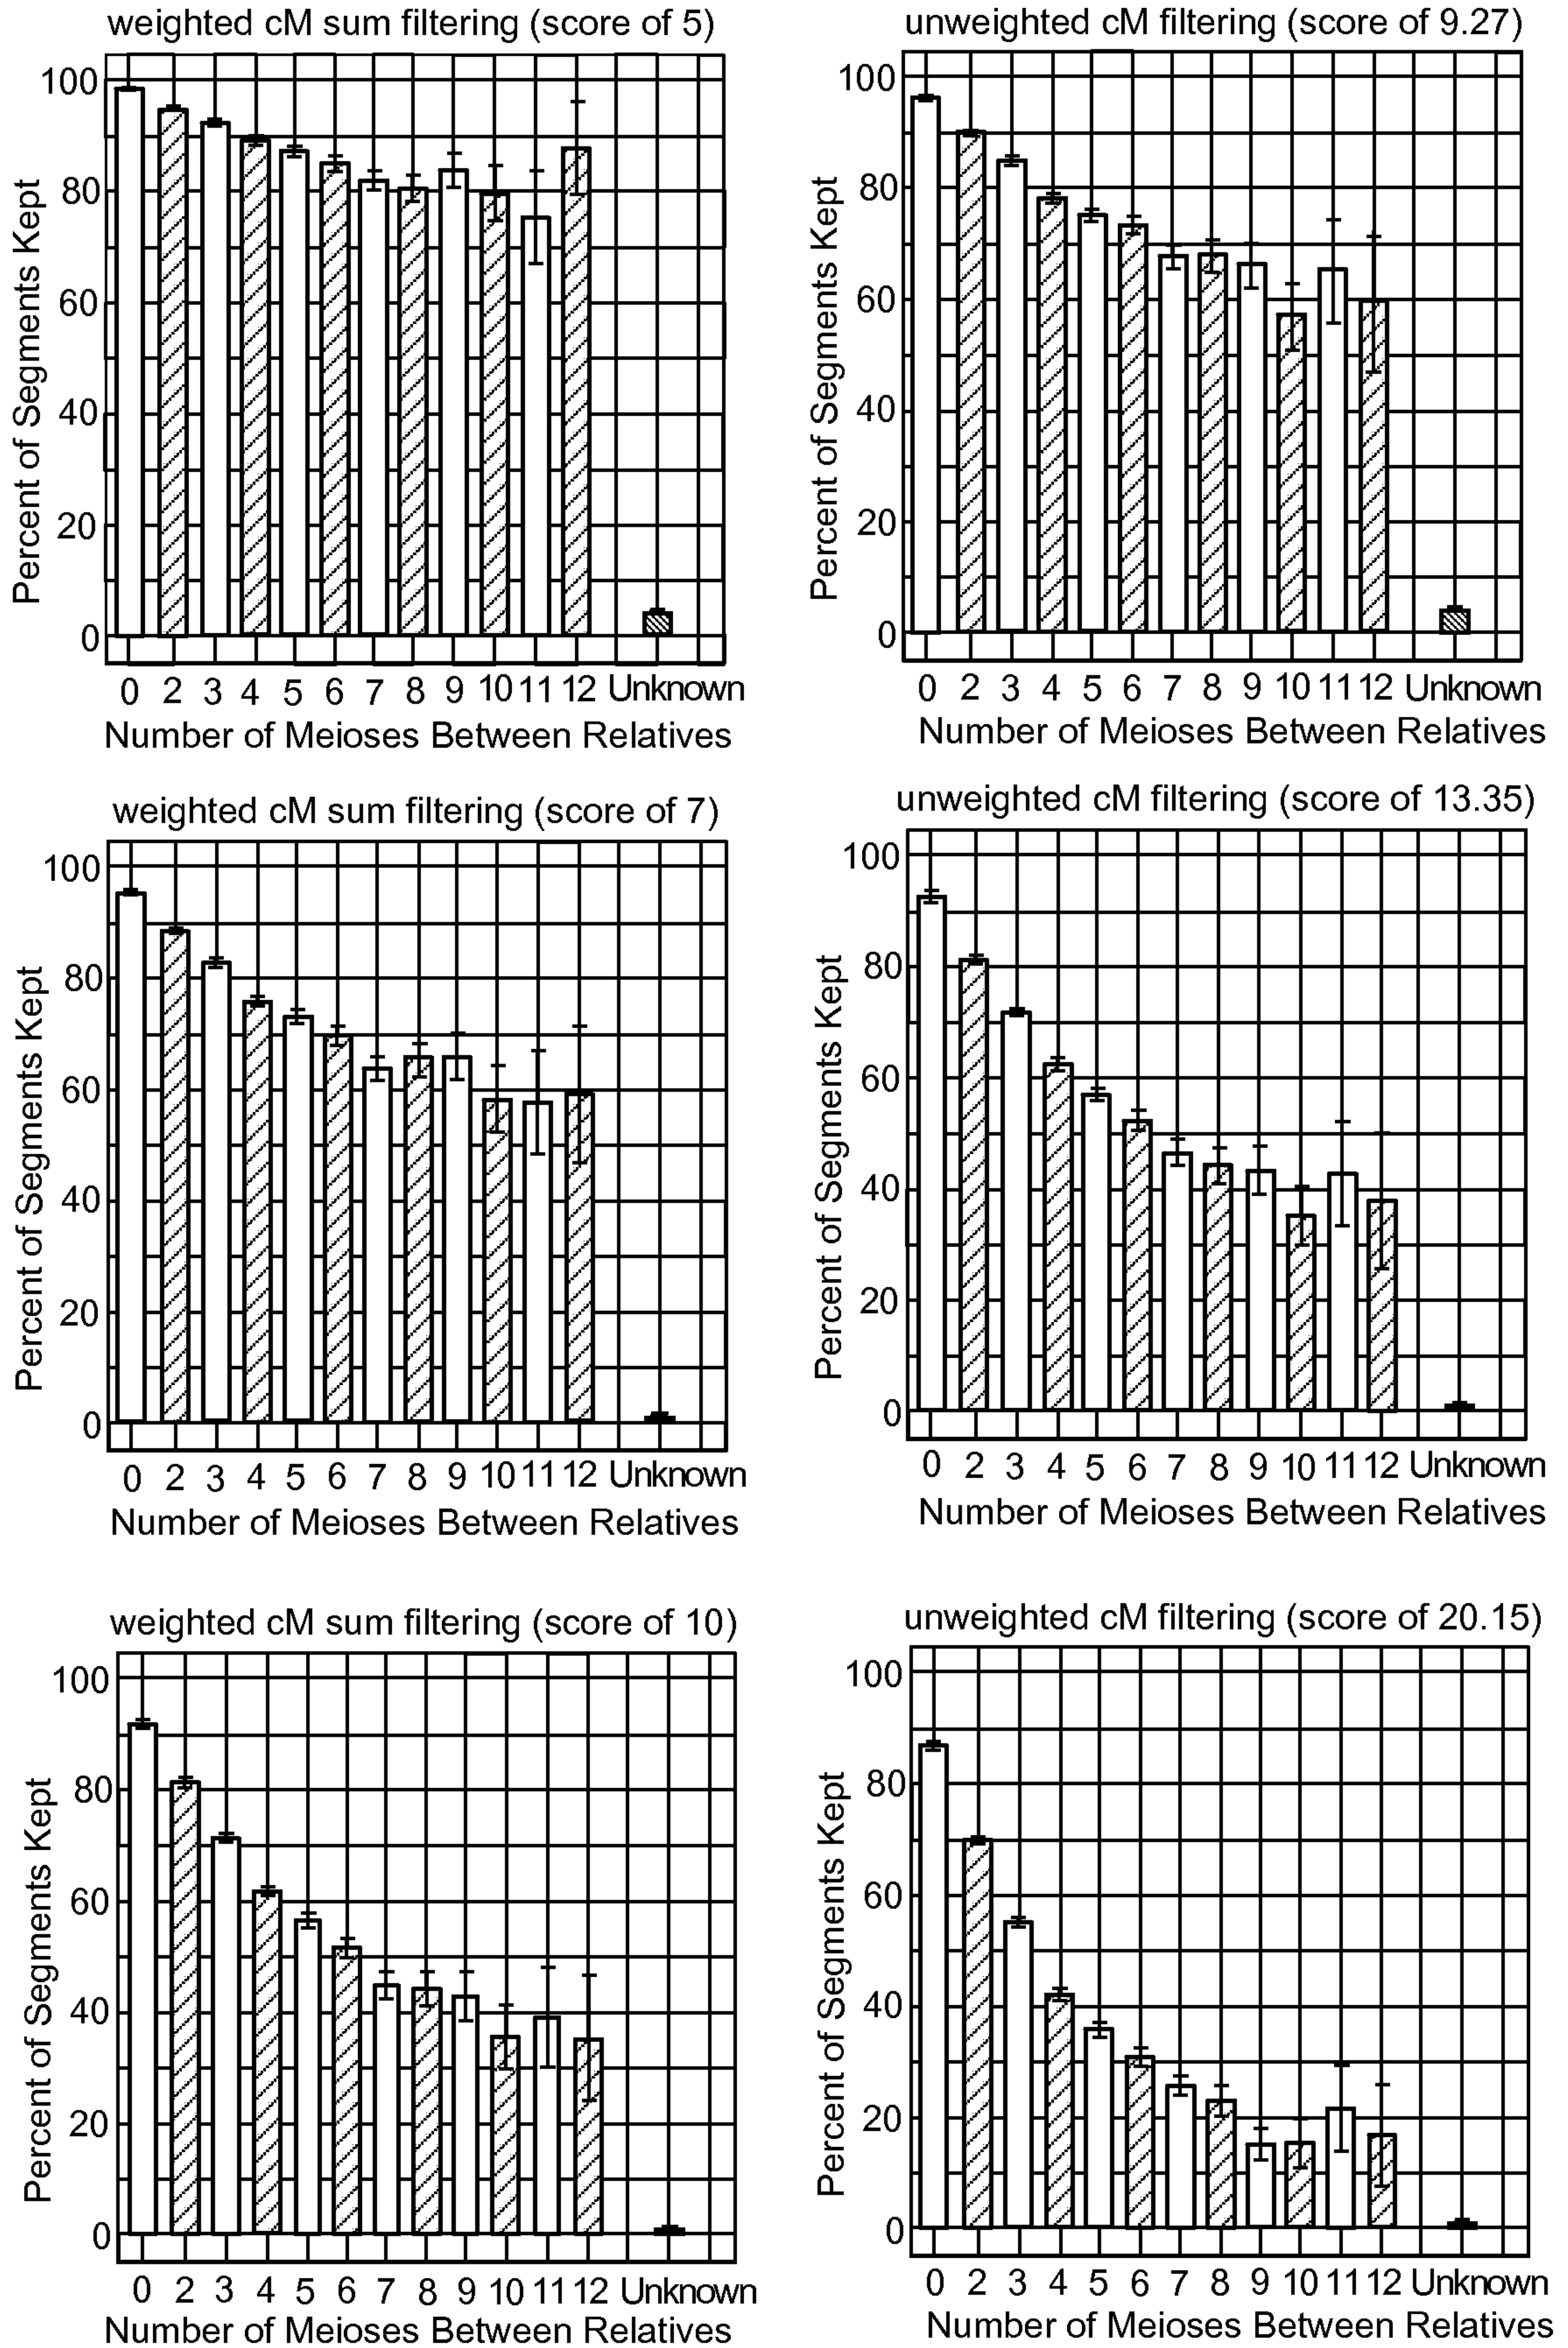


FIG. 10

100

Receive haplotype data for a population of individuals, the haplotype data including a plurality of genetic markers shared among the individuals

102

Divide the haplotype data into segment windows based on the genetic markers

104

For each individual in the population:

Based on the genetic markers, match segments of the haplotype data that are identical between the individual and any other individual in the population, each matched segment having a first cM width exceeding a threshold cM width and being part of one or more of the segment windows

106

Count the matched segments in each segment window

108

Estimate a weight associated with each segment window based on the count of matched segments in the associated segment window

110

Calculate a weighted sum of per-window cM widths for each matched segment based on the first cM width and the weights associated with the segment windows of the matched segment

112

Estimate a degree of ancestral relatedness between two individuals based on the weighted sum of per-window cM widths of each matched segment between the two individuals

114

FIG. 1A