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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
18 March 2021 (18.03.2021)

WIPO | PCT

(10) International Publication Number
WO 2021/051018 A1

(51) International Patent Classification:
G16B 20/40 (2019.01) *C12Q 1/68* (2018.01)
G16B 20/00 (2019.01) *C12N 15/10* (2006.01)

(21) International Application Number:
PCT/US2020/050582

(22) International Filing Date:
11 September 2020 (11.09.2020)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/900,373 13 September 2019 (13.09.2019) US
62/911,026 04 October 2019 (04.10.2019) US
62/936,267 15 November 2019 (15.11.2019) US

(71) Applicant: **23ANDME, INC.** [US/US]; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US).

(72) Inventors: **JEWETT, Ethan M.**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **SEAMAN, Andrew C.**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **MCMANUS, Kimberly F.**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **FREYMAN, William A.**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **BLAKKAN, Cordell T.**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **AUTON, Adam**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **MOUNTAIN, Joanna L.**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **FUREST, Susan M.**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **LOPATIN, Rachel E.**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **XU, Hang**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US). **VANCE, Hilary M.**; 223 N Mathilda Avenue, Sunnyvale, California 94086 (US).

(74) Agent: **ZHOU, Feng** et al.; Weaver Austin Villeneuve & Sampson LLP, P.O. Box 70250, Oakland, California 94612-0250 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN, KP, KR, KW, 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, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(54) Title: METHODS AND SYSTEMS FOR DETERMINING AND DISPLAYING PEDIGREES

(57) Abstract: The disclosed embodiments concern methods, apparatus, systems and computer program products for determining and displaying pedigrees based on IBD data. Some implementations use a probabilistic relationship model to obtain various likelihoods of various potential relationships based on pairwise IBD data and pairwise age data. Some implementations build large pedigrees by combining smaller pedigrees. Some implementations display pedigree graphs with various features that are informative and easy to understand.

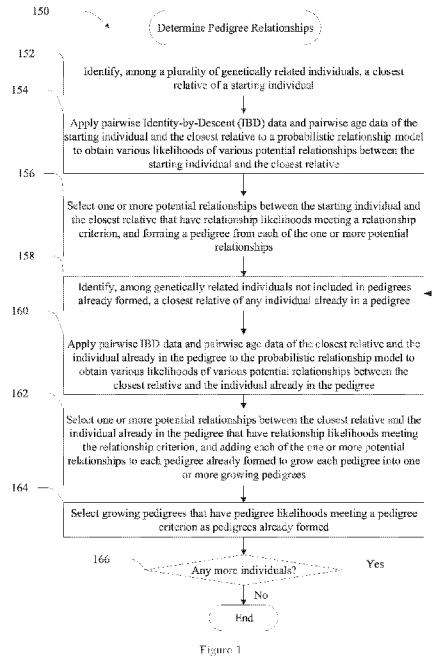

Figure 1

WO 2021/051018 A1

[Continued on next page]

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

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

METHODS AND SYSTEMS FOR DETERMINING AND DISPLAYING PEDIGREES

INCORPORATION BY REFERENCE

[0001] A PCT Request Form is filed concurrently with this specification as part of the present application. Each application that the present application claims benefit of or priority to as identified in the concurrently filed PCT Request Form is incorporated by reference herein in its entirety and for all purposes.

BACKGROUND

[0002] A pedigree refers to the genetic relationships among a group of genetically related individuals. Pedigrees can be used to produce family trees for consumers or genealogists. They can also be used to determine the heritability and genetic models for traits and disorders. Pedigree structure can be used to enable or improve genetic-analysis tools such as linkage, family-based association, pedigree-aware imputation, and pedigree-aware phasing.

[0003] However, there are many technical challenges in determining pedigrees using genetic data. Manually reconstructing an unknown pedigree with pairwise relationship comparisons requires arduous, error-prone labor. For example, Pemberton et al. manually reconstructed cryptic HapMap3 pedigrees, but the authors encountered inconsistencies they could not resolve by hand. Pemberton, et al. (2010). Inference of unexpected genetic relatedness among individuals in HapMap Phase III. Am. J. Hum. Genet. 87, 457–464. These problems become even more impractical or impossible to solve when the pedigrees are large and numerous.

[0004] Computer tools using identity-by-descent (IBD) genetic data to construct pedigrees have been developed to address some of these problems. However, the accuracy,

qualities, and efficiencies of available computer tools have many limitations. In various implementations, methods and systems disclosed herein for determining, constructing and visualizing pedigrees provide various advantages and improvements over conventional approaches.

SUMMARY

[0005] The disclosed implementations concern methods, apparatus, systems, and computer program products for determining and displaying pedigrees among genetically individuals based on IBD data

[0006] A first aspect of the disclosure provides computer-implemented methods for determining pedigree relationships among a plurality of genetically related individuals.

[0007] Another aspect of the disclosure provides systems for determining pedigree relationships among a plurality of genetically related individuals. In some implementations, the system involves a processor and one or more computer-readable storage media having stored thereon instructions for execution on said processor to determine pedigree relationships among a plurality of genetically related individuals.

[0008] Another aspect of the disclosure provides a computer program product including a non-transitory machine readable medium storing program code that, when executed by one or more processors of a computer system, causes the computer system to implement the methods above for determining pedigree relationships among a plurality of genetically related individuals.

[0009] Although the examples herein concern humans and the language is primarily directed to human concerns, the concepts described herein are applicable to genomes from any plant or animal. These and other objects and features of the present disclosure will become more fully apparent from the following description and appended claims, or may be learned by the practice of the disclosure as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] For a more complete understanding of the invention, reference is made to the following description and accompanying drawings, in which:

[0011] Figure 1 shows a flowchart illustrating a process for determining pedigree relationships using IBD data and age data according to some implementations;

[0012] Figure 2 shows parts 1-4 of a flow chart illustrating process 170 that can be used to combine smaller pedigrees into a larger one according to some implementations;

[0013] Figure 3 shows the use of closely-related individuals for identifying background IBD. Genotyped individuals are represented by large shaded squares and circles. Vertical solid lines indicate regions where genotyped individuals in the left pedigree share IBD with individuals in the right pedigree. A most recent common ancestor, G, has transmitted genetic material to both pedigrees, some of which is shared as IBD among the left and right pedigrees. Note that there can be either one or two most recent common ancestors, G, and that the observed segments are the union of all segments inherited from these ancestors. White dots on genotyped individuals indicate the set of genotyped individuals with no direct genotyped ancestors. White dots on ungenotyped individuals indicate the most recent common ancestors transmitting the segments to the genotyped individuals. Dashed lines in the pedigree indicate the induced tree whose branch lengths correspond to degrees of relationship among individuals. All observed IBD segments in the left pedigree were inherited through their common ancestor A_1 and all IBD segments in the right pedigree were inherited through their common ancestor A_2. The number of meioses separating A_1 and A_2 from a common ancestor, G, are $d_{A_1,G}$ and $d_{A_2,G}$. The arrow and horizontal bar indicate a specific position along the genome at which one considers the pattern of presence and absence of IBD;

[0014] Figure 4 illustrates the utility of considering IBD segments among groups of individuals rather than pairwise IBD. Genotyped individuals are represented by large

shaded squares and circles. Vertical solid lines indicate IBD segments shared between the genotyped descendants of A_1 and the genotyped descendants of A_2;

[0015] Figure 5 Determining the parental side of distant relatives. Individual 1 in the cyan pedigree shares segment A IBD with individuals 2 and 5 in the purple pedigree and they share segment IBD B with individuals 3 and 11 in the red pedigree. If the lineage connecting individual 1 to the purple pedigree passes through ancestor 8 and the lineage connecting individual 1 to the red pedigree passes through individual 9, then the ranges of segments A and B cannot overlap because individual 10 only transmits one recombined haplotype to individual 1. Observing abutting segments A and B is evidence that the cyan pedigree is connected to the purple and red pedigrees through the same parent. Observing spatially overlapping segments A and B is evidence that the purple and red pedigrees are connected through different parents of individual 1. In the absence of segment overlaps and splicing information, the orange dashed lines indicate equally reasonable ways to connect the purple and cyan pedigrees;

[0016] Figure 6 includes data comparing pedigree estimates between a likelihood method (Equation 21) according to some implementations and a novel generalized version of the existing DRUID method (Equation 24). The figure shows that both the likelihood method and the generalized DRUID method perform accurately. Pedigrees were simulated in which two smaller pedigrees were connected by a path of degree d between their respective common ancestors. The accuracy of the methods for inferring degree d is shown for four different tolerances: (A) exactly equal to the true degree, (B) within one degree of the true degree, (C) within two degrees of the true degree, and (D) within three degrees of the true degree;

[0017] Figure 7 shows a flowchart illustrating a process for training a probabilistic relationship model according to some implementations;

[0018] Figure 8 schematically illustrates an example of how to train the machine learning relationship model using multiple training sets;

[0019] Figure 9 shows a flowchart illustrating a process for generating pedigree graphs

according to some implementations;

[0020] Figure 10 illustrates three entries that can be used in the process for generating the pedigree graph;

[0021] Figure 11 shows a subtree that can be formed from a root node according to some implementations;

[0022] Figure 12 illustrates shifting one branch of a subtree;

[0023] Figure 13 shows a pedigree graph obtained by merging pedigree subtrees;

[0024] Figure 14 illustrates positioning of a row of four parent nodes according to some implementations;

[0025] Figure 15 shows an example of a pedigree graph in two views according to some implementations;

[0026] Figure 16 shows a close-up of a grandparent tree;

[0027] Figure 17 shows a functional diagram illustrating a programmed computer system for performing processes for determining or displaying pedigrees in accordance with some implementations;

[0028] Figure 18 shows a block diagram illustrating an implementation of an IBD-based pedigree determination and display system;

[0029] Figure 19 shows results comparing computer runtime between a method according to some implementation and an existing method;

[0030] Figure 20 shows accuracies of inferred relationships for pedigrees with 10 percent of individuals sampled according to some implementations and the existing method;

[0031] Figure 21 shows accuracies of inferred relationships for pedigrees with 20 percent of individuals sampled according to some implementations and the existing method;

[0032] Figure 22 shows accuracies of inferred relationships for pedigrees with 30 percent of individuals sampled according to some implementations and the existing method;

[0033] Figure 23 shows accuracies of inferred relationships for pedigrees with 40 percent of individuals sampled according to some implementations and the existing method;

[0034] Figure 24 shows accuracies of inferred relationships for pedigrees with 50 percent of individuals sampled according to some implementations and the existing method;

[0035] Figure 25 shows a flowchart illustrating a process for re-annotating pedigree graphs according to some implementations;

[0036] Figure 26 illustrates another process for re-annotating pedigree graph according to some implementations;

[0037] Figure 27 shows an example implementation of displaying a pedigree graph and receiving user input for annotating un-genotyped nodes of the pedigree graph;

[0038] Figure 28 shows an example implementation of displaying a pedigree graph and receiving user input for annotating un-genotyped nodes of the pedigree graph;

[0039] Figure 29 shows an example implementation of displaying a pedigree graph and receiving user input for annotating un-genotyped nodes of the pedigree graph;

[0040] Figure 30 shows an example implementation of displaying a pedigree graph and receiving user input for annotating un-genotyped nodes of the pedigree graph;

[0041] Figure 31 shows an example of a first pedigree graph and a second pedigree graph respectively labeled as Tree 1 and Tree 2, illustrating how ungenotyped nodes can be matched according to some implementations.

[0042] Figure 32 shows an example of a pedigree graph with health related information.

[0043] Figure 33 shows two GUIs activated from the pedigree graph with health related information.

DETAILED DESCRIPTION

[0044] The disclosure concerns methods, apparatus, systems, and computer program

products for determining pedigree relationships among a plurality of genetically related individuals. Various implementations operate on IBD data to perform the disclosed functions. IBD data may be provided in different formats or obtained by various methods. For example, U.S. Patent Application No.: 16/947,107, entitled: PHASE-AWARE DETERMINATION OF IDENTITY-BY-DESCENT DNA SEGMENTS, filed on July 17, 2020, which is incorporated by reference in its entirety, discloses suitable methods for determining IBD using genotype data.

[**0045**] Numeric ranges are inclusive of the numbers defining the range. It is intended that every maximum numerical limitation given throughout this specification includes every lower numerical limitation, as if such lower numerical limitations were expressly written herein. Every minimum numerical limitation given throughout this specification will include every higher numerical limitation, as if such higher numerical limitations were expressly written herein. Every numerical range given throughout this specification will include every narrower numerical range that falls within such broader numerical range, as if such narrower numerical ranges were all expressly written herein.

[**0046**] The headings provided herein are not intended to limit the disclosure.

[**0047**] Unless defined otherwise herein, all technical and scientific terms used herein have the same meaning as commonly understood by one of ordinary skill in the art. Various scientific dictionaries that include the terms included herein are well known and available to those in the art. Although any methods and materials similar or equivalent to those described herein find use in the practice or testing of the embodiments disclosed herein, some methods and materials are described.

[**0048**] The terms defined immediately below are more fully described by reference to the Specification as a whole. It is to be understood that this disclosure is not limited to the particular methodology, protocols, and reagents described, as these may vary, depending upon the context they are used by those of skill in the art.

Definitions

[0049] As used herein, the singular terms "a," "an," and "the" include the plural reference unless the context clearly indicates otherwise.

[0050] The term "plurality" refers to more than one element. For example, the term is used herein in reference to a number of nucleic acid molecules or sequence reads that is sufficient to identify significant differences in repeat expansions in test samples and control samples using the methods disclosed herein.

[0051] A DNA segment is identical by state (IBS) in two or more individuals if they have identical nucleotide sequences in this segment. An IBS segment is identical by descent (IBD) in two or more individuals if they have inherited it from a common ancestor without recombination, that is, the segment has the same ancestral origin in these individuals. DNA segments that are IBD are IBS per definition, but segments that are not IBD can still be IBS due to the same mutations in different individuals or recombinations that do not alter the segment.

[0052] The terms "nucleic acid" and "nucleic acid molecules" are used interchangeably and refer to a covalently linked sequence of nucleotides (i.e., ribonucleotides for RNA and deoxyribonucleotides for DNA) in which the 3' position of the pentose of one nucleotide is joined by a phosphodiester group to the 5' position of the pentose of the next. The nucleotides include sequences of any form of nucleic acid, including, but not limited to RNA and DNA molecules such as cell-free DNA (cfDNA) molecules.

[0053] The term "parameter" herein refers to a numerical value that characterizes a physical property. Frequently, a parameter numerically characterizes a quantitative data set and/or a numerical relationship between quantitative data sets. For example, the maximum degree of genetic distance between two genotyped individuals in a pedigree is a parameter of a genetic pedigree model.

[0054] The term "based on," when used in the context of obtaining a specific quantitative value, herein refers to using another quantity as input to calculate the specific quantitative value as an output.

[0055] As used herein the term "chromosome" refers to the heredity-bearing gene carrier

of a living cell, which is derived from chromatin strands comprising DNA and protein components (especially histones). The conventional internationally recognized individual human genome chromosome numbering system is employed herein.

Introduction and Overview

[**0056**] Some existing computer implemented methods use identity-by-descent (IBD) data to estimate pedigrees. One such method for determining a pedigree is PRIMUS. Staples et al. (2014), PRIMUS: Rapid Reconstruction of Pedigrees from Genome-wide Estimates of Identity by Descent, The American Journal of Human Genetics 95, 553–564. PRIMUS uses the total lengths of half and full IBD between pairs of individuals to obtain likelihoods of different relationship types. It then attempts to construct a pedigree for which the product of all pairwise likelihoods induced by the pedigree is greatest. PRIMUS does not use the count of IBD fragments for determining a pedigree and it uses age information only to resolve apparent discrepancies, such as an inferred grandparent being younger than a grandchild or an inferred nephew-uncle pair having an age difference greater than a specified threshold. In general, previous methods do not use age information as part of the likelihood of relationship in constructing pedigrees. In contrast, some implementations disclosed herein use the count of IBD segments and pairwise age differences directly in modeling the likelihoods of relationships. These aspects of the implementations improve the accuracy of relationship estimate.

[**0057**] PRIMUS uses a kernel density estimation to estimate IBD-length distributions. Kernel density estimation is a non-parametric technique that can result in over-fitting data. In contrast, some implementations disclosed herein model distributions of IBD data (*e.g.*, IBD length, number of IBD segments) and age difference data as parametric probability distributions. In various implementations, the probability distributions are modeled as Gaussian distributions, exponential distributions, or Poisson distributions. The parametric approach can provide more reliable estimates by avoiding overfitting the data.

[0058] PRIMUS computes likelihoods for only six general categories corresponding to different degrees of relationship:

- Parent-child
- Full-sibling
- Half-sibling, avuncular, grandparental
- First-cousin, great-grandparental, great-avuncular, half-avuncular
- Distantly related
- Unrelated

[0059] PRIMUS does not compute separate estimates for different relationships within each category (*e.g.*, half siblings, avuncular, and grandparental). This was because relationships within a category could be difficult to distinguish from one another based on genetic data alone. In contrast, methods disclosed herein estimate relationship likelihood for each specific relationship, including each relationship in a same category above. Methods disclosed herein use age data to determine likelihoods of relationships, which is especially helpful for distinguishing relationships with the same coefficient of relationship (*e.g.* half siblings vs. grandparents).

[0060] The coefficient of relationship is a measure of the degree of consanguinity (or biological relationship) between two individuals. With a simplifying assumption of non-consanguineous common ancestors, it can be calculated as:

$$r_{BC} = \sum_{p} e^{-L(p)},$$

[0061] where p enumerates all paths connecting B and C to unique common ancestors, and $L(p)$ is the length of a common-ancestor path p, which may be expressed in generations or meioses through the path. The coefficient of relationship sometimes is also referred to as "average fraction of DNA shared." Table 1 lists various relationships and corresponding coefficients of relationships.

Table 1: Coefficient of Relationship and Degree of Relationship for Various Relationships

Degree of Relationship	Relationship	Coefficient of Relationship (r)
0	identical twins; clones	100%(1)
1	parent-offspring	50% (2^{-1})
2	full siblings	50% ($2^{-2} + 2^{-2}$)
2	3/4 siblings or sibling-cousins	37.5% ($2^{-2} + 2^{-3}$)
2	grandparent-grandchild	25% (2^{-2})
2	half siblings	25% (2^{-2})
3	aunt/uncle-nephew/niece	25% ($2 \cdot 2^{-3}$)
4	double first cousins	25% ($4 \cdot 2^{-4}$)
3	great grandparent-great grandchild	12.5% (2^{-3})
4	first cousins	12.5% ($2 \cdot 2^{-4}$)
6	quadruple second cousins	12.5% ($8 \cdot 2^{-6}$)
6	triple second cousins	9.38% ($6 \cdot 2^{-6}$)
4	half-first cousins	6.25% (2^{-4})
5	first cousins once removed	6.25% ($2 \cdot 2^{-5}$)
6	double second cousins	6.25% ($4 \cdot 2^{-6}$)
6	second cousins	3.13% ($2 \cdot 2^{-6}$)
8	third cousins	0.78% ($2 \cdot 2^{-8}$)
10	fourth cousins	0.20% ($2 \cdot 2^{-10}$)

[**0062**] PRIMUS lumps all relationships having the same coefficient of relationship into one category, and relationships with coefficients smaller than 25% cannot be used to generate a pedigree. The disclosed methods herein provide pairwise relationship estimates at many levels beyond a 25% coefficient of relationship. In many applications, relationships up to 15th degree are estimated. This makes it possible to build very large pedigrees with many degrees of relationship.

[**0063**] Conventional methods for determining pedigrees using IBD data do not properly address noise caused by background IBD. For example, some individuals from Ashkenazi Jewish, Mexican, Puerto Rican, and other populations share IBD due to historical bottlenecks, rather than true recent relationships. Such shared IBD constitutes noise in the IBD data for determining recent relationships in pedigrees. Some implementations herein estimate the level of background IBD by computing the amount of IBD that each person in a group shares with him or herself between two chromosomes. Their IBD data can then be adjusted to remove the background IBD noise. This approach can help to improve the accuracy of pedigree estimates in various populations.

[0064] Some implementations described herein also statistically infer whether the IBD carried by an individual in a pedigree is due simply to background IBD. These approaches leverage previously inferred close relatives of such individuals to make these inferences. The methods then exclude such individuals from consideration when inferring degrees of relationship.

[0065] When adding a person to a pedigree, PRIMUS checks that the person's maximum likelihood estimated relationship with any person in the pedigree exceeds an initial threshold of 0.3, although this threshold can be adjusted downward over time if a pedigree fails to build properly on the first attempt. The methods disclosed herein do not have such a restriction. Being free of this restriction makes it possible to build larger pedigrees without having to progressively reduce this threshold, saving considerable computational time.

[0066] In the process of building pedigrees, PRIMUS does not distinguish among individuals within the same category of relationship. The PRIMUS method builds in stages, first combining all siblings and parent-child pairs, then second degree relatives (half-sibling, avuncular, and grandparent), then third degree relationships (first cousin, half-avuncular, great-avuncular, and great-grandparent). The goal is to combine high-confidence classes of relatives first, although confidence in an estimate can vary within a degree class. In contrast, methods disclosed herein first include a person that is most closely related to any individuals in the pedigree. By giving priority to the close relationship with the highest confidence, the disclosed methods can improve the accuracy of the pedigree estimate.

[0067] In adding a person to pedigree, PRIMUS adds the person in all possible relationships regardless of the likelihoods of the relationships. In contrast, the methods disclosed herein add a person to a pedigree in potential relationships that are highly likely. Moreover, the methods disclosed herein exclude low likelihood pedigrees in the process of building pedigrees. These likelihood-based techniques can greatly reduce the pedigree space that needs to be explored without significantly sacrificing accuracy. As a result, it can greatly improve computational speed and efficiency, and reduce memory

and CPU loads.

[0068] All previous pedigree inference methods, including PRIMUS, attempt to search the full pedigree space. The full search of all possible pedigrees quickly becomes computationally intractable when the number of individuals in the pedigree is moderate or large. This is so even using modern computers. In contrast the methods disclosed herein use a two stage approach. In the first stage, small pedigrees are inferred using approaches that thoroughly search the space of possible pedigrees. In the second stage, small pedigrees are combined into large pedigrees using heuristic methods that greatly reduce the number of pedigrees that must be searched. Without this heuristic second step, it is computationally intractable to build large pedigrees. The methods disclosed herein are the only known methods to use such heuristic approaches and are the only computer-implemented methods capable of building very large pedigrees.

Process for Determining a Pedigree

[0069] Figure 1 shows a flowchart illustrating a process 150 for determining pedigree relationships using IBD data and age data. The process is applicable for determining pedigrees for individual humans or animals, as well as to some other organisms having relevant genetic mechanisms that are analogous to humans and animals. Although persons, people, or humans are referred to in some descriptions, such descriptions are often applicable to other organisms with suitable modifications.

[0070] The illustrated process is implemented using a computer system that includes one or more processors and system memory. In many real-world applications, it is not practical or possible to implement these methods in a person's mind or using pen and paper. For pedigrees including a large number of individuals, many types or levels of relationship, or ambiguous data, the computation involved in the process would be too complex to be performed in the human mind.

[0071] The methods illustrated here apply IBD data and age data to a probabilistic relationship model to obtain likelihoods of many potential relationships. In each iteration of

adding an individual to a pedigree, the number of possible pedigrees grows exponentially. The computational task of determining the likelihood of a single pairwise relationship given IBD data and age data is time-consuming. This computation needs to be performed for tens of relationships in each iteration of growing pedigrees. As the pedigrees grow larger, the computation of a pedigree likelihood becomes impractical to perform by hand.

[**0072**] Process 150 starts by identifying, among the plurality of genetically related individuals, the closest relative of a starting individual. See block 152. In some applications, the starting individual is an individual of interest, such as a consumer who wants to obtain a pedigree or pedigree graph with herself as a focal person. In some implementations, the starting individual is an individual meeting certain genetic relationship criteria, such as a person who has a high average degree of relationship with other individuals being considered. Various methods may be used to determine how closely two individuals are related or the relationship distance between them. For example, IBD data may be used to calculate a coefficient of relationship as explained above.

[**0073**] In various implementations, the plurality of genetically related individuals includes at least 20, 50, 100, 200, 300, 400, or 500 individuals. In some implementations the pedigree can include both genetically related individuals that have been genotyped and those individuals where genotype data is unknown or not available. In some implementations, every pair of individuals in the plurality of genetically related individuals has a total IBD length larger than an IBD threshold. In various implementations, the IBD threshold is 1 centimorgan (cM), 2 cM, 3 cM, 4 cM, 5 cM, 6 cM, 7 cM, 8 cM, 9 cM, 10 cM, 15 cM, 20 cM, 25 cM, 50 cM, 75 cM, 100 cM, 200 cM, or 500 cM. In some implementations, the total IBD length is adjusted by subtracting background IBD from the pairwise IBD data.

[**0074**] Various methods may be used to determine background IBD for a group of individuals or a population of individuals. In some implementations two chromosomes in each pair of one or more pairs of the 22 pairs of somatic chromosomes of the same individual can be compared to identify IBD regions. Two corresponding fragments on a pair of chromosomes are respectively inherited from two parents. Assuming that an

individual's two parents are not more consanguineous than unrelated individuals in the population, the IBD amount between the two chromosomes of a pair in the individual provides a good estimate of population background IBD.

[**0075**] In some implementations, the level of background IBD can be inferred by estimating IBD between pairs of individuals assumed to be non-consanguineous.

[**0076**] In some implementations, IBD lengths are adjusted for the background IBD before being used to model or determine the relationship likelihood or pedigree likelihood. In some implementations, IBD lengths are adjusted before being compared to an IBD threshold to determine whether individuals should be included for consideration in a pedigree. In other implementations, pairs of individuals whose IBD sharing levels are inferred to be significantly lower than expected by chance are removed from consideration.

[**0077**] When selecting a next individual to be added to a pedigree, the process considers how closely individuals already included in the pedigrees are related to individuals not yet included. In some implementations, pairwise IBD data between two individuals are used to determine how closely related the two individuals are, or the relationship distance between the two individuals. In some implementations, the relatedness or relationship distance between individuals may be inferred from IBD data using a likelihood expression for the degree of relationship derived using a probabilistic recombination model. Other genetic information and methods may also be used to determine relatedness or relationship distance. In some implementations, relatedness or relationship distance may be measured by meioses on a common ancestor path. In some implementations, relatedness or relationship distance may be expressed as or measured by coefficient of relationship.

[**0078**] Process 150 proceeds to apply pairwise identity by descent (IBD) data and pairwise age data of the starting individual and the closest relative to the probabilistic relationship model to obtain various likelihoods of various potential relationships between the starting individual and the closest relative. In various implementations, the pairwise age data reflect the age difference between two individuals. In some implementations, the pairwise age data are obtained by simple subtraction. In other implementations, other operations may be performed on ages of two individuals, such as division (*e.g.*, to obtain

a ratio of two ages) or normalization (*e.g.*, to obtain a z-score).

[**0079**] It is also possible to extrapolate from empirical distributions of age differences between different types of relatives to obtain distributions for relationships that are unobserved empirically. In particular, pairs of relatives sharing third great-grandparental relationships (5 generations) may be unobserved, and therefore, it is not possible to obtain the distribution of the age difference of a third great-grandparental pair empirically. However, the age difference distribution for third great-grandparents can be estimated by computing the mean (μ_{PC}) and variance σ^2_{PC} of the age differences among observed parent-child pairs. Then, noting that a third great-grandparental relationship is a string of five statistically independent parent-child relationships, we find that the mean and variance of the age difference distribution for third great-grandparental relationships are $\mu_{5GGP} = 5\mu_{PC}$ and $\sigma^2_{5GGP} = 5\sigma^2_{PC}$, respectively. This result is obtained by using properties of the means and variances of sums of independently distributed random variables.

[**0080**] Given the IBD data and the pairwise age data of the two individuals, the probabilities or likelihoods of different relationships between the two individuals can be determined using the probabilistic relationship model. Various probabilistic relationship models are further described herein after. See block 154. In some implementations, the pairwise IBD data include the lengths of IBD segments, such as the total or summed length of the IBD segments. In some implementations, the lengths of IBD segments include the length of full IBD segments (IBD2) and/or length of half IBD segments (IBD1). In some implementations, the two types of IBD lengths may be combined. In other implementations, the two IBD segment lengths are kept separate and are modeled by the probabilistic relationship model to have different probability distributions. In some implementations the lengths of half IBD segments (IBD1) are summed and the sum is used to compute the likelihood. Similarly, in some implementations the lengths of full IBD segments (IBD2) are summed and the sum is used to compute the likelihood. In some implementations, the pairwise IBD data also include numbers or counts of IBD segments. Similar to lengths of the two types of IBD segments, the numbers of the two types of IBD segments may be combined or modeled separately.

[0081] Given the IBD data and the pairwise age data of the two individuals, the probabilities or likelihoods of different relationships between the two individuals can be determined using the probabilistic relationship model. Various probabilistic relationship models are further described herein after. See block 154. In some implementations, the pairwise IBD data include the lengths of IBD segments, such as the total or summed length of the IBD segments. In some implementations, the lengths of IBD segments include the length of full IBD segments (IBD2) and/or length of half IBD segments (IBD1). In some implementations, the two types of IBD lengths may be combined. In other implementations, the two IBD segment lengths are kept separate and are modeled by the probabilistic relationship model to have different probability distributions. In some implementations the lengths of half IBD segments (IBD1) are summed and the sum is used to compute the likelihood. Similarly, in some implementations the lengths of full IBD segments (IBD2) are summed and the sum is used to compute the likelihood. In some implementations, the pairwise IBD data also include numbers or counts of IBD segments. Similar to lengths of the two types of IBD segments, the numbers of the two types of IBD segments may be combined or modeled separately.

[0082] In some implementations, the probabilistic relationship model is a machine learning model obtained by training the model using training data to determine a plurality of parameters of the model, including parameters of probability distributions for various independent/input variables and various relationships. In some implementations, the probabilistic relationship model models the probability distribution of the pairwise IBD as a Gaussian distribution, a Poisson distribution, an exponential distribution, a binomial distribution, a beta binomial distribution, or other distributions suitably determined from prior information. In some implementations, the probabilistic relationship model also models the probability distribution of the pairwise age data for each relationship using one or more of said forms of distributions.

[0083] In some implementations, the various potential relationships include more than 10, 20, 30, 40, or 50 different relationships. In various implementations, the various relationships include relationships of the 0th, $1^{st}$and 2nd, 3rd, 4th, 5th, 6th, 7th, 8th, 9th, 10th, 11th, 12th, 13th, 14th, or 15^{th}degree or further. In some implementations, the various

relationships include relationships of at least 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16 or more meioses on a common-ancestor path between the two individuals through a common ancestor. In some implementations, the various relationships include two or more different relationships of the same degree or of the same coefficient of relationship (*e.g.*, half sibling, grandparent, avuncular have a coefficient of relationship of 0.25). So in some implementations, the various relationships include these three relationships as different relationships instead of as a same category of relationship.

[**0084**] Process 150 involves selecting the one or more potential relationships between the starting individual and the closest relative that have relationship likelihoods meeting a relationship criterion, and forming a pedigree from each of the one or more potential relationships. See block 156. In various implementations, different relationship criteria may be used. For example, the relationship criterion may be determined by likelihood ranks or percentile. There may simply be a number of the most likely relationships, *e.g.*, the top 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, or 15 most likely relationships. In other implementations, the relationship criterion is based on a ratio of the candidate relationship likelihood over the maximum relationship likelihood. In some implementations, the ratio is a log likelihood ratio, and the criterion is for the ratio to be larger than a threshold c. In general, the larger the parameter c, the fewer potential relationships are included. By reducing the potential relationships to be used to construct different pedigrees, the process can reduce the number of relationships to be processed. This can increase computational speed and reduce computational load.

[**0085**] Process 150 proceeds to identify, among genetically related individuals not yet included in pedigrees already formed, a closest relative of any individual in the formed pedigrees. See block 158.

[**0086**] Process 150 further involves applying pairwise IBD data and pairwise age data of the closest relative and the individual already in the formed pedigrees to the probabilistic relationship model to obtain various likelihoods of various potential relationships between the closest relative and the individual already in the pedigrees. See block 160.

[**0087**] Process 150 then proceeds to select one or more potential relationships between

the closest relative and the individual already in the pedigrees that have relationship likelihoods meeting the relationship criterion. The process also adds each of the one or more potential relationships with the individual already in the pedigrees to grow each pedigree into one or more growing pedigrees. See block 162.

[0088] Process 150 further involves selecting growing pedigrees that have pedigree likelihoods meeting a pedigree criterion. In some implementations, a pedigree likelihood can be obtained by aggregating the likelihood of all the relationships in a pedigree, such as summing the log likelihoods of the pairwise relationships in a pedigree. In some implementations, the pedigree criterion is met when a ratio of the candidate pedigree likelihood over a maximum pedigree likelihood is larger than or equal to a threshold value d. In various implementations, d = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, or 1.0. In other implementations, d = 1/100,000, 1/500,000, 1/1,000,000, 1/2,000,000, 1/4,000,000, or the like. In some implementations, the pedigree criterion may also be determined by pedigree likelihood ranks or percentile. Similar to the parameter c above, as d gets larger, fewer pedigrees are included for pedigree building. By increasing the value of d, one can increase computational speed and reduce CPU or memory load for exploring potential pedigrees.

[0089] Process 150 then decides whether there are more individuals to be considered for adding to the pedigrees. See block 166. If so, the process loops back to block 158 to identify another closest relative of any individual already in the pedigree. In some implementations, the process continues the loop until all individuals of the plurality of genetically related individuals have been identified as a closest relative or excluded from the pedigrees according to particular exclusion criteria.

[0090] Figure 1 shows that process 150 ends when no more individuals need to be considered. But in some applications, a number of pedigrees having high likelihoods are selected for further downstream processing. In some implementations, the pedigree having the highest likelihood is stored in memory. In some implementations, the data of the pedigree having the highest pedigree likelihood are retrieved and used to generate a pedigree graph, such as those described hereinafter. The pedigree graph then can be displayed on a display device.

Combining Smaller Pedigrees into Larger Pedigrees

[0091] Pedigrees created using process 150 can be combined into even larger pedigrees. Figure 2 shows a flow chart in four parts illustrating process 170 that can be used to combine smaller pedigrees into a larger one according to some implementations. Combining pedigrees has two primary uses. First, combining multiple pedigrees together makes it possible to create very large pedigrees connected by common ancestors many generations in the past. For such large pedigrees, it is not possible to add a single individual at a time to grow the pedigree as is done in process 150 and the PRIMUS method. This is because the amount of IBD shared between two distantly-related people decreases quickly in their degree of relationship. For example, relatives with more than 10 degrees of separation have a high probability of sharing no detectable IBD segments, especially if segments with lengths similar to background IBD are ignored or removed from the analysis.

[0092] Combining smaller pedigrees into larger pedigrees makes it possible to leverage all IBD segments observed between the smaller pedigrees when inferring the degree of relationship between individuals. This is because it is more likely that some individual in a small pedigree shares IBD with some individual in a related pedigree, even if not all cross-pedigree pairs of individuals share IBD. Figure 2 show four parts of process 170 of using IBD segments in close relatives to improve the inference of relationship degrees.

[0093] Methods that combine pedigrees are also computationally much faster than methods that add one individual at a time. The reason for this is that, the amount of IBD shared between a single unplaced individual and a set of genotyped individuals in a pedigree is often consistent with many possible relationships. Consequently, it is necessary to consider many ways of placing the new individual, which is computationally slow. In contrast, when combining two pedigrees, the IBD shared between individuals in the first pedigree and individuals in the second pedigree is larger than the amount shared with any single individual, providing additional information about the way in which the two pedigrees are related. As a result, there are fewer highly likely ways in which the pedigrees can be related, reducing the number of combinations that must be explored

and considerably increasing the speed of computation. Combining pedigrees makes it computationally possible to infer pedigrees that are much larger than pedigrees that are computationally tractable for PRIMUS or methods that must search many possible pedigree configurations.

[0094] Another reason to combine pedigrees is that background IBD can be detected more effectively. When comparing a single pair of individuals, it is difficult to detect whether the IBD they share is due to background or to a recent relationship. However, by examining the amount of IBD shared among all genotyped or sequenced close relatives of a pair of individuals, it becomes easier to determine when observed IBD is background IBD. Figure 3 shows the use of closely-related individuals for identifying background IBD.

[0095] Referring to Figure 2, process 170 for combining smaller pedigrees into larger pedigrees begins by first considering a set of pedigrees inferred using various methods such as process 150 shown in Figure 1. In some implementations the set of pedigrees can be obtained from another appropriate source besides process 150. Process 170 then proceeds by computing the total amount of IBD shared between each pair of pedigrees. The total amount of IBD between two pedigrees, P_1 and P_2, is found by merging the IBD segments between all pairs of individuals i_1 and i_2 such that i_1 is in P_1 and i_2 is in P_2.

[0096] Process 170 proceeds by identifying the two smaller pedigrees, P_1 and P_2, that share the greatest amount of IBD. See Box 1. These will be the next pair of pedigrees that will be combined. To combine pedigrees P_1 and P_2, the set S_1 of individuals in P_1 who share IBD with individuals in P_2 are then identified. Conversely, the set S_2 individuals in P_2 who share IBD with individuals in P_1 are identified. See box 2. Process 170 then proceeds by identifying a common ancestor A_1 of the set S_1 (box 3 and box 4) and a common ancestor A_2 of the set S_2 (box 5 and box 6). These common ancestors are identified using the small pedigree structures that were previously inferred using methods such as shown in process 150.

[0097] The degree of relationship between common ancestors A_1 and A_2 is then inferred. See box 7. In some implementations of the method, the degree of relationship between A_1 and A_2 is inferred by considering a degree of relationship between A_1 and A_2 and attaching

A_1 and A_2 by a chain of dummy nodes reflecting this degree to create a combined pedigree P comprising P_1, P_2, and the newly-formed chain of dummy nodes. The log likelihood of this pedigree can then be computed as the sum over all pairwise log likelihoods among genotyped individuals in P. Process 170 considers many different possible degrees between A_1 and A_2 and forms a new pedigree P for each degree. The degree between A_1 and A_2 is then inferred as the degree that yields the pedigree P with the highest sum of pairwise log likelihoods.

[**0098**] In other implementations, the degree of relationship between A_1 and A_2 is inferred using a version of the DRUID estimator (M.D. Ramstetter, S.A. Shenoy, T.D. Dyer, D.M. Lehman, J.E. Curran, R. Duggirala, J. Blangero, J.G. Mezey, and A.L. Williams. Inferring identical-by-descent sharing of sample ancestors promotes high-resolution relative detection. Am. J. Hum. Genet., 103:30–44, 2018) that we generalize to the case of pedigrees with arbitrary outbred topologies. This generalized estimator of the degree between A_1 and A_2 is discussed in Section Distant Relatives Likelihood.

[**0099**] In some implementations, process 170 involves identifying individuals in the sets S_1 and S_2 whose observed IBD is likely due to background IBD. The IBD observed in these individuals can lead to biased estimates of the degree of relatedness between ancestors A_1 and A_2 and, more importantly, it can lead to the incorrect identification of A_1 and A_2, themselves.

[**0100**] The way in which background IBD can contribute to the mis-identification of A_1 and A_2 is shown in Figure 3.

[**0101**] In Figure 3, individuals 1 and 2 share a small amount of IBD with individuals 5 and 6. Because individuals 1 and 2 have approximately the same degree of relationship to 5 and 6 as individuals 3 and 4, it appears that the amount of IBD that 1 and 2 share with 5 and 6 is much lower than that expected by chance. Thus, one can ignore the IBD in 1 and 2 when connecting pedigrees P_1 and P_2.

[**0102**] Ignoring the IBD in individuals 1 and 2 will not only lead to a different inferred degree between P_1 and P_2, it also affects the choice of A_1. In particular, if individuals 1

and 2 are unrelated to individuals 5 and 6, then the correct common ancestor in pedigree P_1 to whom one will connect A_2 is individual 8. In some implementations, process cycles over nodes descended from A_1 and A_2 and identifies nodes whose descendants share significantly less IBD than expected, conditional on the current estimate of the pedigrees P_1 and P_2, the choice of ancestors A_1 and A_2, and the degree of relationship between A_1 and A_2. See box 8. In other implementations, the processing of these nodes is optional.

[**0103**] To determine whether the amount of observed IBD in the descendants of node n_1 below A_1 is statistically significantly lower than that expected by chance, some implementations consider the set N_1 of genotyped descendants of n_1 and compute the total merged amount of IBD shared between individuals in N_1 and all nodes in S_2. The process then uses the likelihood described hereinafter in the Distant Relatives Likelihood section to evaluate whether the total merged length of IBD is significantly lower than that expected by chance. If the amount of IBD is lower than expected, the descendants N_1 of node n_1 are removed from S_1. The aforementioned approach is then used to identify and remove nodes below of A_2 whose descendants have an amount of IBD that is significantly lower than that expected by chance. Some implementations cycle through nodes descended from A_1 and A_2 by repeat consideration of shared IBD of the nodes described above until the amount of IBD in all the remaining descendant nodes of A_1 and A_2 is not significantly different from that expected by chance.

[**0104**] In some implementations, process 170 identifies a new pair of common ancestors A_1' and A_2' of the reduced sets S_1' and S_2', where S_1' consists of the individuals in S_1 who remain after removing individuals whose IBD is inferred to be due to background IBD. Similarly, S_2' consists of the individuals in S_2 who remain after removing individuals whose IBD is inferred to be due to background IBD. See box 9. In some implementations, the operation shown in box 9 is optional, and downstream processes are performed on A_1 and A_2. Process 170 then computes the degree of relatedness between A_1' and A_2' using the likelihood described in the Distant Relatives Likelihood section. See box 10. The process then attaches A_1' to A_2' by a string of dummy ancestral nodes. This step yields a new pedigree P comprised of small pedigrees P_1 and P_2, and the dummy nodes connecting A_1' and A_2'. See box 11.

[0105] Process 170 can repeat operations covered by box 1 through box 11 and interim boxes until all small pedigrees have been combined into a single pedigree, or until no two small pedigrees share an amount of merged IBD greater than a predetermined threshold.

Distant Relatives Likelihood

[0106] Likelihoods computed among pairs of individuals provide high accuracy for inferring the degree of relatedness when the degree is relatively small. However, the amount of IBD shared between two individuals decreases exponentially in their degree of relatedness, resulting in very little information for inferring degrees between distant relatives. In fact, there can be a sizable probability that distant relatives will share no IBD segments at all, especially if IBD segments below a threshold are discarded to reduce the rate of observation of false positives.

[0107] When inferring the degree of relatedness between two distant relatives, it is helpful to leverage information from IBD segments shared among close relatives of these two individuals. Figure 4 illustrates the utility of considering IBD segments among groups of individuals rather than pairwise IBD when the degree of relatedness is not small. In particular, individuals 3 and 4 in Figure 4 share no IBD segments. Thus, one cannot infer their degree of relatedness without prior knowledge. However, if close relatives of 3 and close relatives of 4 have been genotyped and local pedigree structures have been previously inferred, one can use the IBD in close relatives of 3 and close relatives of 4 to estimate their degree of relationship.

[0108] Practitioners have developed a likelihood estimator of the pairwise degree of relatedness between the common ancestors A_1 and A_2 of two sets of genotyped individuals. To do this, practitioners derive the probability of the observed pattern of IBD shared among descendants of A_1 and A_2, given the degree $d = d_{A_1,G} + d_{A_2,G}$ separating A_1 and A_2 from their most recent common ancestor, G. Note that there can be more than one most recent common ancestor, G. There are two such individuals, G, if A_1 and A_2 are descended from a single ancestral couple and there is one most recent common ancestors

if A_1 and A_2 are descended from a pair of half siblings.

[0109] Figure 4 shows an IBD segment shared among genotyped individuals in two small pedigrees. If one considers a fixed position on the genome, the segment will be present in some individuals at the position and absent in others. The probability of the observed presence and absence pattern can be computed recursively by conditioning on whether the segment was observed in the ancestor of each individual using an approach similar to that of Felsenstein's (1973) tree pruning algorithm. Felsenstein, J. (1981). Evolutionary trees from DNA sequences: A maximum likelihood approach. J. Mol. Evol.17, 368–376.

[0110] Consider one of ancestor G's two alleles at a single locus and let O_i be a random variable describing the event that a copy of the allele is transmitted to descendant i and is observed. One sets $O_i = 1$ if the allele is observed in individual i and $O_i = 0$ if it is not observed. The probability $Pr(O_i = 1)$ can be computed by conditioning on whether G's allele was observed in a recent ancestor of individual i.

[0111] Consider the tree relating a set of genotyped individuals with no genotyped direct ancestors and their respective most recent common ancestors (dashed orange lines and red dots in Figure 4). Let $a(i)$ denote the parent node of node i in this tree. For example, in the tree in Figure 4, $a(1) = A_1$, $a(6) = A_1$, $a(2) = 6$, $a(3) = 6$, $a(4) = A_2$, $a(5) = A_2$, $a(A_1) = G$ and $a(A_2) = G$. It follows:

$$\begin{aligned} Pr(O_i = 1) &= Pr(O_i = 1|O_{a(i)} = 1)Pr(O_{a(i)} = 1) + Pr(O_i = 1|O_{a(i)} = 0)Pr(O_{a(i)} = 0) \\ &= Pr(O_i = 1|O_{a(i)} = 1)Pr(O_{a(i)} = 1) \\ &= 2^{-d_{i,a(i)}}Pr(O_{a(i)} = 1), \end{aligned} \tag{1}$$

and

$$\begin{aligned} Pr(O_i = 0) &= Pr(O_i = 0|O_{a(i)} = 1)Pr(O_{a(i)} = 1) + Pr(O_i = 0|O_{a(i)} = 0)Pr(O_{a(i)} = 0) \\ &= [1 - 2^{-d_{i,a(i)}}]Pr(O_{a(i)} = 1) + Pr(O_{a(i)} = 0), \end{aligned} \tag{2}$$

where $d_{i,a(i)}$ is the number of meioses separating individual i from their ancestor $a(i)$. In the final lines of Equations (1) and (2), one has used the fact that the probability that an allelic copy is transmitted in one meiosis is 1/2.

[0112] Equations (1) and (2) establish a recursion for computing the probability of an observed presence and absence pattern from a given ancestral allelic copy at a single base of the genome. Defining

$$p_{i,0} \equiv Pr(O_i = 0), p_{i,1} \equiv Pr(O_i = 1), \tag{3}$$

one can express the recursion compactly as

$$\begin{aligned} p_{i,0} &= [1 - 2^{-d_{i,a(i)}}]p_{a(i),1} + p_{a(i),0}, \\ p_{i,1} &= 2^{-d_{i,a(i)}}p_{a(i),1}, \end{aligned} \tag{4}$$

with the base conditions $p_{g,0} = 0$ and $p_{g,1} = 1$ for $g \in G$. The probability of an observed IBD sharing pattern $\{O_1, ..., O_k\}$ across k leaf nodes can then be computed recursively using Equation (4).

[0113] Equation (4) allows one to compute the expected total length $T_{1,2}$ of the genome that is covered by an IBD segment between some descendant of A_1 and some descendant of A_2. In other words, this is the length of IBD one would obtain if one merged all observed IBD segments between descendants of A_1 and A_2. The expected fraction of the genome that is observed IBD between descendants of A_1 and A_2 is given by the probability that an ancestral allele copy at a locus in G is passed down to at least one descendant of A_1 and at least one descendant of A_2.

[0114] Let $\mathcal{N}_1$ be a set of nodes descended from A_1 and let $\mathcal{N}_2$ be a set of nodes descended from A_2. In some implementations, $\mathcal{N}_1$ and $\mathcal{N}_2$ are the sets of genotyped nodes below A_1 and A_2. Let D_1 denote the event that a copy of the allele from G is observed in at least one descendant in $\mathcal{N}_1$. Then, given that the ancestral allele copy was passed to A_1, the probability of the event D_1^c that no copy was passed to any node in $\mathcal{N}_1$ is

$$Pr(D_1^c|O_{A_1} = 1) = Pr(O_n = 0 \text{ for } n \in \mathcal{N}_1|O_{A_1} = 1), \tag{5}$$

which can be computed using the recursion in Equation (4) with the base conditions $p_{A_1,0} = 0$ and $p_{A_1,1} = 1$. The equivalent probability $Pr(D_2^c|O_{A_2} = 1)$ that a particular allelic copy from G is not observed in any node in $\mathcal{N}_2$, given that A_2 inherited the copy is computed in the same way.

[0115] The probability that the allelic copy was observed in some member of $\mathcal{N}_1$ and in some member of $\mathcal{N}_2$ is then

$$\begin{aligned} Pr(D_1, D_2) &= Pr(D_1, D_2 | O_{A_1} = 1, O_{A_2} = 1) Pr(O_{A_1} = 1) Pr(O_{A_2} = 1) \\ &= Pr(D_1 | O_{A_1} = 1) Pr(D_2 | O_{A_1} = 1) Pr(O_{A_1} = 1) Pr(O_{A_2} = 1) \\ &= [1 - Pr(D_1^c | O_{A_1} = 1)][1 - Pr(D_2^c | O_{A_2} = 1)] Pr(O_{A_1} = 1) Pr(O_{A_2} = 1) \\ &= [1 - Pr(D_1^c | O_{A_1} = 1)][1 - Pr(D_2^c | O_{A_2} = 1)] 2^{-(d_{A_1,G} + d_{A_2,G})}, \end{aligned} \tag{6}$$

where $Pr(D_1^c | O_{A_1} = 1)$ and $Pr(D_2^c | O_{A_2} = 1)$ are computed using the recursion (4) and Equation (5).

[0116] If A_1 and A_2 had exactly one common ancestor with one allele to transmit, then Equation (6) would be the fraction of the genome in which we expect to find some IBD segment shared between some member of $\mathcal{N}_1$ and some member of $\mathcal{N}_2$. However, we must now account for the fact that each common ancestor of A_1 and A_2 in G carries two allelic copies and that there can be either one or two such common ancestors.

[0117] Let $|G|$ denote the number of common ancestors of A_1 and A_2, each of which carries two alleles at the locus of interest. The probability that a specific one of these $2|G|$ alleles is not observed IBD between the descendants of A_1 and A_2 is $1 - Pr(D_1, D_2)$ and the probability that none of them results in an observed IBD segment is $[1 - Pr(D_1, D_2)]^{2|G|}$. Therefore, the probability $Pr(\mathcal{I}_{1,2})$ that one of the $2|G|$ ancestral alleles results in an observed IBD segment between some descendant of A_1 and some descendant of A_2 is

$$Pr(\mathcal{I}_{1,2}) = 1 - [1 - Pr(D_1, D_2)]^{2|G|}. \tag{7}$$

[0118] One can use the probability $Pr(\mathcal{I}_{1,2})$ to obtain an approximate likelihood of the total length $T_{1,2}$ of IBD observed between descendants of A_1 and A_2. The mean of this distribution is simply the expected length of the genome in a state of IBD between the two clades, which is

$$E[T_{1,2}] = Pr(\mathcal{I}_{1,2}) L_{genome}. \tag{8}$$

An approximation of the variance of $T_{1,2}$ is derived by noting that the length of a patch of IBD can be approximated as the maximum length of $|\mathcal{N}_1| \times |\mathcal{N}_2|$ different IBD segments,

where $\mathcal{N}_i$ is the set of genotyped nodes below ancestor A_i at locus m in which the IBD segment is observed. This approximation comes from conceptualizing IBD shared among the $|\mathcal{N}_1|$ IBD segment carrying descendants of A_1 and the $|\mathcal{N}_2|$ IBD segment carrying descendants of A_2 as $|\mathcal{N}_1| \times |\mathcal{N}_2|$ independent segments with a single point at which all segments overlap. The length of the merged segment to one side of this focal point then has a distribution given by the maximum of $|\mathcal{N}_1| \times |\mathcal{N}_2|$ exponential random variables whose means depend on the degree of separation between the corresponding pairs of leaf individuals.

[**0119**] This approximation is a simplification of the IBD sharing pattern because the segments are not truly independent and need not overlap at a single point. Moreover, under this approximation, the length of the merged segment would actually be the maximum over sums of identically distributed random variables, representing the sum of the length of a segment to the right of the center point and the length of the segment to the left. However, one need not be overly concerned with these drawbacks of the conceptualization because the goal is to obtain an accurate, yet simple approximation of the variance of the distribution. One may also assume that no member of $\mathcal{N}_i$ is the direct ancestor of another member of the set, which holds in practice if we drop all individuals from $\mathcal{N}_i$ who are descended from others.

[**0120**] The length, $\ell_{i,j}$, of an IBD segment between leaf nodes i and j is can be modeled as an exponentially distributed random variable with mean length $\mu_{ij} = L_{genome}/d_{i,j}R$, where $d_{i,j}$ is the degree of relationship between them and R is the expected number of recombination events, genome wide, in one meiosis. This approximation is due to Huff, C.D., Witherspoon, D.J., Simonson, T.S., Xing, J., Watkins, W.S., Zhang, Y., Tuohy, T.M., Neklason, D.W., Burt, R.W., Guthery, S.L., Woodward, S.R., and Jorde, L.B. (2011). Maximum-likelihood estimation of recent shared ancestry (ERSA). Genome Research, 21, 768–774. When the length of the genome is expressed in centimorgans (cM), the expected number of recombination events in the genome is $L_{genome}/100$. Thus, the expected length in cM of an IBD segment between individuals i and j separated by $d_{i,j}$ meioses is $\mu_{ij} = 100/d_{i,j}$.

[0121] Let $L_{1,2}$ denote a random variable describing the length of the segment formed by merging all segments at a given locus m between descendants of A_1 and A_2. If the lengths of all segments at this locus were independent, their merged length would have a distribution given (approximately) by the maximum over independent exponentially distributed random variables with means $\{\mu_{i,j}\}_{i\in\mathcal{N}_1,j\in\mathcal{N}_2}$.

[0122] Then we have $L_{1,2} = \max(\{\ell_{i,j}\}_{i\in\mathcal{N}_1,j\in\mathcal{N}_2})$. Under this condition, the cumulative density function (CDF) $F_L(\ell;\mathcal{N}_1,\mathcal{N}_2)$ of L is

$$
\begin{aligned}
&F_{L_{1,2}}(\ell;\mathcal{N}_1,\mathcal{N}_2) \\
&\quad = Pr(L_{1,2} < \ell;\mathcal{N}_1,\mathcal{N}_2) \\
&\quad = Pr(\ell_{i,j} < \ell, \text{ for } i \in \mathcal{N}_1, j \in \mathcal{N}_2) \\
&\quad = \prod_{i\in\mathcal{N}_1}\prod_{j\in\mathcal{N}_2} Pr(\ell_{i,j} < \ell) \\
&\quad = \prod_{i\in\mathcal{N}_1}\prod_{j\in\mathcal{N}_2} (1 - e^{-\lambda_{i,j}\ell}) \\
&\quad = 1 - \sum_{i\in\mathcal{N}_1,j\in\mathcal{N}_2} e^{-\lambda_{i,j}\ell} + \sum_{i,u\in\mathcal{N}_1,j,v\in\mathcal{N}_2} e^{-(\lambda_{i,j}+\lambda_{u,v})\ell}(1-\delta_{(i,j),(u,v)}) \\
&\qquad - \sum_{i,u,w\in\mathcal{N}_1,j,v,z\in\mathcal{N}_2} e^{-(\lambda_{i,j}+\lambda_{u,v}+\lambda_{z,w})\ell}(1-\delta_{(i,j),(u,v)})(1-\delta_{(i,j),(z,w)})(1-\delta_{(u,v),(z,w)}) + \cdots,
\end{aligned}
\tag{9}
$$

where $\delta_{(a,b),(c,d)}$ is the Kronecker delta between tuples (a,b) and (c,d), which is equal to one when $(a,b) = (c,d)$ and zero, otherwise.

[0123] The sets $\mathcal{N}_1$ and $\mathcal{N}_2$ are, themselves, random variables. Summing over all sets $\mathcal{N}_1$ and $\mathcal{N}_2$, one obtains

$$
F_{L_{1,2}}(\ell) = \sum_{\mathcal{N}_1,\mathcal{N}_1} F_{L_{1,2}}(\ell;\mathcal{N}_1,\mathcal{N}_2)Pr(\mathcal{N}_1)Pr(\mathcal{N}_2), \tag{10}
$$

where the probabilities $Pr(\mathcal{N}_1)$ and $Pr(\mathcal{N}_2)$ are probabilities of observing IBD in the sets of leaf nodes below A_1 and A_2 and can be computed using the recursion in Equation (4).

[0124] Over the length of the genome, the number $N_{1,2}$ of IBD segments between the descendants of A_1 and A_2 is approximately Poisson distributed with mean $Pr(\mathcal{I}_{1,2})L_{genome}/E[L_{1,2}]$, where $\mathcal{I}_{1,2}$ is the event that IBD is observed between some individual in $\mathcal{N}_1$ and some individual in $\mathcal{N}_2$. This rate comes from the fact that the average total amount of the genome

in a patch of IBD is $Pr(\mathcal{I}_{1,2})L_{genome}$ while the average length of any given segment is $E[L_{1,2}]$. Thus, there are approximately $Pr(\mathcal{I}_{1,2})L_{genome}/E[L_{1,2}]$ patches of IBD in the genome, on average. When the lengths of IBD are relatively short and far apart, which they are when the degree between A_1 and A_2 is large, this is a reasonable approximation. This is precisely the regime in which the distribution in Equation (20) is most useful.

[**0125**] The total length $T_{1,2}$ of merged IBD among the descendants of A_1 and A_2 is then

$$T_{1,2} = \sum_{n=1}^{N_{1,2}} L_{1,2}. \tag{11}$$

One can derive the variance of $T_{1,2}$ using the law of total variance as

$$\begin{aligned} \mathrm{Var}(T_{1,2}) &= \mathrm{E}[\mathrm{Var}(T_{1,2}|N_{1,2})] + \mathrm{Var}(\mathrm{E}[T_{1,2}|N_{1,2}]) \\ &= \mathrm{E}[N_{1,2}\mathrm{Var}(L_{1,2})] + \mathrm{Var}(N_{1,2}\mathrm{E}[L_{1,2}]) \\ &= \mathrm{E}[N_{1,2}]\mathrm{Var}(L_{1,2}) + \mathrm{Var}(N_{1,2})\mathrm{E}[L_{1,2}]^2. \end{aligned} \tag{12}$$

Note that because $N_{1,2} \sim \mathrm{Poisson}(Pr(\mathcal{I}_{1,2})L_{genome}/\mathrm{E}[L_{1,2}])$, one obtains

$$\mathrm{E}[N_{1,2}] = \mathrm{Var}(N_{1,2}) = Pr(\mathcal{I}_{1,2})L_{genome}/\mathrm{E}[L_{1,2}]. \tag{13}$$

So Equation (12) simplifies to

$$\begin{aligned} \mathrm{Var}(T_{1,2}) &= \frac{Pr(\mathcal{I}_{1,2})L_{genome}}{\mathrm{E}[L_{1,2}]}[\mathrm{Var}(L_{1,2}) + \mathrm{E}[L_{1,2}]^2] \\ &= Pr(\mathcal{I}_{1,2})L_{genome}\frac{\mathrm{E}[L_{1,2}^2]}{\mathrm{E}[L_{1,2}]}, \end{aligned} \tag{14}$$

where the fact that $\mathrm{Var}(X) = \mathrm{E}[X^2] - \mathrm{E}[X]^2$ has been used.

[**0126**] It remains to find $\mathrm{E}[L_{1,2}]$ and $\mathrm{E}[L_{1,2}^2]$. Using the cumulative density function (CDF) of $L_{1,2}$ in Equation (10) and the fact that $\mathrm{E}[X^m] = m! \int_{\mathbb{R}} x^{m-1}[1 - F_X(x)]dx$, one

obtains

$$\begin{aligned}
\mathrm{E}_{\mathcal{N}_1,\mathcal{N}_2}[L_{1,2}^m] &= m!\int_{\ell=0}^{\infty} x^{m-1}[1-F_{L_{1,2}}(\ell;\mathcal{N}_1,\mathcal{N}_2)]d\ell \\
&= \sum_{i\in\mathcal{N}_1,j\in\mathcal{N}_2}\int_{\ell=0}^{\infty} m!\ell^{m-1}e^{-\lambda_{i,j}\ell}d\ell - \sum_{i,u\in\mathcal{N}_1,j,v\in\mathcal{N}_2}\int_{\ell=0}^{\infty} m!\ell^{m-1}e^{-(\lambda_{i,j}+\lambda_{u,v})\ell}d\ell \\
&\quad + \sum_{i,u,w\in\mathcal{N}_1,j,v,z\in\mathcal{N}_2}\int_{\ell=0}^{\infty} m!\ell^{m-1}e^{-(\lambda_{i,j}+\lambda_{u,v}+\lambda_{z,w})\ell}d\ell + \cdots \\
&= \sum_{i\in\mathcal{N}_1,j\in\mathcal{N}_2}\frac{m!}{\lambda_{i,j}^m} - \sum_{i,u\in\mathcal{N}_1,j,v\in\mathcal{N}_2}\frac{m!}{(\lambda_{i,j}+\lambda_{u,v})^m} \\
&\quad + \sum_{i,u,w\in\mathcal{N}_1,j,v,z\in\mathcal{N}_2}\frac{m!}{(\lambda_{i,j}+\lambda_{u,v}+\lambda_{z,w})^m} + \cdots
\end{aligned} \tag{15}$$

where the integrals in Equation (15) can be evaluated by noting that they are essentially expressions for the moments of exponential random variables with parameters λ_i, $(\lambda_i+\lambda_j)$, $(\lambda_i+\lambda_j+\lambda_k)$, etc.

[**0127**] Thus, one can use Equation (15) to compute

$$\mathrm{E}[L_{1,2}^m] = \sum_{\mathcal{N}_1,\mathcal{N}_1}\mathrm{E}_{\mathcal{N}_1,\mathcal{N}_2}[L_{1,2}^m]Pr(\mathcal{N}_1,\mathcal{N}_2), \tag{16}$$

where $Pr(\mathcal{N}_1,\mathcal{N}_2)$ is the probability of observing IBD segments at the leaves $\mathcal{N}_1$ and $\mathcal{N}_2$, and is obtained using the recursion in Equation (4). Equation (16) is then used in to Equation (14) to obtain the variance of $T_{1,2}$.

[**0128**] In practice, it is too computationally demanding to compute the sums in Equation (16) because the terms $\mathrm{E}_{\mathcal{N}_1,\mathcal{N}_2}[L_{1,2}]$, $\mathrm{E}_{\mathcal{N}_1,\mathcal{N}_2}[L_{1,2}^2]$, and $Pr(\mathcal{N}_1,\mathcal{N}_2)$ are not fast to compute in large quantities. However, the probabilities $Pr(\mathcal{N}_1,\mathcal{N}_2)$ can be computed quickly, making it possible to find the most likely sets of leaf nodes, $\hat{\mathcal{N}}_1$ and $\hat{\mathcal{N}}_2$, with observed IBD. Thus, in some implementations one can use an approximation in which it is assumed that the most likely IBD pattern has been observed and one computes

$$\mathrm{E}[L_{1,2}^m] \approx \mathrm{E}_{\hat{\mathcal{N}}_1,\hat{\mathcal{N}}_2}[L_{1,2}^m]. \tag{17}$$

The assumption used in this approximation is that most patterns of observed IBD at the leaves are unlikely compared with the most likely pattern and that most high-likelihood patterns of IBD will yield similar moments $E[L_{1,2}^m]$.

[0129] Equation (17) can then be used to obtain an approximation of the variance of $T_{1,2}$ as

$$\mathrm{Var}(T_{1,2}) \approx Pr(\mathcal{I}_{1,2}) L_{genome} \frac{\mathrm{E}[L_{1,2}^2]}{\mathrm{E}[L_{1,2}]}, \tag{18}$$

where $L_{1,2}$ is the length of any given IBD segment between A_1 and A_2 formed by merging all IBD segments between leaf nodes in A_1 and A_2 that overlap one another.

[0130] If the segments, $L_{1,2}$ were each exponentially distributed, then $T_{1,2}$ would have a gamma distribution. In practice, a gamma distribution is an accurate approximation for the distribution of $T_{1,2}$, given that the length $T_{1,2}$ is greater than zero, so one can approximate the distribution of $T_{1,2}$ by

$$T_{1,2}|T_{1,2} > 0 \sim \mathrm{Gamma}(k_{1,2}, \theta_{1,2}),$$

where $k_{1,2}$ and $\theta_{1,2}$ are found by matching the mean and variance of the gamma distribution with $\mathrm{E}[T_{1,2}]$ and $\mathrm{Var}(T_{1,2})$. Thus, one obtains

$$T_{1,2}|T_{1,2} > 0 \sim \mathrm{Gamma}(\frac{E[L_{1,2}]^2}{\mathrm{Var}(L_{1,2})}, \frac{\mathrm{Var}(L_{1,2})}{E[L_{1,2}]}), \tag{19}$$

where $\mathrm{E}[L_{1,2}]$ and $\mathrm{E}[L_{1,2}^2]$ are given by Equation (15).

[0131] If every IBD segment has some length, one can assume that $T_{1,2}$ is only identically zero when there are no IBD segments. The distribution of the number of segments can be modeled as a Poisson random variable with mean $\mathrm{E}[N_{1,2}]$ equal to the expected number $N_{1,2}$ of merged segments shared between $\mathcal{N}_1$ and $\mathcal{N}_2$. The probability that there are no segments is then $e^{-\mathrm{E}[N_{1,2}]}$. Thus, one has the approximation

$$f_{T_{1,2}}(t_{1,2}) \approx \begin{cases} \frac{t_{1,2}^{k-1}}{\Gamma(k)\theta^k} e^{-t_{1,2}/\theta}(1 - e^{-\mathrm{E}[N_{1,2}]}) & \text{if } t_{1,2} > 0 \\ e^{-\mathrm{E}[N_{1,2}]} & \text{if } t_{1,2} = 0. \end{cases}, \tag{20}$$

where $k = E[L_{1,2}]^2/\mathrm{Var}(L_{1,2})$ and $\theta = \mathrm{Var}(L_{1,2})/E[L_{1,2}]$.

Estimators of Distant Relationships

[0132] A maximum likelihood estimator of the degree between A_1 and A_2 can be obtained by determining the degree $d_L(A_1, A_2)$ between A_1 and A_2 for which value of the

distribution in Equation (20) is maximized. This gives the likelihood estimator

$$d_L(A_1, A_2) = \arg\max_d f_{T_{1,2}}(t_{1,2}; d). \tag{21}$$

[**0133**] One can also use Equation (5) to obtain a generalized version of the DRUID estimator of Ramstetter, M.D., Shenoy, S.A., Dyer, T.D., Lehman, D.M., Curran, J.E., Duggirala, R., Blangero, J., Mezey, J.G., and Williams, A.L. (2018). Inferring identical-by-descent sharing of sample ancestors promotes high-resolution relative detection. Am. J. Hum. Genet. 103, 30–44. The generalized estiamtor can provide fast estimates of the degree without the need to evaluate Equation (21). Using the approach of Ramstetter et al., the total length of IBD shared between A_1 and A_2 can be estimated as the total length of IBD shared between $\mathcal{N}_1$ and $\mathcal{N}_2$, divided by the total fraction of genetic material A_1 and A_2 are expected to pass to their descendants. The fraction f_i of the genome of A_i passed to their descendants $\mathcal{N}_i$ is given by

$$f_i = 1 - Pr(O_n = 0 \text{ for } n \in \mathcal{N}_i | O_{A_i} = 1) \tag{22}$$

Thus, an estimate $\widehat{IBD}(A_1, A_2)$ of the amount of IBD shared between A_1 and A_2 is

$$\widehat{IBD}(A_1, A_2) = \frac{IBD(\mathcal{N}_1, \mathcal{N}_2)}{f_1 f_2} \tag{23}$$

Using the expression $\hat{\phi} = \widehat{IBD}(A_1, A_2)/4L_{genome}$ for the kinship coefficient when all IBD is of type 1, one obtains the generalized DRUID estimator

$$d_D(A_1, A_2) = d : \frac{1}{2^{d+3/2}} \leq \frac{IBD(\mathcal{N}_1, \mathcal{N}_2)}{4 f_1 f_2 L_{genome}} < \frac{1}{2^{d+1/2}}, \tag{24}$$

are the ones used for the DRUID estimator presented in Ramstetter et al. Thus, one obtains a version of the DRUID estimator that can be applied to general outbred pedigrees.

Likelihood for Identifying Background IBD

[**0134**] Individuals with no recent relationship can share small segments of IBD by chance, especially in populations with recent or severe bottlenecks. This kind of IBD is referred to as background IBD and it poses a considerable challenge to accurate pedigree

inference. Previous methods have addressed background IBD by various approaches. For example, the authors of the ERSA method present an approach for modeling the distribution of background IBD among unrelated individuals and then performing a likelihood ratio test to determine whether the IBD shared between a new pair of individuals is significantly different from background; Huff, C.D., Witherspoon, D.J., Simonson, T.S., Xing, J., Watkins, W.S., Zhang, Y., Tuohy, T.M., Neklason, D.W., Burt, R.W., Guthery, S.L., Woodward, S.R., and Jorde, L.B. (2011). Maximum-likelihood estimation of recent shared ancestry (ERSA). Genome Research, 21, 768–774. This approach requires a background distribution of IBD and it requires testing each pair of individuals separately. The difficulty with detecting background IBD between each pair of individuals separately is that it can result in throwing out many pairs of individuals whose levels of IBD sharing are near background, even when those pairs are truly related. Improved power for detecting background IBD can be obtained by leveraging the information inherent in previously-inferred pedigree structures to infer background IBD for sets of multiple individuals at the same time.

[**0135**] Practitioners take an approach to identifying background IBD in which they consider the information contained in IBD sharing patterns across multiple individuals to determine when IBD is background and when it is due to true recent ancestry. In particular, we consider the problem in which all of the IBD observed in an individual is either background IBD, or true IBD due to a recent relationship.

[**0136**] To illustrate the approach, consider the IBD sharing pattern shown in Figure 3. Individuals 3 and 4 share relatively large amounts of IBD with 5 and 6, compared with the amount shared between $\{1, 2\}$ and $\{5, 6\}$. If 1 and 2 were much more distantly related to 5 and 6 than are 3 and 4, one might not consider the amount of IBD they share with 5 and 6 to be unusually small. However, because 1, 2, 3, and 4 have similar degrees of relatedness to 5 and 6, the amount of IBD shared by 1 and 2 appears to be unusually low. If one can say that the amount of IBD shared below node 7 is smaller than expected by chance, then one can assume that the IBD observed in 1 and 2 is background IBD and remove these nodes from consideration when connecting the left and right pedigrees.

[0137] One can test for background IBD through a series of hypothesis tests. Given that IBD is observed between two sets of nodes, $\mathcal{N}_1$ and $\mathcal{N}_2$, suppose that the putative common ancestors A_1 and A_2 through which the IBD was inherited are the most recent common ancestors of $\mathcal{N}_1$ and $\mathcal{N}_2$, respectively. One can then consider each of the descendant nodes immediately below A_1 in turn (*e.g.*, 7 and 8 in Figure 3) and ask whether the amount of observed IBD below the node is much lower or higher than expected by chance, given the degree between A_1 and A_2 inferred using all the descendant nodes below A_1.

[0138] All nodes that reject the null hypothesis of this test are dropped and the ancestral node is reset to be the common ancestor of all remaining IBD-carrying nodes. For example, if one detected that the clade below node 7 in Figure 3 had much lower IBD than expected by chance, one would drop node 7 and its descendants from consideration and set the true common ancestor relating the two pedigrees to be node 8. One would then iteratively repeat this procedure until no nodes are dropped. One would then then repeat the procedure for the nodes immediately below A_2.

[0139] Let $\mathcal{C}_n$ denote the set of children of node n. To test whether the IBD observed below a child node $c \in \mathcal{C}_{A_1}$ is background IBD, consider the null hypothesis H_0 that the observed IBD below the node is real, and ask whether this hypothesis is rejected in favor of the alternative hypothesis H_1 that the IBD is background. Background IBD can either be lower than the expected true amount of IBD (as in the example in Figure 3), or it can be higher than expected if, for example, the true branch between nodes 7 and A_1 in Figure 3 were very long and the expectation is that no true IBD is observed. Therefore, the hypothesis test must be two-tailed to account for both scenarios.

[0140] Under the null hypothesis H_0, it is assumed that the IBD observed is real and we assume that the degree $d_{H_0}(A_1, A_2)$ between A_1 and A_2 is the maximum likelihood estimate: $d_{H_0}(A_1, A_2) = d_L(A_1, A_2)$, or the generalized DRUID estimate: $d_{H_0}(A_1, A_2) =$

$d_D(A_1, A_2)$. One can then perform the following test

Reject H_0 at level α if:

$$Pr(T_{c,A_2} \leq t_{c,A_2}; d_{H_0}(A_1, A_2)) < \alpha/2, \text{ or}$$
$$Pr(T_{c,A_2} \geq t_{c,A_2}; d_{H_0}(A_1, A_2)) < \alpha/2, \quad (25)$$

where T_{c,A_2} is the random variable describing the observed amount of IBD between descendants of c and descendants of A_2 with observed value t_{c,A_2}. The distribution of T_{c,A_2} is given by Equation (20). It is reasonable to be conservative when dropping background IBD so that true relationships are called as background IBD only a small fraction of the time. Thus, in practice, we take α to be small, such as $\alpha = 10^{-4}$.

Determining When Ancestral Branches are Unrelated

[0141] One difficulty in constructing large pedigrees is determining the ancestors through which two sets of gentoyped individuals are related. A simple fundamental question is whether two lineages are both on the maternal side of an individual, both on the paternal side, or on opposite parental sides. Without genotyped parents, the side through which a lineage passes can be difficult to determine, although sex chromosomes and mitochondrial haplotypes can be used to resolve the parent of origin in some cases.

[0142] Practitioners consider the problem of inferring whether two distant sets of relatives are related through the same parent of a focal individual, or through different parents. The scenario is shown in Figure 5. Even if pedigrees 1 and 2 in Figure 5 share no IBD, they could still be related to individual 1 through the same parent by passing through different grandparents.

[0143] The amount of IBD shared among pedigrees 1 and 2 is uninformative about whether they are related through the same parent. However, if pedigrees 1 and 2 are related to the focal individual 1 through the same parent, the IBD segments pedigree 1 shares with individual 1 cannot spatially overlap with the segments pedigree 2 shares with individual 1. This is because two overlapping segments would have undergone recombi-

nation in the parent (*e.g.*, 10). The result will either be a spliced segment (Figure 5), or the replacement of one segment by the other with possible reduction in segment size.

[**0144**] In the Bonsai method, when there are multiple possible grandparents through which we can connect two pedigrees $\mathcal{P}_1$ and $\mathcal{P}_2$ to a focal set of nodes $\mathcal{N}$ in a focal pedigree $\mathcal{P}$, we examine whether the IBD segments between $\mathcal{P}_1$ and $\mathcal{N}$ overlap with the IBD segments between $\mathcal{P}_2$ and $\mathcal{N}$.

Training the Probabilistic Relationship Model

[**0145**] Figure 7 shows a flowchart illustrating process 200 for training a probabilistic relationship model using a computer system including one or more processors and system memory. In some implementations, the probabilistic model is a machine-learning model. The trained probabilistic relationship model is configured to predict genetic relationships based on IBD data and age data.

[**0146**] Process 200 starts by receiving IBD data of a plurality of training sets for a plurality of relationships. See block 202. All individuals in each training set are genetically related with the same relationship in a pairwise manner. Each training set is associated with a unique relationship. The IBD data of each training set include pairwise IBD data of individuals in the training set. Figure 8 schematically illustrates an example of how to train the machine learning relationship model using multiple training sets. In some implementations, each training set includes hundreds, thousands, tens of thousands of individuals or more. In some implementations, the training set includes data of actual individuals having a particular genetic relationship. In some implementations, each training set includes data of simulated individuals. Such simulated data can be obtained from recombining reference individuals' data according to recombination models to generate related individuals' data for various types of relationships.

[**0147**] Process 200 further involves obtaining age data of the plurality of training sets for the plurality of relationships. The age data of each training set includes age data for individuals in the training set. See block 204.

[0148] Process 200 further involves training the probabilistic relationship model using the IBD data of the plurality of training sets and the age data of the plurality of training sets. See block 206. The trained probabilistic relationship model is configured to take as input pairwise IBD data and age data for two test individuals and provide as output various likelihoods of various potential relationships for the two test individuals.

[0149] Figure 8 schematically illustrates further details on how to train the probabilistic relationship model according to process 200 shown in Figure 7. Boxes 302, 304 and 306 represent training sets for parent-offspring relationship, full siblings relationship, and avuncular relationship, respectively. The figure shows only three training sets for three relationships, but many more training sets that are used are not shown here. In some implementations, the various relationships include relationships of the 0th, 1st, 2nd, 3rd, 4th, 5th, 6th, 7th, 8th, 9th, 10th, 11th, 12th, 13th, 14th, or 15^{th}degree or further. In some implementations, the various relationships include relationships of at least 0, 1, 2, 4, 6, 8, 10, 12, 14, or 16 meioses on a common-ancestor path between the two individuals through a common ancestor. In the training set of 302 for parent-offspring relations, the IBD data and age difference data for each pair of individuals is shown in a box. Here, the training data include half IBD (IBD1) length, number of IBD segments, and age difference. Other data may also be used to train relationship models. For example, the pairwise IBD data may include the length of full IBD (IBD2) segments, and/or the length of half IBD (IBD1) segments. The two types of lengths may be modeled as two separate probability distributions. Alternatively, the different types of IBD lengths may be combined, and the combined length may be modeled as a probability distribution.

[0150] As shown here the number of IBD segments is also used to train the probabilistic relationship model. In some implementations, the numbers of the two types of IBD segments may be modeled separately. In other implementations, the two numbers may be combined and modeled to have one probability distribution. Age difference between the two individuals in a pair is also used to train the probabilistic relationship model.

[0151] Three pairs of individuals are shown for parent-offspring training set 302. Each individual pair provides a data point of the length of half IBD, a data point of the number

of IBD, and a data point of age difference. Although only three individuals are illustrated, the training set may include hundreds, thousands, or tens of thousands of individuals or more.

[0152] The data points from the individuals in the training set for each variable (IBD1, number of IBD segments, age difference) are used to train a probabilistic relationship model. The probabilistic relationship model models the probability distribution of each variable as a Gaussian distribution in this example. In other implementations, the probability distribution of each variable may be modeled as an exponential distribution, a Poisson distribution, a binomial distribution, a beta binomial distribution, and other suitable distributions based on prior knowledge of the variable.

[0153] The data points from the individuals in the training set for each variable (IBD1, number of IBD, age difference) are used to train a probabilistic relationship model. The probabilistic relationship model models the probability distribution of each variable as a Gaussian distribution in this example. In other implementations, the probability distribution of each variable may be modeled as an exponential distribution, a Poisson distribution, a binomial distribution, a beta binomial distribution, and other suitable distributions based on prior knowledge of the variable.

[0154] The data from the training set 302 for the three variables (IBD1 length, number of IBD segments, and age difference) are used to train the probabilistic model. In some implementations, training involves using various techniques to fit the probability distribution to the training data. In some implementations, method-of-moments techniques are used to fit the Gaussian distribution of each variable to the training data. The probability distributions for the three variables are shown in box 312 for parent-offspring relationship. The data and the distributions in the figure are for illustrative purposes only, and they do not reflect biological or mathematical reality. In the same manner, full siblings training set data 304 are used to train the probabilistic relationship model to obtain the probability distributions for the three variables as shown in box 314. The training data of avuncular relationship in box 306 are used to train the probabilistic relationship model to obtain the probability distributions for the three variables for the avuncular relationship.

[0155] After the model is trained, it can be applied to estimate relationship likelihoods between individuals based on IBD data and age differences between the individuals. To apply the probabilistic relationship model, test data of two test individuals are provided to the trained model. The model provides as output the probability of each variable for each relationship. Multiple probabilities for multiple independent variables are aggregated in a relationship to provide a likelihood of the relationship. For example, likelihoods for a relationship as functions of each of the three variables may be summed to provide a composite likelihood indicating how likely the relationship is given the two test individuals' IBD and age data.

[0156] Regarding training, various techniques may be used to fit the model to the data. In some implementations, maximum likelihood methods may be used to obtain parameters that maximize the likelihood of the model given the data. In some implementations, method-of-moments techniques may be used to calculate distribution parameters from the training data. Other model fitting techniques such as kernel density estimation may also be employed, although such techniques tend to provide less accurate estimates in some applications.

Generating Pedigree Graphs

[0157] Another aspect of the disclosure provides methods for generating pedigree graphs that are informative, user-friendly, easy to understand or intuitive.

[0158] Figure 9 shows a flowchart illustrating process 400 for generating pedigree graphs using a computer system that includes one or more processors and system memory. Process 400 involves receiving from a database pedigree relationship data of a plurality of genetically related individuals. See block 402. The pedigree relationship data include a plurality of entries representing the plurality of genetically related individuals. Each entry includes links relating each child entry to its parent entries in a dataset. A child entry and its parent entries represent a child and its biological parents respectively.

[0159] Figure 10 illustrates three such entries that can be used in the process for gen-

erating the pedigree graph. The first entry has an identifier of value 1. It has a null link to its parent entries. This indicates that this entry does not have known parents in the plurality of individuals in a dataset. Such an entry is considered a root entry that may be used to generate a root node of a graph having a tree structure. The entry in the middle has an entry identifier of value 2, and a null link to its parent entries, indicating that the individual represented by the entry does not have parents in the dataset. The entry at the bottom has an identifier of value 12, and a link indicating is parent entries, showing that entry 1 and entry 2 are its parent entries.

[**0160**] In some implementations not shown in Figure 9, various preprocessing steps may be performed after the pedigree relationship data are received. One of the preprocessing steps involves performing passes over entries to determine various properties on the entry. For example, determining for each entry its child entries, sibling entries, partner entries, etc. based on parent-child relations among these entries. Another preprocessing step involves determining ideal ordering of partners when there are multiple partners. For example, an entry having the largest number of partners is placed towards the center of a row of partners. This can help to reduce crossing of lines that connect different nodes. Another preprocessing step involves marking entries for important nodes, such as the focal node, nodes in the nuclear family, core nodes, etc. Another preprocessing step involves determining the chain of nodes linking each node back to the focal node.

[**0161**] Process 400 then involves determining a minimal set of root entries from which all of the plurality of entries are reachable by starting from the minimal set of root entries and traversing from parent entries to child entries and traversing between partner entries. See block 404. As mentioned above, a root entry is an entry having no parent entries. Partner entries are two parent entries of the same child entry. In some implementations, two partner entries may be generated based on other information, such as non-genetic information indicating marriage or partnership that does not yield children. One can travel from one partner to another partner through the partner's children.

[**0162**] Process 400 further involves forming a subtree, for each root entry of the minimal set of root entries, using the root entry as a root node and entries reachable from the root

entry as additional nodes. This operation obtains a plurality of subtrees. See block 406. Root nodes do not have parent nodes, and they tend to be at the top of a subtree. On the contrary, leaf nodes do not have child nodes, and they tend to be at the bottom of a subtree.

[0163] Process 400 further involves positioning nodes in each of the plurality of subtrees. See block 408. In some implementations, positioning the nodes involves starting from the root node and recursively going through nodes in a defined order to reach nodes to be positioned. In some implementations, the defined order is as follows.

1. children of partners to the left;
2. children of partners to the right;
3. partners to the left;
4. self; and
5. partners to the right

[0164] In some implementations, positioning nodes in each of the plurality of subtrees involves placing a leaf node without any siblings on the left at an origin. It also involves placing a leaf node with a sibling immediately to its left at a position immediately to the right of said sibling. In some implementations, it also involves positioning parent nodes relative to their child nodes so that they are further from the horizontal center of a row of nodes than its child nodes are. In some implementations, positioning nodes also involves positioning parent nodes relative to their child nodes' partner nodes, so that they are further from the horizontal center of a row of nodes than their child nodes' partner nodes are.

[0165] Figure 11 shows a subtree that can be formed from a root node (node 6). By recursively going through nodes in the defined order described above, one gets to node 1, which is a leaf node without any siblings on the left. So it can be placed at an origin (0, 0). The process of positioning nodes next would reach node 2, which is a leaf node

with a sibling immediately to its left (node 1). The process places node 2 at a position immediately to the right of node 1 at (1, 0). Following the same rules to reach different nodes and positioning the nodes, the left branch of the subtree can be placed as shown at the top of Figure 12A.

[**0166**] Similarly, the right branch of the subtree may be placed in positions as shown in Figure 12A at the bottom. Note that even though the left branch and the right branch are illustrated as the top half and the bottom half of Figure 12A, they actually overlap, because, *e.g.*, node 1 and node 10 are both placed at the origin (0, 0). The process maintains positions of the right contour of the left branch and the left contour of the right branch. The positions of the contours are obtained from parent nodes that are not on the outside of the subtree and its children nodes not on the outside of the subtree. The right contour of the left branch and the left contour of the right branch may then be used to horizontally shift the two branches so that the two branches do not overlap on non-corresponding nodes as shown in Figures 12A-D.

[**0167**] In some implementations, there are more than two parent nodes in a row. A first parent node is not on either end of the row. A second parent node is immediately to the left of the first parent node. These implementations include a rule of positioning the first parent node relative to the child nodes of the second parent node, so that the first parent node is to the right of the child nodes of the second parent node.

[**0168**] Figure 14 illustrates the positioning of a row of four parent nodes according to the rule. In this illustrative example, there are four parent nodes in a row (702, 704, 706, 708). Node 706 is the partner of nodes 702, 704, and 708. Node 718 is the child of parent nodes 706 and parent nodes 708. Node 716 is the child of nodes 704 and 706. Node 712 and node 714 are children of node 706 and node 702.

[**0169**] Parent node 706 is not on either end of the row. Parent node 704 is immediately to the left of parent node 706. In this case, the rule described above positions the parent node 706 relative to child node 716, which is the child of parent node 704. Node 702 is not a partner of node 704. Nonetheless, the same rule described above applies to node 704 and 702. Node 704 is not on either end of the row. It is placed relative to the two

child nodes of node 702 so that it is to the right of node 714, a child of node 702.

[0170] In some implementations, after positioning the first parent nodes relative to the child nodes of the second parent, the process shifts the child nodes of the first parent node to maintain previous relative relations between positions between the first parent and the child nodes of the first parent.

[0171] Returning to Figure 9, process 400 finally merges the plurality of subtrees to form the pedigree graph for the plurality of genetically related individuals after the nodes are positioned in the subtree. See block 410.

[0172] In some implementations, merging the subtrees to form the pedigree graph involves shifting one or more of the subtrees so that non-corresponding nodes of different subtrees do not overlap. Non-corresponding nodes on different trees are nodes that do not represent the same individual. Figures 12A-D illustrate shifting one branch of a subtree, but the same shifting techniques are also used to shift subtrees. Some implementations also involve merging corresponding nodes on different trees into one node. Corresponding nodes are nodes on different subtrees that represent the same individual.

[0173] In some implementations, merging the subtrees to form the pedigree graph includes identifying core nodes that include a focal node representing a focal individual, any sibling nodes representing siblings of the focal individual, any parent nodes representing parents of the focal individual, any descendant nodes presenting descendants of the focal individual, and any partner nodes presenting partners of the descendants.

[0174] For example, for the pedigree graph in Figure 13, the focal node is node 1. The core nodes include nodes 1-4 shown in box 702. Two subtrees from root node 13 and node 6 merge on the four core nodes. In some implementations, the process further involves merging subtrees that do not intersect with the core nodes into subtrees that do intersect with the core nodes. In Figure 13, the subtree 704 does not overlap with any of the core nodes. However, it does intersect with the subtree including node 6. In such case, the subtree 704 is merged to the subtree including node 6, which intersects with the core nodes.

[0175] In some implementations, merging the plurality of subtrees involves merging each pair of four pairs of subtrees to form four grandparent subtrees. See Figures 15 A and B. Here the four pairs of subtrees respectively correspond to four pairs of great-grandparents of the focal individual. The four pairs of great-grandparents are nodes (821, 822), (823, 824), (825, 826) and (827, 828) as shown in Figure 15B. The focal individual is represented by node 840. The four grandparent trees are shown in box 802 corresponding to grandparent 812, box 804 corresponding to grandparent 814, box 806 corresponding to grandparent 816, and box 808 corresponding to grandparent 818.

[0176] Here, merging each pair of subtrees includes horizontally shifting one subtree so that non-corresponding nodes of the two subtrees do not overlap, and merging two corresponding nodes on the pair of subtrees representing the same grandparent into one node. See Figure 16. Figure 16 is a close-up of the grandparent tree corresponding to grandparent 814 in box 804. This subtree includes two subtrees for two great-grandparents 823 and 824, which are direct ancestors of the focal node. In these two great grandparent trees, the direct ancestors of the focal node 840 are colored as orange. Nodes that are not direct ancestors of the focal node are not connected by colored lines. The colored branch of the great-grandparents subtree above great-grandparent 823 is shown in area 852. The colored branch of the great-grandparents tree above great-grandparent 824 is shown in area 854. These two branches include non-corresponding nodes because they do not represent same individuals on the two great-grandparent subtrees. Merging these two great-grandparents trees involves shifting one of the two trees so that non-corresponding nodes on them do not overlap.

[0177] In some implementations, the process further involves merging the four grandparent subtrees to form the pedigree graph. The merging involves horizontally shifting one or more of the grandparents subtrees so that non-corresponding nodes of different grandparent subtrees do not overlap and merging corresponding nodes on different grandparent subtrees presenting the same individual into one node. In the pedigree tree graph in Figure 15B, node 832, 834, and 840 are obtained by merging corresponding nodes from the four grandparent subtrees.

[0178] In some implementations, the subtree corresponding to a great grandparent is obtained by merging two or more subtrees. See, *e.g.*, Figure 15B, the great-grandparent subtree corresponding to great grandparent 818 shown in box 808, which is obtained by merging two or more subtrees.

[0179] In some implementations, the two grandparent subtrees in the middle are smaller than the two grandparent subtrees on the outside. See Figures 15 A and B for example. In some implementations, generating pedigree graphs further involves swapping positions of two partner nodes whose parental subtrees are on the opposite side from the partner nodes.

[0180] Some implementations further involve removing empty spaces in the pedigree graph. In some implementations, this involves removing a column of empty spaces in the pedigree graph when the removal of the empty spaces does not cause any non-corresponding nodes to overlap.

[0181] In some implementations, the process further includes applying force directed graph drawing techniques to redraw one or more nodes and lines connecting them. In some implementations, the one or more nodes include leaf nodes and their parent nodes.

[0182] In some implementations each pair of two or more pairs of parent nodes in the pedigree graph includes two nodes rendered in different colors. In some implementations, the lines connecting each child node to its parent nodes includes curved lines.

Graphical User Interface for Pedigree Graphs

[0183] Another aspect of the disclosure relates to methods for displaying the pedigree graph for a plurality of genetically related individuals on a graphical user interface (GUI). The method is implemented using a computer system including a processor, system memory, and a display device. The method includes using the display device to display a pedigree graph including a plurality of nodes representing a plurality of genetically related individuals and lines connecting each child node to its pair of parent nodes.

The child node and its pair of parent nodes present a child and its pair of parents. Each pair of two or more pairs of parent nodes includes two nodes rendered in different colors. The lines connecting each child node to its pair of parent nodes include curved lines.

[**0184**] Figure 15B shows an example pedigree graph that is displayed according to some implementations. In the pedigree graph, the two or more pairs of parent nodes include (832, 834), (812, 814), (816, 818), (821, 822), (823, 824), (825, 826), (827, 828). For these pairs, two nodes of each pair are rendered in different colors.

[**0185**] In some implementations, *e.g.,* Figure 15B, all nodes of the two or more pairs of parent nodes are rendered in different colors. In other words, all of the nodes listed above have colors that are different from each other.

[**0186**] In some implementations, the two or more pairs of parent nodes are direct ancestors of a focal node. In the example shown in Figure 15, all of the above listed parent nodes are direct ancestors of focal node 840.

[**0187**] In some implementations, the relative nodes that are not direct ancestors of the focal node have the same coloring as a direct ancestor that is on the same family side and at the same generational level as the relative nodes.

[**0188**] In some implementations, generation levels at and above great-grandparents are rendered in the same color on the same side of the family.

[**0189**] In some implementations, nodes at the pedigree graph include core nodes that include the focal node representing a focal individual, any sibling nodes resenting siblings of the focal individual, any parent nodes representing parents of the focal individual, any descendant nodes resenting descendants of the focal individual, and any nodes resenting partners of the descendants. In the example in Figure 15B, the core nodes include focal node 840, sibling nodes 836, 838, and parent nodes 832 and 834.

[**0190**] In some implementations, the pedigree graph includes lines indicating direct ancestry of the focal individual. The lines indicating the direct ancestry of the focal individual are rendered in a color or shade that is different from lines not indicating the direct ancestry of the focal individual. In the example in Figure 15B, the lines indicating

the direct ancestry of the focal individual 840 are rendered in yellow, which is different from other lines not indicating direct ancestry of the focal individual.

[**0191**] In some implementations, at least one pair of parent nodes has an off-center alignment relative to their child nodes. See, *e.g.*, parent nodes 827 and 828 relative to their child node 818.

[**0192**] In some implementations, two subtrees having the same topology in the pedigree graph are represented in different forms.

[**0193**] In some implementations, at least one pair of parent nodes has an inter-pair physical distance of larger than the smallest possible distance. See, *e.g.*, parent nodes 814 and 812.

[**0194**] In some implementations, one or more nodes and lines connecting them are drawn using force directed graph drawing techniques. In some implementations, the one or more nodes include leaf nodes and their parent nodes. See *e.g.*, leaf nodes 836, 840, and 838, and parent nodes 832 and 834.

[**0195**] In some implementations, each child node is connected through a curved line to a straight line connecting the child node's parent nodes. See most of the direct ancestry lines on pedigree graph in Figure 15B.

[**0196**] In some implementations, the GUI can interactively display and update the pedigree graph. A user may provide user input relating to genealogy information on any individuals represented by the nodes of the pedigree graph. In some implementations, the user input is provided in a way that involves an interaction with an input text field of the GUI and/or an interaction with a graphical element of the pedigree graph. For example, the user may point and click at a node on the pedigree graph, which activates an editing mode of the node. Then the user may provide genealogical or other information about the individual represented by the node. The information may include age, gender, partnership, ethnicity, nationality, relative relationship, name, photos, etc. A computer processor can then use data provided by or derived from the user input data to update the pedigree information underlying the pedigree graph or update the pedigree graph directly. Some

implementations store and/or propagate the updated pedigree information to generate other pedigree graphs involving said information.

[0197] In some implementations, two or more different users may provide input to two or more different pedigree graphs. In some implementations, a computer system uses the input from one user to update the pedigree graph of another user, and vice versa. In such implementations, two or more different users can collaboratively update their pedigree graphs in real time.

[0198] In various embodiments, the pedigree graph is interactive. Namely, the pedigree graph is designed or configured to receive user input, modify information associated with the pedigree graph, and update the pedigree graph using the modified information. In some implementations, the user input is received via a user interaction with the pedigree graph in a GUI. In some implementations, the user interaction includes clicking an interactive node in the pedigree. An interactive node is one that is configured to receive user input and present information, sometimes the presented information being updated by the user input.

[0199] In some implementations, updating the pedigree graph automatically updates one or more display elements of the pedigree graph or in the GUI. In some implementations, the user interaction includes entering data using a window activated by clicking the interactive node. In some implementations, updating one or more elements of the pedigree graph includes changing a relationship between the interactive node and at least one other node in the pedigree graph. In some implementations, at least one interactive node in the pedigree graph is associated with information of health, traits, diseases, physical conditions, or phenotypes of an individual represented by the at least one interactive node.

[0200] In some implementations, the at least one interactive node is associated with a graphical element representing the information of health traits, diseases, physical conditions, or phenotypes of the individual. In some implementations, the user input includes clicking the interactive node. The user provides input by clicking the interactive node. In some implementations, the user input includes entering information of individuals in

a window activated by clicking the interactive node. In some implementations, updating the one or more elements of the pedigree graph includes changing at least one relationship between two nodes. In some implementations, updating the one or more elements of the pedigree graph includes changing the graphical element representing information of traits, diseases, physical conditions, or phenotypes of the individual.

[**0201**] In some implementations, the nodes of the pedigree graph, or associated with information relating to diseases, physical conditions, traits, or phenotypes. In some implementations, the information of diseases, traits, physical conditions, or phenotypes is represented by a graphical icon by graphical elements positioned next to the node. The graphical elements reflect conditions of the individual represented by the node.

[**0202**] Figure 32 shows a pedigree graph as an example according to some implementations. The pedigree graph 1602 has a focal node 1604, a sibling node 1606, a father node 1613, a mother node 1614, a grandfather node 1617, a grandmother node 1618, etc. Node 1604, node 1614, and node 1618 are associated with health information of the individuals represented by the nodes.

[**0203**] Graphical icons 1608 and 1609 respectively represent a heart condition and a lung condition of a focal individual represented by node 1604. Graphical icons 1608 and 1609 are positioned with respect to node 1604. Similarly, graphical icons 1620 and 1621 respectively represent an eye color and a heart condition of the mother represented by node 1614. Graphical icons 1620 and 1621 are positioned with respect to mother node 1614. Similarly, graphical icons 1622 and 1623 respectively represent an eye color and a heart condition of the grandmother represented by 1618. Graphical icons 1622 and 1623 are positioned with respect to the node 1618. The juxtaposition of the nodes and icons in the pedigree graph visualizes the heritability of diseases and traits.

[**0204**] In this example, the user interacts with the pedigree graph by inputting information related to the pedigree or individuals in the pedigree. In some implementations, the information comprises traits, diseases, physical conditions, or phenotypes. In some implementations, the user may provide input regarding a node, such as node 1606 by clicking the node 1606, which brings up a graphical window or GUI 1610 for displaying

and receiving information about Jodi Neville represented by node 1606. In GUI 1610, user may interact with elements in the window to display further information and/or input information related to Jodi. In this example, the user may further view Jodi's medical history information by clicking a link or button 1612, which brings up another GUI 1702 in Figure 33 showing Jodi's medical information or ways to enter information about Jodi.

[**0205**] Figure 33 shows GUI 1702 and GUI 1708 for viewing and inputting data about Jodi. Jodi has hypertension as indicated by display element 1704 in GUI 1702. GUI 1702 also allows the user to input other conditions of the individual via display element 1706. The user may click on display element 1704 to bring up GUI 1708 that includes more information of the hypertension condition. GUI 1708 includes community and relative information regarding hypertension in display element 1710. In GUI 1708, the user may also edit the condition in area 1712 and area 1714. A user may also save the information using display element 1716, or delete the information using display element 1718.

Re-annotating Pedigree Graphs

[**0206**] In many implementations and applications, a pedigree graph includes nodes for both genotyped individuals and ungenotyped individuals. Pedigree graphs can also be referred to as family trees. A genotyped node represents an individual whose genetic data have been used to determine the pedigree relationships depicted by the pedigree graph. An ungenotyped node represents an individual whose genetic data have not been used to determine the pedigree relationships depicted by the pedigree graph. Since ungenotyped nodes are inferred from the pedigree relationships, information about the individual is limited to the inference from the pedigree relationships.

[**0207**] For example, Figure 27 shows a pedigree graph including genotyped nodes and ungenotyped nodes. Each genotyped node is labeled with two letters. For example, the focal node (node 1402) labeled as FZ has a genotyped sibling labeled as LZ. It also has two genotyped parents 1403 and 1404. It can be inferred that each of the parents has two parents and four grandparents. For instance, the parent node 1404 is inferred to have two

parents 1405 and 1406. These inferred individuals are also shown in the pedigree graph. They are not associated with data beyond the inferred relationships among them.

[0208] A user may annotate ungenotyped nodes using annotation information such as name, gender, date of birth, etc. Such information and annotation are helpful for understanding individuals and the relationships represented by the pedigree graph, making the graph more informative. When a pedigree graph is updated with new relationships or additional individuals based on genotyped data, the identities of genotyped nodes are known, and their matching between an old graph and a new graph is straightforward. However, because the identities of the unannotated, ungenotyped nodes in a new graph are unknown, matching them to annotated, ungenotyped nodes in an old graph is not as straightforward. The matching requires using relationships between ungenotyped nodes and genotyped nodes. But such relationships in the old graph and the new graph may not be the same, making it difficult to re-annotate ungenotyped nodes in the new graph using annotation data of corresponding nodes in the old graph. Some implementations provide methods and systems for re-annotating un-genotyped nodes in pedigree graphs using annotation data of prior graphs.

[0209] Figure 25 shows a flowchart illustrating process 1300 for re-annotating pedigree graphs. Process 1300 involves generating and/or displaying a first pedigree graph. See box 1302. The first pedigree graph in some implementations can be rendered as the graph displayed in Figure 27. The first pedigree graph depicts relationships among a first plurality of individuals. It includes a plurality of genotyped nodes and one or more ungenotyped nodes. Each node represents an individual. Each genotyped node represents an individual whose genetic data have been used to determine the pedigree relationships depicted by the pedigree graph. Each ungenotyped node represents an individual whose genetic data have not been used to determine the pedigree relationships depicted by the pedigree graph. The ungenotyped individual is inferred from the pedigree relationships determined from genetic data.

[0210] Process 1300 also involves receiving annotation data to annotate one or more un-genotyped nodes of the first pedigree graph. See box 1304. Process 1300 also involves

displaying the first pedigree graph with the one or more un-genotyped nodes annotated. See box 1306.

[0211] Figures 27 – 30 show example implementations of displaying a pedigree graph and receiving user input for annotating un-genotyped nodes of the pedigree graph. A user may provide user input to annotate any one of the ungenotyped nodes. Figure 28 shows that in some implementations, the user may select an ungenotyped node 1406 in a graphical user interface. After the ungenotyped node 1406 is selected, some implementations automatically highlight the focal node, the selected node and intermediate nodes through which the first two are related. In this example, when grandparent node 1406 is selected, focal node 1402, parent node 1404, and grandparent node 1406 are highlighted.

[0212] After annotation data are provided by the user, the first pedigree graph is displayed with grandparent node 1406 annotated as shown in Figure 29. The annotation shown here includes the name of the grandparent and an indication that the grandparent is a grandfather (*e.g.*, based on gender information). In some implementations, the annotation data can be provided through a graphical user interface such as the one shown in Figure 30.

[0213] Returning to Figure 25, process 1300 further involves generating a second pedigree graph. See box 1308. The second pedigree graph also includes a plurality of genotyped nodes and one or more ungenotyped nodes. The individuals represented by the first pedigree graph and the second pedigree graph overlap.

[0214] Process 1300 further involves matching one or more annotated, ungenotyped nodes of the first pedigree graph respectively with one or more corresponding nodes of the second pedigree graph. See box 1310. Process 1300 involves annotating one or more corresponding nodes of the second pedigree graph respectively using annotation data of their matching nodes of the first pedigree graph. This annotation can be referred to as re-annotating nodes in a second pedigree graph using annotation data of corresponding nodes in the first pedigree graph. See box 1312.

[0215] Process 1300 further involves displaying the second pedigree graph with the

annotated one or more corresponding nodes. See box 1314. Figure 31 includes a pedigree graph labeled as "Tree 2," which is an example of a second pedigree graph with one or more genotyped nodes annotated using annotation data from corresponding nodes in a first pedigree graph labeled as "Tree 1." Tree 2 includes a branch with node 5 that is not in Tree 1.

[**0216**] Figure 26 illustrates process 1350 for re-annotating pedigree graph. Process 1350 includes a subset of steps of process 1300. The operations illustrated in boxes 1352, 1354, 1356, and 1358 of process 1350 correspond to operations in boxes 1304, 1308, 1310, and 1312 in process 1300 in Figure 25. Process 1350 includes further details and limitations than process 1300.

[**0217**] Process 1350 for re-annotating pedigree graph involves receiving annotation data to annotate one or more ungenotyped nodes of the first pedigree graph that depicts relationships among a first plurality of individuals. See box 1352. In some implementations, the annotation data include name, maiden name, gender, date of birth, year of birth, place of birth, ethnicity, living or deceased state, date of birth, date of death, place of death, and/or photographic data.

[**0218**] The first pedigree graph includes a plurality of genotyped nodes and one or more ungenotyped nodes, each node representing an individual. The genotyped node represents an individual whose genetic data have been used to determine the pedigree relationships depicted by the pedigree graph. An ungenotyped node represents an individual whose genetic data have not been used to determine the pedigree relationships depicted by the pedigree graph.

[**0219**] In some implementations, the first pedigree graph is displayed before annotation data is received. The annotation data may be provided by the user with reference to the displayed first pedigree graph.

[**0220**] An example of the first pedigree graph is shown in Figure 31 as Tree 1 in the top half of the figure. In this example, genotyped nodes are labeled by numbers and rendered as squares or circles. Ungenotyped nodes are rendered as diamonds. The ungenotyped

nodes that are annotated are labeled with capital letters.

[0221] Process 1350 also involves generating a second pedigree graph that depicts relationships among the second plurality of individuals. The second pedigree graph includes a plurality of genotyped nodes and one or more ungenotyped nodes, each node representing an individual. See box 1354. In some implementations, the second plurality of individuals includes at least one individual who is not among the first plurality of individuals. In other implementations, the first plurality of individuals includes at least one individual who is not among the second plurality of individuals. In various implementations, the first plurality of individuals and the second plurality of individuals overlap. In some implementations, the second plurality of individuals is identical to the first plurality of individuals, but the relationships among the second plurality of individuals are not identical to the relationships among the first plurality of individuals.

[0222] Process 1350 also involves matching one or more annotated, ungenotyped nodes of the first pedigree graph respectively with one or more corresponding nodes of the second pedigree graph. See box 1356.

[0223] Process 1350 further involves annotating the one or more corresponding nodes of the second pedigree graph respectively using annotation data of their matching nodes of the first pedigree graph. See box 1358. In some implementations, the process further involves displaying the annotated second pedigree graph using a display device.

[0224] In some implementations, matching nodes of the first pedigree graph with nodes of the second pedigree graph includes the steps in the following pedigree matching procedure:

Procedure (pedigree matching).

1. Determine that an individual N in the first pedigree graph matches an individual N in the second pedigree graph;

2. identify, among individuals represented by genotyped nodes in the first pedigree graph, relatives of $P(1,1,N)$ and relatives of $P(1,2,N)$, wherein $P(1,1,N)$ is in the first pedigree graph a first parent of N, $P(1,2,N)$ is in the first pedigree graph

a second parent of N, and the relatives are biologically related and exclude any common direct descendants;

3. identify, among individuals represented by genotyped nodes in the second pedigree graph, relatives of $P(2,1,N)$ and relatives of $P(2,2,N)$, wherein $P(2,1,N)$ is in the second pedigree graph a first parent of N, and $P(2,2,N)$ is in the second pedigree graph a second parent of N, and the relatives are biologically related and exclude any common direct descendants; and

4. a) match node $P(1,1,N)$ with node $P(2,1,N)$ or $P(2,2,N)$ when matching conditions are met, wherein the matching conditions comprise: any identified relatives of $P(1,1,N)$ are also identified relatives of $P(2,1,N)$ or $P(2,2,N)$ respectively, or

   b) match node $P(1,1,N)$ with either node $P(2,1,N)$ or node $P(2,2,N)$ when $P(1,1,N)$, $P(1,2,N)$, $P(2,1,N)$ and $P(2,2,N)$ all have zero identified relatives.

[**0225**] In some implementations, one or more of the following must be met for matching a parent $P(1,1,N)$ of N on Tree 1 to a parent $P(2,1,N)$ of N on Tree 2. In some implementations, all of the following must be met to match.

Conditions (pedigree matching).

1. Any identified relative of $P(1,1,N)$ appearing on both trees is also an identified relative of $P(2,1,N)$;

2. any identified relative of $P(1,2,N)$ appearing on both trees is also an identified relative of $P(2,2,N)$;

3. no identified relative of $P(1,1,N)$ is also an identified relative of $P(2,2,N)$;

4. no identified relative of $P(1,2,N)$ is also an identified relative of $P(2,1,N)$;

5. all shared identified descendants of $P(1,1,N)$ and $P(2,1,N)$ have the same degrees of relationship to $P(1,1,N)$ and $P(2,1,N)$;

6. all shared identified descendants of $P(1,2,N)$ and $P(2,2,N)$ have the same degrees of relationship to $P(1,2,N)$ and $P(2,2,N)$;

7. $P(1,1,N)$ and $P(2,1,N)$ have at least one common identified relative, or $P(1,2,N)$ and $P(2,2,N)$ have at least one common identified relative, or $P(1,1,N)$, $P(1,2,N)$, $P(2,1,N)$, and $P(2,2,N)$ have no identified relatives.

[**0226**] In some implementations, the matching conditions further include: each identified relative of $P(1,1,N)$ who is also an identified relative of $P(2,1,N)$ or $P(2,2,N)$ has a same category of relationships with $P(1,1,N)$ and $P(1,2,N)$ or $P(1,1,N)$ and $P(2,2,N)$ respectively.

[**0227**] In some implementations, each category of relationship is selected from: direct ancestor relationships, direct descendant relationships, and other relationships. In some implementations, each category of relationships corresponds to a degree of relationship or similar degrees of relationships.

[**0228**] In some implementations, a data structure with a relationship dictionary is used to store the relatives of each node of interest, as well as relationships, relationship types, relationship degrees, or relationship categories of the relatives. By querying the dictionary, all relatives of an individual of interest can be determined. In some implementations, the relationship dictionary groups relationships of relatives into three categories: ancestors, descendants, and other relatives. In some implementations, relatives who are related through only marriage are excluded.

[**0229**] In some implementations, both the first pedigree graph and the second pedigree graph include a genotyped node representing a same focal individual, and the individual N is the focal individual.

[**0230**] In some implementations, matching nodes includes repeating steps 1–4 of the pedigree matching procedure one or more times using the matched $P(1,1,N)$ and $P(2,1,N)$ in step 4 as the individual N in step 1.

[**0231**] In some implementations, matching nodes further includes, when matching con-

ditions are not met in step 4 of the pedigree matching procedure, repeating steps 1–4 of the pedigree matching procedure using an individual represented by a genotyped node whose parents have not been matched.

[**0232**] In some implementations, matching nodes further includes, matching a first node on the first pedigree graph with a second node on the second pedigree graph when the partner node of the first node and the partner node of the second node are matched. The partner node of the first node and the partner node of the second node can be matched based on the matching conditions. They can also be matched because they correspond to a same genotyped individual in the database.

[**0233**] A matching process can start by selecting a focal node 2 on both Tree 1 and Tree 2 corresponding to a same focal person that was genotyped. The same genotyped data of the person for the two nodes indicate that the two nodes represent the same individual. The matching process identifies all relatives of node A (a first parent of node 2) of Tree 1, which include node 1 and node 4. The relatives are biologically related and exclude any common direct descendants. The matching process also identifies all relatives of node A (a first parent of node 2) of Tree 2, which include node 1 and node 4. Therefore, pedigree matching conditions 1 and 7 are satisfied.

[**0234**] Moreover, the relatives of E in Tree 1 are node 3, and the relatives of node E in Tree 2 are node 3 and node 5. No relative of node A in Tree 1 (1, 4) is also a relative of node E in Tree 2 (3, 5). Similarly, no relatives of node E in Tree 1 (3) is also a relative of node A in Tree 2 (1, 4). As such, pedigree matching conditions 2, 3, and 4 are satisfied.

[**0235**] Finally, on both Tree 1 and Tree 2, the only descendant of A and E is node 2. As such, pedigree matching conditions 5 and 6 are satisfied. Therefore, because all pedigree matching conditions are satisfied, node A of Tree 1 is matched with node A of Tree 2.

[**0236**] In some implementations, the process matches a first node on the first pedigree graph with a second node on the second pedigree graph when the partner nodes of the first and second nodes are matched. In this example, when the partner of node E in Tree

1 (node A) is matched with the partner of node E in Tree 2 (node A), node E in Tree 1 is also matched with node E of Tree 2.

[**0237**] The matching process in some implementations proceeds to use node A as the focal node and identify the relatives of node B and its partner node as two parent nodes of node A. Because relatives of node B in Tree 1 (1, 4) are also relatives of node B in Tree 2 (1,4), pedigree matching conditions 1 and 7 are satisfied. Because the partner of B has no relatives in either Tree 1 or Tree 2, pedigree matching condition 2 is satisfied. Because no relatives of node B in Tree 1 (1, 4) are also relatives of node B's partner in Tree 2 (none), pedigree matching condition 3 is satisfied. Because no relatives of node B's partner in Tree 1 (none) are also relatives of node B in Tree 2 (1,4), pedigree matching condition 4 is satisfied. Finally, because B and their partner have the same descendants (A and 2) in both Tree 1 and Tree 2, and because the relationships between these descendants and B and B's partner are the same in both trees, pedigree matching conditions 5 and 6 are satisfied. Therefore, node B of Tree 1 is matched with node B of Tree 2.

[**0238**] The matching process in some implementations proceeds to use node B as the focal node and identify the relatives of node C and node D as two parent nodes of node B in Tree 1 and node 1502 and node 1504 as two parent nodes of node B in Tree 2. Because a relative of node C in Tree 1 (1) is one of the relatives of node 1502 in Tree 2 (1,4), pedigree matching condition 7 is satisfied. Because no relative of node C in Tree 1 (1) is also a relative of node 1504 in Tree 2 (none), pedigree matching condition 3 is satisfied. However, because a relative of node D in Tree 1 (4) is also one of the relatives of node 1502 in Tree 2 (1,4), pedigree matching condition 4 is not satisfied. Therefore, node C of Tree 1 cannot be matched with node 1502 of Tree 2. Similarly, node D of Tree 1 cannot be matched with node 1502 of Tree 2. Also, neither node C nor node D of Tree 1 cannot be matched with node 1504 of Tree 2. Therefore, node C and node D of Tree 1 cannot be matched with any nodes of Tree 2.

[**0239**] The matching process in some implementations proceeds to use node E as the focal node and identify node G and node F as two parent nodes of node E. Because the relatives of node G in Tree 1 (none) and the relatives of node G in Tree 2 (none)

are the same, pedigree matching condition 1 is satisfied. Because the relative of node F in Tree 1 (3) is also one of the relatives of node F in Tree 2 (3,5), pedigree matching conditions 2 and 7 are satisfied. Because no relatives of node F in Tree 1 (3) are also relatives of node G in Tree 2 (none), pedigree matching condition 3 is satisfied. Moreover, because no relatives of node G in Tree 1 (none) are also the relatives of node F in Tree 2 (3,5), pedigree matching condition 4 is satisfied. Finally, because E and 2 are the only descendants of F and G in both Tree 1 and Tree 2, and because these descendants have the same degrees of relationship to F and G in both trees, pedigree matching conditions 5 and 6 are satisfied. Therefore, node F of Tree 1 is matched with node F of Tree 2. Also, their partner nodes, node G on Tree 1 and node G on Tree 2 are matched.

[**0240**] In some implementations, when two parent nodes in Tree 1 and two parent nodes in Tree 2 all have zero relatives, either parent node in Tree 1 can be matched to either parent node of Tree 2. In these implementations, node J and node K in Tree 1 and node J and node K in Tree 2 all have zero relatives. Note that node 2 is not a relative, because it is a common descendant that is excluded. Moreover, the identified descendants of J and K are 2, E, and G and these descendants have the same degrees to J and K in both Tree 1 and Tree 2. Thus, all criteria are satisfied and either node J or node K in Tree 1 can be matched to either node J or node K in Tree 2.

[**0241**] However, node I and node H in Tree 1 cannot be matched to nodes 1506 or node 1508 in Tree 2, because node I and node H in Tree 1 have no relatives, but 1509 in Tree 2 has relative node 5. Therefore, pedigree matching condition 7 is not satisfied.

[**0242**] Pseudocode for matching nodes on two pedigree graphs for reannotation is provided below.

Apparatus and Systems

[**0243**] Figure 17 is a functional diagram illustrating a programmed computer system for performing processes for determining or displaying pedigrees in accordance with some implementations. Computer system 100, which includes various subsystems as described

if i_1 or i_2 was inferred to be un-matchable on a previous step **then**
    Return
end if

$p_{1,1}, p_{1,2}$ = Get parent ids of i_1 in pedigree 1
$p_{2,1}, p_{2,2}$ = Get parent ids of i_2 in pedigree 2

relSet($p_{1,1}$) = Get non-descendant relatives of parent 1 of i_1
relSet($p_{1,2}$) = Get non-descendant relatives of parent 2 of i_1
relSet($p_{2,1}$) = Get non-descendant relatives of parent 1 of i_2
relSet($p_{2,2}$) = Get non-descendant relatives of parent 2 of i_2

descSet($p_{1,1}$) = Get descendants of parent 1 of i_1
descSet($p_{1,2}$) = Get descendants of parent 2 of i_1
descSet($p_{2,1}$) = Get descendants of parent 1 of i_2
descSet($p_{1,2}$) = Get descendants of parent 2 of i_2

relCheck($p_{1,1}, p_{2,1}$) = Check conditions 1–4 and 7 of the pedigree matching conditions for matching parent 1 of i_1 and parent 1 of i_2
relCheck($p_{1,2}, p_{2,1}$) = Check conditions 1–4 and 7 of the pedigree matching conditions for matching parent 2 of i_1 and parent 1 of i_2

descCheck($p_{1,1}, p_{2,1}$) = Check conditions 5 and 6 of the pedigree matching conditions for matching parent 1 of i_1 and parent 1 of i_2
descCheck($p_{1,2}, p_{2,1}$) = Check conditions 5 and 6 of the pedigree matching conditions for matching parent 2 of i_1 and parent 1 of i_2

Initialize empty list matchedPairList
if relCheck($p_{1,1}, p_{2,1}$) and descCheck($p_{1,1}, p_{2,1}$) **then**
    Map $p_{1,1}$ to $p_{2,1}$ and map $p_{1,2}$ to $p_{2,2}$
    Append the tuples $(p_{1,1}, p_{2,1})$ and $(p_{1,2}, p_{2,2})$ to matchedPairList
else if relCheck($p_{1,2}, p_{2,1}$) and descCheck($p_{1,2}, p_{2,1}$) **then**
    Map $p_{1,2}$ to $p_{2,1}$ and map $p_{1,1}$ to $p_{2,2}$
    Append the tuples $(p_{1,2}, p_{2,1})$ and $(p_{1,1}, p_{2,2})$ to matchedPairList
else
    Add $p_{1,1}$ and $p_{1,2}$ to mismatch set 1
    Add $p_{2,1}$ and $p_{2,2}$ to mismatch set 2
end if

for (i_1, i_2) in matchedPairList **do**
    Recursively match the parents of i_1 to the parents of i_2
end for

below, includes at least one microprocessor subsystem (also referred to as a processor or a central processing unit (CPU)) 102. For example, processor 102 can be implemented by a single-chip processor or by multiple processors. In some implementations, processor 102 is a general purpose digital processor that controls the operation of the computer system 100. Using instructions retrieved from memory 110, the processor 102 controls the reception and manipulation of input data, and the output and display of data on output devices (e.g., display 118). In some implementations, processor 102 includes and/or is used to provide IBD determination, pedigree determination, pedigree drawing and displaying, etc. as described herein.

[**0244**] Processor 102 is coupled bi-directionally with memory 110, which can include a first primary storage, typically a random access memory (RAM), and a second primary storage area, typically a read-only memory (ROM). As is well known in the art, primary storage can be used as a general storage area and as scratch-pad memory, and can also be used to store input data and processed data. Primary storage can also store programming instructions and data, in the form of data objects and text objects, in addition to other data and instructions for processes operating on processor 102. Also as is well known in the art, primary storage typically includes basic operating instructions, program code, data, and objects used by the processor 102 to perform its functions (e.g., programmed instructions). For example, memory 110 can include any suitable computer-readable storage media, described below, depending on whether, for example, data access needs to be bi-directional or uni-directional. For example, processor 102 can also directly and very rapidly retrieve and store frequently needed data in a cache memory (not shown).

[**0245**] A removable mass storage device 112 provides additional data storage capacity for the computer system 100, and is coupled either bi-directionally (read/write) or uni-directionally (read only) to processor 102. For example, storage 112 can also include computer-readable media such as magnetic tape, flash memory, PC-CARDS, portable mass storage devices, holographic storage devices, and other storage devices. A fixed mass storage 120 can also, for example, provide additional data storage capacity. The most common example of mass storage 120 is a hard disk drive. Mass storage 112, 120 generally store additional programming instructions, data, and the like that typically are

not in active use by the processor 102. It will be appreciated that the information retained within mass storage 112 and 120 can be incorporated, if needed, in standard fashion as part of memory 110 (e.g., RAM) as virtual memory.

[0246] In addition to providing processor 102 access to storage subsystems, bus 114 can also be used to provide access to other subsystems and devices. As shown, these can include a display monitor 118, a network interface 116, a keyboard 104, and a pointing device 106, as well as an auxiliary input/output device interface, a sound card, speakers, and other subsystems as needed. For example, the pointing device 106 can be a mouse, stylus, track ball, or tablet, and is useful for interacting with a graphical user interface.

[0247] The network interface 116 allows processor 102 to be coupled to another computer, computer network, or telecommunications network using a network connection as shown. For example, through the network interface 116, the processor 102 can receive information (e.g., data objects or program instructions) from another network or output information to another network in the course of performing method/process steps. Information, often represented as a sequence of instructions to be executed on a processor, can be received from and outputted to another network. An interface card or similar device and appropriate software implemented by (e.g., executed/performed on) processor 102 can be used to connect the computer system 100 to an external network and transfer data according to standard protocols. For example, various process implementations disclosed herein can be executed on processor 102, or can be performed across a network such as the Internet, intranet networks, or local area networks, in conjunction with a remote processor that shares a portion of the processing. Additional mass storage devices (not shown) can also be connected to processor 102 through network interface 116.

[0248] An auxiliary I/O device interface (not shown) can be used in conjunction with computer system 100. The auxiliary I/O device interface can include general and customized interfaces that allow the processor 102 to send and, more typically, receive data from other devices such as microphones, touch-sensitive displays, transducer card readers, tape readers, voice or handwriting recognizers, biometrics readers, cameras, portable mass storage devices, and other computers.

[0249] In addition, various implementations disclosed herein further relate to computer storage products with a computer readable medium that includes program code for performing various computer-implemented operations. The computer-readable medium is any data storage device that can store data which can thereafter be read by a computer system. Examples of computer-readable media include, but are not limited to, all the media mentioned above: magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as optical disks; and specially configured hardware devices such as application-specific integrated circuits (ASICs), programmable logic devices (PLDs), and ROM and RAM devices. Examples of program code include both machine code, as produced, for example, by a compiler, or files containing higher level code (e.g., script) that can be executed using an interpreter.

[0250] The computer system shown in Figure 17 is but an example of a computer system suitable for use with the various implementations disclosed herein. Other computer systems suitable for such use can include additional or fewer subsystems. In addition, bus 114 is illustrative of any interconnection scheme serving to link the subsystems. Other computer architectures having different configurations of subsystems can also be utilized.

[0251] Figure 18 is a block diagram illustrating an implementation of an IBD-based pedigree services system that provides services based on IBD information, which include but are not limited to relatedness estimation, relative detection, pedigree determination, combination, drawing and display. In this example, a user uses a client device 202 to communicate with an IBD-based pedigree services system 206 via a network 204. Examples of device 202 include a laptop computer, a desktop computer, a smart phone, a mobile device, a tablet device or any other computing device. IBD-based personal genomics services system 206 is used to perform a pipelined process to determine and or display pedigrees based on users IBD information. IBD-based pedigree services system 206 can be implemented on a networked platform (e.g., a server or cloud-based platform, a peer-to-peer platform, etc.) that supports various applications. For example, implementations of the platform perform pedigree determination or display and provide users with access (e.g., via appropriate user interfaces) to their personal genetic information (e.g., genetic sequence information and/or genotype information obtained by assaying genetic

materials such as blood or saliva samples) and predicted ancestry information. In some implementations, the platform also allows users to connect with each other and share information. Device 100 can be used to implement 202 or 206.

[**0252**] In some implementations, DNA samples (e.g., saliva, blood, etc.) are collected from genotyped individuals and analyzed using DNA microarray or other appropriate techniques. The genotype information is obtained (e.g., from genotyping chips directly or from genotyping services that provide assayed results) and stored in database 208 and is used by system 206 to make ancestry predictions or pedigree determination. Reference data, including genotype data of reference individuals, simulated data (e.g., results of machine-based processes that simulate biological processes such as recombination of parents' DNA), pre-computed data (e.g., a precomputed reference haplotype data used in phasing and model training) and the like can also be stored in database 208 or any other appropriate storage unit.

Experimental

Accuracy of the Likelihood and Generalized DRUID Estimators

[**0253**] This experiment shows the accuracy of the likelihood and generalized DRUID estimators (Equations 21 and 24) for inferring the degree of relationship between two distantly-related pedigrees.

[**0254**] Practitioners applied these estimators to infer the degree between common ancestors A_1 and A_2 of two small pedigrees $\mathcal{P}_1$ and $\mathcal{P}_2$. For this analysis, two identical small pedigrees $\mathcal{P}_1$ and $\mathcal{P}_2$ were simulated. Each small pedigree had the same topology comprised of the common ancestor A_1 or A_2, their spouse, their two children, and four grandchildren, where the grandchildren were comprised of two children for each child of A_1 or A_2. The ancestors A_1 and A_2 were then connected by degree d through a pair of common ancestors, where the degree d varied from 1 to 10.

[0255] Figure 6 shows the accuracy of the likelihood estimator d_L (Equation 21) and the generalized DRUID estimator d_D (Equation 24) for inferring the degree d, conditional on the event that any IBD at all is observed between the leaf nodes in $\mathcal{P}_1$ and $\mathcal{P}_2$. From Figure 6 it can be seen that both the maximum likelihood estimator d_L and the generalized DRUID estimator d_D have similar accuracies for inferring the degree d. Moreover, the generalized DRUID estimate is nearly identical to the maximum likelihood estimate, which is important in practice because it implies that connecting two pedigrees through the degree inferred by DRUID results in a pedigree that is approximately the maximum likelihood pedigree. This result can dramatically speed up pedigree inference and, in practice, practitioners use the generalized DRUID estimator for inferring the degree of separation between two small pedigrees.

Accuracy and Runtime

[0256] This experiment compares accuracy and computer runtime between a computer implemented method referred to as Bonsai according to some implementation with a method based on the PRIMUS method.

[0257] The experiment is conducted to evaluate 204 pedigrees of 23andMe customers. Pedigrees were chosen in which each nuclear family had at least two genotyped offspring and two genotyped parents. Pedigrees spanned at least two generations.

[0258] To evaluate the accuracy of pedigree inference, practitioners subsampled these pedigrees, sampling 10, 20, 30, 40, or 50% of their members uniformly at random without replacement. The subsampled individuals were then used to reconstruct the pedigrees using PRIMUS and Bonsai.

[0259] The Bonsai and PRIMUS methods were applied to exactly the same pedigree subsets. The x-axis in Figure 19 is the fraction of sampled individuals in each of the 204 pedigrees that were used for pedigree inference. The y-axis is the average computer runtime in seconds required to reconstruct a pedigree. Bonsai is considerably faster than PRIMUS for pedigrees in which an intermediate number of individuals have been sampled.

[0260] Bonsai and PRIMUS have similar runtimes when few or many individuals are sampled because there are many fewer possible pedigrees to explore. This suggests that Bonsai gains considerable computational efficiency by ignoring very low likelihood pedigrees. The experiment illustrates that due to the computational efficiency of Bonsai, it runs faster on computers and conserves computer resources compared to prior art methods.

[0261] Figures 20 – 24 illustrate accuracies of Bonsai compared with PRIMUS, computed on 204 validated customer pedigrees in which all individuals were consented for research. Pedigrees were formed from quartets of individuals containing at least two genotyped offspring and two genotyped parents that could be used to triangulate and validate the other relationships in the quartet. Pedigrees spanned at least two generations, containing at least one pair of first cousins. Because pedigree inference becomes increasingly difficult as the fraction of sampled individuals in a pedigree decreases, we downsampled individuals in each pedigree to create sparse pedigrees in which relationships were known. We considered pedigrees in which 10, 20, 30, 40, and 50 percent of individuals were sampled. The subsampled individuals were then used to reconstruct the pedigrees using PRIMUS and Bonsai. The Bonsai and PRIMUS methods were applied to the same pedigree subsets and the same set of inferred pairwise likelihoods. Figure 20A shows the accuracy of Bonsai for inferring different relationships when 10% of individuals were sampled from each pedigree. Tuples are of the form (up, down, number of ancestors) where up is the number of meioses from the first individual to their common ancestor(s) with the second, down is the number of meioses from the common ancestor(s) to the second individual, and the number of ancestors is the number of most recent common ancestors. The vertical axis indicates the true degrees of relationship, while the horizontal axis indicates the inferred degrees of the relationship. Different levels of shade indicate portions of a particular actual relationship determined by the model as various inferred relationships. Perfect predictions should have inferred relationships 100% on the identity line.

[0262] Figure 20B shows the accuracy for PRIMUS applied to the same individuals with the same pairwise likelihoods as Figure 20A. Comparing Figures 20A and 20B, one

sees that Bonsai has more accurate predictions than PRIMUS when 10% of individuals are sampled.

[**0263**] Similarly, Figures 21 – 24 compare Bonsai with PRIMUS for increasingly large subsamples. Bonsai consistently has more accurate predictions than PRIMUS at different levels of sampling. When the fraction of sampled individuals was at least 50%, Bonsai had perfect accuracy, whereas PRIMUS continued to infer some relationships incorrectly.

CLAIMS

What is claimed:

1. A method, implemented using a computer system that includes one or more processors and system memory, for determining pedigree relationships among a plurality of genetically related individuals, the method comprising:

a) identifying, among the plurality of genetically related individuals, a closest relative of a starting individual using genetic data of the plurality of genetically related individuals;

b) applying pairwise Identity-by-Descent (IBD) data and pairwise age data of the starting individual and the closest relative to a probabilistic relationship model to obtain various likelihoods of various potential relationships between the starting individual and the closest relative;

c) selecting one or more potential relationships between the starting individual and the closest relative that have relationship likelihoods meeting a relationship criterion, and forming a pedigree from each of the one or more potential relationships;

d) identifying, among genetically related individuals not included in pedigrees already formed, a closest relative of any individual already in a pedigree;

e) applying pairwise IBD data and pairwise age data of the closest relative and the individual already in the pedigree to the probabilistic relationship model to obtain various likelihoods of various potential relationships between the closest relative and the individual already in the pedigree;

f) selecting one or more potential relationships between the closest relative and the individual already in the pedigree that have relationship likelihoods meeting the relationship criterion, and adding each of the one or more potential

relationships to each pedigree already formed to grow each pedigree into one or more growing pedigrees;

g) selecting growing pedigrees that have pedigree likelihoods meeting a pedigree criterion as the pedigrees already formed; and

h) repeating (1d)-(1g) one or more times,

wherein (1a)–(1h) are performed by the computer system.

2. The method of claim 1, wherein the pairwise IBD data comprise a length of IBD segments.

3. The method of claim 2, wherein the lengths of IBD segments comprise a length of full IBD segments and/or a length of half IBD segments.

4. The method of any one of the preceding claims, wherein the pairwise IBD data comprise a number of IBD segments.

5. The method of claim 4, wherein the number of IBD segments comprise a number of full IBD segments and/or a number of half IBD segments.

6. The method of any one of the preceding claims, wherein the probabilistic relationship model is a machine-learning model.

7. The method of any one of the preceding claims, wherein the probabilistic relationship model models the probability distribution of the pairwise IBD data for each relationship and/or the probability distribution of the pairwise age data for each relationship as a Gaussian distribution, a Poisson distribution, or an exponential distribution.

8. The method of any one of the preceding claims, further comprising: storing into a database or retrieving from a database relationship data of a pedigree having the highest pedigree likelihood among the growing pedigrees selected in (1g).

9. The method of claim 8, further comprising:

a) generating a pedigree graph using the relationship data of the pedigree having the highest pedigree likelihood; and

b) displaying the pedigree graph on a display device.

10. The method of any one of the preceding claims, wherein the various relationships comprise relationships of the 0th, 1st, 2nd, 3rd, 4th, 5th, 6th, 7th, 8th, 9th, 10th, 11th, 12th, 13th, 14th, or 15^{th}degree or further.

11. The method of any one of the preceding claims, wherein the various relationships comprise relationships of at least 4 meioses on a common-ancestor path.

12. The method of any one of the preceding claims, wherein the various relationships comprise relationships having a coefficient of relationship smaller than 0.25.

13. The method of any one of the preceding claims, wherein the various relationships comprise two or more relationships of a same degree or of a same coefficient of relationship.

14. The method of any one of the preceding claims, wherein every pair of individuals in the plurality of genetically related individuals has a total IBD length larger than an IBD threshold.

15. The method of claim 14, wherein the IBD threshold is 50 centimorgan (cM).

16. The method of claim 14, wherein the IBD threshold is 10 cM.

17. The method of claim 14, wherein the total IBD length is adjusted by subtracting background IBD from the pairwise IBD data.

18. The method of claim 14, wherein likelihoods of observing IBD shared among one or more individuals are computed and IBD segments that are likely to be background IBD are removed.

19. The method of any one of the preceding claims, wherein the starting individual is an individual of interest.

20. The method of any one of the preceding claims, wherein pairwise IBD data between two individuals are used to determine how closely related the two individuals are.

21. The method of any one of the preceding claims, wherein the relationship criterion is a ratio of an instant relationship likelihood over a maximum relationship likelihood being larger than a value c.

22. The method of any one of the preceding claims, wherein the pedigree criterion is a ratio of an instant pedigree likelihood over a maximum pedigree likelihood being larger than a specific value d.

23. The method of any one of the preceding claims, wherein (1h) comprises repeating (1d)–(1g) until all individuals of the plurality of genetically related individuals have been identified as a closest relative or excluded from the pedigree.

24. The method of any of the preceding claims, further comprising:

    c) identifying two pedigrees from among a plurality of pedigrees constructed using operations (1a)–(1h), the two pedigrees being a genealogically closest pair of pedigrees among all pairs in the plurality pedigrees.

25. The method of claim 24, wherein in operation (24c) a genealogical similarity between two pedigrees that are the genealogically closest pair of pedigrees is measured as a union over all IBD segments shared between an individual in a first pedigree in the genealogically closest pair of pedigrees and an individual in a second pedigree in the genealogically closest pair of pedigrees.

26. The method of claim 24 or 25, further comprising:

    d) combining the two pedigrees identified in operation (24c) into a combined pedigree.

27. The method of claim 26, further comprising:

    e) repeating operations (24c) and (26d) to agglomerate a plurality of pedigrees into a large pedigree.

28. The method of claim 26, wherein in operation (26d) the two pedigrees are combined by:

a) identifying a first set of individuals in a first pedigree who share IBD with individuals in a second pedigree;

b) identifying a second set of individuals in the second pedigree who share IBD with individuals in the first pedigree;

c) identifying a common ancestor of the first set of individuals;

d) identifying a common ancestor of the second set of individuals;

e) inferring a degree of relatedness between the common ancestor of the first set and the common ancestor of the second set; and

f) connecting the two common ancestors by an inferred degree of relatedness between the common ancestors.

29. The method of claim 28, wherein only IBD segments longer than a predetermined threshold are considered when determining whether two individuals share IBD.

30. The method of claim 29, wherein the predetermined threshold is at least 5 cM.

31. The method of claim 28, wherein a likelihood is used to determine whether the observed IBD shared between two individuals is due to recent relatedness or to background IBD.

32. The method of claim 28, wherein the common ancestor of the plurality of individuals in the first pedigree who share IBD with any member of the plurality of individuals in the second pedigree is identified as the most recent common ancestor of these individuals in the pedigree constructed using the steps (1a)-(1h).

33. The method of claim 32, wherein selecting the common ancestor involves:

a) identifying a most recent common ancestor of the first set of individuals;

b) identifying a set relatives of the common ancestor who are not descendants;

c) computing IBD segments between the set of relatives and the first set of individuals (28a);

d) computing IBD segments between the second set of individuals (28b) and the first set of individuals (28a);

e) identifying any overlapping IBD segments between segments (33c) and segments (33d);

f) rejecting the common ancestor if the total overlap of segments identified is greater than a fraction f of the total length of all segments.

34. The method of claim 33, wherein the fraction f is at least 0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, 0.09, 0.10, 0.11, 0.12, 0.13, 0.14, 0.15, or higher.

35. The method of claim 28, wherein the inferred degree of relatedness between the common ancestor of the first set and the common ancestor of the second set is inferred using a likelihood applied to a total amount of IBD shared between the descendants of the common ancestor in the first pedigree and the descendants of the common ancestor in the second pedigree.

36. A method, implemented using a computer system that includes one or more processors and system memory, for combining two or more pedigrees into a larger pedigree, the method comprising:

a) identifying a first pedigree and a second pedigree from among a plurality of pedigrees, the two pedigrees being a genealogically closest pair of pedigrees among all pairs in the plurality pedigrees;

b) identifying a first set of individuals in the first pedigree who share IBD with individuals in the second pedigree;

c) identifying a second set of individuals in the second pedigree who share IBD with individuals in the first pedigree;

d) identifying a common ancestor of the first set of individuals;

e) identifying a common ancestor of the second set of individuals;

f) inferring a degree of relatedness between the common ancestor of the first set and the common ancestor of the second set; and

g) connecting the two common ancestors by an inferred degree of relatedness between the common ancestors to form a larger pedigree.

37. The method of claim 36, further comprising repeating (36a)–(36g) for one or more times for one or more pairs of pedigrees.

38. A method, implemented using a computer system that includes one or more processors and system memory, for training a probabilistic relationship model for predicting genetic relationships based on Identity-by-Descent (IBD) data and age data, the method comprising:

a) receiving, by the computer system from a database, IBD data of a plurality of training sets for a plurality of relationships, wherein each training set comprises one or more pairs of individuals, each pair of individuals being genetically related with a same relationship, each training set is associated with a unique relationship, and the IBD data of each training set comprise pairwise IBD data of individuals in the training set;

b) obtaining, by the computer system, age data of the plurality of training sets for the plurality of relationships, wherein the age data of each training set comprise pairwise age data of individuals in the training set; and

c) training, by the computer system, the probabilistic relationship model using the IBD data of the plurality of training sets and the age data of the plurality of training sets, wherein the trained probabilistic relationship model is configured to take as input pairwise IBD data and pairwise age data for two test individuals and provide as output various likelihoods of various potential relationships for the two test individuals.

39. The method of claim 38, wherein the pairwise IBD data comprise a length of IBD segments.

40. The method of claim 39, wherein the lengths of IBD segments comprise a length of full IBD segments and/or a length of half IBD segments.

41. The method of any one of claims 38–40, wherein the pairwise IBD data comprise a number of IBD segments.

42. The method of claim 41, wherein the number of IBD segments comprise a number of full IBD segments and/or a number of half IBD segments.

43. The method of any one of claims 38–42, wherein the probabilistic relationship model is a machine-learning model.

44. The method of any one of claims 38–43, wherein the probabilistic relationship model models the probability distribution of the pairwise IBD data for each relationship and/or the probability distribution of the pairwise age data for each relationship as a Gaussian distribution, a Poisson distribution, or an exponential distribution.

45. The method of claim 44, wherein training the probabilistic relationship model comprises using a method-of-moments technique to fit the probability distribution to training data.

46. The method of any one of claims 38–45, wherein each likelihood of each potential relationship for the two individuals is a composite likelihood obtained by combining a likelihood given a length of IBD segment, a likelihood given a number of IBD segments, and a likelihood given an age difference between the two individuals.

47. The method of any one of claims 38–46, wherein the plurality of relationships comprise relationships of the 4th, 5th, 6th, 7th, 8th, 9th, 10th, 11th, 12th, 13th, 14th, or 15^{th}degree or further.

48. The method of any one of claims 38–47, wherein the various relationships comprise relationships of at least 4 meioses on a common-ancestor path.

49. The method of any one of claims 38–48, wherein the various relationships comprise relationships having a coefficient of relationship smaller than 0.25.

50. The method of any one of claims 38–49, wherein each training set comprises hundreds of individuals.

51. The method of any one of claims 38–50, wherein each training set comprises thousands of individuals.

52. The method of any one of claims 38–51, wherein each training set comprises tens of thousands of individuals.

53. The method of any one of claims 38–52, wherein each training set comprises data of actual individuals.

54. The method of any one of claims 38–53, wherein each training set comprises data of simulated individuals.

55. A method, implemented using a computer system that includes one or more processors and system memory, for generating a pedigree graph for a plurality of genetically related individuals, the method comprising:

a) receiving from a database pedigree relationship data of the plurality of genetically related individuals, wherein the pedigree relationship data comprise a plurality of entries representing the plurality of genetically related individuals, each entry comprising links relating each child entry to its parent entries, wherein a child entry and its parent entries represent a child and its biological parents;

b) determining a minimal set of root entries from which all of the plurality of entries are reachable by starting from the minimal set of root entries and traversing from parent entries to child entries and traversing between partner entries, wherein partner entries are two parent entries of a same child entry, and wherein a root entry is an entry having no parent entries;

c) for each root entry of the minimal set of root entries, forming a subtree using the root entry as a root node and entries reachable from the root entry as additional nodes, thereby obtaining a plurality of subtrees;

d) positioning nodes in each of the plurality of subtrees; and

e) merging the plurality of subtrees to form the pedigree graph for the plurality of genetically related individuals,

wherein 55a–55e are performed using the one or more processors.

56. The method of claim 55, wherein positioning nodes in each of the subtrees comprises starting from the root node and recursively going through nodes in a defined order to reach nodes to be positioned.

57. The method of claim 56, wherein the defined order is:

a) children of partners to the left;

b) children of partners to the right;

c) partners to the left;

d) self; and

e) partners to the right.

58. The method of any one of claims 55–57, wherein positioning nodes in each of the plurality of subtrees comprises:

a) placing a leaf node without any siblings on the left at an origin;

b) placing a leaf node with a sibling immediately to its left at a position immediately to the right of said sibling;

c) positioning parent nodes relative to their child nodes so that they are further from the horizontal center of a row of nodes than their child nodes are.

59. The method of claim 58, wherein the parent nodes are positioned further from the horizontal center of a row of nodes than their child nodes' partners are.

60. The method of claim 58, wherein

a) there are more than two parent nodes in a row,

b) a first parent node is not on either end of the row,

c) a second parent node is immediately to the left of the first parent node, and

d) the positioning nodes other than leaf nodes comprises positioning the first parent node relative to the child nodes of the second parent node so that the first parent node is to the right of the child nodes of the second parent node.

61. The method of claim 60, further comprises: shifting the child nodes of the first parent node to maintain previous relative positions between the first parent node and the child nodes of the first parent node.

62. The method of any one of claims 55–61, wherein merging the subtrees to form the pedigree graph comprises:

a) horizontally shifting one or more of the subtrees so that non-corresponding nodes of different subtrees do not overlap, wherein non-corresponding nodes are nodes on different trees that do not represent the same individual; and

b) merging corresponding nodes on different subtrees into one node, wherein corresponding nodes are nodes on different subtrees that represent the same individual.

63. The method of claim 62, wherein merging the subtrees to form the pedigree graph comprise:

a) identifying core nodes that comprise a focal node representing a focal individual, any sibling nodes representing siblings of the focal individual, any parent nodes representing parents of the focal individual, any descendant nodes representing descendants of the focal individual, and any nodes representing partners of the descendants;

b) merging each pair of four pairs of subtrees to form four grand-parent subtrees, wherein the four pairs of subtrees respectively correspond to four pairs of great grandparents of the focal individual, and wherein merging each pair of subtrees comprises:

i. horizontally shifting one subtree so that non-corresponding nodes of the two subtrees do not overlap, and

ii. merging two corresponding nodes on the pair of subtrees representing a same grandparent into one node; and

c) merging the four grand-parent subtrees to form the pedigree graph by:

iii. horizontally shifting one or more of the grand-parent subtrees so that non-corresponding nodes of different grand-parent subtrees do not overlap; and

iv. merging corresponding core nodes on different grand-parent subtrees representing a same individual into one node.

64. The method of claim 63, wherein a subtree corresponding to a great grandparent is obtained by merging two or more subtrees.

65. The method of claim 63, wherein the two grand-parent subtrees in the middle are smaller than the two grand-parent subtrees on the outside.

66. The method of claim 63, further comprising merging subtrees that do not intersect with the core nodes into subtrees that do intersect with the core nodes.

67. The method of any one of claims 55–66, further comprising swapping positions of two partner nodes whose parental subtrees are on the opposite sides from the partner nodes.

68. The method of any one of claims 55–67, further comprising removing empty spaces in the pedigree graph.

69. The method of any one of claims 55–68, further comprising applying force-directed graph drawing techniques to redraw one or more nodes and lines connecting them.

70. The method of claim 69, wherein the one or more nodes comprise leaf nodes and their parent nodes.

71. The method of any one of claims 55–70, wherein each pair of two or more pairs of parent nodes comprises two nodes rendered in different colors.

72. The method of any one of claims 55–71, wherein the lines connecting each child node to its parent nodes comprise curved lines.

73. A method, implemented using a computer system comprising a processor, system memory, and a display device, for displaying a pedigree graph for a plurality of genetically related individuals on a graphical user interface (GUI), the method comprising:

a) displaying, using the display device, the pedigree graph comprising a plurality of nodes representing the plurality of genetically related individuals and lines

connecting each child node to its pair of parent nodes, wherein a child node and its pair of parent nodes represent a child and its pair of parents; wherein

i. each pair of two or more pairs of parent nodes comprises two nodes rendered in different colors, and

ii. the lines connecting each child node to its pair of parent nodes comprise curved lines.

74. The method of claim 73, wherein all nodes of the two or more pairs of parent nodes are rendered in different colors.

75. The method of claim 74, wherein the two or more pairs of parent nodes are direct ancestors of a focal node.

76. The method of claim 75, wherein relative nodes that are not direct ancestors of the focal node have the same coloring as a direct ancestor that is on the same family side and at the same generational level as the relative nodes.

77. The method of claim 74, wherein intra-pair color differences are smaller than inter-pair color differences for the two or more pairs of parent nodes.

78. The method of any one of claims 73–77, wherein nodes at one generation level are rendered in one or more colors that are different from colors of nodes at another generation level.

79. The method of any one of claims 73–78, wherein the pedigree graph comprises:

a) core nodes that comprise a focal node representing a focal individual, any sibling nodes representing siblings of the focal individual, any parent nodes representing parents of the focal individual, any descendant nodes representing descendants of the focal individual, and any nodes representing partners of the descendants.

80. The method of claim 79, wherein the pedigree graph comprises lines indicating direct ancestry of the focal individual, and wherein the lines indicating the direct ancestry of the focal individual are rendered in a color or shade that is different from lines not indicating the direct ancestry of the focal individual.

81. The method of any one of claims 73–80, wherein at least one pair of parent nodes has an off-center alignment relative to their child nodes.

82. The method of any one of claims 73–81, wherein two subtrees having the same topology in the pedigree graph are represented in different forms.

83. The method of any one of claims 73–82, wherein at least one pair of parent nodes has an inter-pair physical distance larger than the smallest possible distance.

84. The method of any one of claims 73–83, wherein one or more nodes and lines connecting them are drawn using force-directed graph drawing techniques.

85. The method of claim 84, wherein the one or more nodes comprise leaf nodes and their parent nodes.

86. The method of any one of claims 73–85, wherein each child node is connected through a curved line to a straight line connecting the child node's parent nodes.

87. The method of any one of claims 73–86, further comprising,

a) receiving a user input; and

b) updating one or more elements of the pedigree graph based on the user input.

88. The method of claim 87, wherein the user input is received via a user interaction with the GUI or the pedigree graph in the GUI.

89. The method of claim 88, wherein the user interaction comprises clicking an interactive node in the pedigree, wherein an interactive node is a node that is configured to receive user input and present output data modifiable by the user input.

90. The method of claim 89, wherein the user interaction comprises entering data using a window activated by clicking the interactive node.

91. The method of claim 89, wherein the updating one or more elements of the pedigree graph comprises changing a relationship between the interactive node and at least one other node in the pedigree.

92. The method of claim 87, wherein at least one interactive node in the pedigree is associated with information of health, trait, disease, physical conditions, or phenotypes of an individual represented by the at least one interactive node.

93. The method of claim 92, wherein the at least one interactive node is associated with a graphical element representing the information of health, trait, disease, physical conditions, or phenotypes of the individual.

94. The method of claim 93, wherein the user input comprises clicking the interactive node.

95. The method of claim 94, wherein the user input comprises entering inforamtion of the individual in a window activated by clicking the interactive node.

96. The method of claim 95, wherein the updating the one or more elements of the pedigree graph comprises changing at least one relationship between two nodes.

97. The method of claim 96, wherein the updating the one or more elements of the pedigree graph comprises changing the graphical element representing the information of health, trait, disease, physical conditions, or phenotypes of the individual.

98. A system for determining pedigree relationships among a plurality of genetically related individuals, the system comprising a processor and one or more computer-readable storage media having stored thereon instructions for execution on said processor to:

a) identify, among a plurality of genetically related individuals, a closest relative of a starting individual;

b) apply pairwise Identity-by-Descent (IBD) data and pairwise age data of the starting individual and the closest relative to a probabilistic relationship model to obtain various likelihoods of various potential relationships between the starting individual and the closest relative;

c) select one or more potential relationships between the starting individual and the closest relative that have relationship likelihoods meeting a relationship

criterion, and forming a pedigree from each of the one or more potential relationships;

d) identify, among genetically related individuals not included in pedigrees already formed, a closest relative of any individual already in a pedigree;

e) apply pairwise IBD data and pairwise age data of the closest relative and the individual already in the pedigree to the probabilistic relationship model to obtain various likelihoods of various potential relationships between the closest relative and the individual already in the pedigree;

f) select one or more potential relationships between the closest relative and the individual already in the pedigree that have relationship likelihoods meeting the relationship criterion, and adding each of the one or more potential relationships to each pedigree already formed to grow each pedigree into one or more growing pedigrees;

g) select growing pedigrees that have pedigree likelihoods meeting a pedigree criterion as the pedigrees already formed; and

h) repeat 98d–98g one or more times.

99. A computer program product comprising one or more computer-readable non-transitory storage media having stored thereon computer-executable instructions that, when executed by one or more processors of a computer system, cause the computer system to:

a) identify, among a plurality of genetically related individuals, a closest relative of a starting individual;

b) apply pairwise Identity-by-Descent (IBD) data and pairwise age data of the starting individual and the closest relative to a probabilistic relationship model to obtain various likelihoods of various potential relationships between the starting individual and the closest relative;

c) select one or more potential relationships between the starting individual and the closest relative that have relationship likelihoods meeting a relationship

criterion, and forming a pedigree from each of the one or more potential relationships;

d) identify, among genetically related individuals not included in pedigrees already formed, a closest relative of any individual already in a pedigree;

e) apply pairwise IBD data and pairwise age data of the closest relative and the individual already in the pedigree to the probabilistic relationship model to obtain various likelihoods of various potential relationships between the closest relative and the individual already in the pedigree;

f) select one or more potential relationships between the closest relative and the individual already in the pedigree that have relationship likelihoods meeting the relationship criterion, and adding each of the one or more potential relationships to each pedigree already formed to grow each pedigree into one or more growing pedigrees;

g) select growing pedigrees that have pedigree likelihoods meeting a pedigree criterion as the pedigrees already formed; and

h) repeat 99d–99g one or more times.

100. A method, implemented using a computer system comprising a processor and system memory, of generating pedigree graphs, the method comprising:

a) receiving, by the processor, annotation data to annotate one or more ungenotyped nodes of a first pedigree graph that depicts relationships among a first plurality of individuals, wherein the first pedigree graph comprises a plurality of genotyped nodes and one or more ungenotyped nodes, each node representing an individual, a genotyped node represents an individual whose genetic data have been used to determine the pedigree relationships depicted by the pedigree graph, and an ungenotyped node represents an individual whose genetic data have not been used to determine the pedigree relationships depicted by the pedigree graph,

b) generating, using the processor, a second pedigree graph that depicts relationships among a second plurality of individuals, wherein the second pedigree

graph comprises a plurality of genotyped nodes and one or more ungenotyped nodes, each node representing an individual;

c) matching, using the processor, one or more annotated, ungenotyped nodes of the first pedigree graph respectively with one or more corresponding nodes of the second pedigree graph;

d) annotating, using the processor, the one or more corresponding nodes of the second pedigree graph respectively using annotation data of their matching nodes of the first pedigree graph.

101. The method of claim 100, wherein (100c) comprises:

a) determining that an individual N in the first pedigree graph matches an individual N in the second pedigree graph;

b) identifying, among individuals represented by genotyped nodes in the first pedigree graph, relatives of $P(1,1,N)$ and relatives of $P(1,2,N)$, wherein $P(1,1,N)$ is in the first pedigree graph a first parent of N, $P(1,2,N)$ is in the first pedigree graph a second parent of N, and the relatives are biologically related to $P(1,1,N)$ or $P(1,2,N)$ and exclude any common direct descendants;

c) identifying, among individuals represented by genotyped nodes in the second pedigree graph, relatives of $P(2,1,N)$ and relatives of $P(2,2,N)$, wherein $P(2,1,N)$ is in the second pedigree graph a first parent of N, and $P(2,2,N)$ is in the second pedigree graph a second parent of N, and the relatives are biologically related to $P(1,1,N)$ or $P(1,2,N)$ and exclude any common direct descendants; and

d) i. matching node $P(1,1,N)$ with node $P(2,1,N)$ when matching conditions are met, wherein the matching conditions comprise: any identified relatives of $P(1,1,N)$ are also identified relatives of $P(2,1,N)$, or

ii. matching node $P(1,1,N)$ with node $P(2,1,N)$ when $P(1,1,N)$, $P(1,2,N)$, $P(2,1,N)$ and $P(2,2,N)$ all have zero identified relatives.

102. The method of claim 101, wherein the matching conditions further comprise:

a) no identified relatives of $P(1,1,N)$ are also identified relatives of $P(2,2,N)$; and

b) no identified relatives of $P(1,2,N)$ are also identified relatives of $P(2,1,N)$.

103. The method of claim 101 or 102 wherein the matching conditions further comprise:

a) $P(1,1,N)$ has at lest one identified relative who is also a relative of $P(2,1,N)$ or $P(2,2,N)$ respectively; or

b) $P(1,2,N)$ has at lest one identified relative who is also a relative of $P(2,2,N)$ or $P(2,1,N)$ respectively.

104. The method of claim 101, 102, or 103 wherein the matching conditions further comprise:

a) each identified descendant of $P(1,1,N)$ who is also an identified relative of $P(2,1,N)$, has the same degree of relationship to $P(1,1,N)$ and $P(2,1,N)$; and

b) each identified descendant of $P(1,2,N)$ who is also an identified relative of $P(2,2,N)$, has the same degree of relationship to $P(1,2,N)$ and $P(2,2,N)$.

105. The method of claim 101, 102, 103, or 104, wherein the matching conditions further comprise: each identified relative of $P(1,1,N)$ who is also an identified relative of $P(2,1,N)$ has a same category of relationships with $P(2,1,N)$.

106. The method of claim 105, wherein each category of relationships is selected from: direct ancestor relationships, direct descendant relationships, and other relationships.

107. The method of claim 105, wherein each category of relationships corresponds to a degree of relationship or similar degrees of relationships.

108. The method of claim 101, wherein both the first pedigree graph and the second pedigree graph comprise a genotyped node representing a same focal individual, and the individual N is the focal individual.

109. The method of claim 108, further comprising, repeating 101a–101d one or more times using the matched $P(1, 1, N)$ and $P(2, 1, N)$ in 101d as the individual N in 101a.

110. The method of claim 109, further comprising, when matching conditions are not met in 101d, repeating 101a–101d using an individual represented by a genotyped node whose parents have not been matched.

111. The method of any one of claims 100–110, wherein 100c comprises: matching a first node on the first pedigree graph with a second node on the second pedigree graph when the partner node of the first node and the partner node of the second node are matched.

112. The method of any one of claims 100–111, wherein the annotation data comprise data selected from the group consisting of: name, maiden name, gender, date of birth, year of birth, place of birth, ethnicity, living/diseased, date of death, place of death, photographic data, and any combinations thereof.

113. The method of any one of claims 100–112, wherein the second plurality of individuals comprises at least one individual who is not among the first plurality of individuals.

114. The method of any one of claims 100–113, wherein the first plurality of individuals comprises at least one individual who is not among the second plurality of individuals.

115. The method of any one of claims 100–114, wherein the second plurality of individuals is identical to the first plurality of individuals, but the relationships among the second plurality of individuals are not identical to the relationships among the first plurality of individuals.

116. The method of any one of claims 100–115, further comprising displaying, before 100a, the first pedigree graph using a display device.

117. The method of any one of claims 100–116, further comprising displaying, after (d), the second pedigree graph with the annotated one or more corresponding nodes using a display device.

118. A system comprising a processor and one or more computer-readable storage media having stored thereon instructions for execution on said processor to perform any one of the proceeding methods.

119. A computer program product comprising one or more computer-readable non-transitory storage media having stored thereon computer-executable instructions that, when executed by one or more processors of a computer system, cause the computer system to perform any one of the proceeding methods.

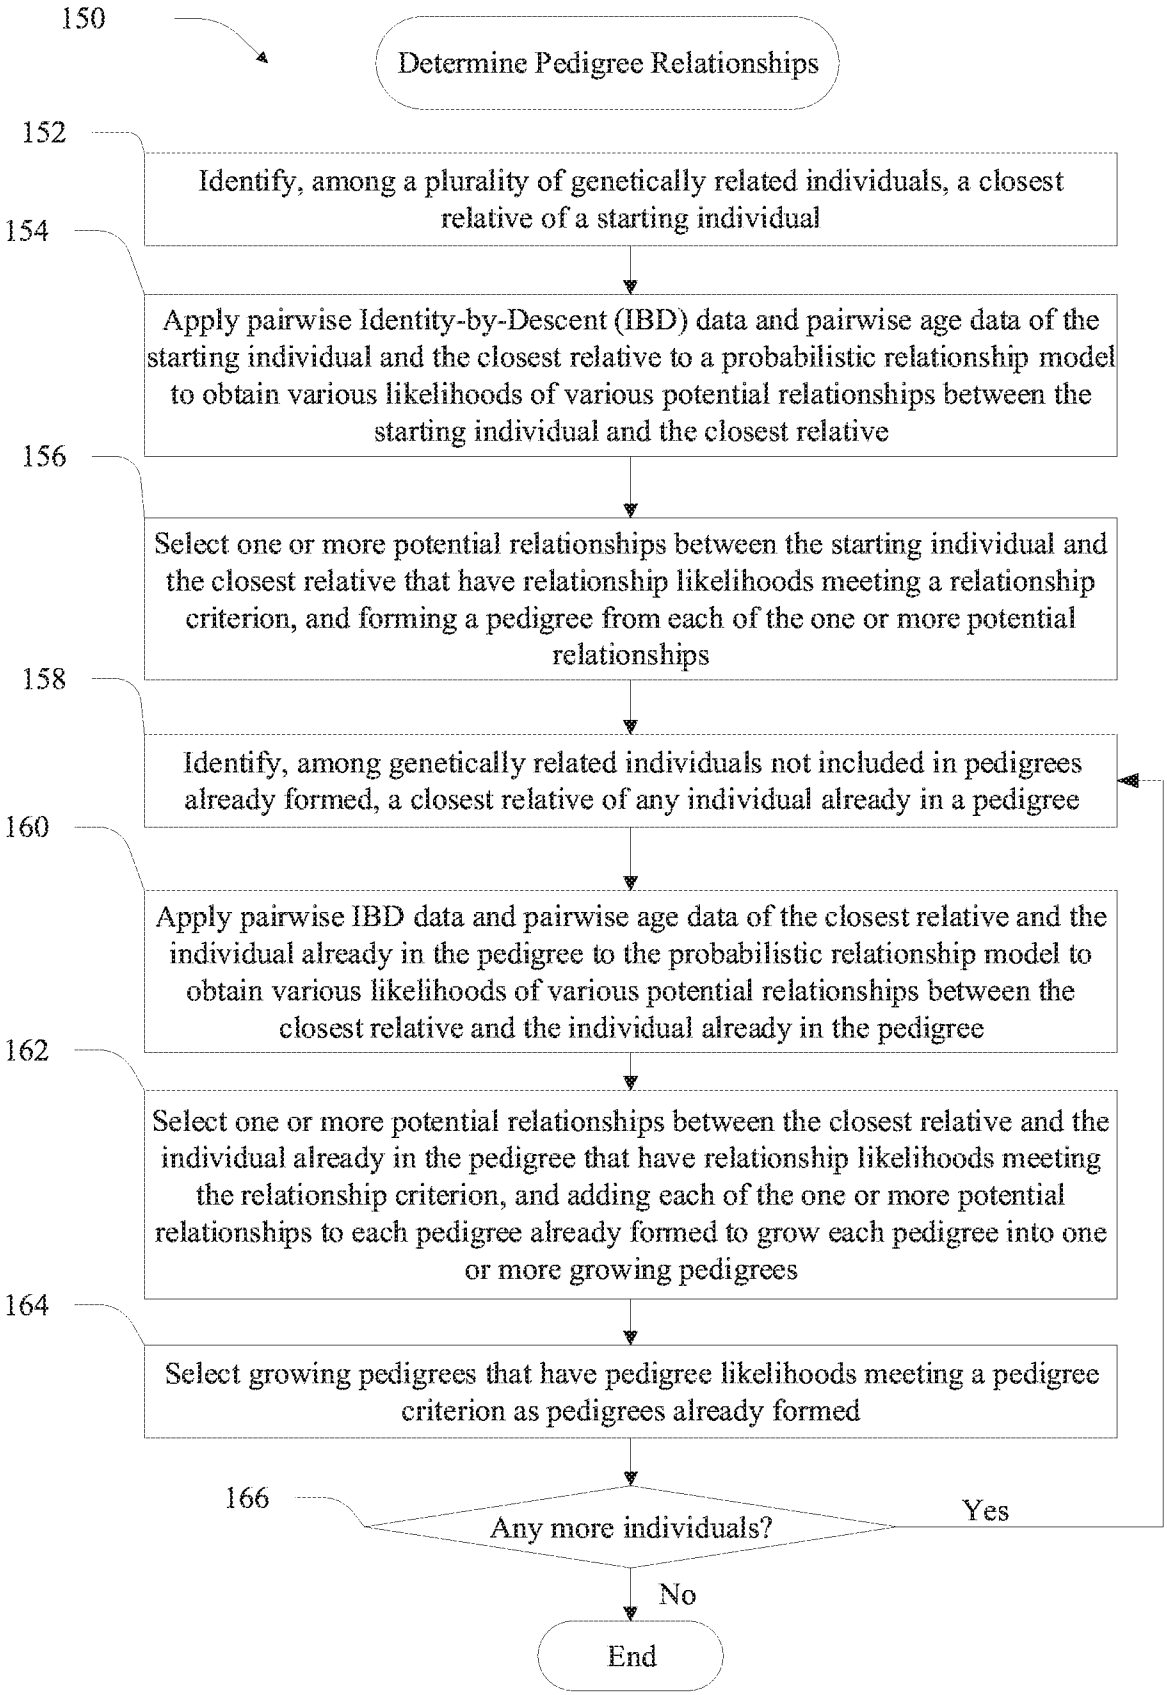
150
Determine Pedigree Relationships
152
Identify, among a plurality of genetically related individuals, a closest relative of a starting individual
154
Apply pairwise Identity-by-Descent (IBD) data and pairwise age data of the starting individual and the closest relative to a probabilistic relationship model to obtain various likelihoods of various potential relationships between the starting individual and the closest relative
156
Select one or more potential relationships between the starting individual and the closest relative that have relationship likelihoods meeting a relationship criterion, and forming a pedigree from each of the one or more potential relationships
158
Identify, among genetically related individuals not included in pedigrees already formed, a closest relative of any individual already in a pedigree
160
Apply pairwise IBD data and pairwise age data of the closest relative and the individual already in the pedigree to the probabilistic relationship model to obtain various likelihoods of various potential relationships between the closest relative and the individual already in the pedigree
162
Select one or more potential relationships between the closest relative and the individual already in the pedigree that have relationship likelihoods meeting the relationship criterion, and adding each of the one or more potential relationships to each pedigree already formed to grow each pedigree into one or more growing pedigrees
164
Select growing pedigrees that have pedigree likelihoods meeting a pedigree criterion as pedigrees already formed
166
Any more individuals?
Yes
No
End

Figure 1

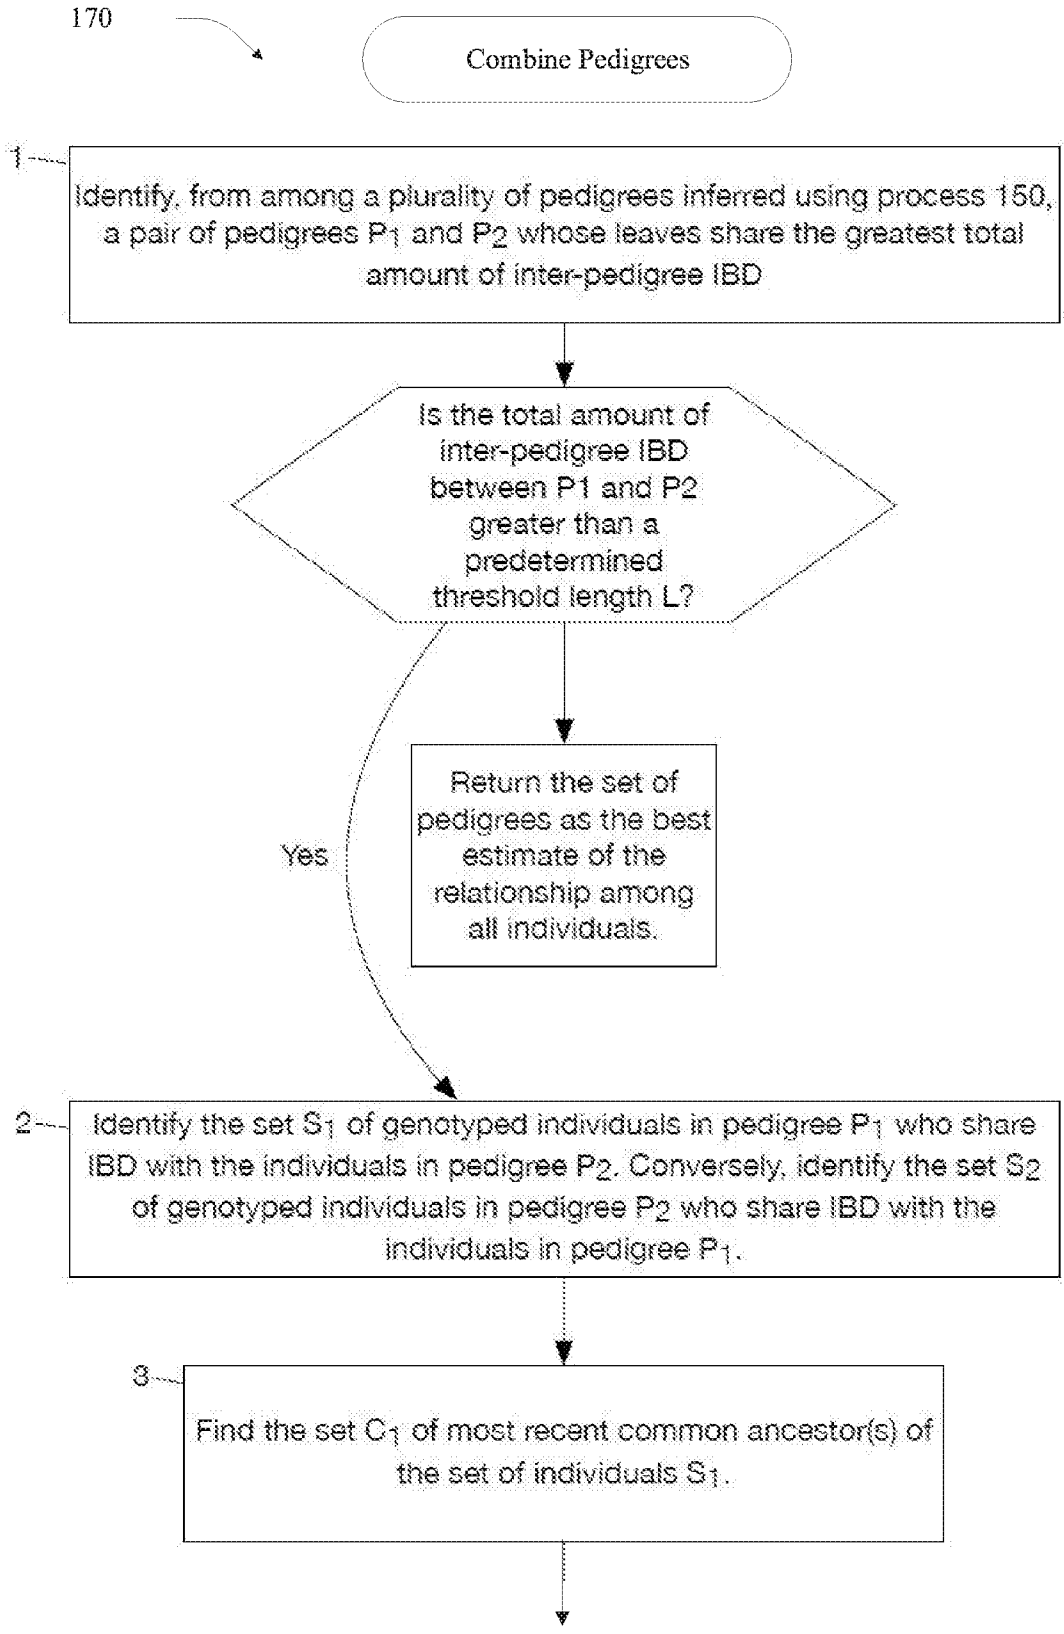
170
Combine Pedigrees
1
Identify, from among a plurality of pedigrees inferred using process 150, a pair of pedigrees P_1 and P_2 whose leaves share the greatest total amount of inter-pedigree IBD
Is the total amount of inter-pedigree IBD between P1 and P2 greater than a predetermined threshold length L?
Yes
Return the set of pedigrees as the best estimate of the relationship among all individuals.
2
Identify the set S_1 of genotyped individuals in pedigree P_1 who share IBD with the individuals in pedigree P_2. Conversely, identify the set S_2 of genotyped individuals in pedigree P_2 who share IBD with the individuals in pedigree P_1.
3
Find the set C_1 of most recent common ancestor(s) of the set of individuals S_1.

Figure 2

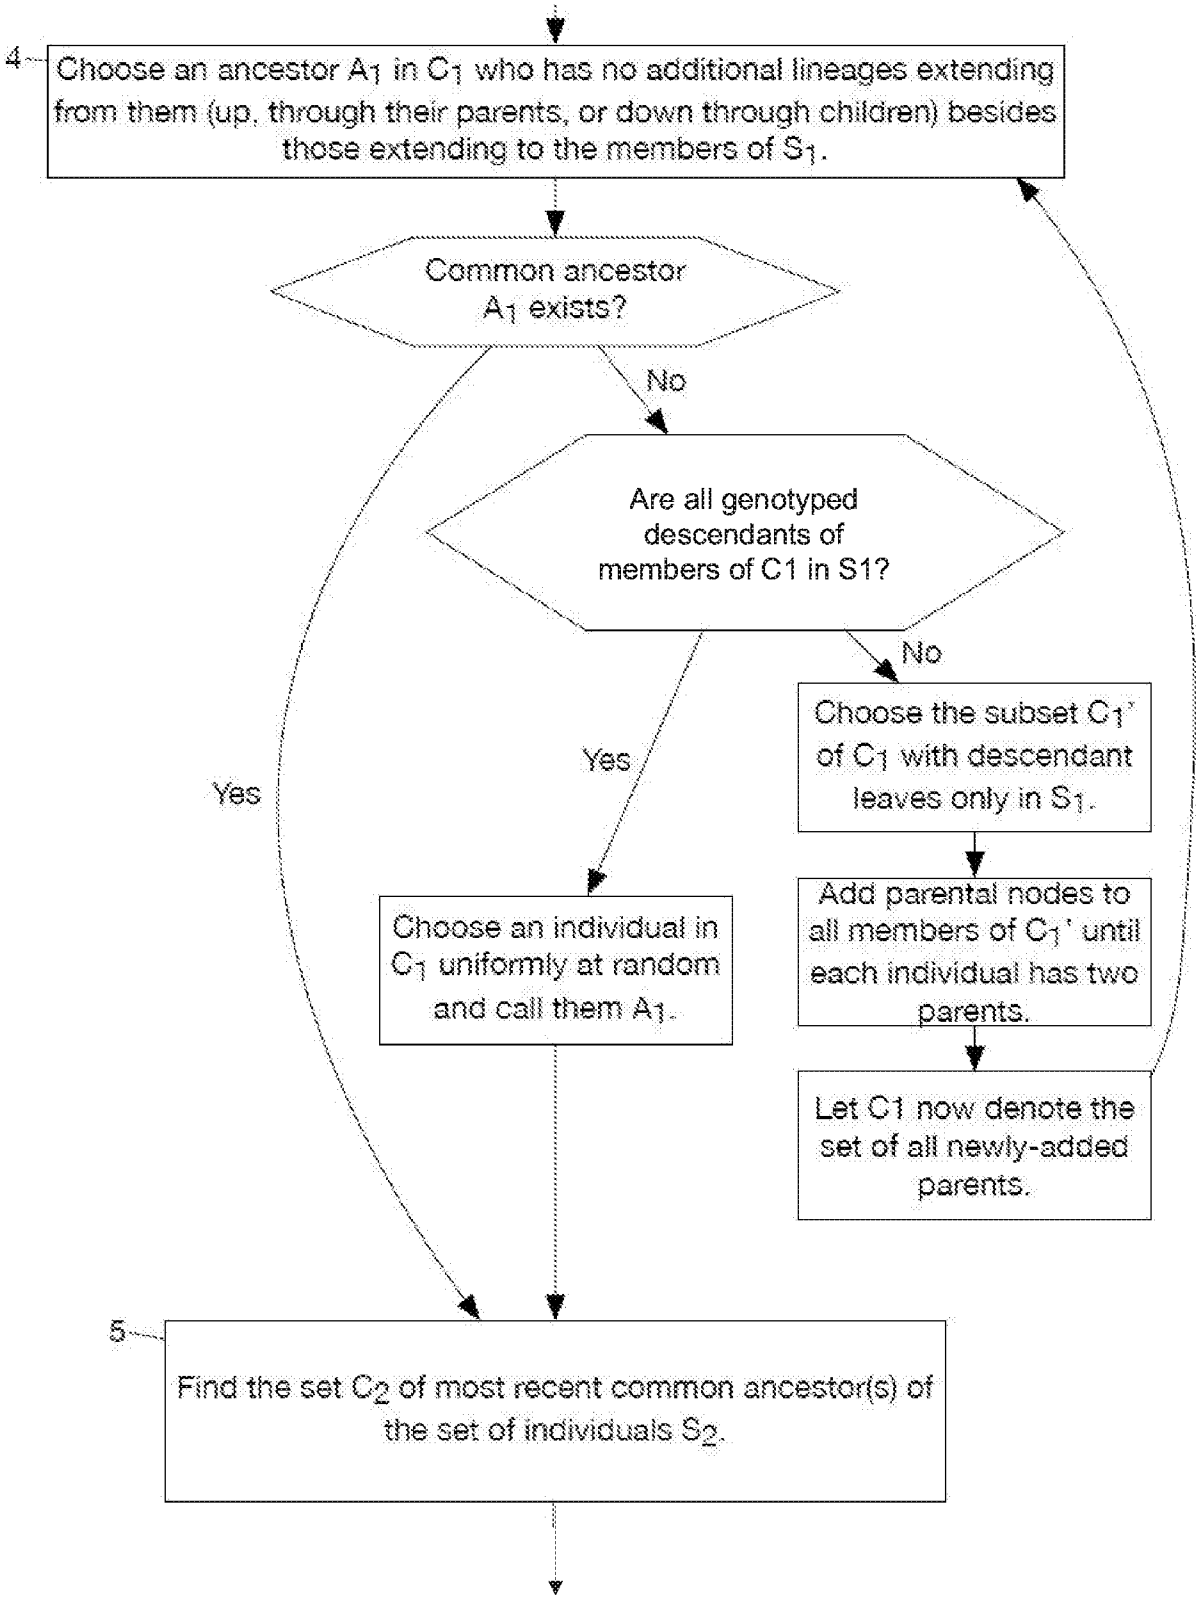

Figure 2 (Continued)

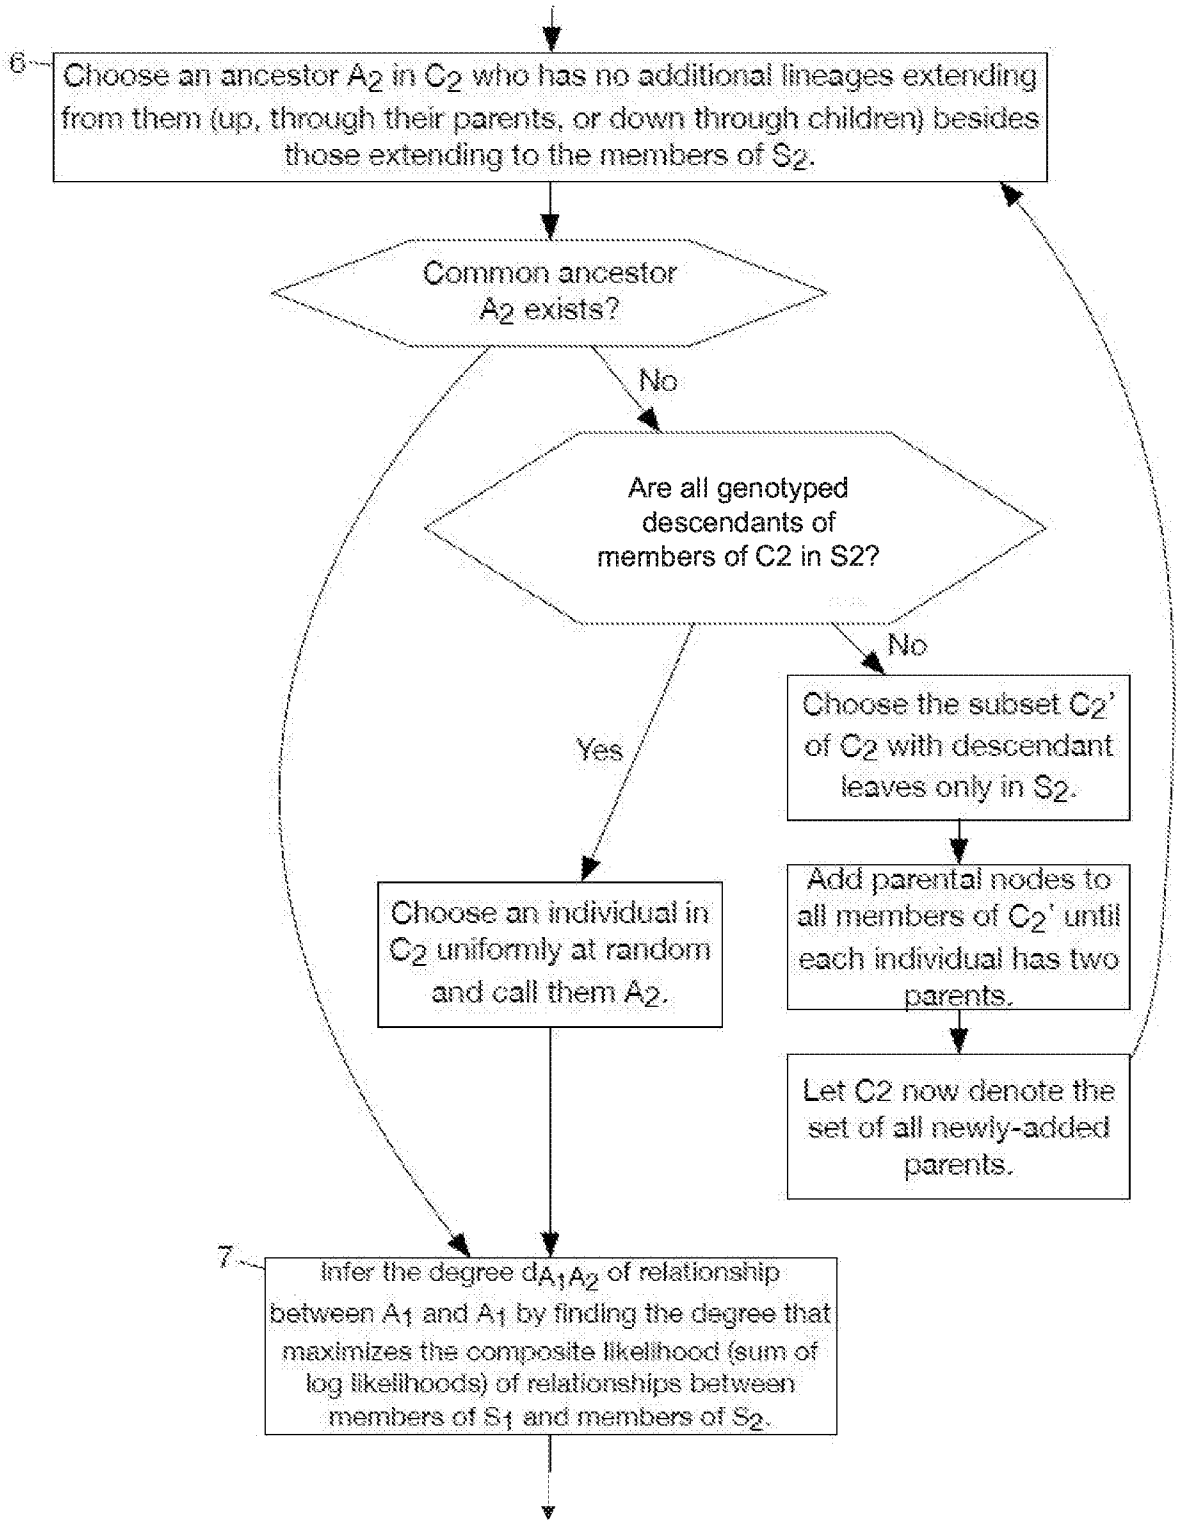

Figure 2 (Continued)

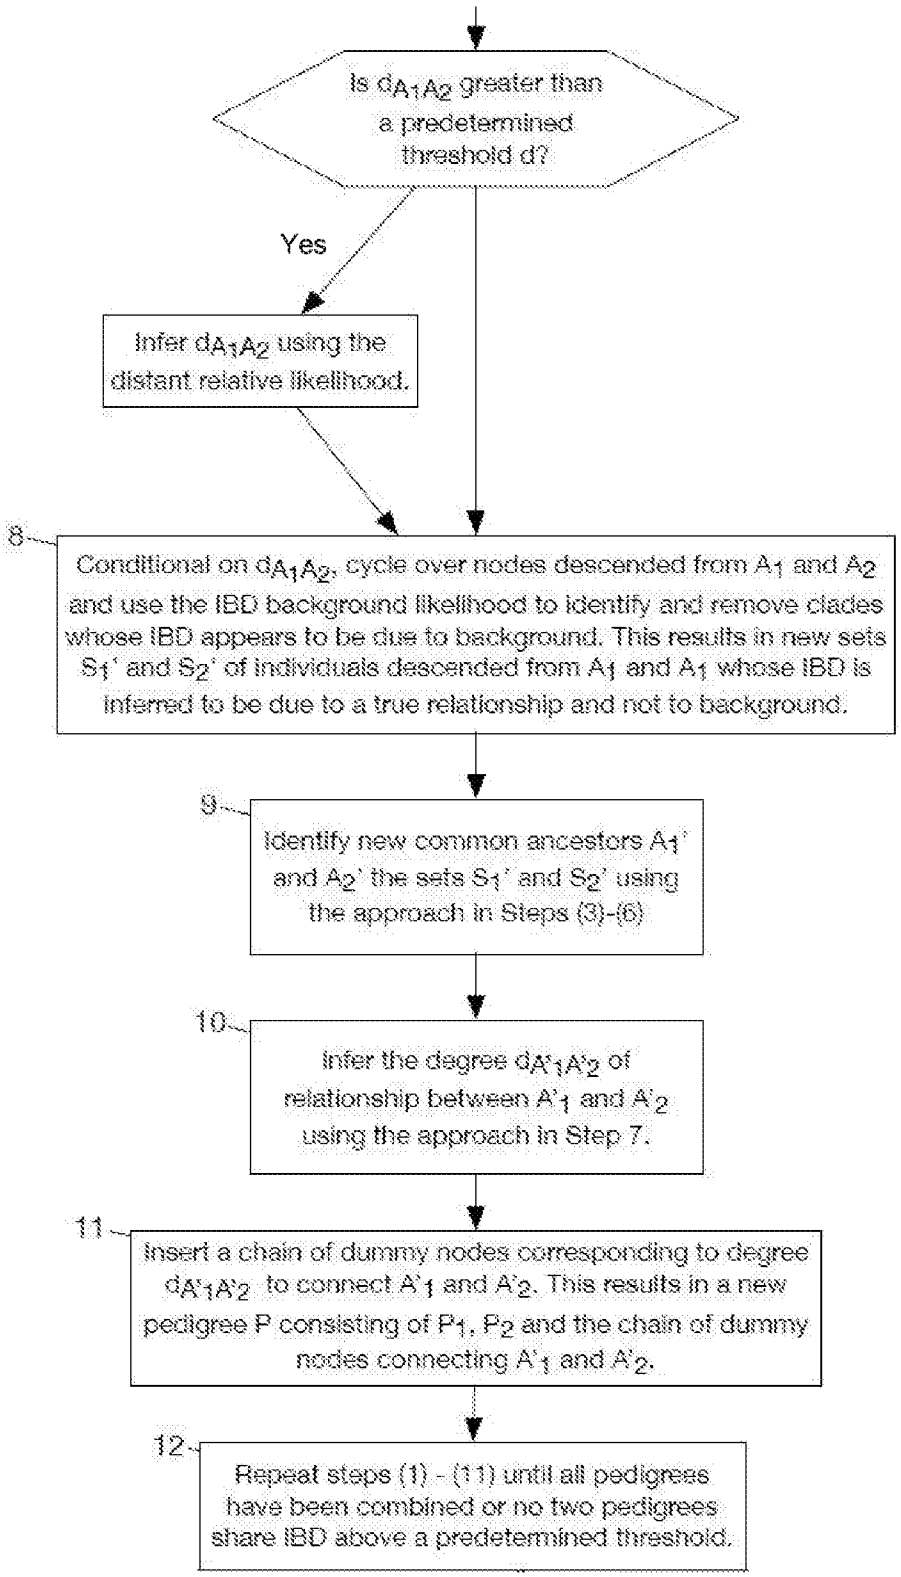

Figure 2 (Continued)

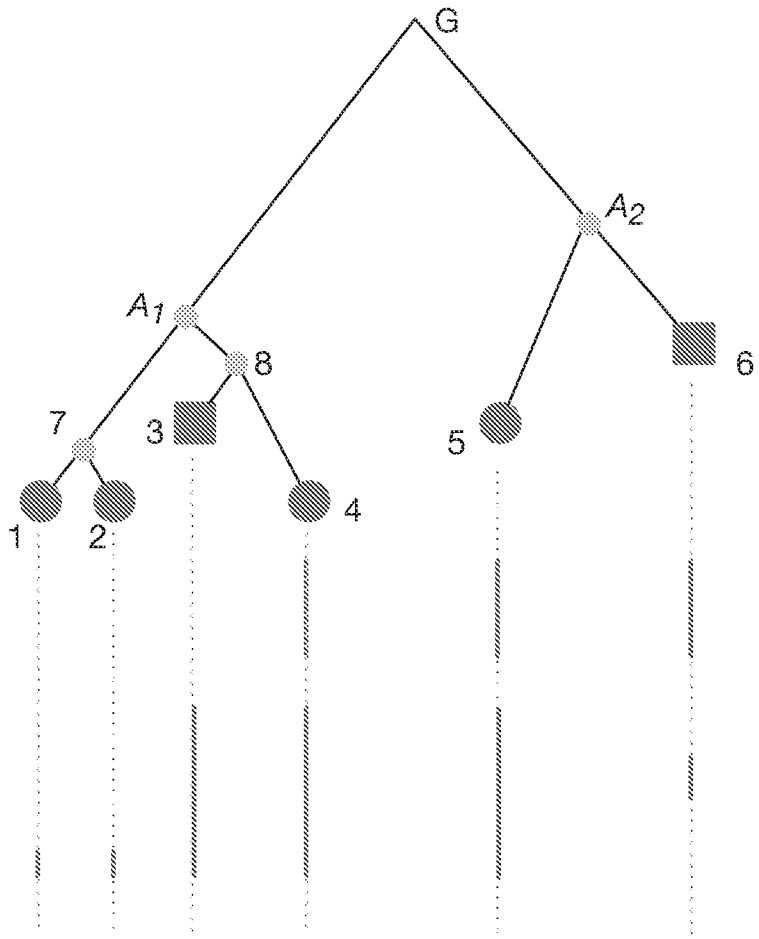

Figure 3

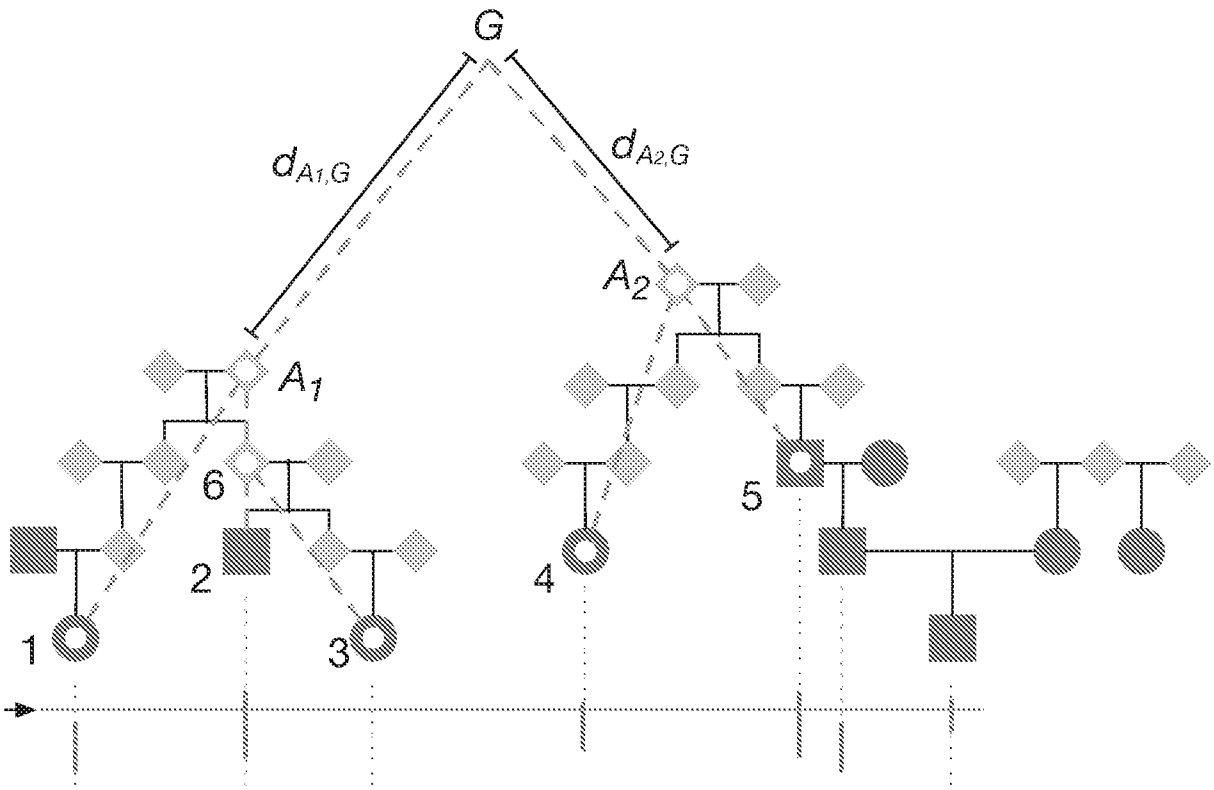

Figure 4

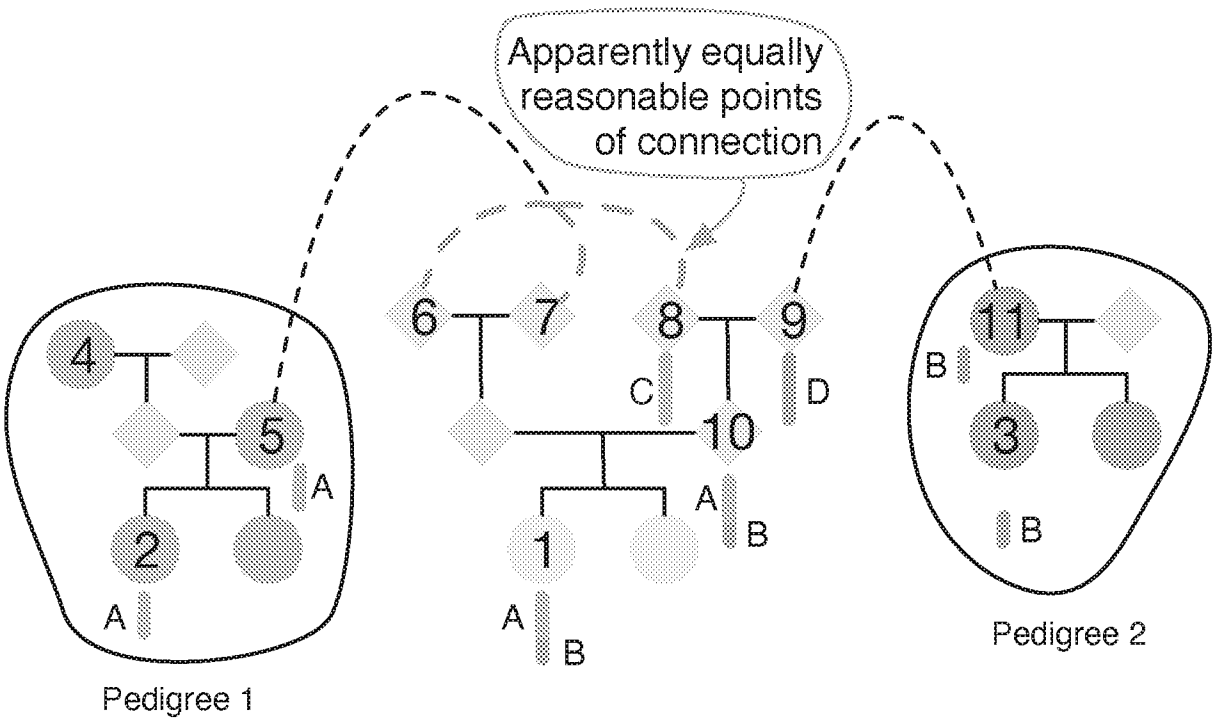

Figure 5

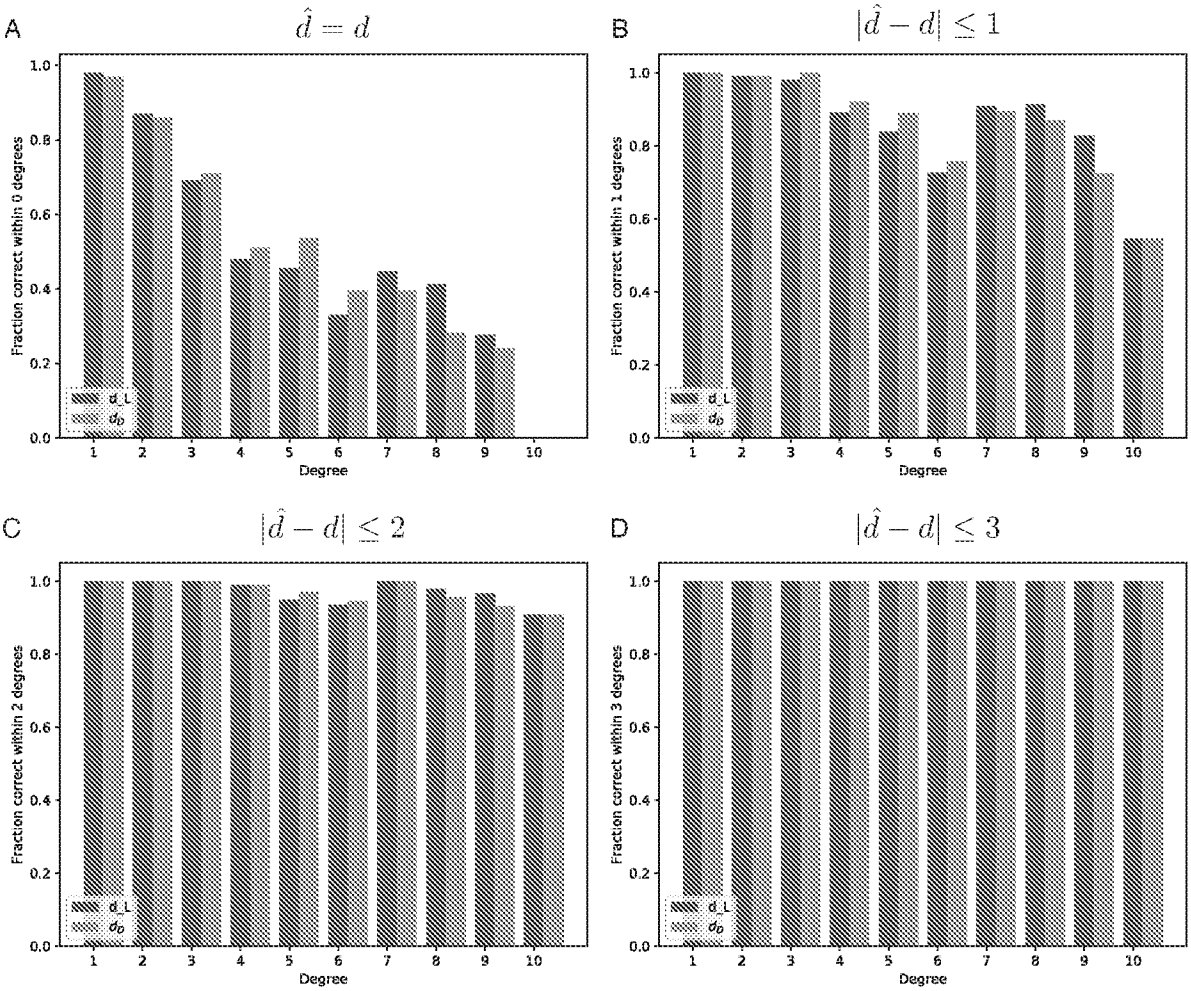
A
$\hat{d} = d$
B
$|\hat{d} - d| \leq 1$
C
$|\hat{d} - d| \leq 2$
D
$|\hat{d} - d| \leq 3$
Fraction correct within 0 degrees
Fraction correct within 1 degrees
Fraction correct within 2 degrees
Fraction correct within 3 degrees
Degree
d_L
d_D

Figure 6

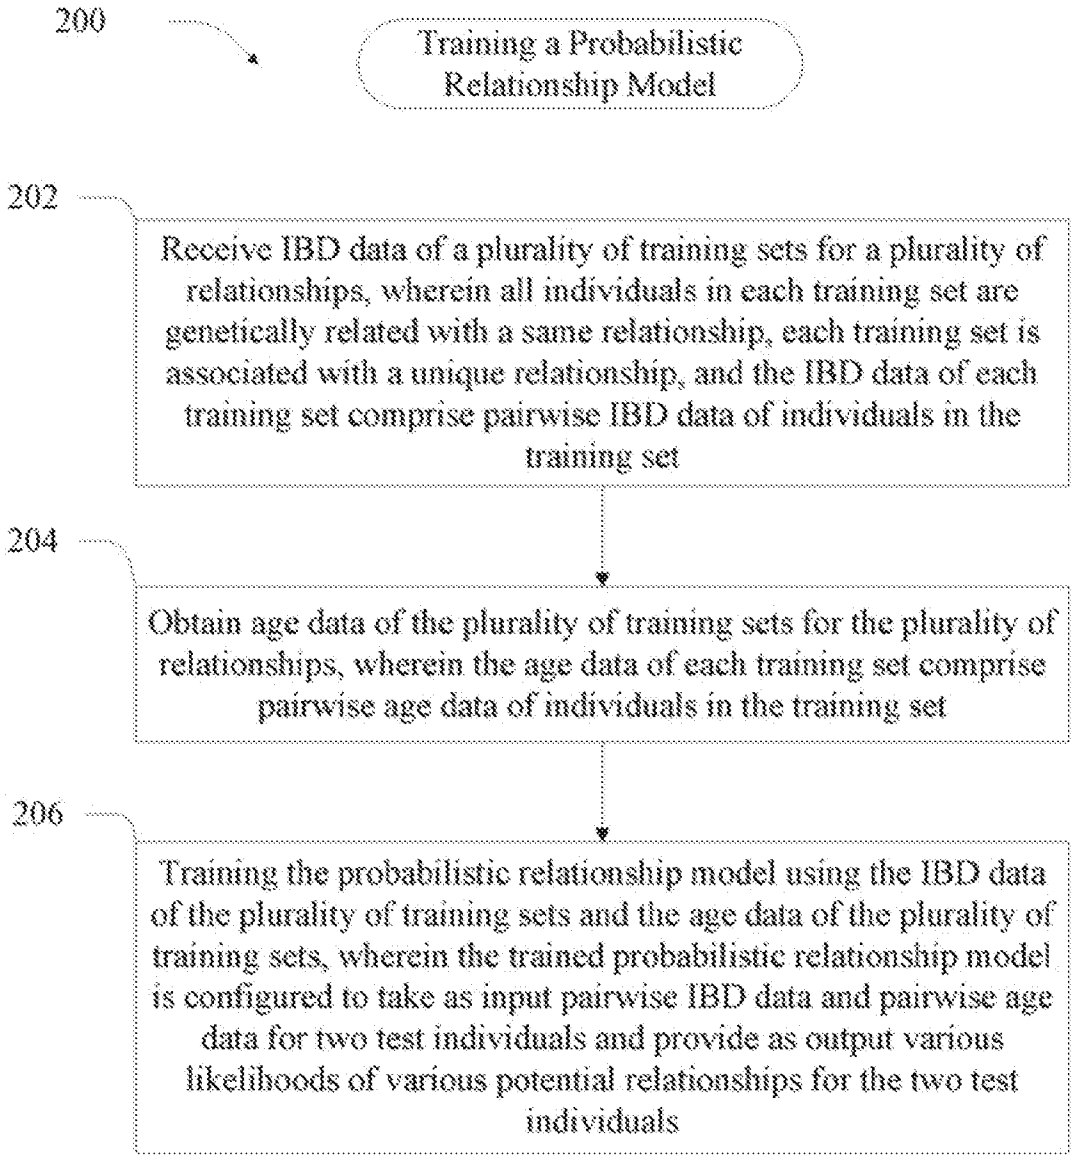
200
Training a Probabilistic Relationship Model
202
Receive IBD data of a plurality of training sets for a plurality of relationships, wherein all individuals in each training set are genetically related with a same relationship, each training set is associated with a unique relationship, and the IBD data of each training set comprise pairwise IBD data of individuals in the training set
204
Obtain age data of the plurality of training sets for the plurality of relationships, wherein the age data of each training set comprise pairwise age data of individuals in the training set
206
Training the probabilistic relationship model using the IBD data of the plurality of training sets and the age data of the plurality of training sets, wherein the trained probabilistic relationship model is configured to take as input pairwise IBD data and pairwise age data for two test individuals and provide as output various likelihoods of various potential relationships for the two test individuals

Figure 7

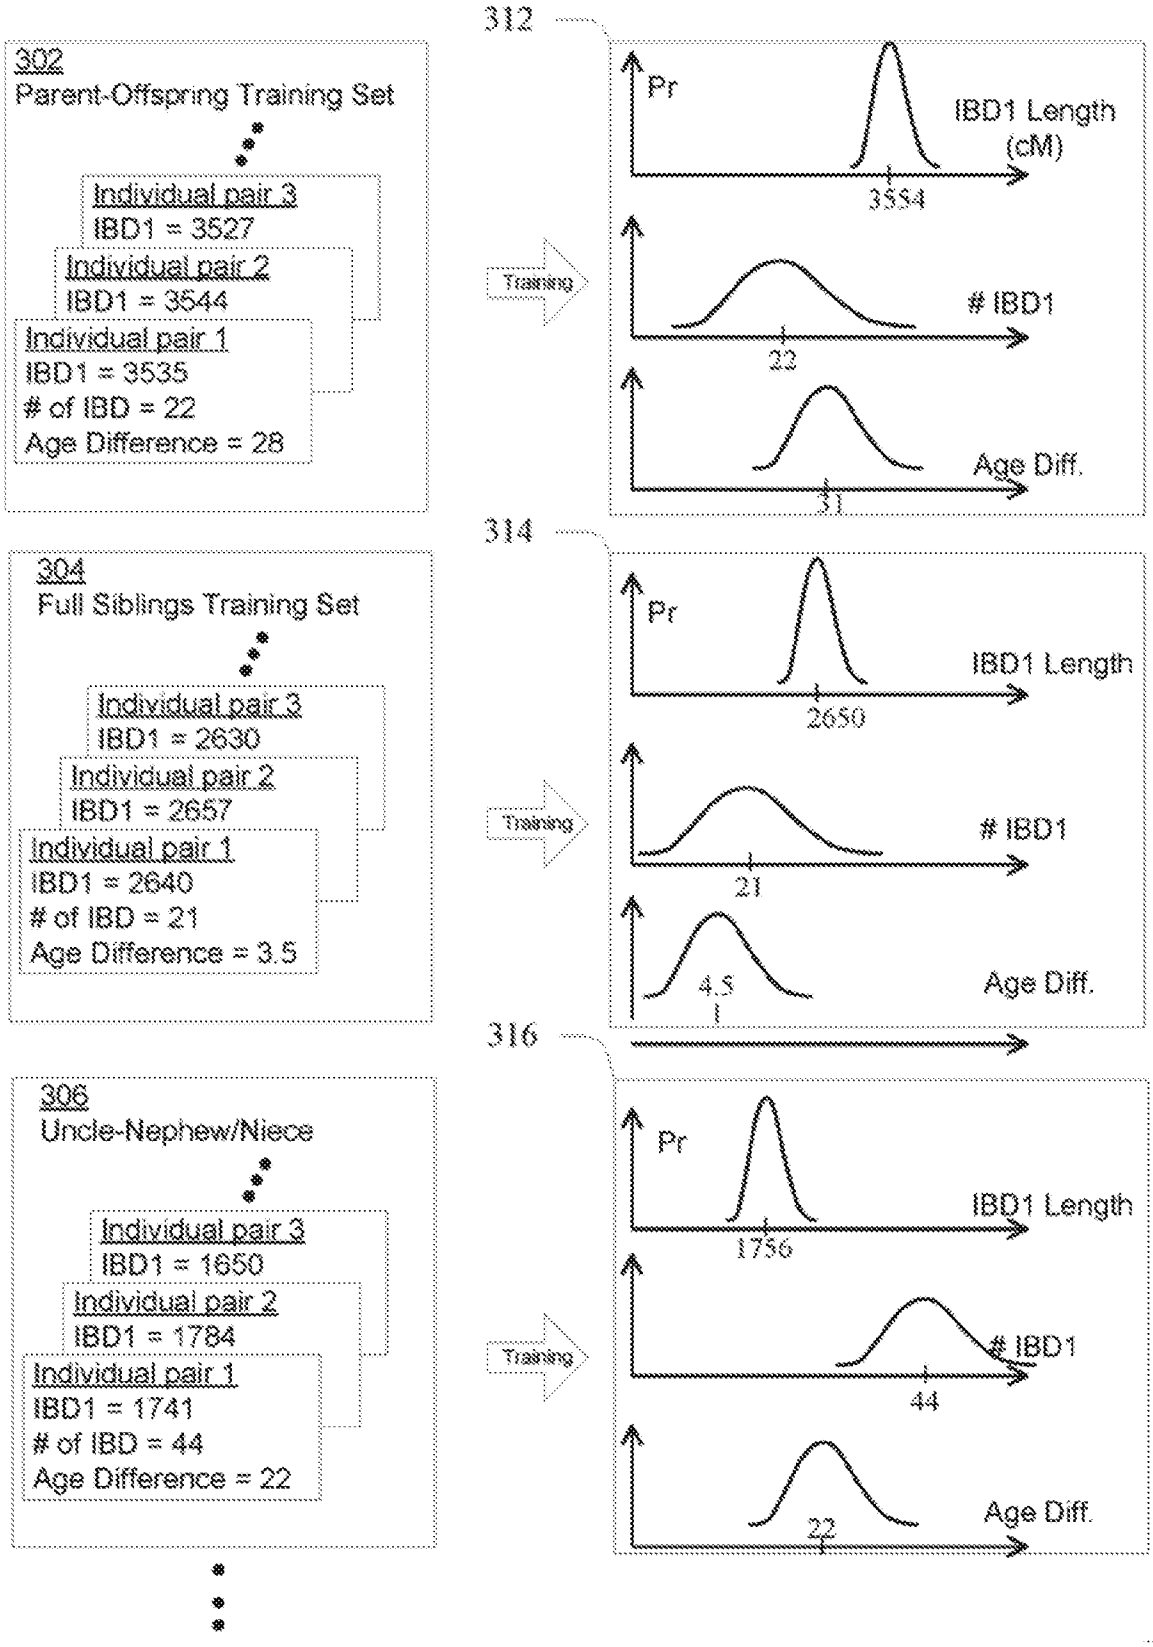
302
Parent-Offspring Training Set
Individual pair 3
IBD1 = 3527
Individual pair 2
IBD1 = 3544
Individual pair 1
IBD1 = 3535
# of IBD = 22
Age Difference = 28
304
Full Siblings Training Set
Individual pair 3
IBD1 = 2630
Individual pair 2
IBD1 = 2657
Individual pair 1
IBD1 = 2640
# of IBD = 21
Age Difference = 3.5
306
Uncle-Nephew/Niece
Individual pair 3
IBD1 = 1650
Individual pair 2
IBD1 = 1784
Individual pair 1
IBD1 = 1741
# of IBD = 44
Age Difference = 22
312
Pr
IBD1 Length (cM)
3554
Training
IBD1
22
Age Diff.
31
314
Pr
IBD1 Length
2650
Training
IBD1
21
4.5
Age Diff.
316
Pr
IBD1 Length
1756
Training
IBD1
44
Age Diff.
22

Figure 8

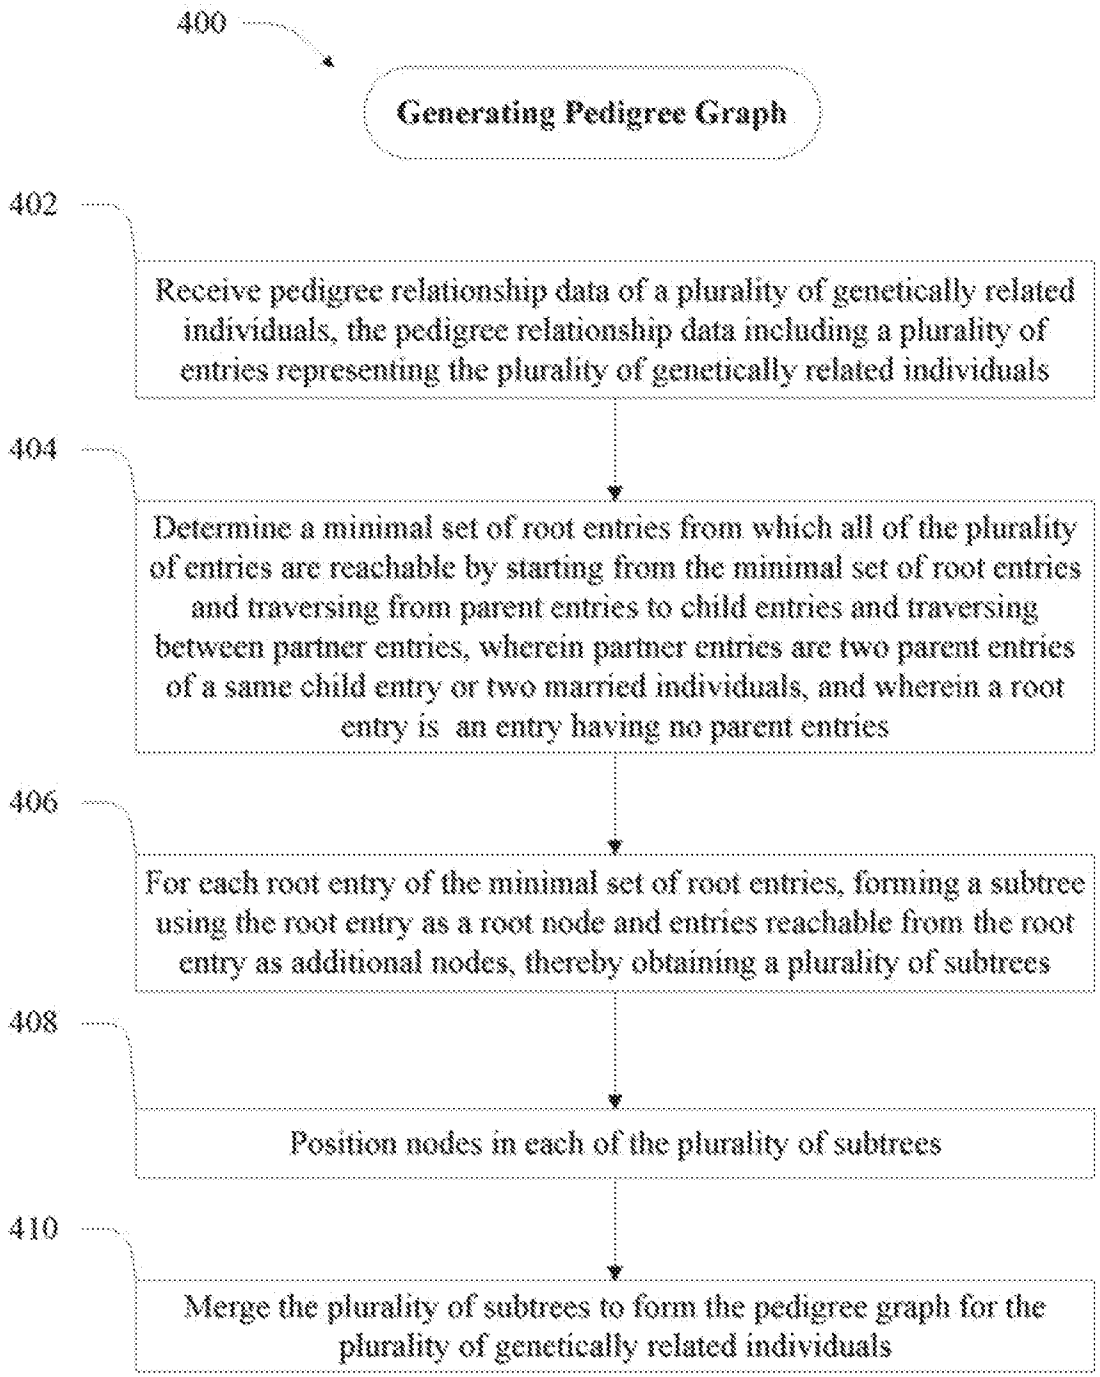
400
Generating Pedigree Graph
402
Receive pedigree relationship data of a plurality of genetically related individuals, the pedigree relationship data including a plurality of entries representing the plurality of genetically related individuals
404
Determine a minimal set of root entries from which all of the plurality of entries are reachable by starting from the minimal set of root entries and traversing from parent entries to child entries and traversing between partner entries, wherein partner entries are two parent entries of a same child entry or two married individuals, and wherein a root entry is an entry having no parent entries
406
For each root entry of the minimal set of root entries, forming a subtree using the root entry as a root node and entries reachable from the root entry as additional nodes, thereby obtaining a plurality of subtrees
408
Position nodes in each of the plurality of subtrees
410
Merge the plurality of subtrees to form the pedigree graph for the plurality of genetically related individuals

Figure 9

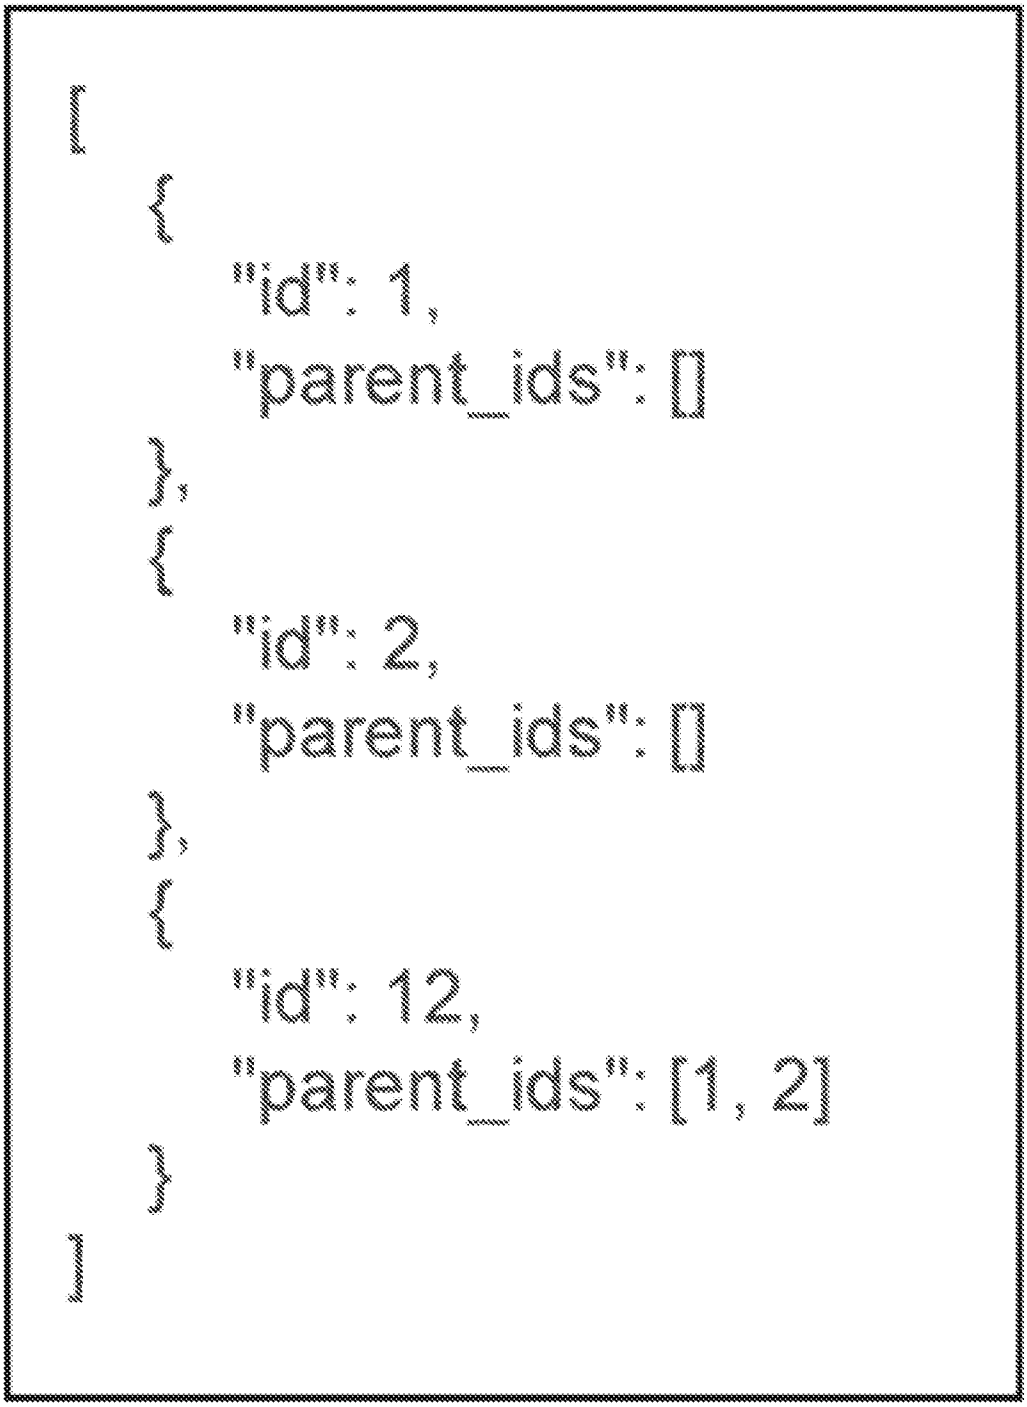

Figure 10

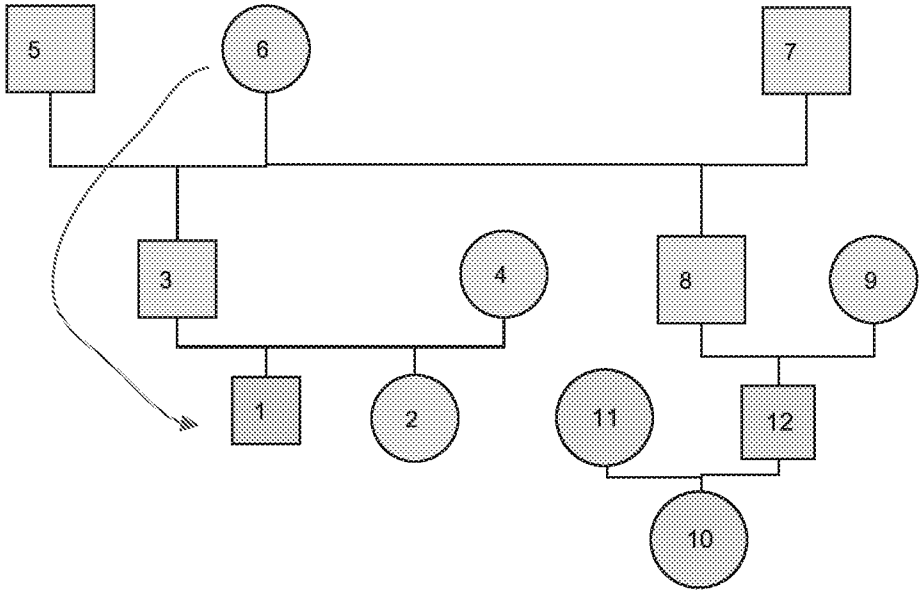

Figure 11

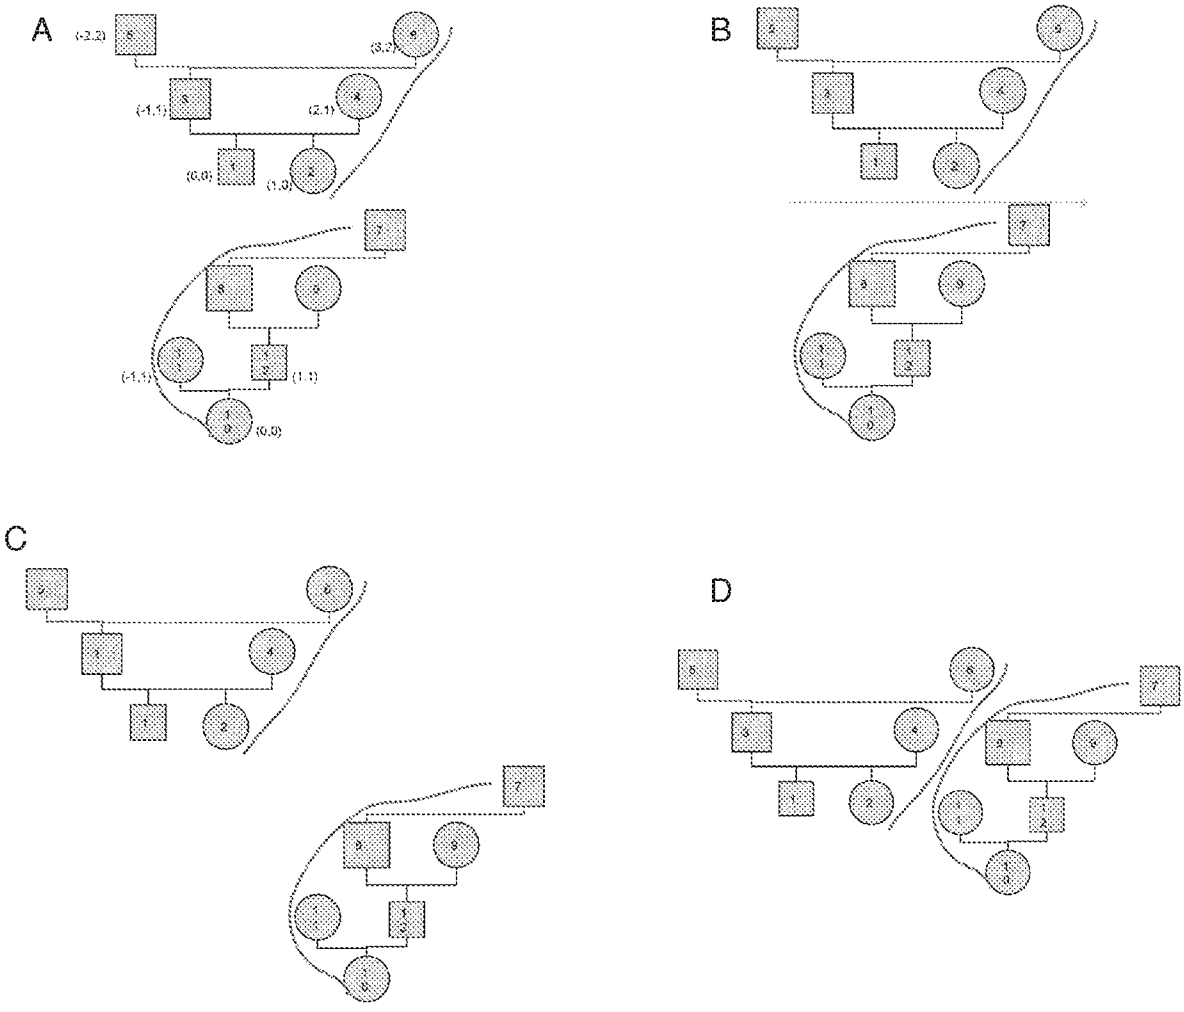

Figure 12

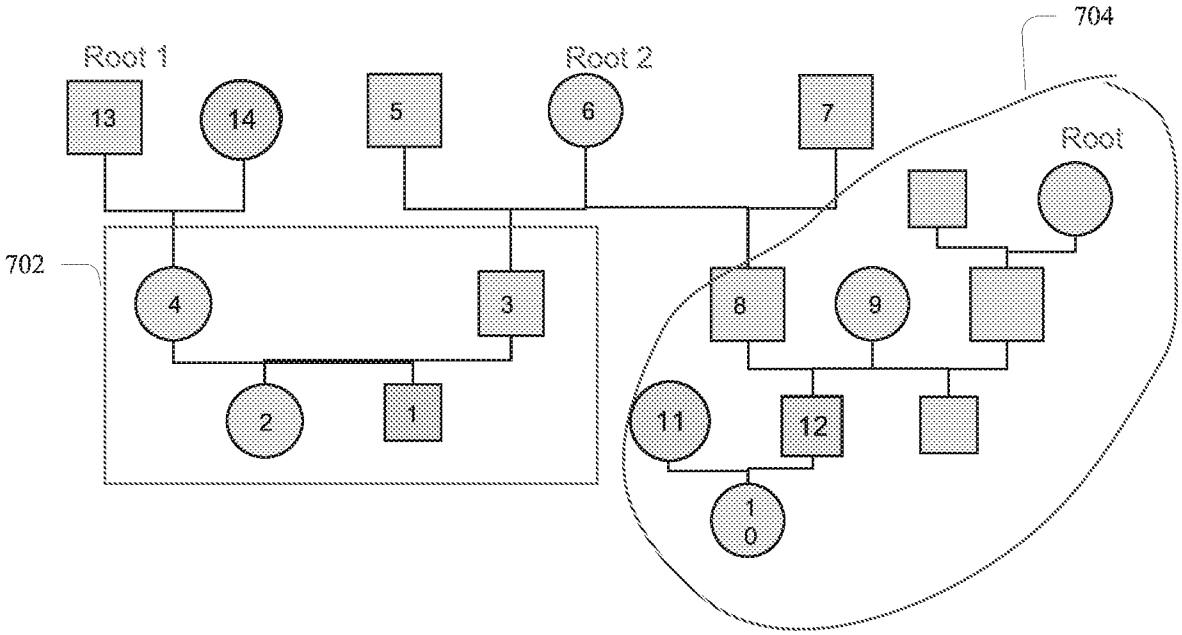

Figure 13

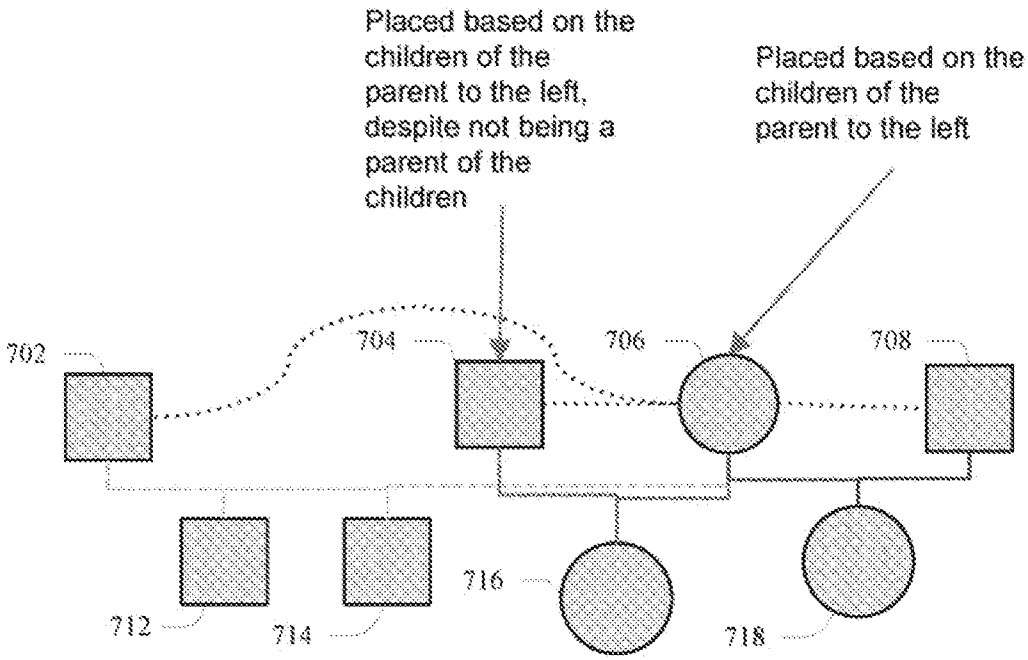

Figure 14

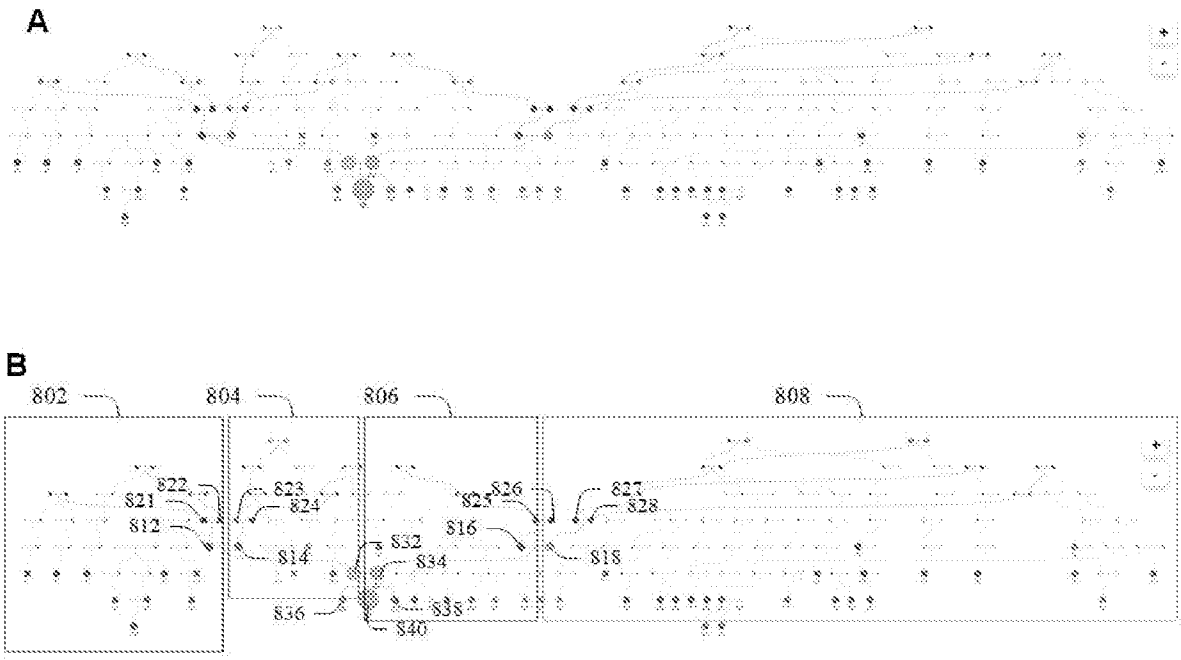
A
B
802
804
806
808
821
822
823
824
812
814
832
834
836
838
840
816
825
826
827
828
818

Figure 15

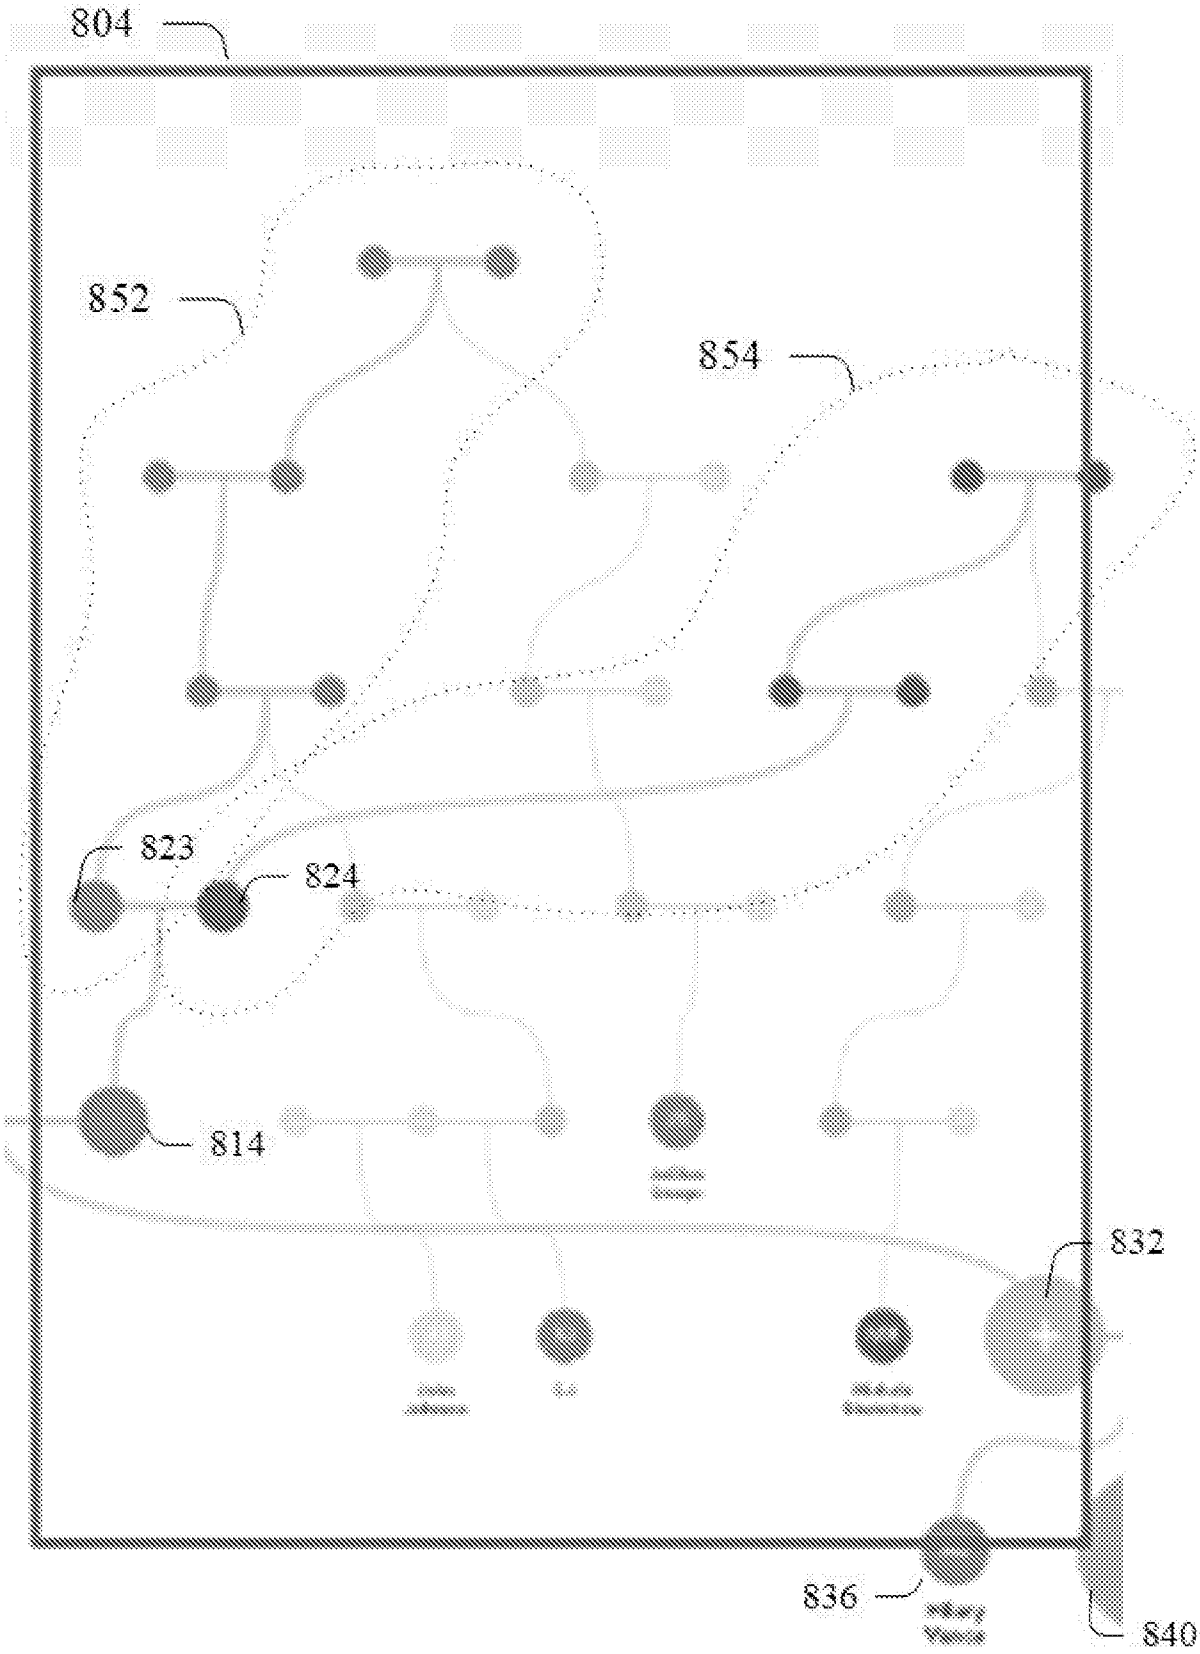

Figure 16

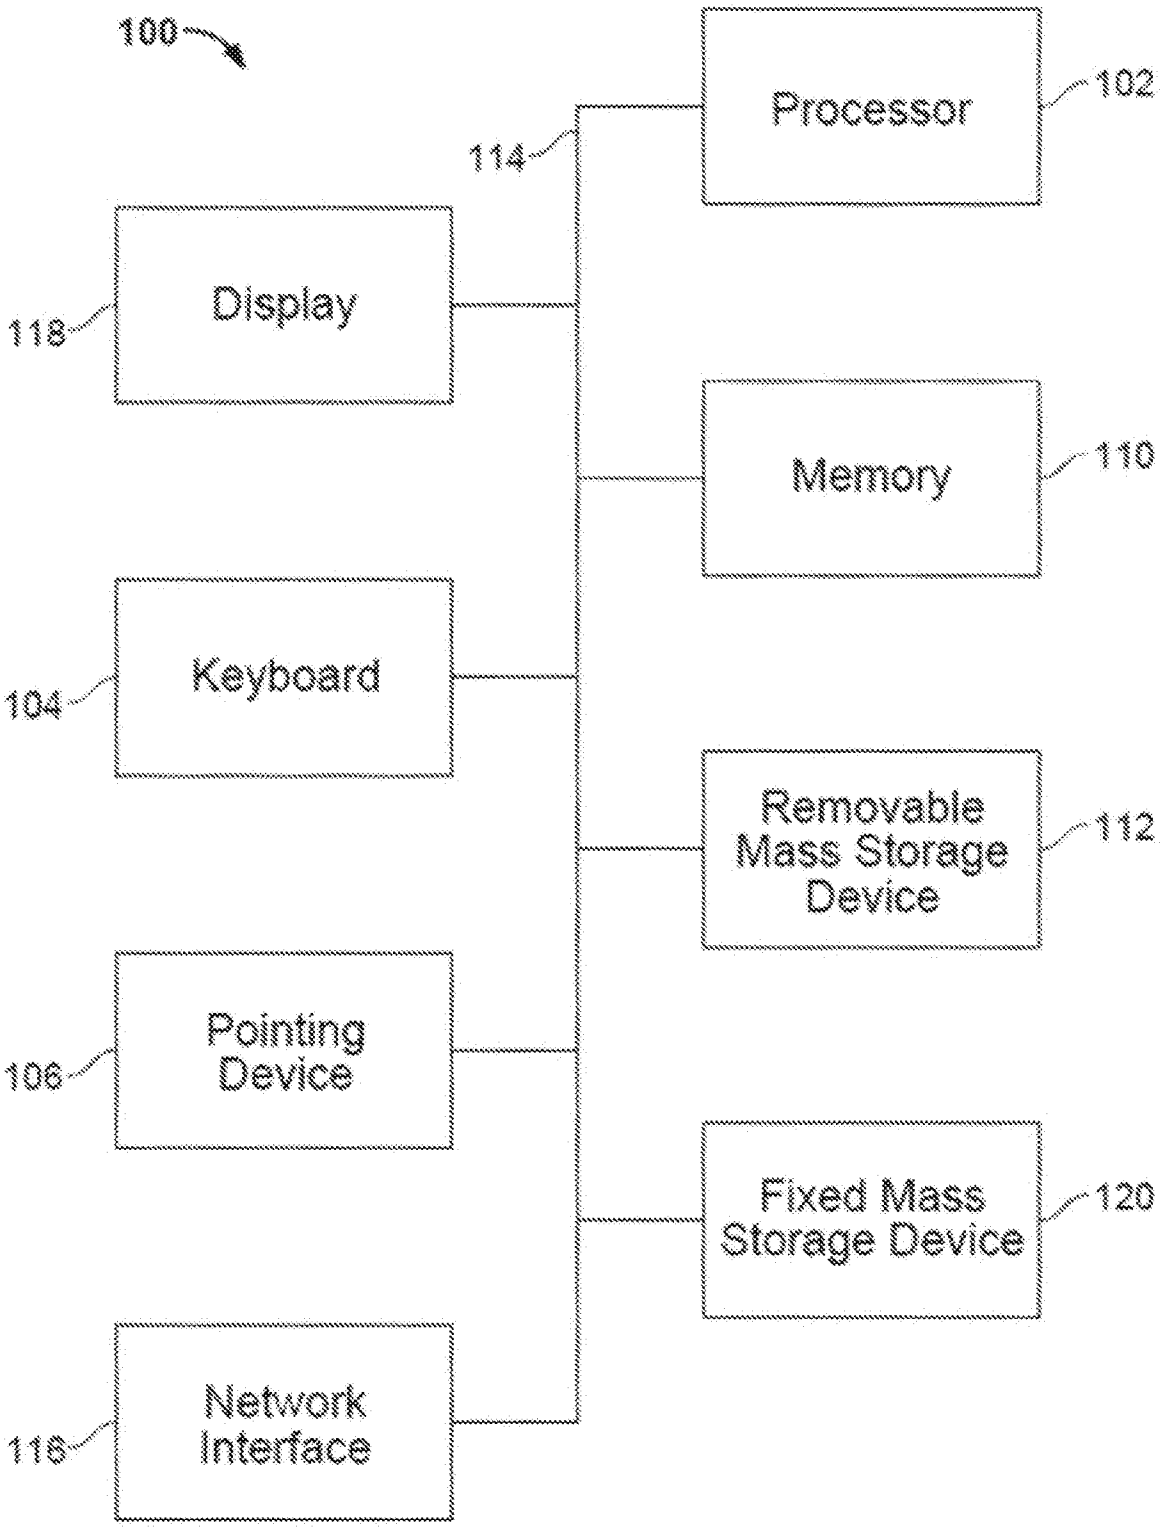
100
Processor
102
114
Display
118
Memory
110
Keyboard
104
Removable Mass Storage Device
112
Pointing Device
106
Fixed Mass Storage Device
120
Network Interface
116

Figure 17

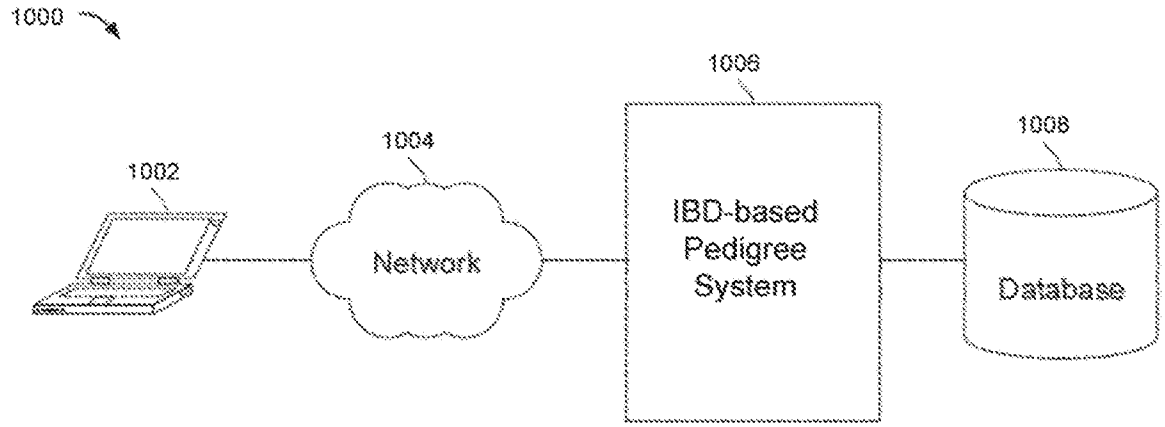
1000
1002
1004
Network
1006
IBD-based
Pedigree
System
1008
Database

Figure 18

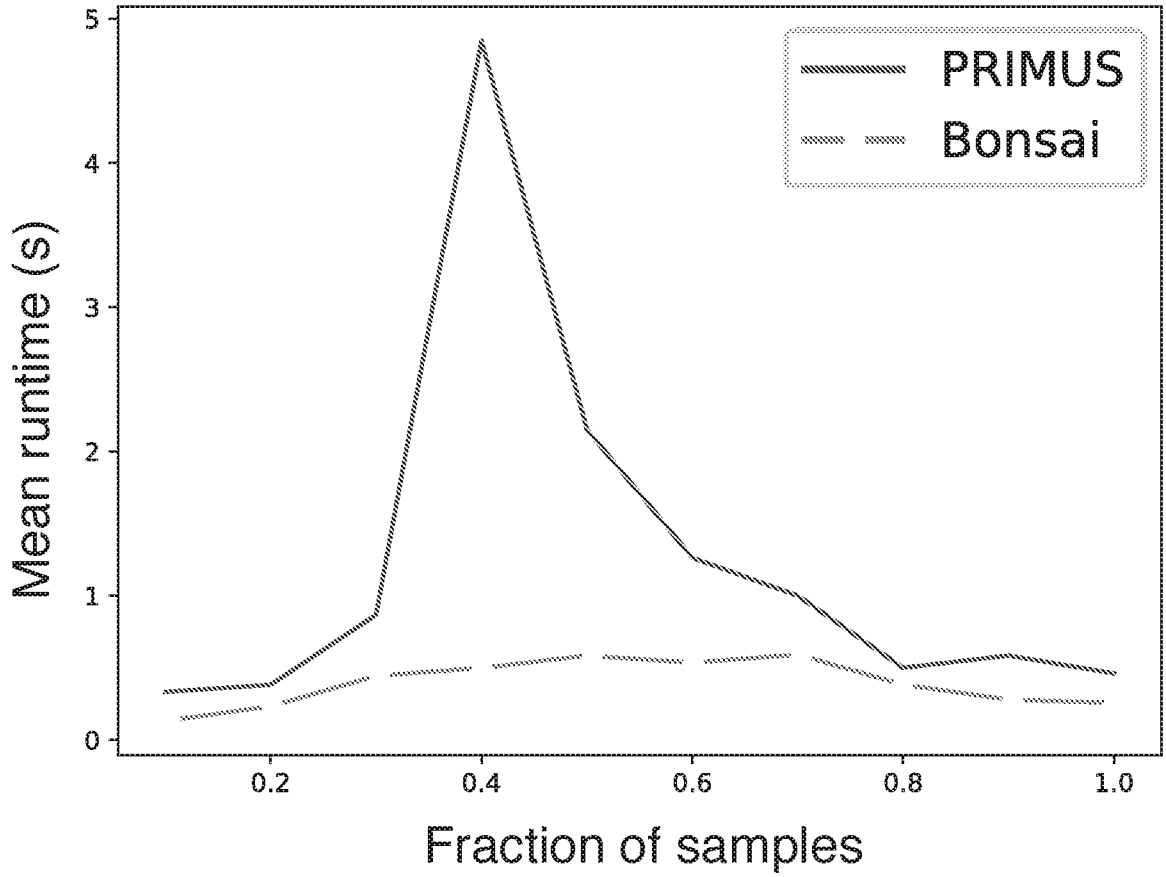
PRIMUS
Bonsai
Mean runtime (s)
0
1
2
3
4
5
0.2
0.4
0.6
0.8
1.0
Fraction of samples

Figure 19

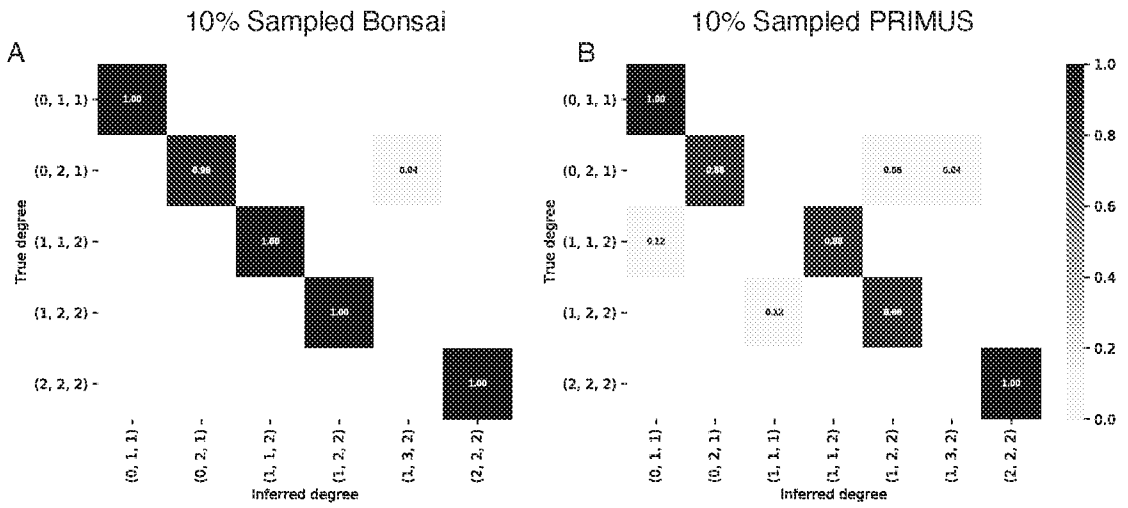
A
10% Sampled Bonsai
True degree
Inferred degree
(0, 1, 1)
(0, 2, 1)
(1, 1, 2)
(1, 2, 2)
(2, 2, 2)
(3, 3, 2)
1.00
0.04
B
10% Sampled PRIMUS
0.06
0.32
1.0
0.8
0.6
0.4
0.2
0.0

Figure 20

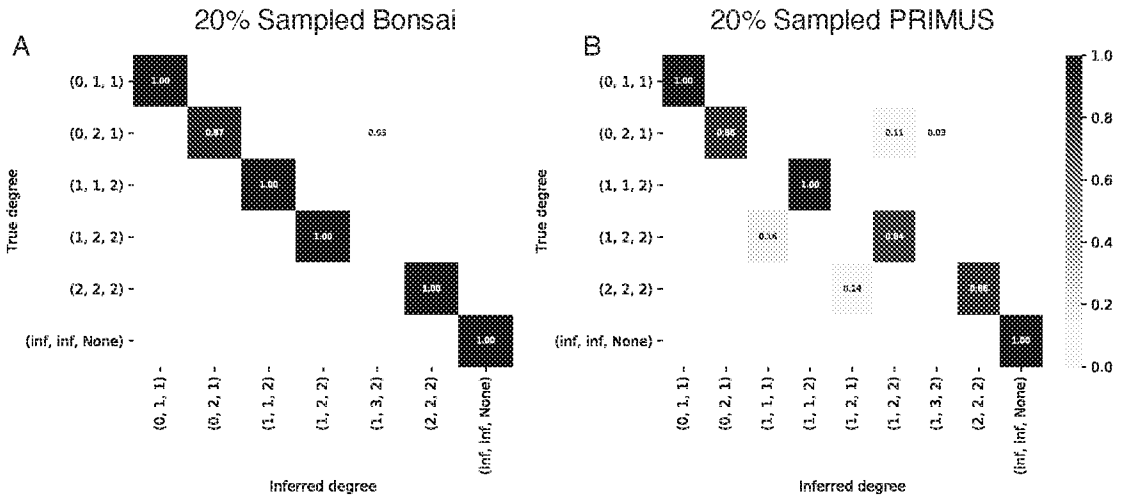
20% Sampled Bonsai
A
20% Sampled PRIMUS
B
True degree
Inferred degree
(0, 1, 1)
(0, 2, 1)
(1, 1, 2)
(1, 2, 2)
(2, 2, 2)
(inf, inf, None)
(1, 3, 2)
(1, 1, 1)
(1, 2, 1)
1.0
0.8
0.6
0.4
0.2
0.0

Figure 21

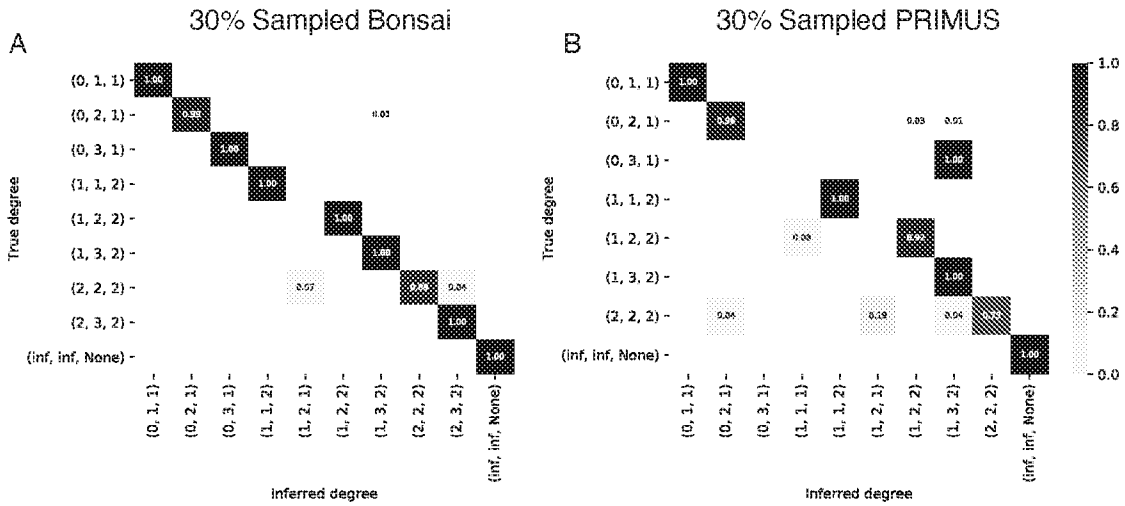
A
30% Sampled Bonsai
True degree
Inferred degree
(0, 1, 1)
(0, 2, 1)
(0, 3, 1)
(1, 1, 2)
(1, 2, 1)
(1, 2, 2)
(1, 3, 2)
(2, 2, 2)
(2, 3, 2)
(inf, inf, None)
B
30% Sampled PRIMUS
1.0
0.8
0.6
0.4
0.2
0.0

Figure 22

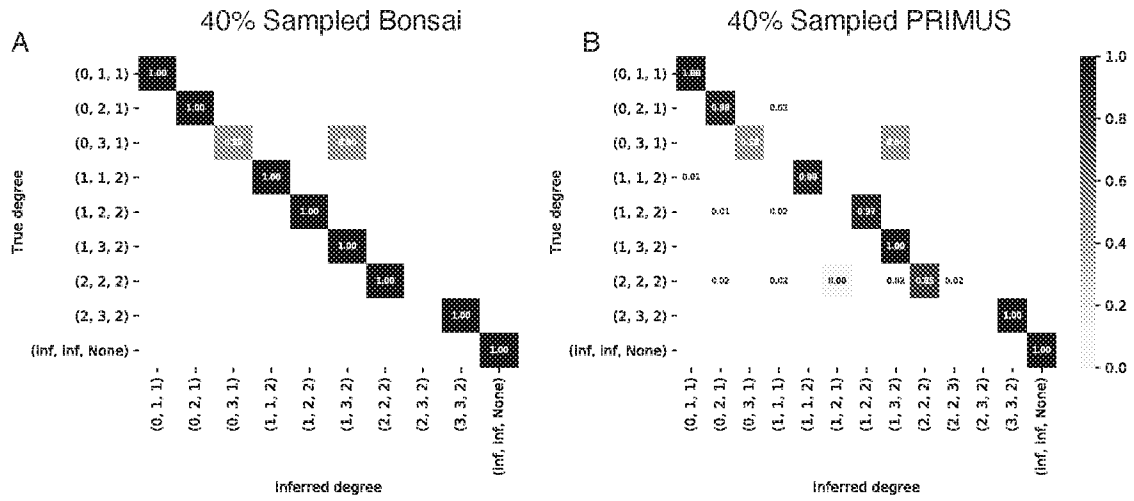
A
40% Sampled Bonsai
True degree
Inferred degree
(0, 1, 1)
(0, 2, 1)
(0, 3, 1)
(1, 1, 2)
(1, 2, 2)
(1, 3, 2)
(2, 2, 2)
(2, 3, 2)
(3, 3, 2)
(inf, inf, None)
B
40% Sampled PRIMUS
1.0
0.8
0.6
0.4
0.2
0.0

Figure 23

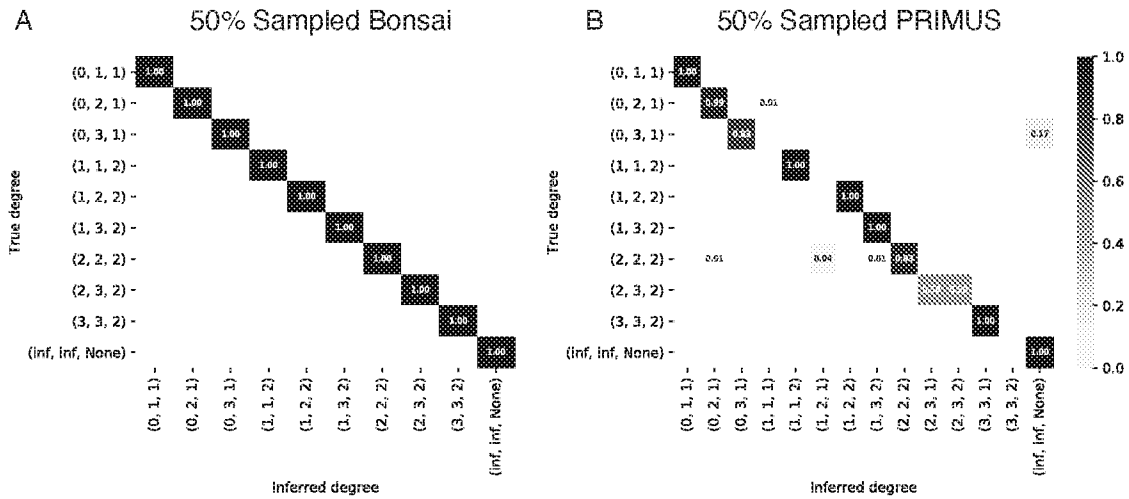
A
50% Sampled Bonsai
B
50% Sampled PRIMUS
True degree
Inferred degree
(0, 1, 1)
(0, 2, 1)
(0, 3, 1)
(1, 1, 2)
(1, 2, 2)
(1, 3, 2)
(2, 2, 2)
(2, 3, 2)
(3, 3, 2)
(inf, inf, None)
1.0
0.8
0.6
0.4
0.2
0.0

Figure 24

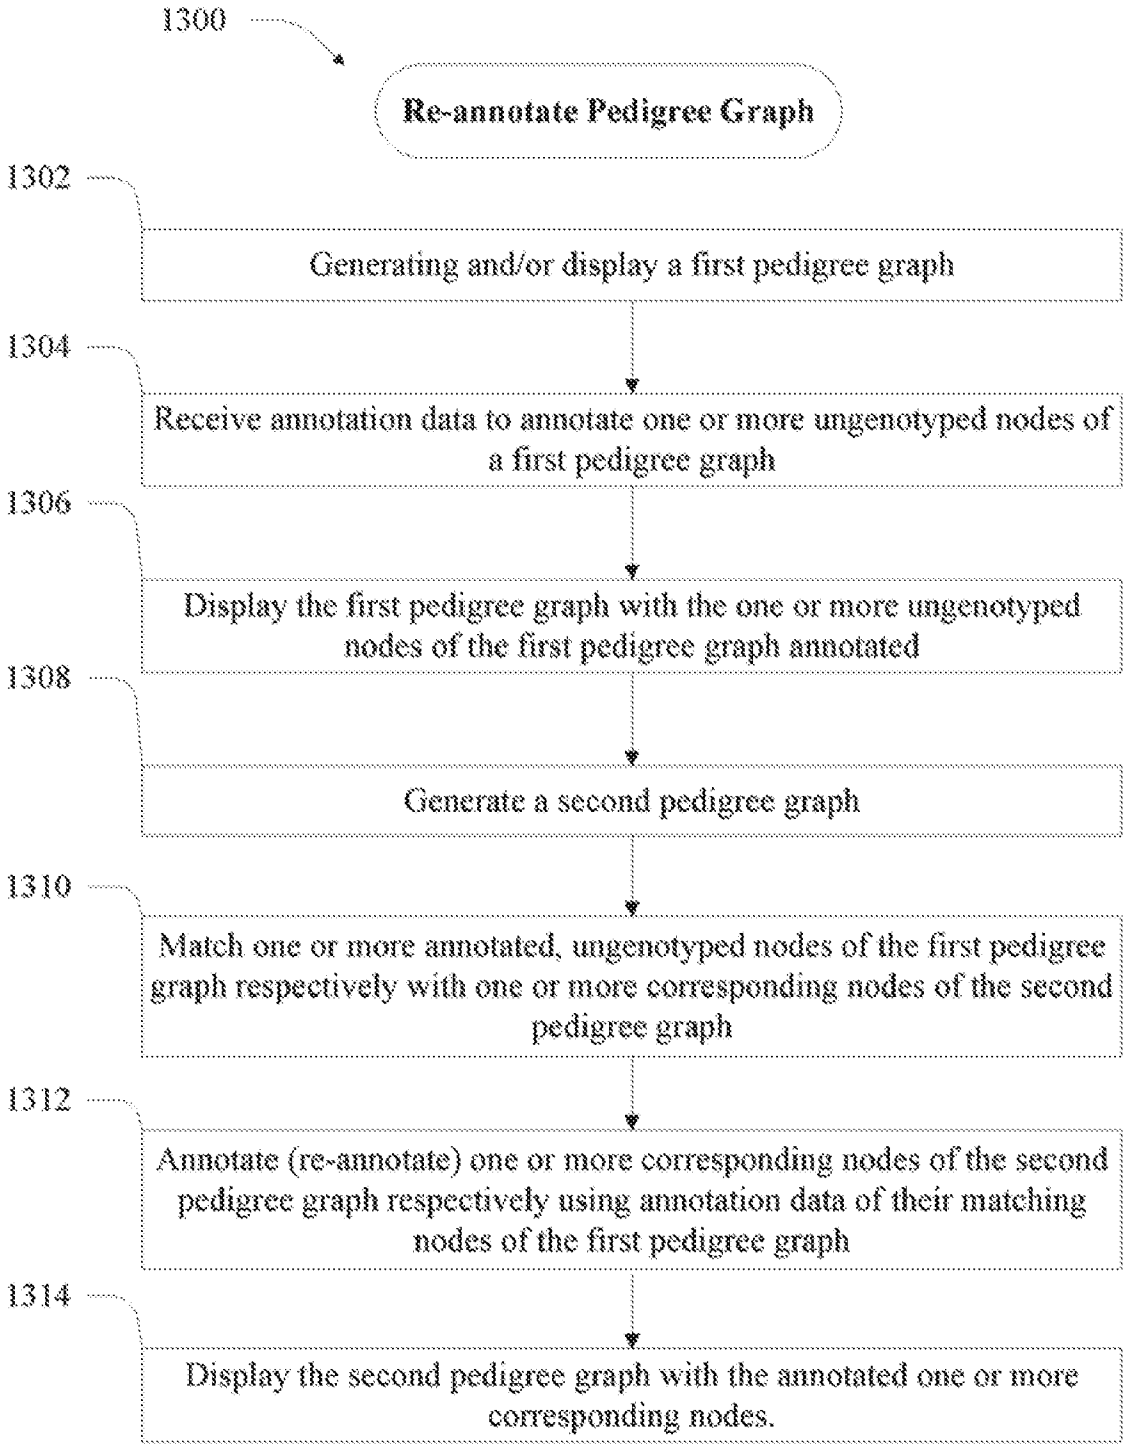
1300
Re-annotate Pedigree Graph
1302
Generating and/or display a first pedigree graph
1304
Receive annotation data to annotate one or more ungenotyped nodes of a first pedigree graph
1306
Display the first pedigree graph with the one or more ungenotyped nodes of the first pedigree graph annotated
1308
Generate a second pedigree graph
1310
Match one or more annotated, ungenotyped nodes of the first pedigree graph respectively with one or more corresponding nodes of the second pedigree graph
1312
Annotate (re-annotate) one or more corresponding nodes of the second pedigree graph respectively using annotation data of their matching nodes of the first pedigree graph
1314
Display the second pedigree graph with the annotated one or more corresponding nodes.

Figure 25

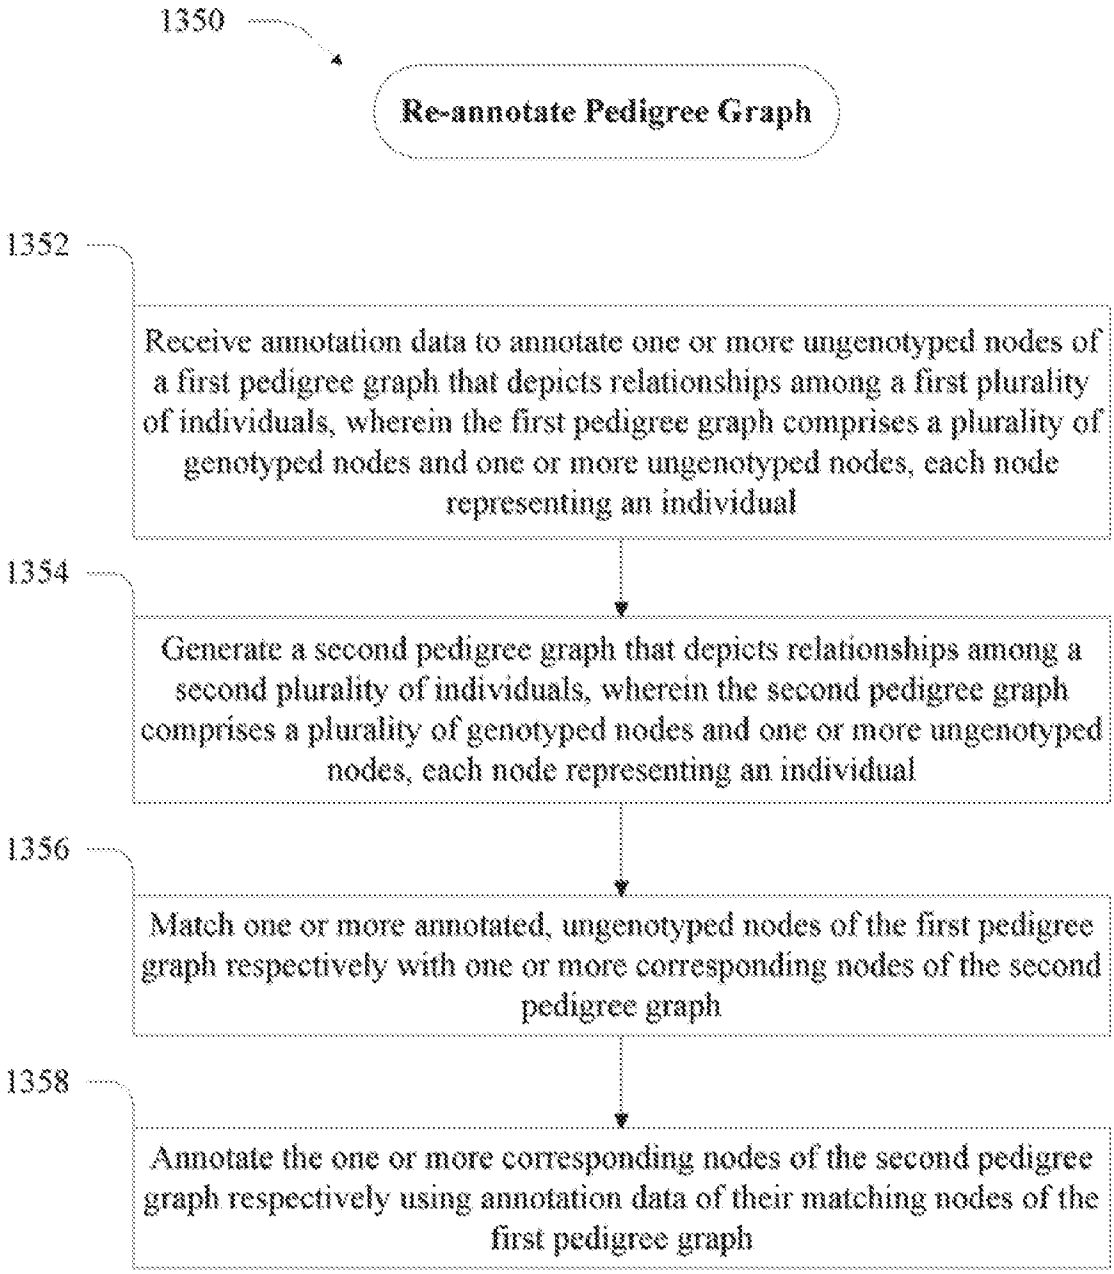
1350
Re-annotate Pedigree Graph
1352
Receive annotation data to annotate one or more ungenotyped nodes of a first pedigree graph that depicts relationships among a first plurality of individuals, wherein the first pedigree graph comprises a plurality of genotyped nodes and one or more ungenotyped nodes, each node representing an individual
1354
Generate a second pedigree graph that depicts relationships among a second plurality of individuals, wherein the second pedigree graph comprises a plurality of genotyped nodes and one or more ungenotyped nodes, each node representing an individual
1356
Match one or more annotated, ungenotyped nodes of the first pedigree graph respectively with one or more corresponding nodes of the second pedigree graph
1358
Annotate the one or more corresponding nodes of the second pedigree graph respectively using annotation data of their matching nodes of the first pedigree graph

Figure 26

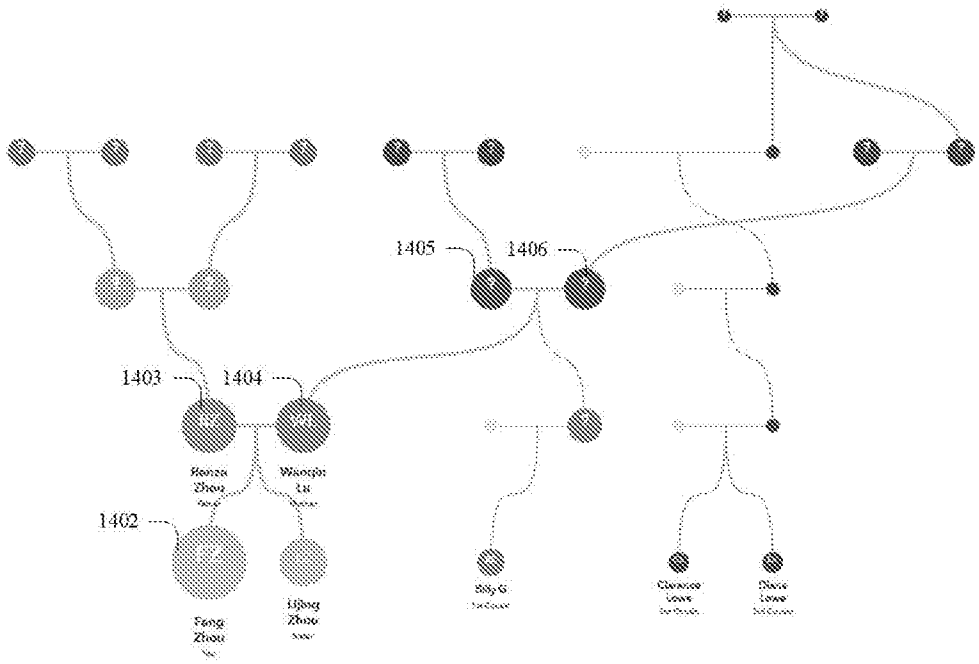

Figure 27

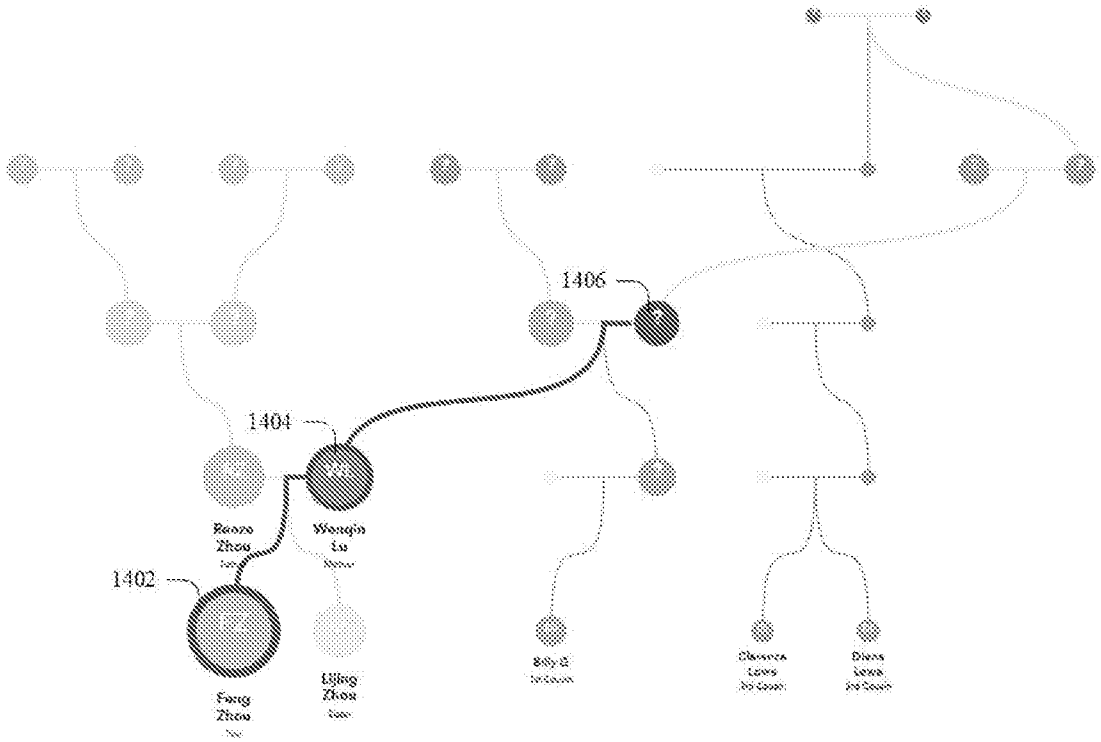

Figure 28

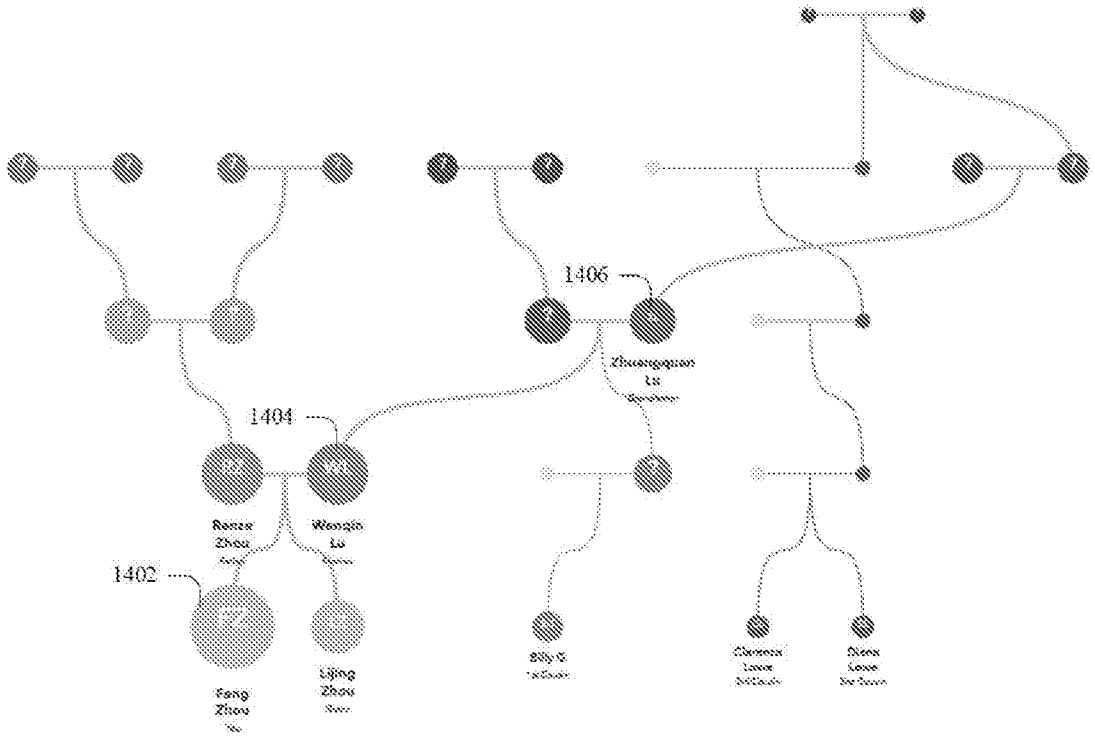

Figure 29

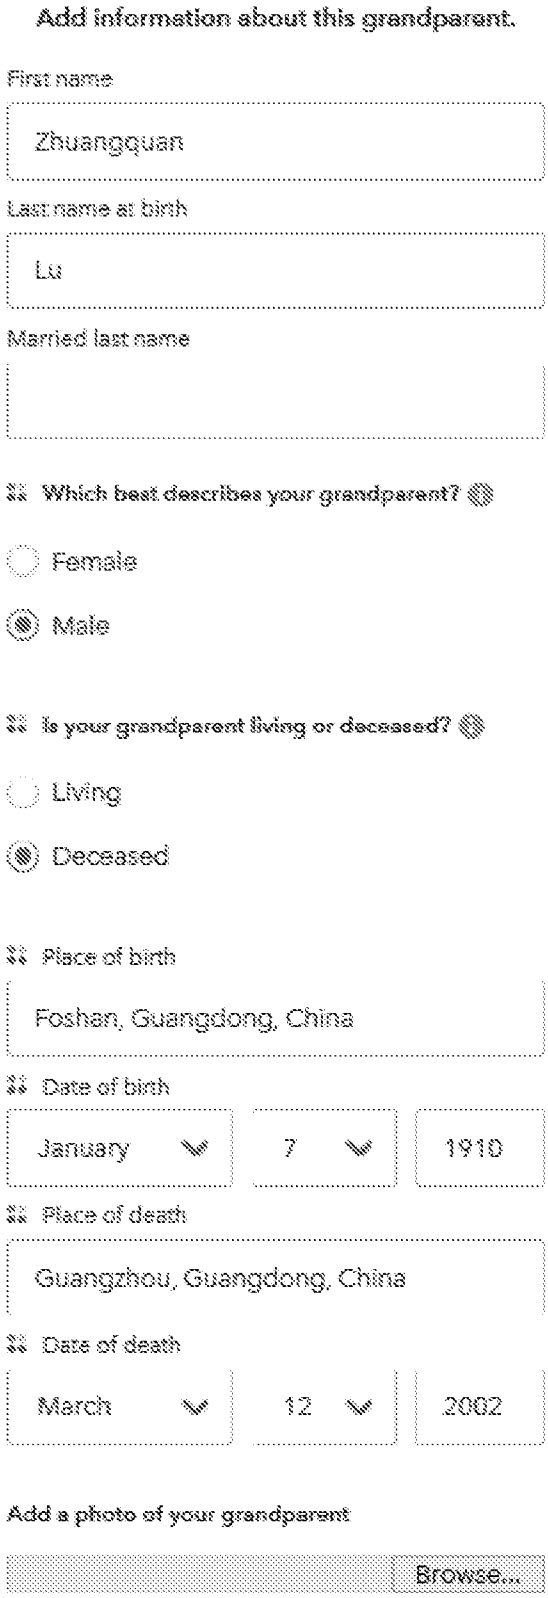
Add information about this grandparent.
First name
Zhuangquan
Last name at birth
Lu
Married last name
Which best describes your grandparent?
Female
Male
Is your grandparent living or deceased?
Living
Deceased
Place of birth
Foshan, Guangdong, China
Date of birth
January
7
1910
Place of death
Guangzhou, Guangdong, China
Date of death
March
12
2002
Add a photo of your grandparent
Browse...

Figure 30

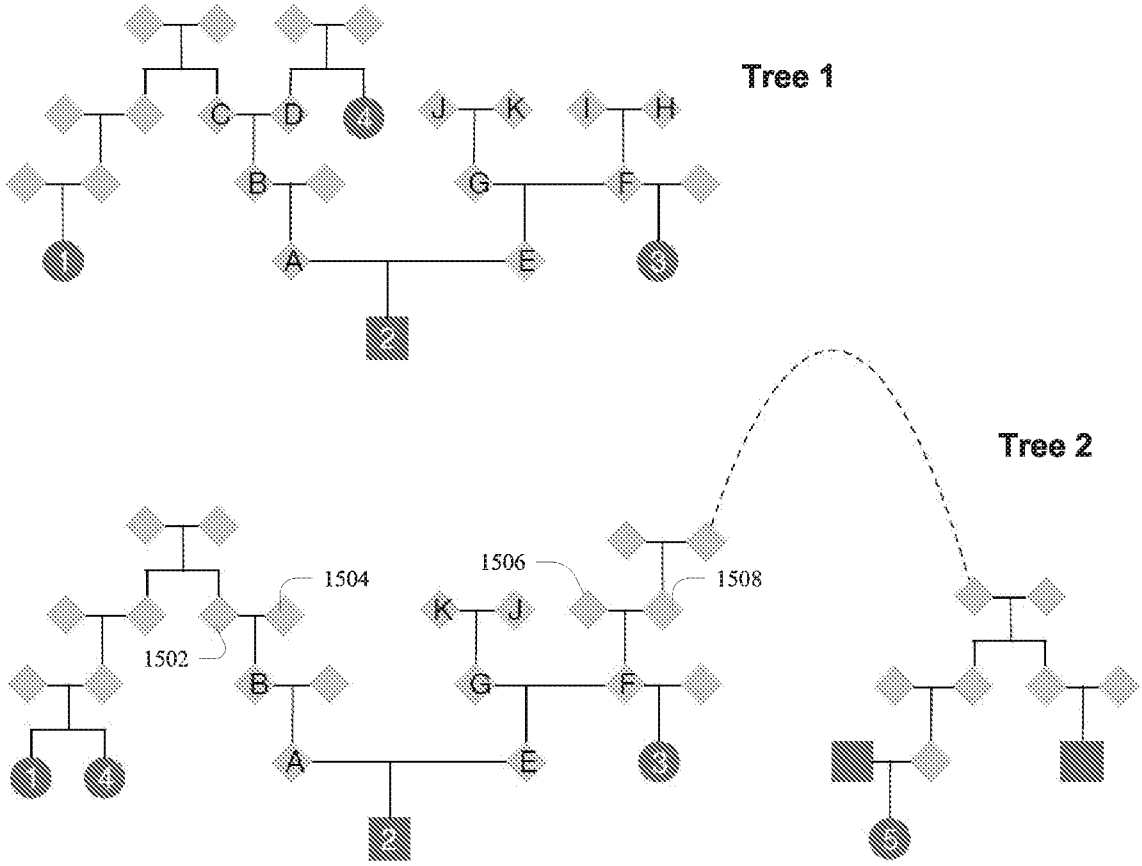
Tree 1
C
D
4
J
K
I
H
B
G
F
1
A
E
5
2
Tree 2
1504
1506
1508
1502
K
J
B
G
F
1
4
A
E
5
2
3

Figure 31

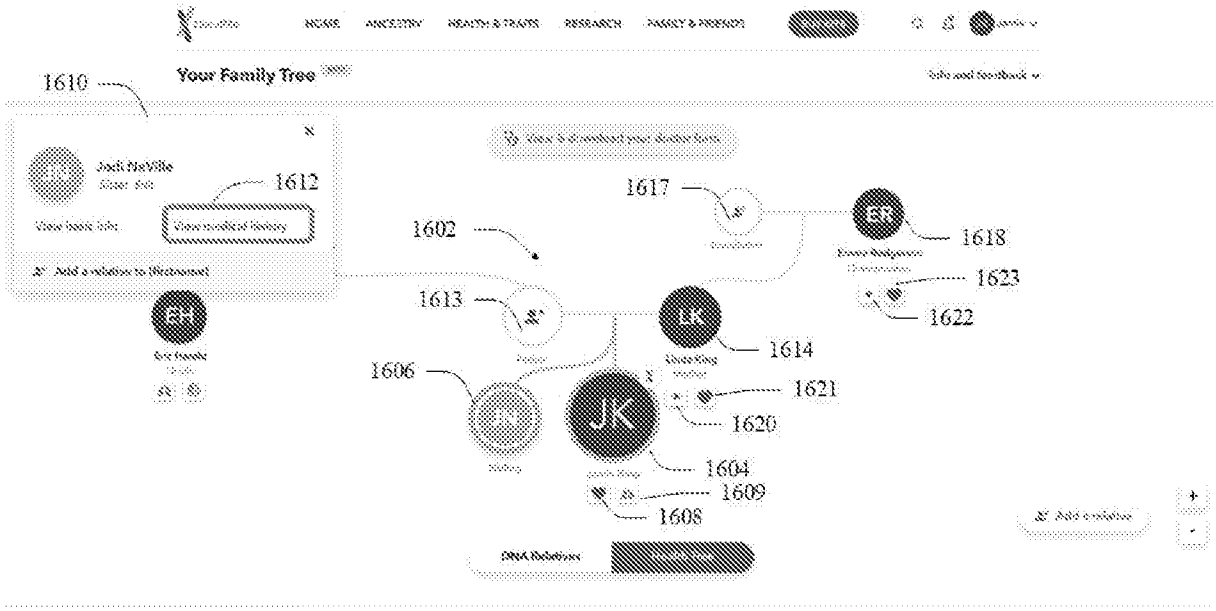
Your Family Tree
1610
1612
1602
1617
1618
1623
1622
1613
1614
1606
1621
1620
1604
1609
1608

Figure 32

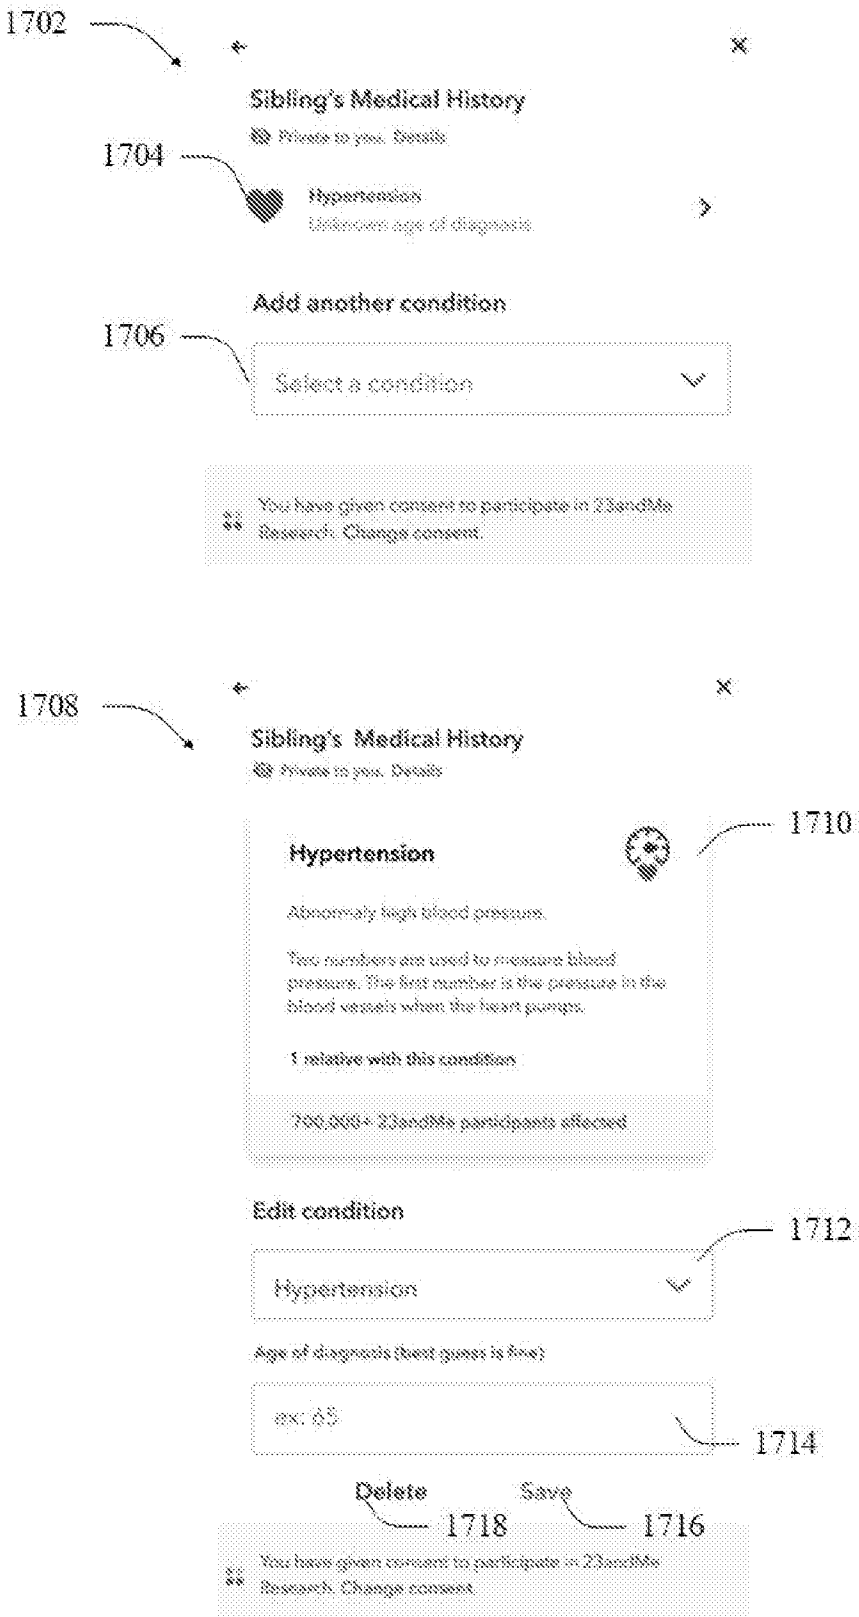
1702
Sibling's Medical History
1704
Hypertension
Add another condition
1706
Select a condition
You have given consent to participate in 23andMe Research. Change consent.
1708
Sibling's Medical History
1710
Hypertension
Abnormally high blood pressure.
1 relative with this condition
Edit condition
1712
Hypertension
Age of diagnosis (best guess is fine)
ex: 65
1714
Delete
Save
1718
1716
You have given consent to participate in 23andMe Research. Change consent.

Figure 33

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US20/50582

A. CLASSIFICATION OF SUBJECT MATTER

IPC - G16B20/40, 20/00; C12Q1/68; C12N15/10 (2021.01)

CPC - G16B20/40, 20/00; C12Q1/6827, 1/6809, 1/6837; C12N15/10, 15/102, 15/1034, 15/1072; C12Q1/6841

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
See Search History document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
See Search History document

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2017/0213127 A1 (DUNCAN, MC) 27 July 2017; see entire document	1-5, 38-42, and 98-99
A	US 2016/0300012 A1 (ANCESTRY.COM DNA, LLC) 13 October 2016; see entire document	1-5, 38-42, and 98-99
A	US 2014/0025308A1 (UNIVERSITY OF UTAH RESEARCH FOUNDATION) 23 January 2014; see entire document	1-5, 38-42, and 98-99
A	US 2005/0176031 A1 (SEARS, CP) 11 August 2005; see entire document	1-5, 38-42, and 98-99

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"D" document cited by the applicant in the international application
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
07 January 2021 (07.01.2021)	08 FEB 2021
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-8300	Authorized officer Shane Thomas Telephone No. PCT Helpdesk: 571-272-4300

Form PCT/ISA/210 (second sheet) (July 2019)

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US20/50582

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☒ Claims Nos.: 6-35, 43-54, 62-72, 79-97, 111-119
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

Group I: Claims 1-5, 38-42, 98-99; Group II: Claims 36-37; Group III: Claims 55-61; Group IV: Claims 73-78; Group V: Claims 100-110

-***-Continued within extra sheet-***-

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-5, 38-42, 98-99

Remark on Protest

☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.

☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.

☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US20/50582

-***-Continued from Box No. III - Observations where unity of invention is lacking -***-

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fee must be paid.

Group I: Claims 1-5, 38-42, and 98-99 are directed towards a method, system, and computer program product for determining pedigree relationships.
Group II: Claims 36-37 are directed towards a method of combining two or more pedigrees into a larger pedigree.
Group III: Claims 55-61 are directed towards a method of generating a pedigree graph by determining root entries.
Group IV: Claims 73-78 are directed towards a method of displaying a pedigree graph.
Group V: Claims 100-110 are directed towards a method of generating a pedigree graph, comprising genotyped nodes.

The inventions listed as Groups I-V do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons:

The special technical features of Group I include at least a probabilistic relationship model; pairwise Identity-by-Descent (IBD) data; age data; and likelihoods of various potential relationships, which are not present in Groups II-V.

The special technical features of Group II include at least identifying a first set of individuals in the first pedigree who share IBD with individuals in the second pedigree; identifying a second set of individuals in the second pedigree who share IBD with individuals in the first pedigree; identifying a common ancestor of the first set of individuals; identifying a common ancestor of the second set of individuals; inferring a degree of relatedness between the common ancestor of the first set and the common ancestor of the second set; and connecting the two common ancestors by an inferred degree of relatedness between the common ancestors to form a larger pedigree, which are not present in Groups I and III-V.

The special technical features of Group III include at least determining a minimal set of root entries from which all of the plurality of entries are reachable by starting from the minimal set of root entries and traversing from parent entries to child entries and traversing between partner entries, wherein partner entries are two parent entries of a same child entry, and wherein a root entry is an entry having no parent entries; for each root entry of the minimal set of root entries, forming a subtree using the root entry as a root node and entries reachable from the root entry as additional nodes, thereby obtaining a plurality of subtrees, which are not present in Groups I-II and IV-V.

The special technical features of Group IV include at least each pair of two or more pairs of parent nodes comprises two nodes rendered in different colors, and the lines connecting each child node to its pair of parent nodes comprise curved lines, which are not present in Groups I-III and V.

The special technical features of Group V include at least receiving, by the processor, annotation data to annotate one or more ungeno• typed nodes; a plurality of genotyped nodes and one or more ungenotyped nodes, which are not present in Groups I-IV.

The common technical features shared by Groups I-V are a method, implemented using a computer system that includes one or more processors and system memory; generating pedigree graphs; and links relating each child entry to its parent entries, wherein a child entry and its parent entries represent a child and its biological parents.

However, these common features are previously disclosed by: US 2017/0213127 A1 (DUNCAN). Duncan discloses a method, implemented using a computer system that includes one or more processors and system memory (computer software and systems for Genomics assisted Genealogy; para [0001]); generating pedigree graphs (a virtual family tree, VFT (pedigree graph), is made; Fig. 11; para [0353], [0525]); and links relating each child entry to its parent entries, wherein a child entry and its parent entries represent a child and its biological parents (nodes and edges form a traditional pedigree view of a family tree and all nodes below the top row (child entries) have mother and father pointers (links); para [0525]).

Since the common technical features are previously disclosed by Duncan, these common features are not special and so Groups I-V lack unity.