Detection of stable community structures within gut microbiota co-occurrence networks from different human populations

Matthew A Jackson[1], Marc Jan Bonder[2], Zhana Kuncheva[3], Jonas Zierer[1,4], Jingyuan Fu[2,5], Alexander Kurilshikov[2], Cisca Wijmenga[2,6], Alexandra Zhernakova[2], Jordana T Bell[1], Tim D Spector[1], and Claire J Steves[1*]

[1]**Department of Twin Research and Genetic Epidemiology, King's College London, London, UK**
[2]**University Medical Center Groningen, Department of Genetics, Groningen, University of Groningen, The Netherlands**
[3]**Department of Mathematics, Imperial College London, London, UK**
[4]**Institute of Bioinformatics and Systems Biology, Helmholtz Zentrum München, Neuherberg, Germany**
[5]**Department of Pediatrics, University Medical Center Groningen, University of Groningen, Groningen, The Netherlands**
[6]**K.G. Jebsen Coeliac Disease Research Centre, Department of Immunology, University of Oslo, Norway**

Corresponding author:
Claire Steves[1]

Email address: claire.j.steves@kcl.ac.uk

ABSTRACT

Microbes form sub-communities of interacting taxa in the the gut microbiome. The interactions defining these communities can be inferred from the co-occurrence of taxa across multiple samples. Here, we present an approach to identify comparable communities within different gut microbiota co-occurrence networks. We demonstrate its use by comparing the gut microbiota community structures of three geographically diverse populations. We combine gut microbiota profiles from 2764 British, 1023 Dutch, and 639 Israeli individuals and derive co-occurrence networks between their operational taxonomic units. Applying our approach, we then detect comparable communities within them. Comparing populations we find community structure is significantly similar between datasets than expected by chance. Mapping communities across the datasets, we also show that communities can have similar associations to host phenotypes in different populations. This study shows that the community structure within the gut microbiota is stable across populations and provides a novel approach that facilitates comparative community-centric microbiome analyses.

INTRODUCTION

The gut microbiome is a complex bacterial community whose structure is determined by many factors including the interactions between its members. Bacteria can interact in numerous ways - in either an actively targeted or passive manner, which can result in both beneficial and detrimental effects for the parties involved (Faust and Raes, 2012). Given the increasing evidence of the importance of the gut microbiome in human health, it is necessary to understand the inter-microbial effects underlying its composition.

Gut microbiota are frequently profiled using marker gene sequencing. Sequencing reads from amplicons of the chosen marker are typically collapsed to operational taxonomic units (OTUs), analytical units used to approximate taxonomic abundances (Schloss and Handelsman, 2005; Navas-Molina et al.,

2013). One approach to infer interactions between bacteria in the gut microbiota is to quantify the co-occurrence of OTUs across multiple samples (Faust and Raes, 2012). High correlation between OTUs can reflect the interactions between their source bacteria and similarities/differences in their responses to environmental conditions (Lozupone et al., 2012; Faust and Raes, 2012). However, OTU counts are relative to the sequencing depth of a sample, which introduces inherent correlations to the data (Friedman and Alm, 2012). As a result, several specialist approaches have been developed to estimate correlations from microbiota data (Faust et al., 2012; Friedman and Alm, 2012; Deng et al., 2012; Gevers et al., 2014; Fang et al., 2015; Kurtz et al., 2015). Whilst these have seen use within the research community (Gevers et al., 2014; Goodrich et al., 2014; Tong et al., 2013; McHardy et al., 2013), correlation metrics for microbiome studies have only recently been compared systematically by Weiss et al. who found that using an ensemble of metrics can improve the precision of co-occurrence detection (Weiss et al., 2016).

Given co-occurrence between OTUs it is possible to generate networks of their inferred interactions. Within this we would not expect all bacteria to interact but rather that subsets taxa are more likely to interact within one another forming identifiable and distinct interacting groups termed communities. The formation of such communities can be driven by factors such as cross-species metabolism and geospatial environmental variation (Faust and Raes, 2012; Levy and Borenstein, 2014; Lozupone et al., 2012). Previous studies have identified communities within microbial co-occurrence networks, often using the WGCNA (weighted gene co-expression network analysis) method developed to identify modules in gene co-expression networks (Lozupone et al., 2012; Tong et al., 2013; Jackson et al., 2016b; Duran-Pinedo et al., 2011; Langfelder and Horvath, 2008). However, whilst they may be more biologically accurate, such approaches allow OTUs to have weighted contributions to multiple communities, which complicates comparison and mapping of equivalent communities between networks. Using a community detection algorithm that assigns OTUs to single communities simplifies such analyses enabling comparison across datasets. This was demonstrated by the use of a modularity maximisation approach to detect and compare community structures between gut microbiome networks of irritable bowel syndrome patients and healthy controls (Baldassano and Bassett, 2016).

It has yet to be established if gut microbiota form similar communities in different human populations. Here, we describe a rational data-driven approach to the creation of ensemble co-occurrence networks and community detection using modularity maximisation. This produces robust community definitions that can be compared and mapped between different networks. Applying this to gut microbiota data from three distinct human populations in the UK, Netherlands, and Israel we find that OTUs form similar community structures across all three, and that these communities have similar associations with the health-related host factors of age and body mass index (BMI) in their respective populations.

MATERIALS AND METHODS

Data aggregation

Given our objective to compare network community structures across data sets (Figure 1A), we required OTUs that would be comparable between them. To this end, we carried out clustering of sequences across combined data from multiple sources. To maximise sequencing similarity, we selected two data sets whose experimental approach best matched that of the gut microbiota profiles obtained for TwinsUK through 16S rRNA gene sequencing of faecal samples. The LLDEEP and Israeli-PN datasets were selected as they carried out gut microbiota profiling by amplifying the V4 region of the 16S rRNA gene using the same PCR primers and used paired-end sequencing with read lengths sufficient to capture the whole V4 region. Notable differences between studies include faecal sampling and DNA extraction techniques. Both TwinsUK and LLDEEP utilised aliquots from faecal samples stored at -80°C, the Israeli-PN study utilised a mixture of faecal swabs stored at -80°C and OMNIgene-GUT stool collection kits stored at -20°C. All three studies used both chemical and mechanical lysis in DNA extraction but employed different protocols: TwinsUK utilised the MoBio PowerSoil HTP extraction kit, the LLDEEP cohort utilised the Qiagen AllPrep kit, and the Israeli-PN DNA was extracted using the MoBio PowerMag Soil DNA extraction kit.

Extraction of DNA from faecal samples, amplification and sequencing of the V4 region of the 16S rRNA gene, and demultiplexing of sequencing reads has previously been described for the TwinsUK cohort (Goodrich et al., 2014). The paired-end demultiplexed reads were joined using join_paired_ends within QIIME with an overlap of at least 200nt to form single reads covering the full V4 region (Navas-Molina et al., 2013). DNA extraction and 16S rRNA gene sequencing within the LLDEEP samples has been described in detail previously (Fu et al., 2015). Data from the LLDEEP cohort was provided in a similar format having
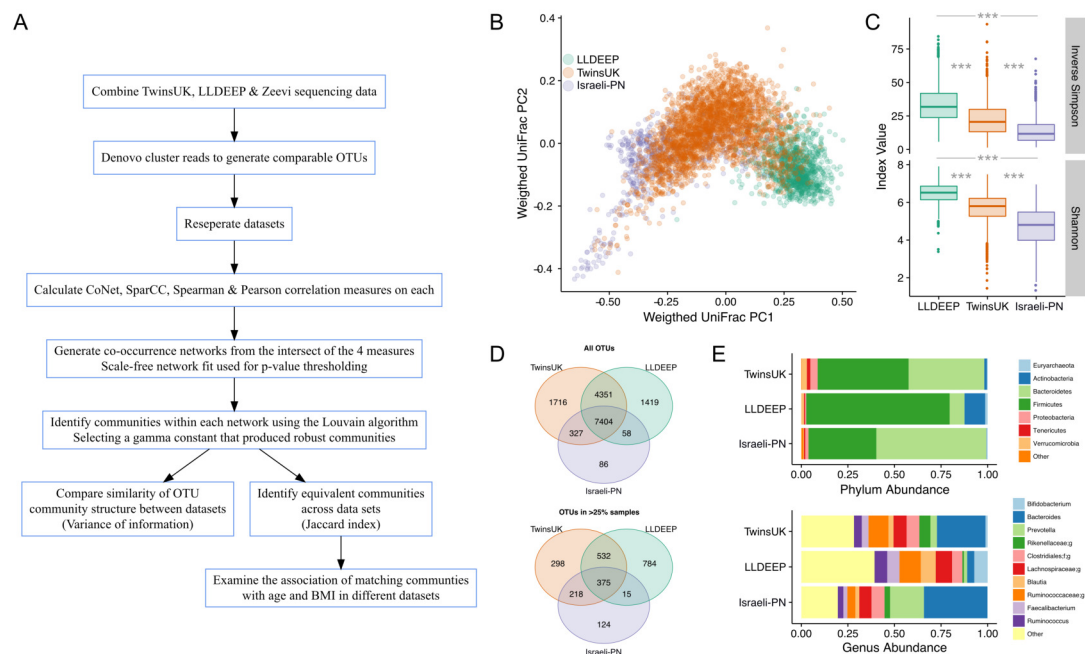
Figure 1. An overview of the study and the data used in analyses. A) An outline of the study design. B) Plot of the first two components of PCoA from weighted UniFrac distance measures between samples in the study, coloured by dataset, shows there is some separation but significant overlap in microbiota compositions by study. C) Comparison of alpha diversity measures in all three data sets. There was a significant difference in all pair-wise comparisons for both measures (Mann-Whitney U p<0.0001). D) Venn diagrams showing the number of OTUs shared across datasets. The top includes all OTUs, the lower diagram only those found in at least 25% of samples in each dataset. E) Comparison of the mean relative abundances of taxonomies at the phylum and genus level across the complete table for each data set. Phyla at less than 1% abundance and genera at less than 5% abundance in all sets are collapsed as Other. In genera f; and g; represent unassigned family and genus names in the Greengenes reference.

used custom scripts to merge the paired end data to full length reads covering the V4 region and split data by individual (Gevers et al., 2014; Fu et al., 2015). Raw data accompanying the Israeli-PN publication was downloaded from the European Nucleotide Archive (ENA) (Accession:PRJEB11532)(Zeevi et al., 2015). This was processed similarly to the TwinsUK data and demultiplexed using published barcode mappings. Samples not listed in the accompanying metadata or with ambiguous read identifiers were removed. The cleaned, joined and demultiplexed data was then concatenated to produce one sequencing file covering all three studies. After quality control, 4426 samples (2764 TwinsUK, 1023 LLDEEP, 639 Israeli-PN) were included in the analysis.

OTU clustering, table filtering, and diversity analyses

The complete sequencing from all three sets contained 381,767,528 reads. These were dereplicated, removing reads occurring only once, resulting in 5,728,288 unique reads. We chose to use a de novo approach to cluster these to OTUs, as de novo OTU clustering is not influenced by the reference database used, can capture novel diversity, and can produce more accurate clustering of OTUs than reference-based approaches (Westcott and Schloss, 2015; Jackson et al., 2016a). De novo clustering was carried out using the cluster_fast command in the VSEARCH package with a 97% similarity threshold. The resultant 94,070 representative sequences were filtered to remove chimeric reads using the uchime_denovo command (within VSEARCH), producing a final set of 17,123 representative sequences. These were used to generate OTU counts across all samples using the VSEARCH usearch_global command. Post-processing, average read counts were 82,695±745 for TwinsUK, 49,962±964 for LLDEEP, and 130,378±5,534 for Israeli-PN (mean±SEM).

A phylogenetic tree was generated from the representative sequences using the default parameters of

120 the make_phylogeny command in QIIME. Taxonomy of OTUs was assigned by matching representative
121 sequences against the Greengenes v13_8 database using the default parameters of the assign_taxonomy
122 command in QIIME. OTUs occurring in only one sample, and samples with less than 10,000 OTU counts
123 were removed. Weighted and unweighted UniFrac beta diversity measures and subsequent principle
124 co-ordinates analysis of them was carried out using the beta_diversity_through_plots script in QIIME. The
125 combined OTU table was then split by data set. For the purposes of alpha diversity calculations, the raw
126 counts tables were rarefied to a depth of 10,000 reads. For each sample, Shannon and inverse Simpson
127 diversity indices were calculated as the mean across ten rarefactions. Significant differences in alpha
128 diversity between datasets were assessed using Mann-Whitney U-tests.

Calculating co-occurrence measures

130 Co-occurrence calculation was carried out on each data set independently. Sub-tables were generated
131 from raw (unrarefied) OTU tables that only contained OTUs found in at least 25% of the samples. All co-
132 occurrence measures were calculated within these subsets as they are less sparse and hence more amenable
133 to correlation measures (Weiss et al., 2016). The mean diversity was assessed on rarefied versions of
134 the subset tables using the inverse Simpson index as for the full tables. OTU sparsity was assessed
135 from the unrarefied table using the biom summarize_table command. Following the recommendations of
136 Weiss et al. (2016), we then used these estimates of the mean inverse Simpson index (TwinsUK=20.2,
137 LLDEEP=29.0, and Israeli-PN=13.1) and OTU table sparsity (0.49 in all three) to select an ensemble
138 approach to co-occurrence detection that combines four different correlation measures: CoNet, SparCC,
139 Pearson's, and Spearman's.

CoNet

141 CoNet is itself an ensemble approach. The package allows a range of co-occurrence measures to be
142 combined with several options for how to combine the weighting of edges and their p-values. CoNet
143 addresses compositionality within the data using the ReBoot procedure, which involves permutation
144 followed by renormalisation of data. Calculating co-occurrence within these renormalized data allows
145 assessment of the levels of correlation expected simply as a result of the compositionality within the data
146 (Faust and Raes, 2012). For this study we used four measures of co-occurrence within CoNet: Spearman's
147 and Pearson's correlations and Kullback-Leibler and Bray-Curtis distance measures (Faust and Raes,
148 2012). Initial correlation thresholds were selected for each of these measures that produced 2000 positive
149 and 2000 negative edges concordant across the four metrics (Weiss et al., 2016); 1000 permutations were
150 then used for renormalisation to account for compositionality, and bootstrapping to identify edge p-values
151 for each metric. The Simes method was chosen to merge p-values across edges by keeping the minimum.
152 Final p-values were adjusted for multiple testing using the Benjamini-Hochberg FDR approach.

SparCC

154 SparCC was developed to calculate correlations between OTU abundances in microbiome data whilst
155 accounting for their inherent sparsity and compossitionality (Friedman and Alm, 2012). It uses the centred
156 log-ratio transformation to address data compositionality. SparCC was used with default parameters to
157 calculate correlations from the raw count OTU tables (Friedman and Alm, 2012). The MakeBoostraps
158 command was used to generate 100 bootstrapped tables, which were in turn used to calculate SparCC
159 correlations. The bootstrapped correlations were then used with the PseudoPvals command to generate
160 two-tailed p-values for the SparCC correlations from the true table.

Pearson's and Spearman's

162 Pearson's and Spearman's correlation metrics do not take data compositionality into account, but were
163 included to follow the approach outline by Weiss et al. (2016). Both measures were calculated using the
164 relative abundance tables. The rcorr function from the Hmisc R package was used to calculate correlations
165 and generate two-tailed p-values pairwise between all OTUs for both Pearson's and Spearman's measures
166 (Harrell Jr et al., 2008). P-values were adjusted for multiple testing using the Bonferroni method in R.
167 The outputs of all four co-occurrence approaches were converted into simplified unweighted edge
168 tables detailing the direction of association (1 or -1) and bootstrapped/adjusted p-values.

P-value thresholding to generate co-occurrence networks

170 Intersected networks were generated by combining the edge tables from the CoNet, SparCC, Pearson's,
171 and Spearman's methods. This was done independently for each data set. Edges were retained in the final

network if the direction of co-occurrence matched and the edge p-values were below a given threshold in all four methods. This was carried out at multiple different p-value thresholds (0.05, then ranging in powers of ten from 0.01 to 10^{-8}), to generate multiple intersect networks for each dataset with a gradient of stringency for edge inclusion. The resultant networks were then assessed to determine how well their structure fit that of a scale-free network to identify the optimum p-value threshold to use.

A scale-free network is a network whose node degree distribution follows a power law, i.e. there are few highly connected nodes and many more less connected nodes, and is the expected distribution of several biological networks (Jeong et al., 2000; Albert, 2005; Zhang and Horvath, 2005). Fit to a scale free network was calculated by first extracting the degree distribution of a network using igraph (Csardi and Nepusz, 2006). The fit of this distribution was then assessed using the scaleFreeFitIndex function from the WGCNA package in R (Langfelder and Horvath, 2008). This provides the R2 of the fit to a scale-free model, which the creators of WGCNA suggest should be >0.8, and the slope of the fit, which they suggest should be close to -2 to indicate a good fit. The optimum p-value threshold was selected based on these criteria across all three data sets (see Results) and the resultant intersect network at this threshold used in all further analyses. Visualisation and generation of descriptive statistics from networks was carried out using Gephi (Bastian et al., 2009).

Detecting communities within co-occurrence networks

Between OTU adjacency matrices were generated that represented the final intersect networks for each dataset. Negative correlations were removed (considered zeroes in the adjacency matrix) to generate unsigned networks as they represented a considerably small proportion of edges ($<1\%$ in all data sets). Community detection was carried out on each dataset's network independently using the genLouvain 2.0 package within MATLAB (Mucha et al., 2010), which implements the Louvain approach to modularity maximisation (Blondel et al., 2008). This defines community partitions by assigning nodes to unique communities, then iteratively combining neighbouring nodes into communities if it results in an increase in modularity across the whole network. Modularity is a measure of the number of edges within communities relative to between communities and a higher value represents better community definition (Newman and Girvan, 2004).

The genLouvain algorithm includes a γ parameter, which controls the size and number of detected communities (Mucha et al., 2010). A smaller gamma value promotes the detection of a small number of larger communities, while larger gamma values promote the detection of a high number of smaller communities. To find an optimal value for the γ parameter, we carried out community detection using a range of γ values for each dataset (0.1-1, increments of 0.01), and assessed the stability and significance of the resultant communities.

Stability

To assess the stability of community definitions at each γ, we carried out community detection 25 times on the real network followed by pairwise comparisons of the similarity of community clustering between the runs. To assess similarity between two community groupings we used the normalised variation of information (variation of information divided by ln(number of nodes in the network)) as a measure of similarity between assignments. A high value for variation of information means OTUs are sub set more differently in the two partitions compared, whereas a value of zero means the two partitions are identical. In figures and text, we report 1-normalised variation of information so a higher value represents more similar community structure.

Significance

To assess the significance of the community definitions, we also carried out community detection at each γ on 100 randomised networks whose nodes followed the same degree distribution as the real network (generated using the randomGraphFromDegreeSequence command in the Octave networks toolbox package). We then compared the mean modularity of the 25 runs on the real network to the 100 randomised networks.

From observation of both the variation of information (stability) and modularity (significance) results we then selected a suitable value for γ that produced both stable and significant community groupings (see results). Once a suitable value for γ was identified, community detection was repeated 100 times at this value, and the community definitions from the run producing the highest modularity were retained as the final community definitions.

5/19

Community properties and host associations in TwinsUK

The relative abundance of each community in a sample was found by summing the read counts of its constituent OTUs and dividing by the total number of reads observed in the sample. Association between the mean abundance of an OTU in a dataset and the number of OTUs in its parent community was assessed using Spearman's correlation. Taxonomy within communities was investigated by counting the number of OTUs assigned to different taxa at each taxonomic level.

The heritability of TwinsUK communities was estimated using the mets package in R. Log transformed relative abundances were used for each community. For each community ACE, CE, E, and AE models were fit using data from complete twin pairs. Co-variates included age, BMI, sample collection method, sex, and sequencing depth. For each community, the best fitting model was determined as that with the lowest Akaike information criterion.

Association analyses between community abundances and BMI and age was carried out using log transformed relative abundances of the communities ($\log10$(relative abundance + 10^{-6})). These phenotypes were selected as they were also available in both the LLDEEP and Israeli-PN data sets. To investigate associations linear models were fitted for each community using the lm function in R. Community abundance was the dependent variable with BMI, age, gender, and sequencing depth as independent variables, as these were the maximal set available across all datasets. The coefficient and significance of associations were extracted for BMI and age. P-values were then adjusted for multiple testing using the FDR method in R.

Comparison of community structure between populations

Overall community structure comparisons

We carried out pairwise comparisons between datasets to assess the similarity of overall community structures between the networks. In each comparison, the two networks were subset to just the OTUs shared by both and the normalised variation of information between their community structures calculated. To find the variation that would be expected by chance we generated randomised community sets for each network by shuffling OTU labels. Each randomised comparison therefore shared the same number OTUs between the two real networks with the same community sizes, but without the biological basis for the linkage of OTUs. We then carried out pairwise comparisons of the variation of information between the randomised communities. Shuffling and comparisons were repeated 1,000 times. The highest score observed (1- normalised variation of information) in any pairwise comparison, in any permutation was 0.41 (mean=0.34, SD=0.014).

Mapping of individual communities between datasets

To map communities between networks we carried out pairwise comparisons between all the communities in all three networks. We used the Jaccard index to quantify the number of OTUs shared between the two communities relative to the number not shared in each comparison. Matches were considered positive with scores >0.25 (range 0-1, no shared OTUs - complete overlap). This was chosen as above this threshold there were no instances of multiple mapping of communities between datasets.

Community-types were defined where matches could be found linking the communities in all three data sets. These were labelled with colour names using the standardColors function from the WGCNA package in R. The log transformed abundances for the 14 community-types that could be mapped across all three data sets were generated for the LLDEEP and Israeli-PN data sets as for TwinsUK. These were analysed for associations with age and BMI using linear regressions as for TwinsUK.

RESULTS

Integrating microbiota data from different populations

The aim of this study was to compare gut microbiota community structures across populations (study overview in Figure 1A). To this end comparable OTUs were generated by combining gut microbiota sequencing data from three geographically diverse populations. These were the British TwinsUK and Dutch Lifelines-DEEP (LLDEEP) cohorts and data from an Israeli personalised nutrition study (Israeli-PN) (Table 1) (Fu et al., 2015; Zeevi et al., 2015). Principle coordinates analysis of overall microbiota composition between samples found that whilst there was some grouping by data set, there was also significant overlap between them (Figure 1B). Comparison of median alpha diversities between the populations found significant differences between all three (Figure 1C).

Table 1. Participant summary statistics for the datasets considered in this study.

Dataset	BMI (Mean±SD)	Age (Mean±SD)	Sex (M/F/Unknown)	No. Samples
TwinsUK	26.1±4.8	59.5±12.3	308 / 2456 / 0	2764
LLDEEP	25.3±4.2	45.3±13.7	445 / 578 / 0	1023
Israeli-PN	26.4±5	43.5±12.9	376 / 251 / 12	639

Table 2. Summary statistics for the final intersect co-occurrence networks ($p<0.01$) for each dataset. Graph density is the percentage of all possible edges represented, mean path length is pairwise between all nodes, mean clustering coefficient is across all nodes in the network and provides a comparative indicator of overall clustering in the networks.

Dataset	Nodes	Edges	Mean Node Degree	Graph Density	Mean Path Length	Mean Clustering Coefficient
TwinsUK	844	2843	3.3	0.008	5.9	0.29
LLDEEP	922	2967	3.2	0.007	5.9	0.28
Israeli-PN	406	2573	6.7	0.033	3.7	0.31

Across the 15361 OTUs detected in the complete data 48% were shared across all three data sets, with 79% being found in at least two datasets (Figure 1D). There was also considerable overlap between TwinsUK and LLDEEP of OTUs not found in the Israeli-PN data. Similar patterns were observed when considering more abundant OTUs (found in >25% of the respective populations). Although LLDEEP and TwinsUK shared more OTUs, examining the mean taxonomic distributions at the phyla level the TwinsUK and Israeli-PN cohorts were most alike (Figure 1E). The LLDEEP cohort contained relatively lower levels of the phylum Bacteroidetes, and a higher abundance of Firmicutes bacteria. Further differences were observed at the genus level, where the Israeli-PN study had a higher average abundance of Prevotella.

Scale-free p-value thresholding of edges in co-occurrence networks

We split the OTU data and generated co-occurrence networks for each dataset using an ensemble approach - taking the intersect of four different correlation measures (TwinsUK example in Figure 2A). This required selection of a p-value threshold at which correlations were considered significant and retained as edges in the networks. Rather than make an arbitrary selection, we opted for a data-driven approach using the networks fit to a scale-free distribution to determine thresholds.

We generated ensemble networks for each dataset at a range of p-value thresholds. We then assessed how well the resulting networks' degree distributions fit a scale-free distribution (Figure 2B). We observed no consistent trend across the p-values, but this might be expected given the differences in the number of OTUs and samples in the datasets and varying levels of p-value precision between the co-occurrence measures used. However, we found that including all edges with a p-value <0.01 produced intersect networks with a good fit to the scale-free model in all three data sets. As such, this threshold was selected to generate the networks used for community detection.

From the 6761 unique edges observed across in all three networks, 166 were observed in all three, 882 in the TwinsUK and LLDEEP, 677 in TwinsUK and Israeli-PN, and 229 in the LLDEEP and Israeli-PN datasets. Summary statistics for the resultant networks are show in Table 2. The TwinsUK and LLDEEP network structures were most alike. The Israeli-PN data, which contained fewer nodes (OTUs), produced a smaller and more connected network.

Detection of communities in microbial co-occurrence networks

After establishing co-occurrence networks we used the Louvain modularity maximisation algorithm to detect communities within them (Mucha et al., 2010; Blondel et al., 2008; Newman and Girvan, 2004). The version used includes a constant (γ) that can be used to manipulate the number and size of resultant communities (see Materials and Methods). To determine an appropriate value for γ, we carried out repeated community detection on the three co-occurrence networks at various γ values and assessed the stability (as variation of information between runs) and significance (mean modularity of real compared to randomised networks) of the community definitions at each γ (Figure 3).
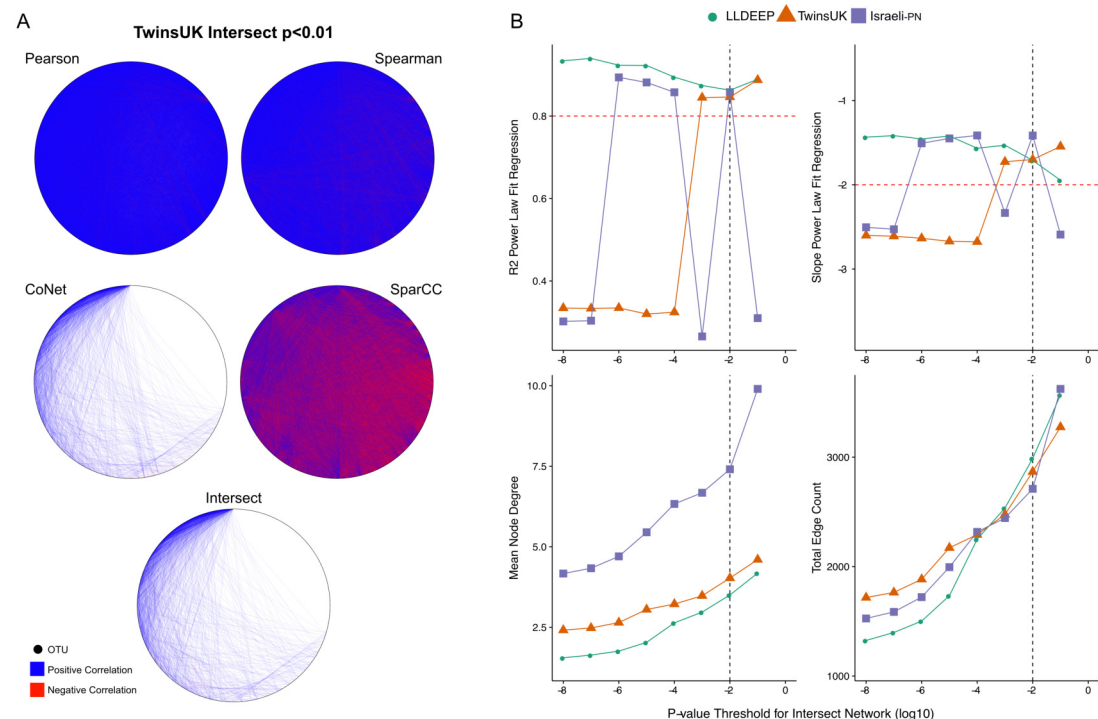
Figure 2. P-value thresholding of co-occurrence networks. A) A visualisation of the ensemble process at the p<0.01 threshold for TwinsUK, the top four networks combine to make the intersect network below. B) Selection of final p-value threshold for generating ensemble networks. Networks for each data set were made by intersecting four correlation measures with different p-value thresholds. Top left, the R2 of the resultant intersect networks fit to a scale-free distribution as assessed using WGCNA. Highlighted is 0.8 which the developers suggest as the minimum value to consider a good fit. Top right, the slope of the regression fit to a scale-free network returned by WGCNA, again highlighted is the developers recommended value. Bottom, summary measures of the resultant networks at each value showing the mean node degree of OTUs in the networks and the total number of edges in each.

Stability of community definitions at different γ values

We found that the community definitions were very stable (low variation of information between runs) at γ values approaching zero. This would be expected as in this instance most OTUs fall into a few large communities. Whilst the mean variation of information fluctuated across the γ range, it should be noted that the lowest estimates (1 – variation of information) were above 0.9 (range 0-1, least to most similar), showing that even the least stable γ produced very similar communities between runs.

Significance of community definitions at different γ values

At low γ values (with large communities) a networks edges are most likely to be within communities, hence we observed high modularity values for the true and randomised networks. As γ increased we found modularity estimates for the random networks dropped significantly lower than those of the true networks (Figure 3). This shows that the real co-occurrence networks have more community structure than would be expected by chance, and provides confidence that the communities identified at these γ values reflect the biological relationships of their member OTUs.

Final community definitions

From Figure 3, we selected a γ value of 0.4 for community detection. This γ provided high modularity estimates (that were significantly higher than the random networks) and good stability in all three data sets. We used the Louvain algorithm with γ=0.4 to define the final OTU communities in the three co-occurrence networks (Figure 4). Communities in each network are hence referred to using arbitrary numbers ranging from 1 to the number of communities in the network.

The number of communities detected within the networks (TwinsUK=96, LLDEEP=105, Israeli-PN=31) reflected the number of OTUs within them, but was more similar when only considering communities containing at least 5 OTUs (TwinsUK=35, LLDEEP=36, Israeli-PN=11). There was a positive correlation between the mean relative abundance of an OTU and the total number of OTUs in the community it was assigned to in the TwinsUK and LLDEEP networks (rs=0.1, p=0.002 and rs=0.27, p<0.0001 respectively). However, overall each community in all three networks contained a range of both high and low abundance OTUs (Supplementary Figure 1).

Taxonomy, heritability, and host associations of communities in TwinsUK

Taxonomic distribution within communities

Investigating taxonomic assignments of OTUs within the TwinsUK communities, we found that at finer taxonomic levels (genus and species) several communities contained OTUs with a mixture of assignments (Supplementary Table 1). Whilst some communities retained a mixture of taxa at the family level, the majority were dominated by one taxon (Figure 4B), and at higher levels nearly all contained a single taxon. A similar pattern was observed in the LLDEEP and Israeli-PN networks (Supplementary Figure 2).

The influence of host genetics on communities

To identify host associations with the microbial communities in TwinsUK, we used an ACE model (which estimates the variance attributable to additive genetics - A, to the environment common within twin pairs - C, and to the environment unique to individuals - E) to estimate the heritability of community abundances within 654 monozygotic and 495 dizygotic pairs. From the 96 communities in the TwinsUK network 52 had some variance attributed to additive genetic effects (Supplementary Table 1). Within these, the variance due to genetic effects ranged from 0.07-0.46 (Figure 4C). The estimates for the most heritable communities are of a similar magnitude to those of the most heritable taxa previously reported within TwinsUK (Goodrich et al., 2014). Indeed, examination of taxonomies within these communities found that they contained heritable taxa (Supplementary Table 1). For instance, the most heritable community (Community 39) contained a *Turicibacter* OTU, and the second most (Community 26) contained numerous Christensenellaceae and Ruminococcaceae OTUs (Goodrich et al., 2014).

Community associations with age and BMI

To further explore host effects, we investigated if communities associated with properties relating to health. We carried out linear regression analysis between community abundances and age and BMI in TwinsUK (Supplementary Table 2). These phenotypes were chosen as they were available for all three datasets enabling replication. In TwinsUK, we found that of the 96 communities 47 were significantly associated with BMI (FDR adjusted p<0.05). The highly heritable Community 26 that contained multiple Christensenellaceae OTUs had a significant negative correlation with BMI (FDR adjusted $p<10^{-10}$,
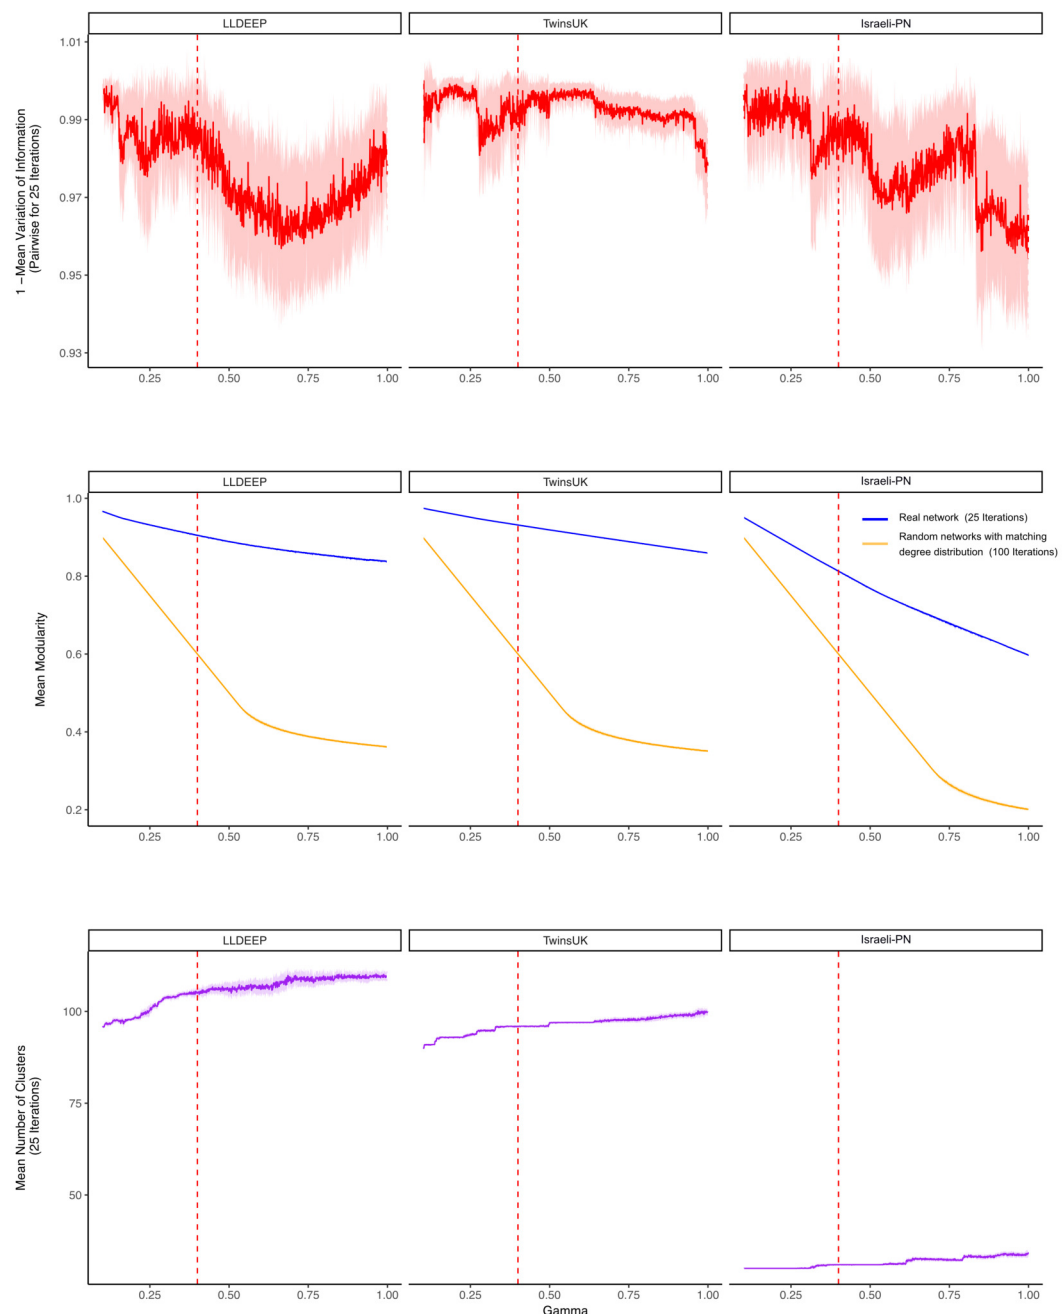
Figure 3. Selecting a γ parameter for Louvain thresholding. Louvain community detection was carried out multiple times on each network at a range of γ values. The top row shows the mean variation of information between the 25 runs at each gamma value. Plotted is 1-variation of information so a higher value means more stability in community definitions between runs. The middle row show the mean modularity of the 25 iterations compared to the mean modularity of 100 randomised networks with the same degree distribution. The lowest row shows the mean number of clusters generated at each value of γ. In all cases the mean line is surrounded by the standard deviation across the interactions. This is also true for the middle row, however modularity estimates had extremely low variance between runs, hence even a small difference between the random and true means is significant. The dashed red line indicates the final selected value for γ.
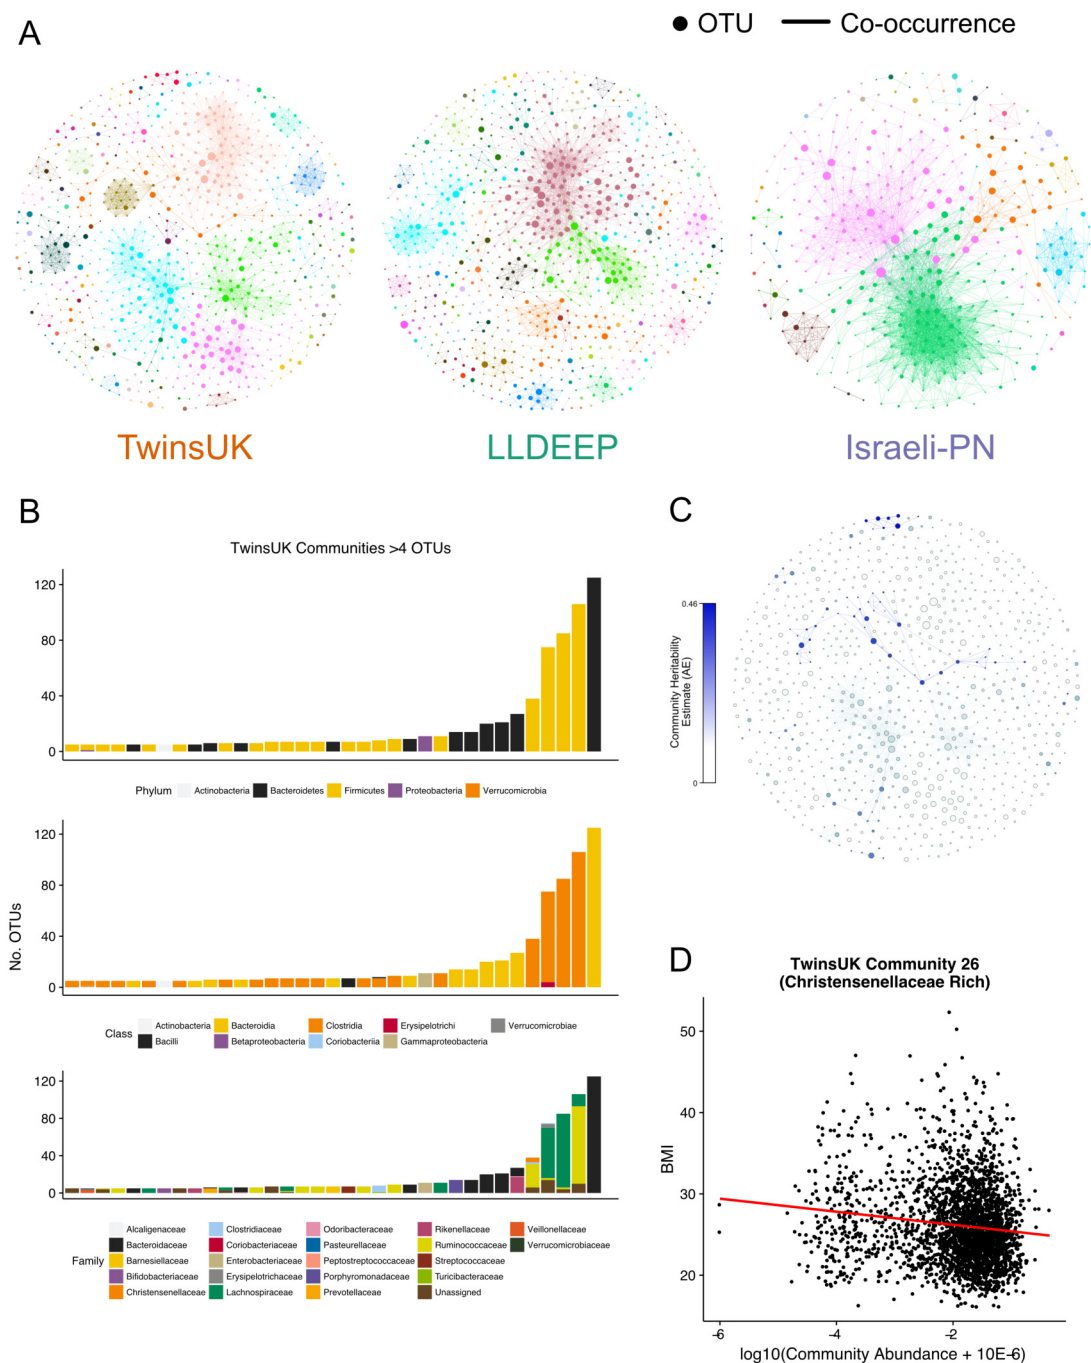
Figure 4. Communities detected within the co-occurrence networks. A) A visualisation of the communities detected by the Louvain algorithm. Communities are arbitrarily coloured. Node size represents an OTUs relative abundance within the dataset. B) Summary of the taxonomic distributions within the TwinsUK communities. Only communities with at least 5 OTUs are shown. Stacked bars represent the number of OTUs in the community assigned to each taxon. C) Community plot for TwinsUK as in A, but coloured by the heritability estimate of the community. D) Highlighting the significant negative association between BMI and the abundance of a Christensenellaceae rich community (linear regression with age, gender, and seq. depth β=-0.13, FDR q=0.001).

β=-0.13)(Figure 4D), reflecting our previous observation that Christensenellaceae in conjunction with correlated taxa are protective against weight gain in mice (Goodrich et al., 2014). There were several communities containing exclusively OTUs belonging to the order Clostridiales that were negatively associated with BMI. The strongest positive association was with Community 5 (FDR adjusted $p<10^{-7}$, β=0.11), which was a large community also dominated by Clostridiales OTUs. However, these were also assigned lower level taxonomies with the majority being various genera from the Lachnospiraceae family.

There were 48 communities significantly associated with age (FDR adjusted p<0.05). The three most significant associations were negative, two of these (Community 54 and 24), were dominated by *Bifidobacterium* OTUs. There was also a significant negative association with Community 27 (FDR adjusted p=0.007, β=-0.05), containing *Faecalibacterium prausnitzii* OTUs. The strongest positive association with age was observed with Community 17 (FDR adjusted $p<10^{-7}$, β=0.11), which contained exclusively Enterobacteriaceae OTUs. There were also several small communities positively associated with age that contained exclusively OTUs assigned to the Ruminococcaceae family.

Comparison of communities across populations
Similarities in overall community structure
Having identified communities within each dataset we aimed to determine how similar the segmentation of OTUs was between them. We again used the normalised variation of information measure to compare groupings and found that the similarity of community assignments was significantly higher in all three pair-wise comparisons than would be expected by chance (Figure 5A). This shows that OTUs are forming similar communities in all three data sets. However, it should be noted that variation of information can only be used to compare matching sets, so this only shows that the OTUs shared between the populations have a similar community structure.

Mapping equivalent communities between datasets
We then mapped the communities between the three networks using the Jaccard index to identify those which were equivalent to one another. Across the 96 TwinsUK, 105 LLDEEP, and 31 Israeli-PN communities, we found 14 instances where communities could be matched in all three data sets (Figure 5B), although there were many more that could be mapped across only two datasets. For distinction, we refer to these 14 matched groups as community-types and label each using arbitrary colour names (Figure 5B and Figure 5C, Supplementary Table 3). For example, the Green community type was community 45 in the LLDEEP network, community 31 in the Israeli-PN network, and community 36 in the TwinsUK network.

Communities have similar associations with age and BMI across populations
To determine if each of the 14 community-types had similar associations with age and BMI in their respective populations, linear regressions were carried out as for TwinsUK. Comparing the results (Figure 5D), we found that seven significant BMI associations within TwinsUK were replicated in the LLDEEP data, with two also significant in the Israeli-PN data. There were six community-types significantly associated with age in both the TwinsUK and LLDEEP data, four of which were also significant in the Israeli-PN data. BMI had more negative associations within the community-types, whereas age had more positive associations. For both BMI and age, the significance of community associations was most similar between TwinsUK and LLDEEP, this is may be due to their higher samples sizes relative to the Israeli-PN dataset.

The Greenyellow and Green community-types had significant negative associations with BMI in all three datasets. Greenyellow communities were dominated by Ruminococcus OTUs in all three, and similarly all the Green communities consisted solely of Ruminococcaceae OTUs. More widely, most community-types with at least one significant negative association with BMI consisted of Clostridiales OTUs. The Pink community type for instance had a significant negative association with BMI in TwinsUK and LLDEEP, and contained multiple Lachnospiraceae OTUs assigned to the species *Coprococcus eutactus* in all three data sets.

The Salmon, Turquoise, Cyan, and Red community-types were significantly positively associated with age across all three data sets. All three contained exclusively Clostridiales OTUs. The strongest positive association with age in TwinsUK was with the Brown module which was not associated with age in the other sets, this contained *Veillonella*, and *Haemophilus parainfluenzae* OTUs. Several associations for both age and BMI were significant in TwinsUK and LLDEEP but not in the smaller Israeli-PN dataset,

12/19
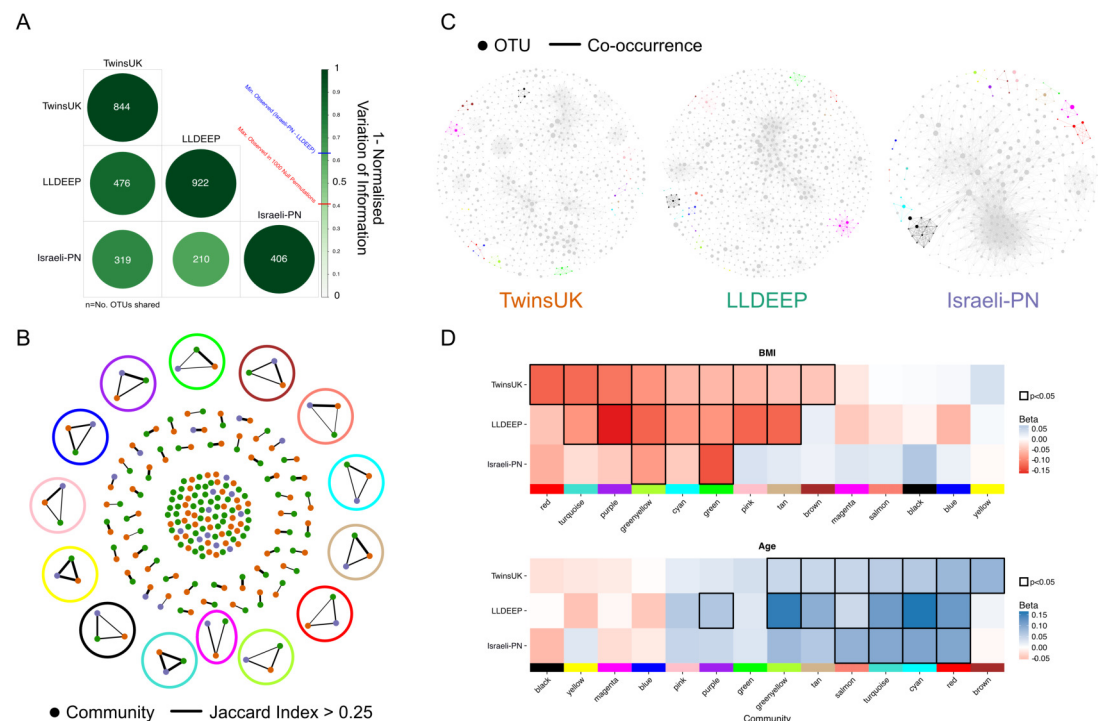
Figure 5. Comparison of communities across populations. A) Pair-wise comparison of variation of information between community definitions for OTUs shared between networks. 1-variation of information is shown where 1 represents identical segmentation of OTUs to communities and 0 indicates no similarity, with the highest score observed in randomised permutations being 0.41. Numbers represent the number of OTUs shared between the sets (OTUs with at least one edge). B) A network showing the matching of communities based on Jaccard index. Edges are considered where the index >0.25 and are weighted by the index 0.25-1. Highlighted are the community-types selected for later analysis, using the colour matching their name. C) Visualisation of networks as in 4A but coloured based on community-types as in B. Communities not in these 14 types are coloured grey. D) Replication of linear regression analysis for each of the 14 mapped communities in the LLDEEP and Israeli-PN comparison. Top comparing the results for BMI and below for age. Squares are coloured by the Beta estimate for the association and nominally significant associations (p<0.05) are highlighted by a black border. Community-types are ordered by their associations within TwinsUK.

but tended to share the same direction of effect. Overall, these results showed that communities associated with age and BMI in the same manner across the thee populations.

DISCUSSION

Here, we have presented a rationalised and data-driven approach to generating comparable co-occurrence networks from 16S rRNA gene sequencing data and identifying community structures within them. Applying this to gut microbiota profiles from three different populations we have found that both co-occurrence networks and community structures are stable across all three. Furthermore, we have shown that genetics, age and BMI can influence the relative abundances of the identified communities, with age and BMI having similar associations with communities in all three populations.

Differences between datasets

In combining these data sets we observed significant differences in the diversity and taxonomic profiles between datasets. These likely reflect a mixture of both the differences in the study populations and the experimental approaches used. Environmental differences between cohorts, in terms of genetics and other lifestyle factors are known to influence the gut microbiome (Goodrich et al., 2014; Falony et al., 2016; David et al., 2014). For instance, low diversity Prevotella dominant microbiomes have been previously observed in cohort studies at differential levels and are potentially linked to dietary intake (Arumugam et al., 2011; Wu et al., 2011). The larger proportion of Prevotella dominance, and associated reduced diversity, in the Israeli-PN gut microbiota might therefore reflect geographical and dietary differences. However, technical disparities such as faecal sampling and extraction method, which were known to be different between the present studies, are also known to influence taxonomic composition and cannot be discounted (Walker et al., 2015; Kennedy et al., 2014). The increased overlap in the TwinsUK and LLDEEP data in comparison to the Israeli-PN dataset in both the OTU and network structures might also result from the higher sample sizes in these studies. However, it is notable that even with the intrinsic population and technical differences between the three studies we still observed consistent elements across them. Most OTUs were found in at least two studies and there was overlap of samples from all three in beta diversity PCoA plots. It is within this common ground that similarities in community structure were observed.

Generating co-occurrence networks

To generate the co-occurrence networks in this study we used an ensemble approach as suggested by Weiss et al. (Weiss et al., 2016). However, we found that CoNet was the main edge filter when intersecting networks and might alone be sufficient. This might be expected given that it is itself an ensemble approach. However, metrics such as SparCC are not implemented within CoNet, and different correlation methods can produce better results depending on the format and properties of the input dataset (Weiss et al., 2016). These observations may therefore only apply to the current study.

We selected edges for inclusion in co-occurrence networks based on their fit to a scale-free topology. It has previously been shown that microbial interaction network structures can approximate a scale-free distribution (Chaffron et al., 2010). Scale-free networks are also well conserved in other aspects of biology across different domains of life (Wolf et al., 2002; Albert, 2005), and this approach is used by existing methods such as WGCNA (Langfelder and Horvath, 2008). The stability of the resultant communities in our study further shows that using scale-free thresholding can provide analytical benefits by generating comparable network structures prior to community detection.

Community detection

Following network creation, we utilised a modularity maximisation community detection algorithm that assigns OTUs to single communities. This may not represent the true nature of interactions, but provides analytical benefits as unambiguous definition of communities facilitates comparison between network structures. This has applications in comparison between disease and control groups (Baldassano and Bassett, 2016) or, as we have shown, replication of community associations across datasets. There is evidence that a community-centric approach to studying the gut microbiome in relation to human health can be more relevant than investigating individual taxa. Several studies have shown that associations between individual gut microbiota and host health can depend on their wider community context(Gevers et al., 2014; Goodrich et al., 2014; Baldassano and Bassett, 2016; Ridaura et al., 2013).

Beyond identifying communities, understanding their formation would require extension of the described approaches to quantification of taxa using metagenomic sequencing. This, especially used in tandem with metatranscriptomics and metabolomics, might provide indications to the mechanisms of interaction. However, cross-sectional approaches are inherently limited to inference of interactions from co-occurrence across samples. Time-series data and in vitro studies will also be required to delineate directional effects and validate individual interactions (Faust et al., 2015). For instance, it has been shown that interactions observed in pair-wise species co-cultures, can be used to predict outcomes in more complex multi-species cultures (Friedman et al., 2017). Combining in vitro observations with community interactions observed in the host environment could be a powerful tool for the design of synthetic bacterial communities for use as gut microbiome targeting medicines Lindemann et al. (2016).

The stability of communities across populations

Although 16S rRNA gene data cannot elucidate interaction mechanisms, we were able to determine several biological phenomena from the microbial approximations provided by OTUs. Most notably, that OTUs formed similar communities with similar host associations in the co-occurrence networks of three different populations. This shows that the variation that exists between the populations (e.g. the OTUs unique to each dataset) does not significantly alter the interactions between the OTUs shared across all three.

Several community associations with age and BMI in TwinsUK replicated in the LLDEEP and Israeli-PN datasets. These associations with age and BMI, and the heritability results within TwinsUK, broadly and reassuringly reflected previously observations from studies investigating taxa in isolation. For instance, in TwinsUK communities negatively associated with BMI were enriched with butyrate producers Ruminococcaceae and *Coprococcus eutactus* (Louis and Flint, 2009). Butryate producing taxa in the gut microbiome have previously been associated with obesity, although at a higher abundance in obese mice (Turnbaugh et al., 2006). This may reflect an increased capacity for energy utilisation by the gut microbiota in leaner individuals, as butyrate has also been shown to have protective effects against metabolic deficits associated with obesity (Qin et al., 2012; Gao et al., 2009). Two Bifidobacterium rich communities were negatively associated with age in line with previous studies (Yatsunenko et al., 2012; Odamaki et al., 2016). We also observed *F.prausnitzii* communities negatively associated and a community of Enterobacteriaceae positively associated with age, similar to previous observations within TwinsUK with frailty (Jackson et al., 2016b). Furthermore, the most heritable communities contained OTUs belonging to heritable taxa Goodrich et al. (2014).

It would be expected that community level associations reflected those of taxa based studies as each community constituted of taxonomically similar OTUs. This is in agreement with a previous study that found that co-occurrence to be higher between genetically similar taxa (Chaffron et al., 2010). This indicates either that interactions might evolve within closely related taxa; or that co-occurrence mainly detects genetically related taxa (more likely to have similar functionality) responding to environmental stimuli in the same manner. Either way, co-occurrence communities provide a useful method to reduce the dimensionality of 16S rRNA gene sequencing data to units reflecting biological phenomena. Further analyses using the described approaches will investigate the influence of host factors such as genetics, diet and health, on the formation of these communities and their reciprocal influences on the host.

CONCLUSION

We have described a data-driven approach to identify robust and comparable community definitions within microbiota co-occurrence networks. This should facilitate comparative, community-centric analyses of microbiome data. Furthermore, our results show that the gut microbiome contains stable communities of bacteria that are similarly associated with host factors across different populations.

DATA AVAILABILITY

TwinsUK 16S rRNA gene sequencing data is available from the ENA (Accession:ERP015317). LLDEEP 16S rRNA gene sequencing data is available from the EGA (Accession: EGAD00001001991). The data from the Israeli-PN study was obtained from the ENA (Accession:PRJEB11532). Code used for the main stages of OTU and co-occurrence network creation and the identification of communities within the networks can be found in Supplementary Materials.
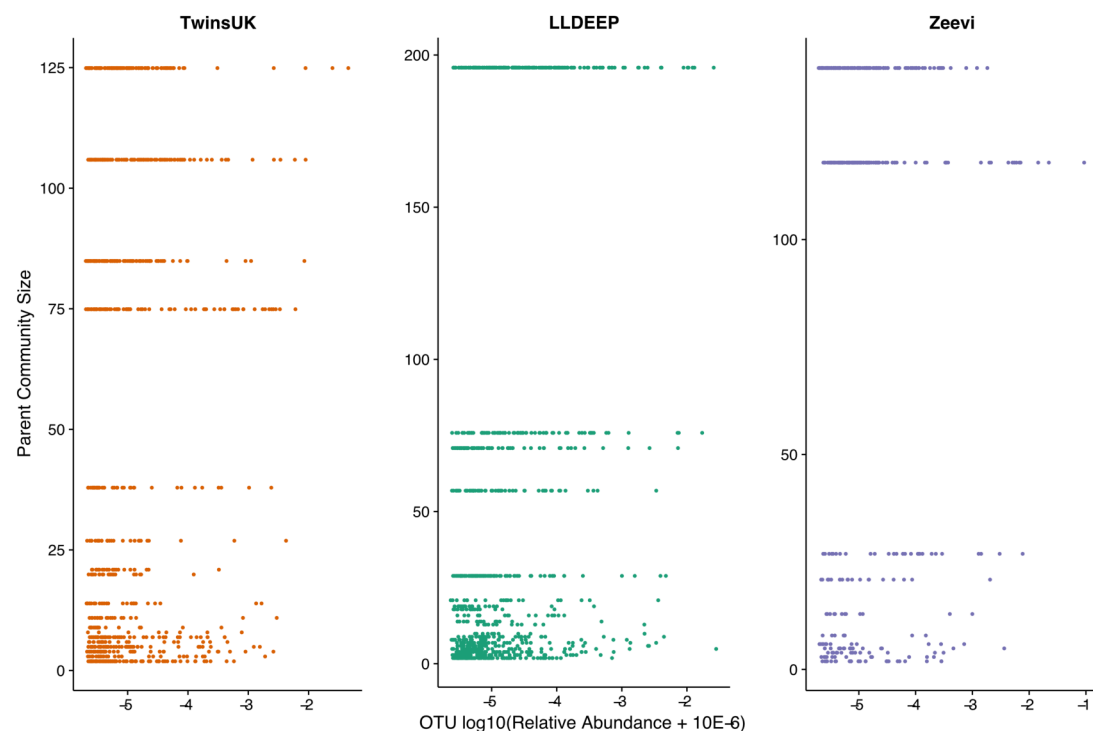
Supplementary Figure 1. A comparison of the relative abundance of an OTU to the number of OTUs in its assigned community for each of the three data sets.

ETHICS APPROVAL

Ethics approval for the TwinsUK study was given by the NRES Committee London - Westminster (REC Reference No. : EC04/015), with all participants providing written consent. The Lifelines-DEEP study was approved by the institutional ethics review boards of the University Medical Center Groningen (ref. M12.113965).

SUPPLEMENTARY MATERIALS

Supplementary Table 1. Heritability analysis results for the communities in TwinsUK.

Supplementary Table 2. Community association results with age and BMI in TwinsUK (Community members can be found in Supplementary Table 1)

Supplementary Table 3. Table summarising the OTUs within each of the community types identified across all three datasets.

Supplementary Methods. An overview of code and scripts used in main stages of network creation and community detection.
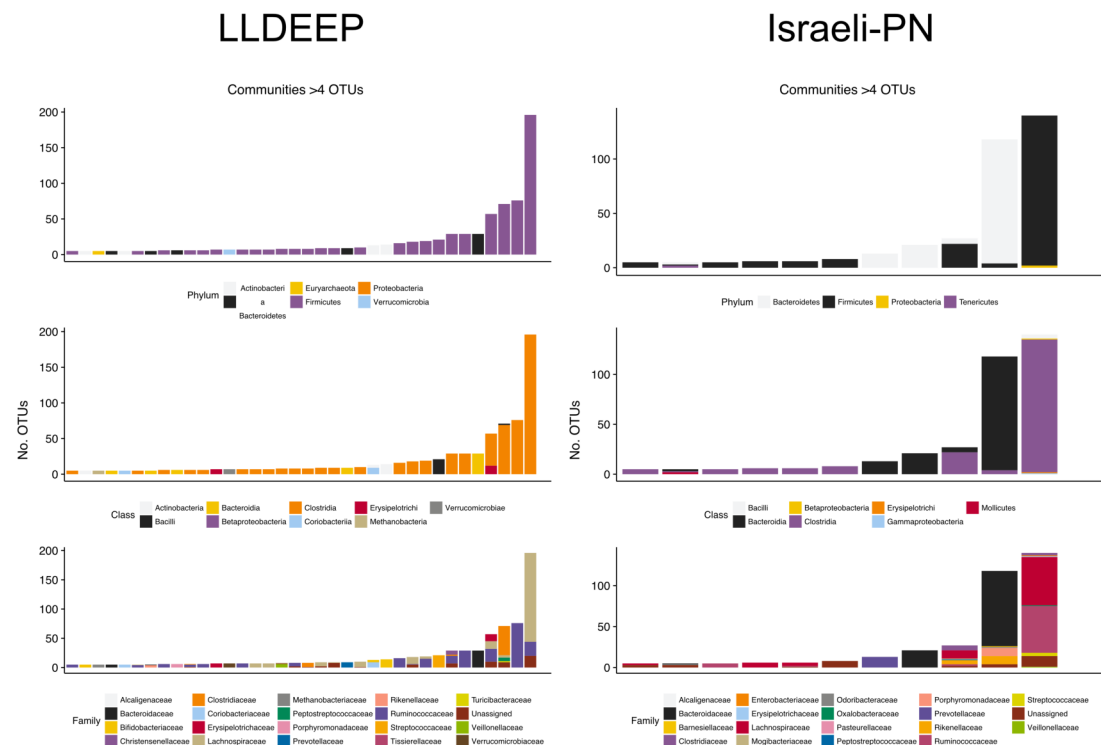
Supplementary Figure 2. Taxonomic summary plots for the communities in the LLDEEP and Israeli-PN data, as for TwinsUK in Figure 4B.

REFERENCES

Albert, R. (2005). Scale-free networks in cell biology. *Journal of Cell Science*, 118(21):4947.

Arumugam, M., Raes, J., Pelletier, E., Le Paslier, D., Yamada, T., Mende, D. R., Fernandes, G. R., Tap, J., Bruls, T., Batto, J.-M., et al. (2011). Enterotypes of the human gut microbiome. *Nature*, 473:174.

Baldassano, S. N. and Bassett, D. S. (2016). Topological distortion and reorganized modular structure of gut microbial co-occurrence networks in inflammatory bowel disease. *Scientific Reports*, 6:26087.

Bastian, M., Heymann, S., Jacomy, M., et al. (2009). Gephi: an open source software for exploring and manipulating networks. *International AAAI Conference on Web and Social Media*, 8:361.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):10008.

Chaffron, S., Rehrauer, H., Pernthaler, J., and von Mering, C. (2010). A global network of coexisting microbes from environmental and whole-genome sequence data. *Genome Research*, 20(7):947.

Csardi, G. and Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal, Complex Systems*, 1695(5):1.

David, L. A., Maurice, C. F., Carmody, R. N., Gootenberg, D. B., Button, J. E., Wolfe, B. E., Ling, A. V., Devlin, A. S., Varma, Y., Fischbach, M. A., et al. (2014). Diet rapidly and reproducibly alters the human gut microbiome. *Nature*, 505:559.

Deng, Y., Jiang, Y.-H., Yang, Y., He, Z., Luo, F., and Zhou, J. (2012). Molecular ecological network analyses. *BMC Bioinformatics*, 13(1):113.

Duran-Pinedo, A. E., Paster, B., Teles, R., and Frias-Lopez, J. (2011). Correlation network analysis applied to complex biofilm communities. *PloS One*, 6(12):e28438.

Falony, G., Joossens, M., Vieira-Silva, S., Wang, J., Darzi, Y., Faust, K., Kurilshikov, A., Bonder, M. J., Valles-Colomer, M., Vandeputte, D., et al. (2016). Population-level analysis of gut microbiome variation. *Science*, 352(6285):560.

Fang, H., Huang, C., Zhao, H., and Deng, M. (2015). Cclasso: correlation inference for compositional data through lasso. *Bioinformatics*, 31(19):3172.

Faust, K., Lahti, L., Gonze, D., de Vos, W. M., and Raes, J. (2015). Metagenomics meets time series analysis: unraveling microbial community dynamics. *Current opinion in microbiology*, 25:56.

Faust, K. and Raes, J. (2012). Microbial interactions: from networks to models. *Nature Reviews Microbiology*, 10(8):538.

Faust, K., Sathirapongsasuti, J. F., Izard, J., Segata, N., Gevers, D., Raes, J., and Huttenhower, C. (2012). Microbial co-occurrence relationships in the human microbiome. *PLoS Computational Biology*, 8(7):e1002606.

Friedman, J. and Alm, E. J. (2012). Inferring correlation networks from genomic survey data. *PLoS Computational Biology*, 8(9):e1002687.

Friedman, J., Higgins, L. M., and Gore, J. (2017). Community structure follows simple assembly rules in microbial microcosms. *Nature Ecology & Evolution*, 1(5).

Fu, J., Bonder, M. J., Cenit, M. C., Tigchelaar, E., Maatman, A., Dekens, J. A., Brandsma, E., Marczynska, J., Imhann, F., Weersma, R. K., et al. (2015). The gut microbiome contributes to a substantial proportion of the variation in blood lipids. *Circulation Research*, page 115.

Gao, Z., Yin, J., Zhang, J., Ward, R. E., Martin, R. J., Lefevre, M., Cefalu, W. T., and Ye, J. (2009). Butyrate improves insulin sensitivity and increases energy expenditure in mice. *Diabetes*, 58(7):1509.

Gevers, D., Kugathasan, S., Denson, L. A., Vázquez-Baeza, Y., Van Treuren, W., Ren, B., Schwager, E., Knights, D., Song, S. J., Yassour, M., et al. (2014). The treatment-naive microbiome in new-onset crohn's disease. *Cell Host & Microbe*, 15(3):382.

Goodrich, J. K., Waters, J. L., Poole, A. C., Sutter, J. L., Koren, O., Blekhman, R., Beaumont, M., Van Treuren, W., Knight, R., Bell, J. T., et al. (2014). Human genetics shape the gut microbiome. *Cell*, 159(4):789.

Harrell Jr, F. E., Dupont, C., et al. (2008). Hmisc: harrell miscellaneous. *R package version*, 3(2).

Jackson, M. A., Bell, J. T., Spector, T. D., and Steves, C. J. (2016a). A heritability-based comparison of methods used to cluster 16s rrna gene sequences into operational taxonomic units. *PeerJ*, 4:e2341.

Jackson, M. A., Jeffery, I. B., Beaumont, M., Bell, J. T., Clark, A. G., Ley, R. E., O'Toole, P. W., Spector, T. D., and Steves, C. J. (2016b). Signatures of early frailty in the gut microbiota. *Genome medicine*, 8(1):8.

Jeong, H., Tombor, B., Albert, R., Oltvai, Z., and Barabasi, A. (2000). The large-scale organization of metabolic networks. *Nature*, 407(6804):651.

Kennedy, N. A., Walker, A. W., Berry, S. H., Duncan, S. H., Farquarson, F. M., Louis, P., and Thomson, J. M. (2014). The impact of different dna extraction kits and laboratories upon the assessment of human gut microbiota composition by 16s rrna gene sequencing. *PloS One*, 9(2):e88982.

Kurtz, Z. D., Müller, C. L., Miraldi, E. R., Littman, D. R., Blaser, M. J., and Bonneau, R. A. (2015). Sparse and compositionally robust inference of microbial ecological networks. *PLoS Computational Biology*, 11(5):e1004226.

Langfelder, P. and Horvath, S. (2008). Wgcna: an r package for weighted correlation network analysis. *BMC Bioinformatics*, 9(1):559.

Levy, R. and Borenstein, E. (2014). Metagenomic systems biology and metabolic modeling of the human microbiome: From species composition to community assembly rules. *Gut Microbes*, 5(2):265.

Lindemann, S. R., Bernstein, H. C., Song, H.-S., Fredrickson, J. K., Fields, M. W., Shou, W., Johnson, D. R., and Beliaev, A. S. (2016). Engineering microbial consortia for controllable outputs. *The ISME journal*, 10(9):2077.

Louis, P. and Flint, H. J. (2009). Diversity, metabolism and microbial ecology of butyrate-producing bacteria from the human large intestine. *FEMS Microbiology Letters*, 294(1):1.

Lozupone, C., Faust, K., Raes, J., Faith, J. J., Frank, D. N., Zaneveld, J., Gordon, J. I., and Knight, R. (2012). Identifying genomic and metabolic features that can underlie early successional and opportunistic lifestyles of human gut symbionts. *Genome Research*, 22(10):1974.

McHardy, I. H., Goudarzi, M., Tong, M., Ruegger, P. M., Schwager, E., Weger, J. R., Graeber, T. G., Sonnenburg, J. L., Horvath, S., Huttenhower, C., et al. (2013). Integrative analysis of the microbiome and metabolome of the human intestinal mucosal surface reveals exquisite inter-relationships. *Microbiome*, 1(1):17.

Mucha, P. J., Richardson, T., Macon, K., Porter, M. A., and Onnela, J.-P. (2010). Community structure in time-dependent, multiscale, and multiplex networks. *Science*, 328(5980):876.

Navas-Molina, J. A., Peralta-Sánchez, J. M., González, A., McMurdie, P. J., Vázquez-Baeza, Y., Xu, Z.,

Ursell, L. K., Lauber, C., Zhou, H., Song, S. J., et al. (2013). Advancing our understanding of the human microbiome using qiime. *Methods in Enzymology*, 531:371.

Newman, M. E. J. and Girvan, M. (2004). Finding and evaluating community structure in networks. *Phys. Rev. E*, 69:026113.

Odamaki, T., Kato, K., Sugahara, H., Hashikura, N., Takahashi, S., Xiao, J.-z., Abe, F., and Osawa, R. (2016). Age-related changes in gut microbiota composition from newborn to centenarian: a cross-sectional study. *BMC microbiology*, 16(1):90.

Qin, J., Li, Y., Cai, Z., Li, S., Zhu, J., Zhang, F., Liang, S., Zhang, W., Guan, Y., Shen, D., et al. (2012). A metagenome-wide association study of gut microbiota in type 2 diabetes. *Nature*, 490:55.

Ridaura, V. K., Faith, J. J., Rey, F. E., Cheng, J., Duncan, A. E., Kau, A. L., Griffin, N. W., Lombard, V., Henrissat, B., Bain, J. R., et al. (2013). Gut microbiota from twins discordant for obesity modulate metabolism in mice. *Science*, 341(6150):1241214.

Schloss, P. D. and Handelsman, J. (2005). Introducing dotur, a computer program for defining operational taxonomic units and estimating species richness. *Applied and environmental microbiology*, 71(3):1501.

Tong, M., Li, X., Parfrey, L. W., Roth, B., Ippoliti, A., Wei, B., Borneman, J., McGovern, D. P., Frank, D. N., Li, E., et al. (2013). A modular organization of the human intestinal mucosal microbiota and its association with inflammatory bowel disease. *PloS One*, 8(11):e80702.

Turnbaugh, P. J., Ley, R. E., Mahowald, M. A., Magrini, V., Mardis, E. R., and Gordon, J. I. (2006). An obesity-associated gut microbiome with increased capacity for energy harvest. *Nature*, 444:1027.

Walker, A. W., Martin, J. C., Scott, P., Parkhill, J., Flint, H. J., and Scott, K. P. (2015). 16s rrna gene-based profiling of the human infant gut microbiota is strongly influenced by sample processing and pcr primer choice. *Microbiome*, 3(1):26.

Weiss, S., Van Treuren, W., Lozupone, C., Faust, K., Friedman, J., Deng, Y., Xia, L. C., Xu, Z. Z., Ursell, L., Alm, E. J., et al. (2016). Correlation detection strategies in microbial data sets vary widely in sensitivity and precision. *The ISME Journal*, 10(7):1669.

Westcott, S. L. and Schloss, P. D. (2015). De novo clustering methods outperform reference-based methods for assigning 16s rrna gene sequences to operational taxonomic units. *PeerJ*, 3:e1487.

Wolf, Y. I., Karev, G., and Koonin, E. V. (2002). Scale-free networks in biology: new insights into the fundamentals of evolution? *BioEssays*, 24(2):105.

Wu, G. D., Chen, J., Hoffmann, C., Bittinger, K., Chen, Y.-Y., Keilbaugh, S. A., Bewtra, M., Knights, D., Walters, W. A., Knight, R., et al. (2011). Linking long-term dietary patterns with gut microbial enterotypes. *Science*, 334(6052):105.

Yatsunenko, T., Rey, F. E., Manary, M. J., Trehan, I., Dominguez-Bello, M. G., Contreras, M., Magris, M., Hidalgo, G., Baldassano, R. N., Anokhin, A. P., et al. (2012). Human gut microbiome viewed across age and geography. *Nature*, 486:222.

Zeevi, D., Korem, T., Zmora, N., Israeli, D., Rothschild, D., Weinberger, A., Ben-Yacov, O., Lador, D., Avnit-Sagi, T., Lotan-Pompan, M., et al. (2015). Personalized nutrition by prediction of glycemic responses. *Cell*, 163(5):1079.

Zhang, B. and Horvath, S. (2005). A general framework for weighted gene co-expression network analysis. *Statistical applications in genetics and molecular biology*, 4(1):1–45.