



A simple statistical test of taxonomic or functional homogeneity using replicated microbiome sequencing samples



Daniel H. Huson^{a,*}, Mike Steel^b, Mohamed El-Hadidi^a, Suparna Mitra^c, Silke Peter^{d,e}, Matthias Willmann^{d,e}

^a Center for Bioinformatics, University of Tübingen, Germany

^b Biomathematics Research Centre, University of Canterbury, Christchurch, New Zealand

^c Faculty of Medicine and Health – University of Leeds, Old Medical School, Leeds General Infirmary, Leeds, UK

^d Institute of Medical Microbiology and Hygiene, University of Tübingen, Germany

^e German Center for Infection Research (DZIF), Partner Site Tübingen, Germany

ARTICLE INFO

Article history:

Received 17 August 2016

Received in revised form 20 October 2016

Accepted 24 October 2016

Available online 27 October 2016

Keywords:

Metagenomics

Taxonomic composition

Functional diversity

Environmental inhomogeneity

Statistical testing

ABSTRACT

One important question in microbiome analysis is how to assess the homogeneity of the microbial composition in a given environment, with respect to a given analysis method. Do different microbial samples taken from the same environment follow the same taxonomic distribution of organisms, or the same distribution of functions? Here we provide a non-parametric statistical “triangulation test” to address this type of question. The test requires that multiple replicates are available for each of the biological samples, and it is based on three-way computational comparisons of samples. To illustrate the application of the test, we collected three biological samples taken from different locations in one piece of human stool, each represented by three replicates, and analyzed them using MEGAN. (Despite its name, the triangulation test does not require that the number of biological samples or replicates be three.) The triangulation test rejects the null hypothesis that the three biological samples exhibit the same distribution of taxa or function (error probability ≤ 0.05), indicating that the microbial composition of the investigated human stool is not homogenous on a macroscopic scale, suggesting that pooling material from multiple locations is a reasonable practice. We provide an implementation of the test in our open source program MEGAN Community Edition.

© 2016 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

There is growing interest in microbiome analysis, the study of microorganisms in a particular environment, such as the human body, water or soil. Taxonomic profiling of microbiome samples is usually based on DNA sequencing, using either an amplicon or a shotgun approach (Handelsman, 2004). While early studies typically involved small numbers of samples, there is now an increased expectation that replicate samples are sequenced (Knight et al., 2012) so as to allow statistical analysis of the significance of obtained results.

One important question is how to assess the homogeneity of the microbial composition in a given environment. Do different samples taken from the same environment follow the same taxonomic distribution of organisms? Do they follow the same distribution of functional features? Presumably, well-mixing environments such

as air or water will display a higher degree of compositional homogeneity than more structured environments such as soil, stool or waste water flocks.

We need to introduce two key concepts. First, we assume that multiple samples are taken from the same given environment. For example, in the case of human stool, one might collect samples from a number of different locations in the same piece of fecal matter. We refer to these samples as *biological samples*. Second, we assume that each biological sample gives rise to two or more *replicates*, which are obtained by constructing multiple DNA libraries for each biological sample.

In addition, we emphasize that the whole chain of steps employed, from sampling, DNA extraction, sequencing, calculation of taxonomic or functional profiles, to the comparison of such profiles using a dissimilarity measure, is part of the input to the problem, and we will use M to denote the complete analysis procedure.

We describe a novel statistical test, which we call the *triangulation test*, that aims at testing the null hypothesis that different biological samples taken from a given environment exhibit the

* Corresponding author.

E-mail address: Daniel.Huson@uni-tuebingen.de (D.H. Huson).

same distribution of taxa, with respect to M . The triangulation test can be applied to a wide range of different sampling regimes involving any number of biological and replicates. This is a nonparametric test and is thus robust against assumptions of population distribution.

In Section 2, we first formulate the triangulation test for the 3×3 case of three biological samples, each represented by three replicates. We then generalize the triangulation test to the case of any number of biological samples, each represented by an arbitrary number of replicates.

In Section 3, we demonstrate the application of the triangulation test in the 3×3 case using human stool samples. In more detail, we collected three samples from different locations (approximately one centimeter apart) in the same piece of stool and sequenced three replicates for each. Application of the triangulation test implies that the microbial composition of the investigated human stool is not homogenous on the macroscopic scale.

We use a series of artificially constructed mixtures of samples to study the performance of the method in lower contrast settings and show that it performs as well as PERMANOVA.

2. Methods

2.1. Triangulation test for 3×3 samples

Assume that we are given three biological samples A , B and C , and for each biological sample $S = A, B$ or C we are given three replicates, S_1, S_2 and S_3 . We thus have nine samples in total: $A_1, A_2, A_3, B_1, B_2, B_3, C_1, C_2, C_3$. Moreover, let M be a specific chain of analysis steps that provides taxonomic profiles for all samples, and a dissimilarity measure.

The null hypothesis and the alternative hypothesis of interest are stated as follows:

- H_0 : the nine samples A_1, \dots, C_3 are all drawn independently according to the same distribution, versus
- H_a : For each of the three biological samples $S = A, B$ or C , the three replicates S_1, S_2, S_3 are drawn from the same distribution for S , but the distributions for each of the biological samples are not all identical. Moreover, replicates within the same biological sample expected to be more similar to each other than any ones that lie in different biological samples.

To address these hypotheses, we define the *triangulation test* as a simple non-parametric significance test for H_0 . It is based on the concept of a *random triangulation* of the set of samples.

For ease of exposition, we first describe this for the special setting used in our practical study, namely three biological samples, each represented by three replicates. Below, we will then present a general version of this test that allows for any number of biological samples, each with its own arbitrary number of replicates.

We define a *triangulation* of the set of nine replicates A_1, A_2, \dots, C_3 to be a graph $G = (V, E)$ with node set $V = \{A_1, A_2, \dots, C_3\}$ and edge set E consisting of undirected edges that form node-disjoint 3-cycles that each involve replicates from exactly two different biological samples. We desire that any such triangulation contains as many triangles as possible, i.e., three triangles as in Fig. 1.

We assume that an analysis of the taxonomic content (or functional content, if desired) of each replicate has been performed using M and we have obtained a dissimilarity measure $d(S_i, T_j)$ between any two replicates $S_i, T_j \in V$.

The triangulation test is performed in two steps. In the first step, we randomly choose a single triangulation that involves all replicates. In Fig. 1 we show one such choice for the case of three biological samples, each represented by three replicates. As




Fig. 1. Triangulation test. The nine replicates A_1, \dots, C_3 (represented by labeled discs) come from three biological samples (indicated by gray ovals). Lines connecting discs represent a triangulation of the data. The null hypothesis is rejected, if for each and every triangle, the distance between replicates from the same sample is less than the distance from either of those replicates to a replicate from a different sample.

required, each depicted triangle contains two replicates from the same biological sample, and a third replicate from a different biological sample.

In the second step, we then ask the following question: For each triangle in the chosen triangulation, is the dissimilarity between the two replicates that are contained in the same biological sample less than the dissimilarity of either of them to the third replicate in the triangle? If the answer is yes for all triangles, then reject H_0 with a significance level that is less than the 0.05.

To see this, note that under the null hypothesis H_0 , all three pairs of replicates in a triangle have equal probability of exhibiting the strictly smallest dissimilarity value, $1/3$ (or less, in the case of a tie). Thus, the probability that it is always the pair of replicates contained in the same biological sample that has the smallest value in all three triangles is at most $(1/3)^3 = (1/27) < 0.05$. Hence, the probability of rejecting the null hypothesis when it is true (i.e., the type 1 error) is less than 0.05.

Whether the triangle test is able to reject H_0 depends not only on the samples, but also on the details of the analysis method M , we would like to emphasize. For example, if the samples are analyzed using a low resolution method M that can only detect the general presence or absence of bacteria, say, and if bacteria are detected in all samples, then the dissimilarities will be constant, in which case the test will fail to reject the null hypothesis. However, it is of course possible that the use of a higher resolution analysis method will lead to rejection of the null hypothesis on the same data.

We would also like to emphasize that application of the test only involves the use of one randomly chosen triangulation. The test does not require that one looks at multiple triangulations. In particular, this makes the test very easy to apply “by hand”.

2.2. Triangulation test in the general setting

We now describe how the simple triangulation significance test described above extends to the general setting where we have m bins (i.e., biological samples) of arbitrary sizes r_1, \dots, r_m . Let $V = \{S_i(j) \mid j = 1, \dots, r_i\}$ be the set of replicates present in bin i , for $i = 1, \dots, m$. (Above we studied the special case of $m = 3$ and $r_i = 3$ for $i = 1, \dots, 3$.) The total number of replicates is $N = \sum_{i=1}^m r_i$. Again,

we assume that a method M is given that provides profiles and dissimilarities.

As above, a *triangulation* of V refers to a set of node-disjoint 3-cycles connecting (all, or a subset of) these N nodes, with the property that each 3-cycle contains exactly two samples from the same bin – we call such 3-cycles *valid*. A simple counting argument shows that the number of valid 3-cycles in any such triangulation can never exceed the largest integer less or equal to $(1/3)N$. Moreover, this bound can be achieved (for example, in the set-up of the previous section where $m = 3; r_1 = r_2 = r_3 = 3$), though not always (for example, if $m = 2, r_1 = 7$ and $r_2 = 2$, then the number of valid 3-cycles is at most 2, while $(1/3)N = 3$).

We now randomly generate (over 100 independent runs, say) triangulations of V , and select a valid triangulation of largest size (i.e., containing the largest number T of valid 3-cycles). Then, under the null hypothesis H_0 that the N samples are all drawn independently according to the same distribution, if X denotes the number of 3-cycles in the sample for which a strictly smallest dissimilarity is between the pair within the bin, then X has a binomial distribution consisting of T trials and success probability (at most) $1/3$. For any $\alpha > 0$ (above we used $\alpha = 0.05$) let k_α denote the α -critical value of this distribution, in other words the smallest integer k for which

$$\sum_{i \geq k} \binom{T}{i} \left(\frac{1}{3}\right)^i \left(\frac{2}{3}\right)^{T-i} < \alpha.$$

Then the triangulation test will reject H_0 in favor of H_a when $X \geq k_\alpha$. Here H_a is the alternative hypothesis that the samples within each

bin are drawn from the same distribution but these m distributions are different, with reads expected to be more similar within each bin than between bins. The proof that this test has type-1 error at most α proceeds in parallel fashion to the earlier argument.

2.3. Illustration of the method using human stool samples

Three samples A, B, C were collected from one piece of human stool provided by a healthy male subject, from three different locations separated by about two centimeters. DNA was extracted within three hours using a modified version of the Human Microbiome Project protocol as described elsewhere (Willmann et al., 2015). For each of the biological samples, A, B and C , three libraries were produced. For this purpose, genomic DNA was sheared by Covaris M220 (Covaris, Woburn, USA) to obtain 550 bp fragments. DNA libraries were prepared with TruSeq Nano DNA LT Kit (Illumina, San Diego, USA) using the standard protocol. Barcoded libraries were analyzed on the QIAXcel Advanced Instrument (Qiagen, Hilden, Germany). All libraries were sequenced at 2×250 bp on an Illumina MiSeq (Illumina, San Diego, USA), obtaining nine shotgun metagenomic datasets $A_1, A_2, A_3, B_1, B_2, B_3, C_1, C_2, C_3$. Quality control was performed using FastQC (<http://www.bioinformatics.babraham.ac.uk/projects/fastqc/>). The total number of reads was 39 million, on average 4.3 million reads per sample (range 3.6–4.8 million).

We used Metascope (Buchfink et al., 2015) to identify human sequences and 0.16% of all reads were removed. We used cutadapt to perform adaptor trimming (Martin, 2011). Quality filtering was performed using the USearch package (Edgar, 2010), with the




Fig. 2. Taxonomic analysis of nine samples. Samples are normalized to the smallest number of aligned reads, 1.9 million each. Each taxonomic node is drawn as a bar chart indicating the number of reads assigned to the taxon for each sample, for each of the nine samples, using a square-root scale so as to make small assignments more visible. To avoid clutter, nodes that have only one child are not labeled.

maximum expected error rate set to 0.01. This removed 19.9% of all reads. The remaining reads (31 million) were aligned against NCBI-nr (Benson et al., 2005) using DIAMOND (Buchfink et al., 2015) (default settings), finding alignments for 24 million reads. We then performed taxonomic and functional binning using MEGAN (Huson et al., 2016), with the following settings: minScore = 50, maxExpected = 0.01, topPercent = 10, minSupportPercent = 0.01. The resulting assignment of reads to taxa is summarized in Fig. 2.

The input of the triangulation test is a dissimilarity matrix on the given set of samples. To compute such a dissimilarity matrix in the case of taxonomic analysis, we summarized all taxonomic read assignments at a fixed taxonomic rank (species, genus, . . . , phylum) and then applied the square-root Jensen–Shannon diversity calculation (Arumugam et al., 2011) to compute pairwise dissimilarities between the samples. We call the resulting matrix a taxonomic dissimilarity matrix.

Similarly, a functional dissimilarity matrix can be obtained by using one of the functional classifications implemented in MEGAN such as the new InterPro2GO classification (Hunter et al., 2014; Mitchell et al., 2015), KEGG (Kanehisa and Goto, 2000) or SEED (Overbeek et al., 2013). Read counts are summarized at a chosen level of the classification and then again the Jensen–Shannon diversity calculation (Arumugam et al., 2011) is used to compute pairwise dissimilarities between the samples.

3. Results

We have implemented the general triangulation test in MEGAN 6, where it can be applied both to taxonomic and functional profiles (Cluster Analysis Viewer → Options menu → Triangulation Test).

3.1. Application to human stool samples

For the nine metagenomic datasets $A_1, A_2, A_3, B_1, B_2, B_3, C_1, C_2, C_3$ described above, together with the described analysis method M , application of the triangulation test to a taxonomic dissimilarity matrix always rejects the null hypothesis that the samples are based on the same distribution of taxa, for all taxonomic ranks from Species up to Phylum. PCoA analysis of the nine samples illustrates that replicates from the same biological sample are much closer to each other than they are to replicates from other samples (see Fig. 3).

The triangulation test applied to a functional dissimilarity matrix based on InterPro2GO, SEED or KEGG, rejects the null hypothesis that the samples are based on the same distribution of function, both when read counts are summarized at a high level (e.g., carbohydrate metabolism, energy metabolism, lipid metabolism, etc. in the case of KEGG), or a low level (e.g., using KO groups in KEGG).




Fig. 3. Application of the triangulation test by hand to a PCoA analysis of stool samples. Choose a random triangulation such that each triangle connects two replicates from the same biological sample (e.g., B_2 and B_3) with one from a different biological sample (e.g., A_3), as shown here. If, for each triangle, the length of the edge connecting the replicates from the same biological sample is the shortest edge in the triangle, then one can reject the null hypothesis that all samples are drawn from the same distribution, with $\alpha = 0.05$. This is the case here.

So, in both terms of taxonomic composition and in terms of functional capacity, the test establishes significant inhomogeneity between biological samples. This result provides statistical justification for the practice of pooling material from multiple locations in a piece of stool, as performed in Willmann et al. (2015).

When using any of the other eight beta-diversity calculations provided by MEGAN, such as Bray–Curtis (Bray and Curtis, 1957), we find that the null hypothesis is always rejected, except for some of the high-level functional comparisons. The latter is consistent with the observation that high-level metabolic modules show much less variation across human microbiome samples than taxonomic phyla do (The Human Microbiome Project Consortium, 2012).

Note that after randomly reassigning replicates to biological samples, the null hypothesis was no longer rejected, as expected.

All non-human sequencing reads, a MEGAN file containing the taxonomic and functional profiles for all nine samples, and the program MEGAN 6 (including source code), are available here: <http://ab.inf.uni-tuebingen.de/data/external/tritest/>.

3.2. Comparison with PERMANOVA on lower contrast data

How does the triangle test compare against standard nonparametric analysis such as PERMANOVA (Anderson, 2001)? To investigate this, we performed permutational multivariate analysis of variance (PERMANOVA) using Bray–Curtis dissimilarities and 999 permutations, employing the Vegan package in R (Oksanen et al., 2016; R Core Team, 2013). The result is that the three groups of samples A, B and C are significantly different with $p = 0.004$.

Fig. 3 suggests that the three groups of samples A, B, C have such high contrast that a significant difference is a foregone conclusion. To evaluate performance of the triangulation test where samples are more similar, we generated a series of artificially constructed mixtures of samples that show increasingly less contrast.

First, we generated a background distribution H of taxa based on the union of all nine original samples. Then, for a given percentage $x = 90, 80, \dots, 10, 7, 5, 3$, we produced a group of nine artificial samples $A_1^x, A_2^x, \dots, A_9^x$ by defining the taxonomic profile associated with each such sample P_i^x to be a linear mixture of P_i (x percent “signal”) and H (the remaining $100 - x$ percent of P_i^x , “background”).

Both the triangle test and PERMANOVA analysis were applied to all groups of artificial samples. For all groups with $x > 3\%$, both the triangle test and PERMANOVA rejected the null hypothesis with a significance level of 0.05 (triangle test) and 0.01 (PERMANOVA). For $x = 3\%$, both tests failed to robustly reject the null hypothesis. This suggests that the triangle test performs as well as PERMANOVA analysis, even in quite low contrast settings.

4. Discussion

To address the problem of assessing microbial homogeneity in a given environment, we suggest that multiple replicates are collected for each biological sample. The non-parametric triangulation test presented here allows one to test the null hypothesis that all samples are drawn independently according to the same distribution. This test is performed relative to a specific analysis pipeline M that is used to calculate profiles and dissimilarities. The implementation of the test that we provide in MEGAN 6 (Community Edition) can be applied to both taxonomic and functional profiles, at different levels of resolution.

The aim of this work was to develop a new statistical tool for testing homogeneity. This test is attractive because it is easily applied by hand (see Fig. 3); moreover, it provides a precise statistical underpinning for the intuition that, if replicates cluster by biological sample, then there is inhomogeneity on the macroscopic

scale. With this tool in hand, it will be possible to systematically investigate the level of homogeneity in different microbial environments, at different macroscopic scales and for different analysis methods.

Author contributions

D.H.H., M.S., S.P. and M.W. designed the study. M.S. developed the statistical test. D.H.H. implemented the test. D.H.H. and M.S. wrote the manuscript. S.P. and M.W. produced the human stool data. M.H. and S.M. performed the data analysis.

References

- Anderson, M., 2001. A new method for non-parametric multivariate analysis of variance. *Austral Ecol.* 26, 32–46.
- Arumugam, M., Raes, J., Pelletier, E., Le Paslier, D., Yamada, T., Mende, D.R., Fernandes, G.R., Tap, J., Bruls, T., Batto, J.-M., Bertalan, M., Borruel, N., Casellas, F., Fernandez, L., Gautier, L., Hansen, T., Hattori, M., Hayashi, T., Kleerebezem, M., Kurokawa, K., Leclerc, M., Levenez, F., Manichanh, C., Nielsen, H.B., Nielsen, T., Pons, N., Poulain, J., Qin, J., Sicheritz-Ponten, T., Tims, S., Torrents, D., Ugarte, E., Zoetendal, E.G., Wang, J., Guarner, F., Pedersen, O., de Vos, W.M., Brunak, S., Dore, J., Weissenbach, J., Ehrlich, S.D., Bork, P., 2011. Enterotypes of the human gut microbiome. *Nature* 473 (7346), 174–180. <http://dx.doi.org/10.1038/nature09944>.
- Benson, D., Karsch-Mizrachi, I., Lipman, D., Ostell, J., Wheeler, D., 2005. *Genbank Nucleic Acids Res.* 1 (33), D34–38.
- Bray, R.J., Curtis, J.T., 1957. An ordination of the upland forest communities of southern Wisconsin. *Ecol. Monogr.* 27, 325–349.
- Buchfink, B., Xie, C., Huson, D., 2015. Fast and sensitive protein alignment using DIAMOND. *Nat. Methods* 12, 59–60.
- Buchfink, B., Huson, D.H., Xie, C., 2015. *Metascope – Fast and Accurate Identification of Microbes in Metagenomic Sequencing Data*. Tech. Rep., arXiv:1511.08753.
- Edgar, R., 2010. Search and clustering orders of magnitude faster than BLAST. *Bioinformatics (Oxford, England)* 26 (19), 2460–2461. <http://dx.doi.org/10.1093/bioinformatics/btq461>.
- Handelsman, J., 2004. Metagenomics: application of genomics to uncultured microorganisms. *Microbiol. Mol. Biol. Rev.* 68 (4), 669–685.
- Hunter, S., Corbett, M., Denise, H., Fraser, M., Gonzalez-Beltran, A., Hunter, C., Jones, P., Leinonen, R., McAnulla, C., Maguire, E., Maslen, J., Mitchell, A., Nuka, G., Oisel, A., Pesseat, S., Radhakrishnan, R., Rocca-Serra, P., Scheremetjew, M., Sterk, P., Vaughan, D., Cochrane, G., Field, D., Sansone, S.-A., 2014. EBI metagenomics – a new resource for the analysis and archiving of metagenomic data. *Nucleic Acids Res.* 42 (Database issue), D600–D606. <http://dx.doi.org/10.1093/nar/gkt961>. URL: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3965009/>.
- Huson, D., Beier, S., Flade, I., Górka, A., El-Hadidi, M., Mitra, S., Ruscheweyh, H.-J., Tappu, R., 2016. MEGAN Community Edition – interactive exploration and analysis of large-scale microbiome sequencing data. *PLoS Comput. Biol.* 12 (6), e1004957. <http://dx.doi.org/10.1371/journal.pcbi.1004957>.
- Kanehisa, M., Goto, S., 2000. KEGG: Kyoto encyclopedia of genes and genomes. *Nucleic Acids Res.* 28 (1), 27–30.
- Knight, R., Jansson, J., Field, D., Fierer, N., Desai, N., Fuhrman, J.A., Hugenholtz, P., van der Lelie, D., Meyer, F., Stevens, R., Bailey, M.J., Gordon, J.I., Kowalchuk, G.A., Gilbert, J.A., 2012. Unlocking the potential of metagenomics through replicated experimental design. *Nat. Biotech.* 30 (6), 513–520. <http://dx.doi.org/10.1038/nbt.2235>.
- Martin, M., 2011. Cutadapt removes adapter sequences from high-throughput sequencing reads. *EMBnet journal* 17 (1), URL: <http://journal.embnet.org/index.php/embnetjournal/article/view/200>.
- Mitchell, A., Chang, H.-Y., Daugherty, L., Fraser, M., Hunter, S., Lopez, R., McAnulla, C., McMenamin, C., Nuka, G., Pesseat, S., Sangrador-Vegas, A., Scheremetjew, M., Rato, C., Yong, S.-Y., Bateman, A., Punta, M., Attwood, T.K., Sigrist, C.J., Redaschi, N., Rivoire, C., Xenarios, I., Kahn, D., Guyot, D., Bork, P., Letunic, I., Gough, J., Oates, M., Haft, D., Huang, H., Natale, D.A., Wu, C.H., Orengo, C., Sillitoe, I., Mi, H., Thomas, P.D., Finn, R.D., 2015. The interpro protein families database: the classification resource after 15 years. *Nucleic Acids Res.* 43 (D1), D213–D221. <http://dx.doi.org/10.1093/nar/gku1243>. URL: <http://nar.oxfordjournals.org/content/43/D1/D213.abstract>. <http://nar.oxfordjournals.org/content/43/D1/D213.full.pdf+html>.
- Oksanen, J., Blanchet, F.G., Friendly, M., Kindt, R., Legendre, P., McGlinn, D., Minchin, P.R., O'Hara, R.B., Simpson, G.L., Solymos, P., Stevens, M.H.H., Szoecs, E., Wagner, H., 2016. *Vegan: Community Ecology Package*. URL: <http://cran.r-project.org/web/packages>.
- Overbeek, R., Olson, R., Pusch, G.D., Olsen, G.J., Davis, J.J., Disz, T., Edwards, R.A., Gerdes, S., Parrello, B., Shukla, M., Vonstein, V., Wattam, A.R., Xia, F., Stevens, R., 2014. The SEED and the rapid annotation of microbial genomes using subsystems technology (RAST). *Nucleic Acids Res.* 42. <http://dx.doi.org/10.1093/nar/gkt1226>. URL: <http://nar.oxfordjournals.org/content/early/2013/11/29/nar.gkt1226.full.pdf+html>. <http://nar.oxfordjournals.org/content/early/2013/11/29/nar.gkt1226.abstract>.

R Core Team, 2013. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

The Human Microbiome Project Consortium, 2012. Structure, function and diversity of the healthy human microbiome. *Nature* 486 (7402), 207–214, <http://dx.doi.org/10.1038/nature11234>.

Willmann, M., El-Hadidi, M., Huson, D., Schütz, M., Weidenmaier, C., Autenrieth, I., Peter, S., 2015. Antibiotic selection pressure determination through sequence-based metagenomics. *Antimicrob. Agents Chemother.* 59 (12), 7335–7345.