

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
Gautreau *et al.*

METHOD

PPanGGOLiN: Depicting microbial diversity via a Partitioned Pangenome Graph

Guillaume Gautreau¹, Adelme Bazin¹, Mathieu Gachet¹, Rémi Planel¹, Laura Burlot¹, Mathieu Dubois¹, Amandine Perrin^{2,5}, Claudine Médigue¹, Alexandra Calteau¹, Stéphane Cruveiller¹, Catherine Matias³, Christophe Ambroise⁴, Eduardo PC Rocha² and David Vallenet^{1*}

*Correspondence:
vallenet@genoscope.cns.fr
¹ LABGeM, Génomique Métabolique, CEA, Genoscope, Institut François Jacob, Université d'Évry, Université Paris-Saclay, CNRS, Evry, France
Full list of author information is available at the end of the article

Abstract

The use of comparative genomics for functional, evolutionary, and epidemiological studies requires methods to classify gene families in terms of occurrence in a given species. These methods usually lack multivariate statistical models to infer the partitions and the optimal number of classes and don't account for genome organization. We introduce a graph structure to model pan genomes in which nodes represent gene families and edges represent genomic neighborhood. Our method, named PPanGGOLiN, partitions nodes using an Expectation-Maximization algorithm based on multivariate Bernoulli Mixture Model coupled with a Markov Random Field. This approach takes into account the topology of the graph and the presence/absence of genes in pan genomes to classify gene families into persistent, cloud, and one or several shell partitions. By analyzing the partitioned pangenome graphs of isolate genomes from 439 species and metagenome-assembled genomes from 78 species, we demonstrate that our method is effective in estimating the persistent genome. Interestingly, it shows that the shell genome is a key element to understand genome dynamics, presumably because it reflects how genes present at intermediate frequencies drive adaptation of species, and its proportion in genomes is independent of genome size. The graph-based approach proposed by PPanGGOLiN is useful to depict the overall genomic diversity of thousands of strains in a compact structure and provides an effective basis for very large scale comparative genomics. The software is freely available at <https://github.com/labgem/PPanGGOLiN>.

Keywords: pangenomics; pangenome graph; microbial genomics; comparative genomics; Bernoulli mixture model; Markov Random Field; Maximization-Expectation algorithm

Background

The analyses of the gene repertoire diversity of species - their pangenome - have many applications in functional, evolutionary, and epidemiological studies regarding both core and accessory genes [1, 2]. The core genome is defined as the set of genes shared by all the genomes of a taxonomic unit (generally a species) whereas the accessory (or variable) genome contains genes that are only present in some genomes. The latter is crucial to understand bacterial adaptation as it contains a large repertoire of genes that may confer distinct traits and explain many of the phenotypic differences across species. Most of these genes are acquired by horizontal gene transfer (HGT) [3]. This usual dichotomy between core and accessory genomes does not consider the diverse ranges of gene frequencies in a pangenome. The main

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

problem in using a strict definition of the core genome is that its size decreases as more genomes are added to the analysis [4] due to gene loss events and technical artifacts (i.e. sequencing, assembly or annotation issues). As a consequence, it was proposed in the field of synthetic biology to focus on persistent genes, i.e. those conserved in a large majority of genomes [5]. The persistent genome is also called the soft core [6], the extended core [7, 8] or the stabilome [9]. These definitions advocate for the use of a threshold on the frequency of observation of a gene among the species of a clade, above which the gene is considered as *de facto* core gene. Persistent gene families are usually defined as those present in a range comprised between 90% [10] and 99% [11] of the strains in the species. This approach addresses some problems of the original definition of core genome but requires the setting of an appropriate threshold. The gene frequency distribution in pangenomes is extensively documented [12, 7, 13, 14, 15, 16, 8]. Due to the variation in the rates of gene loss and gain of genes, the gene frequencies tend to show an asymmetric U-shaped distribution regardless of the phylogenetic level and the clade considered (with the exception of few species having non-homogeneous distributions as described in [17]). Thereby, as proposed by [12] and formally modeled by [14], the pan genome can be split into 3 classes: (1) persistent genome, for the gene families present in almost all genomes; (2) shell genome, for gene families present at intermediate frequencies in the species; (3) cloud genome, for gene families present at low frequency in the species.

The study of pangenomes in microbiology now relies on the comparison of hundreds to thousands of genomes of a single species. The analysis of this massive amount of data raises computational and algorithmic challenges that can be tackled because genomes within a species have many homologous genes and it is possible to design new compact ways of representing and manipulating this information. As suggested by [18], a consensus representation of multiple genomes would provide a better analytical framework than using individual reference genomes. Among others, this proposition has led to a paradigm shift from the usual linear representation of reference genomes to a representation as pan genome graphs bringing together all the different known variations as multiple alternative paths. Methods [19, 20, 21] have been developed aiming at factorizing pangenomes at the genome sequence-level to capture all the nucleotide variations in a graph that enables variant calling and improves the sensitivity of the read mapping (summarized in [22]).

The method presented here, named PPanGGOLiN (Partitioned PanGenome Graph Of Linked Neighbors), introduces a new representation of the gene repertoire variation as a graph, where each node is a family of homologous genes and each edge is a relation of genetic contiguity. PPanGGOLiN fills the gap between the standard pan genomic approach (that uses a set of independent and isolated gene families) and sequence-level pan genome graph (as reviewed in [23]). The interest of a gene-level graph compared to a sequence graph is that it provides a much more compact structure in clades where gene gains and losses are the major drivers of adaptation. This comes at the cost of disregarding polymorphism in genes and ignoring variation in intergenic regions and introns. However, the genomes of prokaryotes have very small intergenic regions and are almost devoid of introns justifying a focus on the variation of gene repertoires [12], which can be complemented

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

by analysis of intergenic and intragenic polymorphism. PPanGGOLiN uses a new statistical model to classify gene families into persistent, cloud, and one or several shell partitions. Unlike the few statistical methods available [24, 25, 26] that partitions gene families using only their frequency, our method combines the information of occurrence of gene families and the pangenome graph to make the classification. In the following sections we present an overview of the method, an illustration of a pangenome graph and then the partitioning of a large set of prokaryotic species from GenBank. We evaluate the relevance of the persistent genome computed by PPanGGOLiN in comparison to the classical soft core genome. Next, we illustrate the importance of the shell structure and dynamics in the study of the evolution of microbial genomes. Finally, we compare GenBank results to the ones obtained with Metagenome-Assembled Genomes (MAGs) to validate the use of PPanGGOLiN for metagenomic applications.

Results and discussion

Overview of the PPanGGOLiN method

PPanGGOLiN builds pangenomes for large sets of prokaryotic genomes (i.e. several thousands) through a graphical model and a statistical method to classify gene families into three classes: persistent, cloud, and one or several shell partitions. It uses as input a set of annotated genomes with their coding regions classified in homologous gene families. As depicted in Figure 1, PPanGGOLiN integrates information on protein-coding genes and their genomic neighborhood to build a graph where each node is a gene family and each edge is a relation of genetic contiguity (two families are linked in the graph if they contain genes that are neighbors in the genomes). Thanks to this graphical model, the structure of the pangenome is resilient to fragmented assemblies: an assembly gap in one genome can be offset by information from other genomes, thus maintaining the link in the graph. To partition this graph, we established a statistical model taking into consideration that persistent genes share conserved genomic organizations along genomes [27] and that horizontally transferred genes (i.e. shell and cloud genes) tend to insert preferentially in a few chromosomal regions (hotspots) [28]. Thereby, PPanGGOLiN assumes that two gene families that are consistent neighbors in the graph are more likely to belong to the same partition. This is achieved by a hidden Markov Random Field (MRF) whose network is given by the pangenome graph. In parallel, the pangenome is also represented as a binary Presence/Absence (P/A) matrix where the rows correspond to gene families and the columns to genomes. Values are 1 for the presence of at least one member of the gene family and 0 otherwise. This P/A matrix is modeled by a multivariate Bernoulli Mixture Model (BMM). Its parameters are estimated via an Expectation-Maximization (EM) algorithm taking into account the constraints imposed by the MRF. Each gene family is then associated to its closest partition according to the BMM. This results in a partitioned pangenome graph made of nodes that are classified as either persistent, shell or cloud. The strength of the MRF constraints increases according to a parameter called β (if $\beta = 0$, the effect of the MRF is disabled and the partitioning only relies on the P/A matrix) and it depends on the weight of the edges of the pangenome graph which represents the number of gene pairs sharing the neighborhood. Another originality of our method

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

is that, even if the number of partitions (K) is estimated to be equal to 3 (persistent, shell, cloud) in most cases (see ‘Analyses of the most represented species in databanks’ section), more partitions can be used if the pangenome matrix contains several contrasted patterns of P/A. These additional partitions are considered to belong to the shell genome and reflect a heterogeneous structure of the shell (see ‘Shell structure and dynamics’ section).

14 Illustration of a Partitioned Pangenome Graph depicting the *Acinetobacter baumannii* 15 species

To illustrate the results of the PPanGGOLiN method, a pangenome graph was built from 3 117 *Acinetobacter baumannii* genomes from GenBank (Figure 2). The gene families classified as persistent (3 080 families, orange nodes) correspond to the conserved paths. Many islands composed of shell (1 534 families, green nodes) and cloud genomes (64 871 families, blue nodes) interrupt the persistent genome. These islands appear to be frequently inserted in hotspots of the persistent genome thus pinpointing regions of genomic plasticity. The average node degree, i.e. the average number of adjacent edges between gene families of the same partition, is 2.80 for the persistent genome while the shell genome has a higher average degree (3.95, $P=5.0\text{e-}6$ with bilateral unpaired 2-sample Student’s t test) and the cloud a lower one (1.97, $P=3.3\text{e-}40$ with the same test). The shell genome is the most diversified in terms of network topology with many interconnections between families reflecting a mosaic composition of regions from different HGT events [28]. The major part of the cloud has a shell-like graph topology with a large connected component containing 60% of the nodes. In addition, the cloud also contains isolated components that are nearly linear (3 606 components having on average 4.25 nodes) and singletons (10 575 nodes), presumably because it includes very recently acquired genetic material. Finally, large families of mobile genes, mostly transposable elements, can be easily detected because they constitute hubs (i.e. highly connected nodes) in the graph. They vary rapidly their genetic neighborhoods and can be found in multiple loci. As an example of the more detailed analysis that can be done using the graph, a zoom on a region containing the genes required for the synthesis of capsular polysaccharides is highlighted in Figure 2. *A. baumannii* strains are involved in numerous nosocomial infections and their capsule plays key roles in the overall fitness and pathogenicity. Indeed, it protects the bacteria against environmental stresses, host immune responses and can confer resistance to some antimicrobial compounds [29]. Over one hundred distinct capsule types and their corresponding genomic organization have been reported in *A. baumannii* [30]. A zoom on this region of the graph shows a wide variety of combinations of genes for the synthesis of capsular polysaccharides. Based on the *A. baumannii* 3 117 genomes available in GenBank, we detected 229 different paths, sharing many common portions, but only a few are conserved in the species (only 24 paths are covered by more than 10 genomes). Among them, two alternative shell paths seem to be particularly conserved (from the *gnaA* to the *weeH* genes in the figure 3 of [30]). Based on the nomenclature of [30], one (colored in khaki green in the Figure 2) corresponds to the serovar called PSgc12, contains 14 gene families of the shell genome and is fully conserved in 581 genomes. The other (colored in fluo green in the Figure 2)

1
2
3
4
5
6
7
8
9
10
11
12
13
14

corresponds to the serovar PSgc9 (equivalent to PSgc7), contains 11 gene families of the shell genome and is fully conserved in 408 genomes. This analysis illustrates how the partitioned pangenome graph of PPanGGOLiN can be useful to study the plasticity of genomic regions. Thanks to its compact structure in which genes are grouped into families while preserving their genomic neighborhood information, it summarizes the diversity of thousands of genomes in a single picture and allows effective exploration of the different paths among regions or genes of interest.

Analyses of the most represented species in databanks

We used PPanGGOLiN to analyze all prokaryotic species of GenBank for which at least 15 genomes were available. This is the minimal number of genomes we recommend to ensure a relevant partitioning. The quality of the genomes was evaluated before their integration in the graph to avoid poor quality assemblies and taxonomic assignation errors (see Methods). This resulted in a dataset of 439 species pan-genomes, whose metrics are available in Additional file 1. We focused our analysis on the 88 species containing at least 100 genomes (Figure 3). This data was used for in-depth analysis of persistent and shell genomes (see the two next sections). Proteobacteria, Firmicutes and Actinobacteria are the most represented phyla in this dataset and comprise a variety of species, genome sizes and environments. In contrast, Spirochaetes, Bacteroidetes and Chlamydiae phyla are represented by only one or two species (*Leptospira interrogans*, *Bacteroides fragilis*, *Flavobacterium psychrophilum* and *Chlamydia trachomatis*). For each species, we computed the median and interquartile range of persistent, shell and cloud families in the genomes. As expected, we observed a large variation in the range of these values: from pathogens with reduced genomes such as *Bordetella pertussis* or *C. trachomatis* which contain only a small fraction of variable gene families (less than $\approx 5\%$ of shell and cloud genomes) to commensal or environmental bacteria such as *Bifidobacterium longum* and *Burkholderia cenocepacia* whose shell represents more than $\approx 35\%$ of the genome. Furthermore, for a few species the number of estimated partitions (K) is greater than 3 (11 out of 88 species), especially for those with a higher fraction of shell genome. Hence, our method provides a statistical justification for the use of three partitions as a default in pangenome analyses, while indicating that species with large shell content might be best modeled using more partitions (see ‘Shell structure and dynamics’ section).

Estimation of the persistent genome in comparison to the soft core approach

We compared a classical approach to identify the persistent genome to the PPanGGOLiN statistical method to demonstrate its added value. For each species, we performed multiple resamplings of the genome dataset in order to measure the variability of the pangenome metrics according to an increasing number of genomes considered in the analyses (hereafter called rarefaction curves) (see Methods and Additional file 2: Figure S1 as an example for *Lactobacillus plantarum*). These rarefaction curves indicate whether the number of families tends to stabilize, increase or decrease. To this end, the curves were fit with the Heaps’ law where γ represents the growth tendency [31] (hereafter called γ -tendency). The persistent component of a pangenome is supposed to stabilize after the inclusion of a certain number of

6
7
8
9
10
11
12
13
14

15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

genomes, which means it has a γ -tendency close to 0. In addition, interquartile range (IQR) areas along the rarefaction curves were computed to estimate the variability of the predictions compared to the sampling. Small IQR areas mean that the predictions are stable and resilient to sampling. Using these metrics, the PPanGGOLiN predictions of the persistent genome were evaluated for the 88 species previously described and compared to a conventional method where persistent genes are those present in at least 95% of the genomes (generally called the soft core approach). This threshold is so well admitted in pangenomic studies that it is the default parameter in Roary [32] which is to date the most cited software to build bacterial pangenomes.

We observed that the γ -tendency of the PPanGGOLiN persistent is closer to 0 than that of the soft core approach (mean of absolute γ -tendency=9.1e-3 versus 2.5e-2, $P=1.5e-9$ with one-sided paired 2-sample Student's t-test) with a lower standard deviation error too (mean=5.3e-04 versus 2.1e-03, $P=9.5e-11$ with one-sided paired 2-sample Student's t-test) (see Figure 4 and Additional file 1). A major problem of the soft core approach is that the γ -tendency is high for many species (32 species have a γ -tendency above 0.025), suggesting that the size of the persistent genome is not stabilized and tends to be underestimated. Indeed, the persistent genome predicted by PPanGGOLiN is larger or almost equal to the soft core in all species (Figure 3). Besides, the IQR area of the PPanGGOLiN prediction is far below the one of the soft core genome (mean=4906.6 versus 11645.9, $P=8.9e-07$ with unilateral paired 2-sample Student's t test). It can be partially explained because the threshold used in the soft core method induces a 'stair-step effect' along the rarefaction curves depending on the number of genomes sampled. This is illustrated on Figure S1 (see Additional file 2) showing a step every 20 genomes (i.e. corresponding to $20 = \frac{100}{100-95}$ where 95% is the threshold of presence used) on the soft core curve of *L. plantarum*. We found a total of 20 species having atypical values of γ -tendency (absolute value above 0.05) and/or IQR area (above 15 000) for the soft core and only 2 species for the persistent genome of PPanGGOLiN, which are *Bacillus anthracis* and *Burkholderia cenocepacia*. For *B. cenocepacia*, it could be explained by the high heterogeneity of its shell (see next section), which is made of several partitions and complicates its distinction from the persistent genome during the process of partitioning. For *Bacillus anthracis*, the source of variability to define the persistent genome is a result of an incorrect taxonomic assignation in GenBank of about 17% of the genomes that are, according to the Genome Taxonomy DataBase (GTDB) [33], actually *B. cereus* or *B. thuringiensis*. This issue was not detected by our taxonomy control procedure because these species are at the boundary of the conspecific genomic distance threshold used (see Methods). Some of persistent gene families of *bona fide* *B. anthracis* may therefore shift between persistent or shell partitions depending on the resampling. Excluding these misclassified genomes, we predicted a larger persistent genome than the one of the initial set of genomes (about a thousand gene families more) with a γ -tendency much closer to 0 (-0.017 versus a γ -tendency of 0.036 for the soft core genome) and a lower IQR area (8367.0 vs 32167.1). Altogether, these results suggest that our approach provides a better and more robust partitioning of gene families in the persistent genome than the use of arbitrary thresholds. Indeed, the statistical method behind

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

PPanGGOLiN uses directly the information of the gene family P/A whereas the soft core is based only on frequency values. PPanGGOLiN can then classify families with similar frequencies in different partitions by distinguishing them according to their pattern of P/A in the matrix and their genomic neighborhood. The main drawback of using family frequency to partition pangenomes is that even if it was possible to determine the best threshold for each species it would still not take into account that some persistent gene families may have atypically low frequency. This may be due to high gene losses in the population or technical reasons like belonging to a genomic region that is difficult to assemble (i.e. genes that are missing or fragmented in draft genome assemblies).

Shell structure and dynamics

Two types of pangenome evolution dynamics are generally distinguished: open pangenomes and closed ones [1, 2, 31]. From rarefaction curves, the dynamics of pangenomes can be assessed using the γ -tendency of a Heaps' law (see Methods) fitting. A low γ -tendency means a rather closed pangenome whereas a higher γ -tendency means a rather open pangenome. A closed pangenome rigorously means a stabilized pangenome and we found no species obeying this strict criterion (that is to say $\gamma = 0$). This suggests that instead of using binary classifications for pangenomes, it is more useful to quantify the degree of openness of pangenomes given the flux of horizontal gene transfer and gene loss [7]. We computed rarefaction curves for the 88 studied species and determined the γ -tendency for different pangenome components (see Additional file 1 and Additional file 2: Figure S2). The distribution of γ values of the PPanGGOLiN shell genome shows a greater amplitude of values than the other components of the pangenome such as the whole pangenome or the accessory component. This indicates that the main differences in terms of genome dynamics between species seem to reside in the shell genome.

As expected, we found a positive correlation (Spearman's $\rho=0.46$, $P=8.2\text{e-}06$) between the total number of shell gene families in a species and the γ -tendency of the shell (Additional file 2: Figure S3). This means that species with high γ -tendency do accumulate genes that are maintained and exchanged in the population at relatively low frequencies, suggesting they may be locally adaptive. More surprisingly, although one could expect that larger genomes have a larger fraction of variable gene repertoires, the fraction of shell and cloud genes per genome does not correlate with the genome size (Spearman's $\rho=0.007$, $P=0.95$, Figure 5). The results remain qualitatively similar when analyzing the shell or the cloud separately (Additional file 2: Figure S4 and S5). During this analysis, we noticed that, among host-associated bacteria with relatively small genomes (between ≈ 2000 and ≈ 3000 genes), three species (*Bifidobacterium longum*, *Enterococcus faecium* and *Streptococcus suis*) have a high fraction of shell genes ($> 28\%$) but low shell γ -tendency. Two of them (*B. longum* and *E. faecium*) are found in the gut of mammals and the third (*S. suis*) in the upper respiratory tract of pigs. They differ from other host-associated species in our dataset that are mainly human pathogens (e.g. bacteria of the genus *Corynebacterium*, *Neisseria*, *Streptococcus*, *Staphylococcus*) and have a low fraction of shell genes ($< 20\%$). It is possible that these three species have specialized in their ecological niches while maintaining a large and stable pool of

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

shell genes for their adaptation to environmental stress. Further analysis would be required to confirm this hypothesis.

We then investigated the importance of the phylogeny of the species on the patterns of P/A of the shell gene families (shell structure). To this end, Spearman's rank correlations were computed between a Jaccard distance matrix generated on the basis of patterns of P/A of the shell gene families and a genomic distance obtained by Mash pairwise comparisons between genomes [34]. Mash distances were shown to be a good estimate of evolutionary distances for closely related genomes [35]. This correlation was examined in relation to the fraction of gene families that are part of the shell genome for each species (Figure 6). We observed that species with a high fraction of shell ($> 20\%$ of their genome) have a shell structure that is mainly explained by the species phylogeny (i.e. shell P/A are highly correlated with genomic distances, Spearman's $\rho > 0.75$). In addition, PPanGGOLiN predicts a number of partitions (K) for these species often greater than 3. Hence, their shell is more heterogeneous between subclades and becomes structured in several partitions whereas for species with a single shell partition the shell is less structured, possibly indicating many gene exchanges between strains from different lineages. Among the nine species with a large shell genome (excluding *B. anthracis* due to taxonomic assignation errors), only two of them (*Shigella sonnei* and *Lactobacillus reuteri*) showed a relatively low correlation of their shell structure with the phylogeny (Figure 6). For *S. sonnei*, this could be explained by a high number of gene losses in the shell of this species that result from convergent gene loss mediated by insertion sequences (preprint: [36]). For *L. reuteri*, these bacteria colonize the gastrointestinal tract of a wide variety of vertebrate species and have diversified into distinct phylogenetic clades that reflect the host where the strains were isolated, but not their geographical provenance [37]. As illustrated in Figure S6 (see Additional file 2), the shell of *L. reuteri* contains several patterns of P/A that are only partially explained by the species phylogeny. Indeed, we observed clusters of families present across strains from distinct lineages that could contain factors for adaptation to the same host. In contrast, the shell structure of *B. longum* strongly depends on phylogenetic distances showing a clear delineation of adult and infant strains that have specialized into two subspecies (see Additional file 2: Figure S7).

We would like to stress the importance of the shell in the study of the evolutionary dynamics of bacteria. The shell content reflects the adaptive capacities of species through the acquisition of new genes that are maintained in the population. We found that the proportion of shell genes does not increase with the genome size. Instead, the shell accounts for a large fraction of the genomes of species when it is structured in several partitions. We can assume that those species are made of non-homogeneous subclades harboring specific shell genes which contribute to the specialization of the latter. Finally, it could be of interest to associate phenotypes to patterns of shell families that co-occur in different lineages independently of the phylogeny.

Analysis of Metagenome-Assembled Genomes in comparison with isolate genomes

The graph approach should make our tool robust to gaps in genome data, making it a useful tool to analyze pangenomes obtained from MAGs. To test this hypothesis,

3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

we built the pangenomes of the Species-level Genome Bins (SGBs, clusters of MAGs that span a 5% genetic diversity and are assumed to belong to the same species) from the recent paper of Pasolli *et al.* [38]. This study agglomerated and consistently built 4 930 SGBs (154 723 MAGs) from 13 studies focussed on the composition of the human microbiome. We skipped the quality control step (already performed by the authors), and computed the pangenomes following the procedure we used for the GenBank species. The only parameter which differs is the K value which is set to 3 as the detection of several shell partitions is inappropriate for MAGs because of their incompleteness. To make the comparison with GenBank species, SGBs were grouped according to their estimated species taxonomy (provided by the supplementary table S4 of [38]). In this table, we noticed potential errors in the taxonomic assignation of two species (*Blautia obeum* and *Chlamydia trachomatis* corresponding to SGBs 4844 and 6877, respectively) and thus excluded them from the analysis. Keeping the same constraint as previously, only species with at least 15 genomes in both MAGs and GenBank were used for the comparison. A total of just 78 species (corresponding to 151 SGBs) could be analyzed as a lot of microbiome species are laborious to cultivate and thus less represented in databanks (see Additional file 3). Then, we compared the MAG pangenome partitions predicted by PPanGGOLiN with those obtained with GenBank genomes. To perform this, we aligned MAG and GenBank families for each species and computed the percentage of common families for each partition (see Methods and Additional file 3).

We observed that the size of the estimated persistent genome of MAGs is similar to the one of GenBank genomes for most species (Figure 7). In 55 out of the 78 species, the absolute fold change of persistent size is less than 1.2 and $\approx 90\%$ ($SD=5\%$) of its content is common between MAGs and GenBank genomes. The 23 other species with more important differences showed smaller persistent genomes with only 60% ($SD=15\%$) of the persistent genome of GenBank being found in MAGs. For these species, the PPanGGOLiN method missed a fraction of the persistent genome due to the incompleteness of MAGs. Indeed in such cases, the missing gene families are mostly classified in the shell of the MAGs which contains 32% ($SD=11\%$) of the GenBank persistent families. Nevertheless, 89% ($SD=9\%$) of the MAG persistent families match the GenBank ones, meaning that PPanGGOLiN correctly assigned persistent families for MAGs even if the persistent genome of these 23 species is incomplete.

However, two species, *Bifidobacterium longum* and *Faecalibacterium prausnitzii*, have less than 75% of their MAG persistent families in common with GenBank ones. For *B. longum*, this could be explained by the fact that the MAGs were obtained mostly from human adult samples while this species in databanks are from a broader host range (infants and pigs). It means that the MAG persistent might contain additional genes related to host-specificity. As a matter of fact, 412 gene families from the MAG persistent (25% of the total MAG persistent) are found in the GenBank shell which supports our hypothesis. For *F. prausnitzii*, the differences might be explained by a poor estimation of the persistent using GenBank data due to the low number of considered genomes (17 genomes versus 4232 MAGs). As expected, the soft core (based on the usual threshold of 95% presence) is unrealistically low in the MAG species with only ≈ 98 gene families

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

on average and only 4 species out of 78 having more than 500 families classified in the soft core (see Additional file 3). Hence, the soft core approach is not well adapted to the analysis of MAGs. Furthermore, using lower thresholds of presence is not adequate because defining a unique threshold for all the families misses the heterogeneity of gene family presence in MAGs.

To explore the diversity within pangenomes, we compared the shell of GenBank genomes and MAGs for the 55 species with similar persistent genomes. Interestingly, we observed for all the 55 species only a partial overlap between the MAGs and GenBank shells (see Additional file 2: Figure S8). Indeed, as the MAGs are obtained only from a specific environment (i.e. the human microbiome), the diversity of GenBank is not fully captured by MAGs. It is especially the case for most of the Firmicutes and Proteobacteria. Conversely, most of the MAGs of Bacteroidetes phylum cover more than half of GenBank diversity while containing a large fraction of shell genes that are lacking in the shell of isolate genomes (i.e. less than 45% of the families are represented in the shell of GenBank). As already reported by Pasolli *et al.* [38], this confirms that the MAGs considerably improve the estimate of the genetic diversity of Bacteroidetes which are key players in the gut microbiome.

In summary, we have shown that PPanGGOLiN is able to provide an estimation of the persistent genome even using MAGs that may miss significant numbers of genes. Indeed, the incompleteness of MAGs and potential contaminations by fragments from other genomes complicates the analysis of their pangenome. This is especially the case for the accessory genome because its assembly coverage and nucleotide composition generally differ from those of the persistent genome making the binning of these regions more difficult. Nevertheless, PPanGGOLiN is able to find shell gene families in MAGs bringing new genes that may be important for species adaptation in the microbiome. Hence, it enables further analyses, even for uncultured species lacking reference genomes, such as the reconstruction of the core metabolism from the persistent genome to predict culture media or the study of the landscape of horizontally transferred genes within species.

Conclusion

We have presented here the PPanGGOLiN method that enables the partitioning of pangenomes in persistent, shell and cloud genomes using a gene family graph approach. This compact structure is useful to depict the overall genomic diversity of thousands of strains highlighting variable paths made of shell and cloud genes within the persistent backbone. The statistical model behind PPanGGOLiN makes a better estimate of the persistent genome than classical approaches based on gene family frequencies in isolate genomes and also in MAGs. The definition of shell partitions based on statistical criteria allowed us to understand genome dynamics within species. We observed different patterns of shell with regard to phylogeny that may suggest different adaptive paths for the diversification of the species.

Future applications of PPanGGOLiN could include the prediction of genomic islands within the shell and cloud genomes. A first version of this application (Bazin *et al.*, in preparation) is already integrated in the MicroScope genome analysis platform [39]. Next, it would be interesting to determine the architecture of these variable regions by predicting conserved gene modules using information on the

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

occurrence of families and their genomic neighborhood in the pangenome graph. Regarding metagenomics, pangenome graphs of PPanGGOLiN could be used as a reference (i.e. instead of individual genomes) for species quantification by mapping short or long reads on the graph to compute the coverage of the persistent genome. Indeed, each gene families node of the partitioned pangenome graph could embed a variation graph as an alignment template. Moreover, coverage variation in the shell or cloud genomes could allow the detection of strain-specific paths in the graph that are signatures of distinctive traits within microbiomes.

To conclude, the graph-based approach proposed by PPanGGOLiN provides an effective basis for very large scale comparative genomics and we hope that drawing genomes on rails like a subway map may help biologists navigate the great diversity of microbial life.

Methods

To explain the partitioning of pangenomes, we first need to describe the method based on the P/A matrix only (BinEM) and then the method built upon it that uses the pangenome graph to improve the partitioning (NEM).

Modeling the P/A matrix via a Multivariate Bernoulli Mixture Model

PPanGGOLiN aims to classify patterns of P/A of gene families into K partitions ($K \in \mathbb{N}; K \geq 3$). Input data consists of a binary matrix X in which a $X_{i,j}$ entry is 1 if family i is present in genome j and 0 otherwise (Figure 1) where $1 \leq i \leq F$ in each of the F gene families and $1 \leq j \leq N$ in each of the N genomes. A first approach for partitioning the data relies on a multivariate Bernoulli Mixture Model (BMM) estimated through the Expectation-Maximization (EM) algorithm [40] (named the BinEM method). The number of partitions K may be larger than 3 (persistent, shell and cloud) due to the possible presence of antagonist P/A patterns in the different strains of a species. Therefore, two of the partitions will correspond to the persistent and cloud genome and a number of $K - 2$ partitions will correspond to the shell genome. The value of K can be either provided by the user or determined automatically (see next section).

In the BMM, the matrix comprises data vectors $X_i = (X_{ij})_{1 \leq j \leq N}$ describing P/A of families, which are assumed to be independent and identically distributed with a mixture distribution given by:

$$P(X_i = (x_{ij})_{1 \leq j \leq N}) = \sum_{k=1}^K \pi_k \prod_{j=1}^N \epsilon_{kj}^{|x_{ij} - \mu_{kj}|} (1 - \epsilon_{kj})^{1 - |x_{ij} - \mu_{kj}|}$$

where $\pi = (\pi_1, \dots, \pi_k, \dots, \pi_K)$ denotes the mixing proportions satisfying $\pi_k \in]0, 1[$; $(\sum_{k=1}^Q \pi_k) = 1$ and where π_k is the unknown proportion of gene families belonging to the k^{th} partition. Moreover, $\mu_k = (\mu_{kj})_{1 \leq j \leq N} \in \{0; 1\}^N$ are the centroid vectors of P/A of the k^{th} partition representing the most probable binary state and $\epsilon_k = (\epsilon_{kj})_{1 \leq j \leq N} \in [0, \frac{1}{2}]^N$ are the unknown vectors of dispersion around μ_k . The parameters of this model, as well as corresponding partitions are estimated by the EM algorithm. To speed up the computation of the EM algorithm, a heuristic is used to initialize the BMM parameters in order to converge to a relevant partitioning using

fewer EM-steps. This heuristics consists in setting π_k with equiprobable proportions equal to $1/K$ while the ϵ_{kj} and μ_{kj} parameters are triangularly initialized.

Given $s = 1/\lceil K/2 \rceil$, the triangular initialization consists of:

$$\begin{aligned}\{\mu_{kj}\}_{1 \leq k \leq K/2, 1 \leq j \leq N} &= 1 \\ \{\mu_{kj}\}_{K/2 < k \leq K, 1 \leq j \leq N} &= 0 \\ \{\epsilon_{kj}\}_{1 \leq k \leq K/2, 1 \leq j \leq N} &= s \cdot k \\ \{\epsilon_{kj}\}_{K/2 < k \leq K, 1 \leq j \leq N} &= s \cdot (K - k + 1)\end{aligned}$$

An interesting consequence of this initialization is that the persistent genome will be the first partition ($k = 1$) while the cloud genome will correspond to the last partition ($k = K$). This particular initialization solves the classical label switching problem in our context.

Partitioning of the P/A matrix

To perform the partitioning, each gene family i must be allocated in one partition only and the variables $\{Z_i\}_{1 \leq i \leq F}$ with state space $\{1, \dots, K\}$ indicate the partition to which each gene family i belongs. Therefore, once the parameters are optimized, the method automatically assigns the gene families to their most probable partition z_i according to the model if its estimated posterior probability is above 0.5 and to the shell otherwise. The default values of the dispersion vector ϵ_k associated to each centroid vector μ_k are constrained to be identical for all the ϵ_{kj} of a specific k partition (for all the genomes of a specific partition) in order to avoid over-fitting but it is possible to release this constraint.

Selection of the optimal number of partitions (K)

To determine the optimal K , named \hat{K} , the algorithm runs multiple partitionings increasing K . After the first initial steps of the EM algorithm (10 steps by default), the Integrated Completed Likelihood (ICL) [41] associated with each K is computed. The ICL corresponds to Bayesian Information Criterion (BIC) [42] penalized by the estimated mean entropy and is calculated as:

$$ICL(K) = BIC(K) - \sum_{k=1}^K \sum_{i=1}^F p(z_i | X, \hat{\theta}, k) \log(p(z_i | X, \hat{\theta}, k)); \forall p(z_i | X, \hat{\theta}, k) > 0$$

and

$$BIC(K) = \log \mathbb{P}_K(X | \hat{\theta}) - 1/2 \dim(K) \log F$$

where $\log \mathbb{P}_K(X | \theta)$ is the data log-likelihood under a multivariate BMM with K partitions and $\theta = (\{\pi_k\}_{1 \leq k \leq K}, \{\mu_{kj}\}_{1 \leq k \leq K, 1 \leq j \leq N}, \{\epsilon_{kj}\}_{1 \leq k \leq K, 1 \leq j \leq N})$. This log-likelihood can be calculated as follows:

$$\log \mathbb{P}_K(X | \theta) = \sum_{i=1}^F \log \left(\sum_{k=1}^K \pi_k \prod_{j=1}^N \epsilon_{kj}^{|x_{ij} - \mu_{kj}|} (1 - \epsilon_{kj})^{1 - |x_{ij} - \mu_{kj}|} \right)$$

Moreover, $\hat{\theta}$ is the maximum likelihood estimator (approximated through the BinEM algorithm) and $\dim(K)$ is the dimension of the parameter space for this

1
2
3
4
5
6
7
8
9
10
11
12
13
14

model. Here, $\dim(K) = K - 1 + 2NK$ if the dispersion vector ϵ_k associated to each centroid vector μ_k is constrained to be identical for all the ϵ_{kj} of a specific k partition and $\dim(K) = K - 1 + NK$ if the dispersion vector ϵ_k is free. Relying on this criterion, one selects the best number of partitions as $\hat{K} = \arg \min_K ((1 - \delta_{ICL}) ICL(K))$ where δ_{ICL} is a small enough margin to avoid obtaining a too high K value which would provide no significant gain regarding a lower K (by default $\delta_{ICL} = 0.05 \times (\max(ICL) - \min(ICL))$).

15 Generation of the pangenome graph

16 Pangome graphs consist in a graph-based representation to store and visualize
17 pangomes. In the PPanGGOLiN representation, the nodes correspond to gene
18 families and the edges to genomic neighborhood information. In this respect, two
19 nodes are connected if two gene families contain at least one pair of genes that are
20 adjacent in a genome. Edges are labeled with corresponding genome identifiers and
21 weighted by the proportion of genomes sharing this link. This process results in a
22 pangome graph (see Figure 2 as an example).

23 Formally, a pangome graph $G = (V, E)$ is a graph having a set of vertices $V =$
24 $\{(v_i)_{(1 \leq i \leq F)}\}$ where F is the number of gene families in the pangome associated
25 with a set of edges $E = \{e_{i \sim i'}\} = \{(v_i, v_{i'})\}, v_i \in V, v_{i'} \in V$ where the couple
26 of vertices $(v_i, v_{i'})$ are gene families having their genes $(v_{i,j}, v_{i',j})$ adjacent on the
27 genome j and where the function $\text{countNeighboringGenes}(v_{i,j}, v_{i',j})$ count these
28 occurrences of adjacencies in the N genomes. Each edge $\{e_{i \sim i'}\}$ has a weight $w_{i \sim i'}$
29 where $w_{i \sim i'} = \frac{1}{N} \sum_{j=1}^N \text{countNeighboringGenes}(v_{i,j}, v_{i',j})$.

30 Partitioning via Neighboring Expectation-Maximization

31 From the graph previously described, neighborhood information on the gene families
32 is used to improve the partitioning results. Indeed, the BinEM approach described
33 above is extended by combining the P/A matrix X with the pangome graph G .
34 This relies on a hidden Markov Random Field (MRF) model whose graph structure
35 is given by G . In this model, each node belongs to some unobserved (hidden)
36 partitions which are distributed among gene families according to a MRF which
37 assumes that two neighbors are more likely to belong to the same partition. Conditional
38 on this hidden structure, the binary vectors of P/A are independent and
39 follow a multivariate Bernoulli distribution, with proportion vectors depending on
40 the associated partition. This approach is called NEM, as it relies on the Neighboring
41 Expectation-Maximization algorithm [43, 44, 45]. As such, NEM will tend
42 to smooth the partitioning by grouping gene families having a weighted majority
43 of neighbors belonging to the same partition. The previously introduced variable
44 $\{Z_i\}_{1 \leq i \leq F}$ is still a latent variable indicating the partition to which each gene family
45 belongs. These random variables are now distributed according to a MRF. More
46 precisely they have the following Gibbs distribution:

$$56 \quad \mathbb{P}(\{Z_i\}_{1 \leq i \leq F}) = W_{\beta}^{-1} \exp\left(\sum_{i=1}^F \sum_{k=1}^K \pi_k 1_{Z_i=k} + \beta \frac{F}{\sum_{i \sim i'} w_{i \sim i'}} \sum_{i \sim i'} w_{i \sim i'} 1_{Z_i=Z_{i'}}\right)$$

57 where 1_A is the indicator function of event A , the second sum concerns every pair
58 $(i \sim i')$ of neighbor gene families. The parameter $\beta \geq 0$ corresponds to the coefficient
59

60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

of spatial regularity. The $\frac{F}{\sum_{i \sim i'}^E w_{i \sim i'}}$ is a corrector term ensuring that the spatial smoothing is balanced by whatever the number of gene families. Indeed when the number of genomes (N) increases, the number of gene families (F) tends to be higher than the sum of the edge weights. Finally,

$$W_\beta = \sum_{\{\tilde{z}_i\} \in \{1 \dots K\}^F} \exp\left(\sum_{i=1}^F \sum_{k=1}^K \pi_k 1_{\tilde{z}_i=k} + \beta \frac{F}{\sum_{i \sim i'} w_{i \sim i'}} \sum_{i \sim i'} w_{i \sim i'} 1_{\tilde{z}_i=\tilde{z}_{i'}}\right)$$

is a normalizing constant. Note that W_β may not be computed, due to a large number of possible configurations. The degree of dependence between elements is controlled by the parameter β . Neighboring elements will be more inclined to belong to the same group with a higher value of this parameter. Here, the data vectors $(X_i)_{1 \leq i \leq F}$ are not independent anymore. However, conditional on the latent groups $(Z_i)_{1 \leq i \leq F}$, they are independent and follow the multivariate Bernoulli distribution:

$$\mathbb{P}(\{X_i\}_{1 \leq i \leq F} | \{Z_i\}_{1 \leq i \leq F}) = \prod_{i=1}^F \prod_{j=1}^N \epsilon_{Z_i, j}^{|x_{ij} - \mu_{Z_i, j}|} (1 - \epsilon_{Z_i, j})^{1 - |x_{ij} - \mu_{Z_i, j}|}.$$

Many different techniques may be used to approximate the maximum likelihood estimator in hidden MRF. NEM relies on a mean-field approximation for the distribution of the latent random variables $Z_{1 \leq i \leq F}$ conditional on the observations. It should be noted that the optimal number of partitions (K) is not determined automatically using NEM and is therefore first estimated using the BinEM approach.

Issues resulting from high-dimensional statistics

As plenty of statistical approaches, NEM is not adapted to high dimensional settings (i.e. whenever the condition $F \gg N$ is not satisfied anymore). Actually, it can be the case in pangenomics as the number of new families added to the pangenome slightly decreases when new genomes are added (see figure 3 in [1]). Mathematical solutions to this issue seem to exist [46, 47, 48] for example via the weighting of features, corresponding to the weighting of genomes in our case. An improved version of NEM should include this improvement and could be perspective of this work. Another drawback of the method is to quadratically scale with the number of genomes leading to heavy computations when thousands of genomes are included in the analysis. Moreover, the method is hard to parallelize as it stands. To date, outside of our analysis the pangenomics field relies on up to 4893 genomes [49] in a single study so the approach must be designed to scale up to thousands of genomes. Our solution to the mentioned issues is to sample the genomes in chunks and to perform multiple partitioning simultaneously. Each family must be involved in at least $N_{total}/N_{samples}$ samplings and will be partitioned only if it is classified in the same partition in at least 50% of all sampling where it is present (absolute majority). If some families do not respect this condition, more samplings are done until all gene families have been partitioned. Chunks have to be large enough to be representative this is why a size of at least 500 genomes is advised.

3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Analysis of isolate genomes and Metagenome-Assembled Genomes

To obtain the set of isolate genomes to be analyzed, we downloaded all archaeal and bacterial genomes (220 561 genomes) of the GenBank database at the date of the 17th of April 2019. Genomes having more than 1000 contigs and a L90 > 100, a content of the “assembly_status.txt” file different from “status=latest” (as another flag indicates poor quality genomes) were discarded. For each species (identified by its NCBI species taxid), a pairwise genomic distance matrix was computed using Mash (version 2.0) [34]. To avoid redundancy, if several genomes share together a Mash distance < 0.0001, only one was kept (the one having the minimum number of contigs). A single linkage clustering using SiLiX [50] was then performed on the adjacency graph of the Mash distance matrix considering only distances below or equal to 0.06, which corresponds to a 94% Average Nucleotide Identity (ANI) cutoff which is a usual value to define species [51]. Genomes that were not in the largest connected component were discarded to remove potential taxonomic assignation errors. Only species having at least 15 remaining genomes were then considered for the analysis. This dataset consists of 439 species encompassing 136 287 genomes (see Additional file 4). MAGs from the Pasolli *et al.* study [38] were downloaded from http://segatalab.cibio.unitn.it/data/Pasolli_et_al.html. In this dataset, the genomes are already grouped in Species Genome Bins. These SGBs do not exactly match the GenBank taxonomy. Thus, SGBs assigned with the same species name (column “estimated taxonomy” in the supplementary table S4 of [38]) were merged to allow comparison with GenBank. SGBs that do not have a taxonomy assigned at the species level were not considered. A total of 583 species encompassing 698 SGBs and 71 766 MAGs were analyzed but only 78 species were finally compared to GenBank species. To avoid introducing a bias in our analysis due to the gene calling, GenBank annotation were not considered as there are obtained using a variety of annotation workflows. Genomes from GenBank and MAGs were consistently annotated using the procedure implemented in PPanGGOLiN. Prodigal (version 2.6.2) [52] is used to detect the coding genes (CDS). tRNA and tmRNA genes are predicted using Aragorn (version 1.2.38) [53] whereas the rRNA are detected using Infernal (version 1.1.2) [54] with HMM models from Rfam [55]. In the case of overlaps between a RNA and a CDS, the overlapping CDS are discarded. Homologous gene families are determined using MMseqs2 (version 8-fac81) [56] with the following parameters: coverage=80% with cov-mode=0, minimal amino acid sequence identity=80% and cluster-mode=0 corresponding to the Greedy Set Cover clustering mode. PPanGGOLiN partitioning was executed on each species using the NEM approach with a parameter $\beta = 2.5$. The nodes having a degree above 10 were not considered to smooth the partitioning in the MRF (parameter “-ms 10”). The number of partitions (K) was determined automatically for each NCBI species using a $\delta_{ICL} = 0.05$. It was fixed at 3 for the MAG analysis. The partitioning was done using chunks of size 500 when there were more than 500 genomes in a species. To compare PPanGGOLiN results between MAGs and GenBank genomes for each species, the representative sequences of each MAG gene family (extracted using the mmseqs2 subcommand: “result2repseq”) are aligned (using mmseqs2 “search”) on those of GenBank genomes. If the best hit of the query has a sequence identity > 80% and a coverage > 80% of the target, the 2 corresponding gene families of each dataset are associated.

3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Rarefaction curves

To represent the pangenome evolution according to the number of sequenced genomes, a multiple resampling approach was used. For each species with at least 100 genomes, 8 rarefaction curves showing the evolution of the persistent, shell, cloud, soft core, soft accessory, exact core and exact accessory components were computed for sizes of 1 to 100 genomes and with 30 samples for each subset size. For each sample, the number of partitions K is automatically determined between 3 and the K obtained on all the genomes of the species. A non-linear Least Squares Regression was performed to fit the rarefaction curves with Heaps' law $F = \kappa N^\gamma$ where F is the number of gene families, N the number of genomes, γ the tendency of the evolution and κ a proportional factor [31]. Subset sizes ≤ 15 were not used for the fitting as they are sometimes too variable to ensure a good fitting. The function “scipy.optimize.curve_fit” of the Python scipy package (version 1.0.0), based on the Levenberg-Marquardt algorithm, was used. For each subset size, the median and quartiles were calculated to obtain a ribbon of interquartile ranges (IQR) along the rarefaction curves. We call the area of this ribbon the IQR area (see Additional file 2: Figure S1 as an example).

PPanGGOLiN software implementation

PPanGGOLiN is designed to be a software suite performing the annotation of the genomic sequences, building the gene families and the pangenome graph before partitioning it. Users can also provide their own annotations (GFF format) and gene families. The application stores its data in a compressed HDF5 file but can also return the graph in GEXF or JSON formats and the P/A matrix with the partitioning in CSV or Rtab files (similarly to the ones provided by Roary [32]). It also generates several illustrative figures, some of which are presented in the article. It was developed in the Python 3 and C languages and is intended to be easily installable on Linux and Mac OS systems via a BioConda package [57] (see <https://bioconda.github.io/recipes/ppanggolin/README.html>). The code is also freely available on the GitHub website at the following address: <https://github.com/labgem/PPanGGOLiN>.

Availability of data and materials

Archaeal and bacterial genomes were downloaded from the NCBI FTP server (<ftp://ftp.ncbi.nlm.nih.gov/genomes/genbank>) at the date of the 17th of April 2019. Metagenome-Assembled Genomes were downloaded from http://segatalab.cibio.unitn.it/data/Pasolli_et_al.html. All analyses described here were run using PPanGGOLiN software (version 1.0). PPanGGOLiN source code is freely available from <https://github.com/labgem/PPanGGOLiN> under a CeCILL license. PPanGGOLiN is written in Python 3 and C languages and has been tested on Linux and Mac OS. It can be installed through conda.

Competing interests

The authors declare that they have no competing interests to declare.

Author's contributions

GG and AB developed the software and performed the analyses. DV proposed the concepts and supervised the work. GG, AB, EPCR and DV wrote the manuscript. GG, CA and CMa developed the statistical approach. AC, MT, AP, SC, CA, CMa and CMb provided knowledge and expertise to analyze the results, improved concepts and the manuscript. GG, AB, RP, MG, LB and MD developed the workflow to analyze the NCBI genomes. All authors read and approved the final manuscript.

Acknowledgements

We acknowledge Alexandre Renaux and Jonathan Mercier for their preliminary insights on pangenome graphs. We thank Mélanie Buy for drawing the PPanGGOLiN logo. Finally, we thank Guilhem Royer, Valentin Sabatet, Johan Rollin, Mohammed-Amin Madoui, Tom Delmont, Nicolas Pons and Pierre Peterlongo for all their advice along this work.

3
4
5
6
7
8
9

Funding

This research was supported in part by the IRTELIS PhD program of the French Alternative Energies and Atomic Energy Commission (CEA) for GG, the French Government 'Investissements d'Avenir' programs namely FRANCE GENOMIQUE [ANR-10-INBS-09-08], the INSTITUT FRANÇAIS DE BIOINFORMATIQUE [ANR-11-INBS-0013] and the Agence Nationale de la Recherche [Projet ANR-16-CE12-29 for EPCR].

10

Consent for publication

11 Not applicable.

12

Ethics approval and consent to participate

13 Not applicable.

14

Author details

15 1 LABGeM, Génomique Métabolique, CEA, Genoscope, Institut François Jacob, Université d'Évry, Université
16 Paris-Saclay, CNRS, Evry, France. 2 Microbial Evolutionary Genomics, Institut Pasteur, CNRS, UMR3525, Paris,
17 France. 3 Laboratoire de Probabilités, Statistique et Modélisation, Sorbonne Université, Université de Paris, Centre
18 National de la Recherche Scientifique, Paris, France. 4 Laboratoire de Mathématiques et Modélisation d'Evry, UMR
19 CNRS 8071, Université d'Evry Val d'Essonne, Evry, France. 5 Sorbonne Université, Collège doctoral, Paris, France.

20

References

1. Tettelin, H., Masignani, V., Cieslewicz, M.J., Donati, C., Medini, D., Ward, N.L., Angiuoli, S.V., Crabtree, J., Jones, A.L., Durkin, A.S., Deboy, R.T., Davidsen, T.M., Mora, M., Scarselli, M., Margarit y Ros, I., Peterson, J.D., Hauser, C.R., Sundaram, J.P., Nelson, W.C., Madupu, R., Brinkac, L.M., Dodson, R.J., Rosovitz, M.J., Sullivan, S.A., Daugherty, S.C., Haft, D.H., Selengut, J., Gwinn, M.L., Zhou, L., Zafar, N., Khouri, H., Radune, D., Dimitrov, G., Watkins, K., O'Connor, K.J., Smith, S., Utterback, T.R., White, O., Rubens, C.E., Grandi, G., Madoff, L.C., Kasper, D.L., Telford, J.L., Wessels, M.R., Rappuoli, R., Fraser, C.M.: Genome analysis of multiple pathogenic isolates of *Streptococcus agalactiae*: implications for the microbial "pan-genome". *Proc. Natl. Acad. Sci. U.S.A.* **102**(39), 13950–13955 (2005)
2. Medini, D., Donati, C., Tettelin, H., Masignani, V., Rappuoli, R.: The microbial pan-genome. *Curr. Opin. Genet. Dev.* **15**(6), 589–594 (2005)
3. Treangen, T.J., Rocha, E.P.C.: Horizontal transfer, not duplication, drives the expansion of protein families in prokaryotes. *PLOS Genetics* **7**(1), 1–12 (2011). doi:[10.1371/journal.pgen.1001284](https://doi.org/10.1371/journal.pgen.1001284)
4. Lukjancenko, O., Wassenaar, T.M., Ussery, D.W.: Comparison of 61 sequenced *Escherichia coli* genomes. *Microb. Ecol.* **60**(4), 708–720 (2010)
5. Acevedo-Rocha, C.G., Fang, G., Schmidt, M., Ussery, D.W., Danchin, A.: From essential to persistent genes: a functional approach to constructing synthetic life. *Trends Genet.* **29**(5), 273–279 (2013)
6. Contreras-Moreira, B., Vinuesa, P.: Get_homologues, a versatile software package for scalable and robust microbial pangenome analysis. *Appl. Environ. Microbiol.* **79**(24), 7696–7701 (2013)
7. Lapierre, P., Gogarten, J.P.: Estimating the size of the bacterial pan-genome. *Trends Genet.* **25**(3), 107–110 (2009)
8. Bolotin, E., Hershberg, R.: Horizontally Acquired Genes Are Often Shared between Closely Related Bacterial Species. *Front Microbiol* **8**, 1536 (2017)
9. Vesth, T., Wassenaar, T.M., Hallin, P.F., Snipen, L., Lagesen, K., Ussery, D.W.: On the origins of a vibrio species. *Microbial Ecology* **59**(1), 1–13 (2010). doi:[10.1007/s00248-009-9596-7](https://doi.org/10.1007/s00248-009-9596-7)
10. Periwal, V., Patowary, A., Vellarikkal, S.K., Gupta, A., Singh, M., Mittal, A., Jeyapaul, S., Chauhan, R.K., Singh, A.V., Singh, P.K., Garg, P., Katoch, V.M., Katoch, K., Chauhan, D.S., Sivasubbu, S., Scaria, V.: Comparative whole-genome analysis of clinical isolates reveals characteristic architecture of *Mycobacterium tuberculosis* pangenome. *PLoS ONE* **10**(4), e0122979 (2015)
11. Livingstone, P.G., Morphew, R.M., Whitworth, D.E.: Genome Sequencing and Pan-Genome Analysis of 23 *Corallococcus* spp. Strains Reveal Unexpected Diversity, With Particular Plasticity of Predatory Gene Sets. *Front Microbiol* **9**, 3187 (2018)
12. Koonin, E.V., Wolf, Y.I.: Genomics of bacteria and archaea: the emerging dynamic view of the prokaryotic world. *Nucleic Acids Res.* **36**(21), 6688–6719 (2008)
13. Baumdicker, F., Hess, W.R., Pfaffelhuber, P.: The infinitely many genes model for the distributed genome of bacteria. *Genome Biol Evol* **4**(4), 443–456 (2012)
14. Collins, R.E., Higgs, P.G.: Testing the infinitely many genes model for the evolution of the bacterial core genome and pangenome. *Mol. Biol. Evol.* **29**(11), 3413–3425 (2012)
15. Lobkovsky, A.E., Wolf, Y.I., Koonin, E.V.: Gene frequency distributions reject a neutral model of genome evolution. *Genome Biology and Evolution* (2013)
16. Bolotin, E., Hershberg, R.: Gene Loss Dominates As a Source of Genetic Variation within Clonal Pathogenic Bacterial Species. *Genome Biol Evol* **7**(8), 2173–2187 (2015)
17. Moldovan, M.A., Gelfand, M.S.: Pangenomic definition of prokaryotic species and the phylogenetic structure of prochlorococcus spp. *Frontiers in Microbiology* **9**, 428 (2018). doi:[10.3389/fmicb.2018.00428](https://doi.org/10.3389/fmicb.2018.00428)
18. Chan, A.P., Sutton, G., DePew, J., Krishnakumar, R., Choi, Y., Huang, X.Z., Beck, E., Harkins, D.M., Kim, M., Lesho, E.P., Nikolich, M.P., Fouts, D.E.: A novel method of consensus pan- chromosome assembly and large-scale comparative analysis reveal the highly flexible pan-genome of *acinetobacter baumannii*. *Genome Biol.* **16**, 143 (2015)
19. Garrison, E., Siren, J., Novak, A.M., Hickey, G., Eizenga, J.M., Dawson, E.T., Jones, W., Garg, S., Markello, C., Lin, M.F., Paten, B., Durbin, R.: Variation graph toolkit improves read mapping by representing genetic variation in the reference. *Nat. Biotechnol.* **36**(9), 875–879 (2018)
20. Kim, D., Paggi, J.M., Park, C., Bennett, C., Salzberg, S.L.: Graph-based genome alignment and genotyping with HISAT2 and HISAT-genotype. *Nat. Biotechnol.* **37**(8), 907–915 (2019)

63
64
65

3

4

5

6 21. Rakocevic, G., Semenyuk, V., Lee, W.P., Spencer, J., Browning, J., Johnson, I.J., Arsenijevic, V., Nadj, J., Ghose, K., Suciu, M.C., Ji, S.G., Demir, G., Li, L., Topa?, B.C., Dolgorobodov, A., Pollex, B., Spulber, I., Glotova, I., Komar, P., Stachyra, A.L., Li, Y., Popovic, M., Kallberg, M., Jain, A., Kural, D.: Fast and accurate genomic analyses using genome graphs. *Nat. Genet.* **51**(2), 354–362 (2019)

7 22. Consortium, T.C.P.-G.: Computational pan-genomics: status, promises and challenges. *Brief. Bioinformatics* (2016)

8 23. Zekic, T., Holley, G., Stoye, J.: Pan-Genome Storage and Analysis Techniques. *Methods Mol. Biol.* **1704**, 29–53 (2018)

9 24. Snipen, L., Almøy, T., Ussery, D.W.: Microbial comparative pan-genomics using binomial mixture models. *BMC Genomics* **10**, 385 (2009)

10 25. van Tonder, A.J., Mistry, S., Bray, J.E., Hill, D.M., Cody, A.J., Farmer, C.L., Klugman, K.P., von Gottberg, A., Bentley, S.D., Parkhill, J., Jolley, K.A., Maiden, M.C., Brueggemann, A.B.: Defining the estimated core genome of bacterial populations using a Bayesian decision model. *PLoS Comput. Biol.* **10**(8), 1003788 (2014)

11 26. Gumi?re, T., Meyer, K., Burns, A.R., Gumi?re, S.J., Bohannan, B.J.M., Andreote, F.D.: A probabilistic model to identify the core microbial community. *bioRxiv* (2018). doi:[10.1101/491183](https://doi.org/10.1101/491183). <https://www.biorxiv.org/content/early/2018/12/10/491183.full.pdf>

12 27. Fang, G., Rocha, E.P., Danchin, A.: Persistence drives gene clustering in bacterial genomes. *BMC Genomics* **9**(1), 4 (2008)

13 28. Oliveira, P.H., Touchon, M., Cury, J., Rocha, E.P.C.: The chromosomal organization of horizontal gene transfer in bacteria. *Nat Commun* **8**(1), 841 (2017)

14 29. Singh, J.K., Adams, F.G., Brown, M.H.: Diversity and Function of Capsular Polysaccharide in *Acinetobacter baumannii*. *Front Microbiol* **9**, 3301 (2018)

15 30. Hu, D., Liu, B., Dijkshoorn, L., Wang, L., Reeves, P.R.: Diversity in the major polysaccharide antigen of *Acinetobacter baumannii* assessed by DNA sequencing, and development of a molecular serotyping scheme. *PLoS ONE* **8**(7), 70329 (2013)

16 31. Tettelin, H., Riley, D., Cattuto, C., Medini, D.: Comparative genomics: the bacterial pan-genome. *Curr. Opin. Microbiol.* **11**(5), 472–477 (2008)

17 32. Page, A.J., Cummins, C.A., Hunt, M., Wong, V.K., Reuter, S., Holden, M.T., Fookes, M., Falush, D., Keane, J.A., Parkhill, J.: Roary: Rapid large-scale prokaryote pan genome analysis. *Bioinformatics* **31**(22), 3691–3693 (2015)

18 33. Parks, D.H., Chuvochina, M., Waite, D.W., Rinke, C., Skarszewski, A., Chaumeil, P.A., Hugenholtz, P.: A standardized bacterial taxonomy based on genome phylogeny substantially revises the tree of life. *Nat. Biotechnol.* **36**(10), 996–1004 (2018)

19 34. Ondov, B.D., Treangen, T.J., Melsted, P., Mallonee, A.B., Bergman, N.H., Koren, S., Phillippy, A.M.: Mash: fast genome and metagenome distance estimation using minhash. *Genome Biology* **17**(1), 132 (2016). doi:[10.1186/s13059-016-0997-x](https://doi.org/10.1186/s13059-016-0997-x)

20 35. Criscuolo, A.: A fast alignment-free bioinformatics procedure to infer accurate distance-based phylogenetic trees from genome assemblies. *Research Ideas and Outcomes* **5**, 36178 (2019). doi:[10.3897/rio.5.e36178](https://doi.org/10.3897/rio.5.e36178). <https://doi.org/10.3897/rio.5.e36178>

21 36. Hawkey, J., Monk, J.M., Billman-Jacobe, H., Palsson, B., Holt, K.E.: Impact of insertion sequences on convergent evolution of shigella species. *bioRxiv* (2019). doi:[10.1101/680777](https://doi.org/10.1101/680777). <https://www.biorxiv.org/content/early/2019/06/24/680777.full.pdf>

22 37. Oh, P.L., Benson, A.K., Peterson, D.A., Patil, P.B., Moriyama, E.N., Roos, S., Walter, J.: Diversification of the gut symbiont *Lactobacillus reuteri* as a result of host-driven evolution. *ISME J* **4**(3), 377–387 (2010)

23 38. Pasolli, E., Asnicar, F., Manara, S., Zolfo, M., Karcher, N., Armanini, F., Beghini, F., Manghi, P., Tett, A., Ghensi, P., Collado, M.C., Rice, B.L., DuLong, C., Morgan, X.C., Golden, C.D., Quince, C., Huttenhower, C., Segata, N.: Extensive Unexplored Human Microbiome Diversity Revealed by Over 150,000 Genomes from Metagenomes Spanning Age, Geography, and Lifestyle. *Cell* **176**(3), 649–662 (2019)

24 39. Vallenet, D., Calteau, A., Dubois, M., Amours, P., Bazin, A., Beuvin, M., Burliot, L., Bussell, X., Fouteau, S., Gautreau, G., Lajus, A., Langlois, J., Planel, R., Roche, D., Rollin, J., Rouy, Z., Sabatet, V., Médigue, C.: MicroScope: an integrated platform for the annotation and exploration of microbial gene functions through genomic, pangenomic and metabolic comparative analysis. *Nucleic Acids Research* (2019). doi:[10.1093/nar/gkz926](https://doi.org/10.1093/nar/gkz926). gkz926. <http://oup.prod.sis.lan/nar/advance-article-pdf/doi/10.1093/nar/gkz926/30294889/gkz926.pdf>

25 40. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the em algorithm. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B* **39**(1), 1–38 (1977)

26 41. Biernacki, C., Celeux, G., Govaert, G.: Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(7), 719–725 (2000). doi:[10.1109/34.865189](https://doi.org/10.1109/34.865189)

27 42. Schwarz, G.: Estimating the dimension of a model. *Ann. Statist.* **6**(2), 461–464 (1978). doi:[10.1214/aos/1176344136](https://doi.org/10.1214/aos/1176344136)

28 43. Ambroise, C., Dang, M., Govaert, G.: Clustering of spatial data by the em algorithm. In: Soares, A., Gómez-Hernández, J., Froidevaux, R. (eds.) *geoENV I — Geostatistics for Environmental Applications*, pp. 493–504. Springer, Dordrecht (1997)

29 44. Ambroise, C., Govaert, G.: Convergence of an em-type algorithm for spatial clustering. *Pattern Recognition Letters* **19**(10), 919–927 (1998). doi:[10.1016/S0167-8655\(98\)00076-2](https://doi.org/10.1016/S0167-8655(98)00076-2)

30 45. Dang, M., Govaert, G.: Spatial fuzzy clustering using em and markov random fields. In: *International Journal of System Research and Information Science*, pp. 183–202 (1998)

31 46. Bouguila, N.: On multivariate binary data clustering and feature weighting. *Computational Statistics and Data Analysis* **54**(1), 120–134 (2010). doi:[10.1016/j.csda.2009.07.013](https://doi.org/10.1016/j.csda.2009.07.013)

32 47. Yamamoto, M., Hayashi, K.: Clustering of multivariate binary data with dimension reduction via l1-regularized likelihood maximization. *Pattern Recognition* **48**(12), 3959–3968 (2015). doi:[10.1016/j.patcog.2015.05.026](https://doi.org/10.1016/j.patcog.2015.05.026).

3

4

5 cited By 6

6 48. Śmieja, M., Hajto, K., Tabor, J.: Efficient mixture model for clustering of sparse high dimensional binary data. Data Mining and Knowledge Discovery (2019). doi:[10.1007/s10618-019-00635-1](https://doi.org/10.1007/s10618-019-00635-1)

7 49. Laing, C.R., Whiteside, M.D., Gannon, V.P.J.: Pan-genome Analyses of the Species *Salmonella enterica*, and Identification of Genomic Markers Predictive for Species, Subspecies, and Serovar. *Front Microbiol* **8**, 1345 (2017)

8 50. Miele, V., Penel, S., Duret, L.: Ultra-fast sequence clustering from similarity networks with SiLiX. *BMC Bioinformatics* **12**, 116 (2011)

9 51. Konstantinidis, K.T., Tiedje, J.M.: Genomic insights that advance the species definition for prokaryotes. *Proc. Natl. Acad. Sci. U.S.A.* **102**(7), 2567–2572 (2005)

10 52. Hyatt, D., Chen, G.L., Locascio, P.F., Land, M.L., Larimer, F.W., Hauser, L.J.: Prodigal: prokaryotic gene recognition and translation initiation site identification. *BMC Bioinformatics* **11**, 119 (2010)

11 53. Laslett, D., Canback, B.: ARAGORN, a program to detect tRNA genes and tmRNA genes in nucleotide sequences. *Nucleic Acids Res.* **32**(1), 11–16 (2004)

12 54. Nawrocki, E.P., Eddy, S.R.: Infernal 1.1: 100-fold faster RNA homology searches. *Bioinformatics* **29**(22), 2933–2935 (2013)

13 55. Kalvari, I., Argasinska, J., Quinones-Olvera, N., Nawrocki, E.P., Rivas, E., Eddy, S.R., Bateman, A., Finn, R.D., Petrov, A.I.: Rfam 13.0: shifting to a genome-centric resource for non-coding RNA families. *Nucleic Acids Res.* **46**(D1), 335–342 (2018)

14 56. Steinegger, M., Soeding, J.: Sensitive protein sequence searching for the analysis of massive data sets. *Nat Biotech* (2017). doi:[10.1038/nbt.3988](https://doi.org/10.1038/nbt.3988)

15 57. Grüning, B., Dale, R., Sjödin, A., Chapman, B.A., Rowe, J., Tomkins-Tinch, C.H., Valieris, K.J. Renan, The Bioconda Team: Bioconda: sustainable and comprehensive software distribution for the life sciences. *Nature methods* **15**(7), 475–476 (2018)

16 58. Bastian, M., Heymann, S., Jacomy, M.: Gephi: An open source software for exploring and manipulating networks. (2009). <http://www.aaai.org/ocs/index.php/ICWSM/09/paper/view/154>

17 59. Jacomy, M., Venturini, T., Heymann, S., Bastian, M.: Forceatlas2, a continuous graph layout algorithm for handy network visualization designed for the gephi software. *PLOS ONE* **9**(6), 1–12 (2014). doi:[10.1371/journal.pone.0098679](https://doi.org/10.1371/journal.pone.0098679)

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

56

57

58

59

60

61

62

63

64

65

Figures

Figure 1 Flowchart of PPanGGOLiN on a toy example of 4 genomes. The method requires annotated genomes of the same species with their genes clustered into homologous gene families. Annotations and gene families can be predicted by PPanGGOLiN or directly provided by the user. Based on this input, a pangenome graph is built by merging homologous genes and their genomic links. Nodes represent gene families and edges represent genomic neighborhood. The edges are labeled by the identifiers of genomes sharing the same gene neighborhood. In parallel, gene families are encoded as a presence/absence matrix indicating in each genome whether or not at least one member is present. The pangenome is then divided into K partitions ($K = 3$ in this example) by estimating the best partitioning parameters through an Expectation-Maximization algorithm. The method involves the maximization of the likelihood of a multivariate Bernoulli Mixture Model smoothed by spreading the resulting partitions along the graph using a hidden Markov Random Field that can penalize inadequately classified families according to the graph. This process is repeated until a trade-off maximizing the overall likelihood is reached. This approach returns a partitioned pangenome graph where persistent, shell and cloud classes are overlaid on the neighborhood graph. In addition, many tables, charts and statistics are provided by the software. The number of partitions (K) can either be provided by the user or determined by the algorithm.

Figure 2 Partitioned pangenome graph of 3 117 *Acinetobacter baumannii* genomes. This partitioned pangenome graph of PPanGGOLiN displays the overall genomic diversity of 3 117 *Acinetobacter baumannii* strains from GenBank. Edges correspond to genomic colocalization and nodes correspond to gene families. The thickness of edges is proportional to the number of genomes covering each edge. The size of nodes is proportional to the total number of genes in each gene family. The edges between persistent, shell and cloud nodes are colored in orange, green and blue, respectively. Nodes are colored in the same way. The edges between gene families belonging to different partitions are shown in mixed colors. For visualization purposes, gene families having less than 20 genes are not shown on this figure although they constitute 84.68% of the nodes (mostly unique genes). The frame on the top shows a zoom on a branching region where multiple alternative shell and cloud paths are present in the species. This region is involved in the synthesis of the major polysaccharide antigen of *A. baumannii*. The two most frequent paths (Sv12/PSgc12 and Sv9/PSgc9) are highlighted in khaki and fluo green. The Gephi software (<https://gephi.org>) [58] with the ForceAtlas2 algorithm [59] was used to compute the graph layout with the following parameters: "Scaling=8000, Stronger Gravity=True, Gravity=4.0, Edge Weight influence=1.3".

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Figure 3 Distribution of PPanGGOLiN partitions in the genomes of the most represented species in GenBank. Each horizontal bar shows the median number of gene families per genome among the different PPanGGOLiN partitions (persistent, shell and cloud) in the 88 most represented species in GenBank (having at least 100 genomes). As the partition projection on the genomes of each species provides variable metrics, the error bars represent the interquartile range around the medians of the number of gene families for each partition. The number of gene families of the persistent genome also in the soft core genome ($\geq 95\%$ of presence) is represented by hatched lines on each persistent genome bar. The species names are colored according to their phylum and sorted by taxonomic order and then by decreasing cumulative bar size. Next to the species names, the number of genomes is indicated in brackets and the number of partitions automatically determined by PPanGGOLiN (K) is also shown.

Figure 4 γ -tendencies and IQR areas of the persistent and the soft core rarefaction curves. Each of the 88 most abundant species in GenBank are represented by two points: orange points correspond to the PPanGGOLiN persistent values and yellow points to the ones of the soft core ($\geq 95\%$ of presence). A dashed line connects the 2 points if either the soft core or the persistent values are not in the range of the grey area. The colored horizontal bars show the standard errors of the fitting of rarefaction curves via the Heaps' law.

Additional Files

Additional file 1: Table compiling all the metrics obtained from the pangenomes of the 439 GenBank species
This is a CSV file.

Additional file 2: Supplementary figures
This is a PDF file.

Additional file 3: Table compiling all the metrics obtained from the comparison of PPanGGOLiN results between MAGs and GenBank genomes in 78 species.
This is a CSV file.

Additional file 4: List of GenBank assembly accessions for the 439 studied species.
This is a TSV file where each line correspond to all the GenBank assembly accession used in this study for each 'species id' in the NCBI taxonomy.

1
2
3
4
5
6
7
8
9
10
11

Figure 5 Fraction of cloud and shell families per genome compared to the number of gene families. Results for the 88 most abundant species in GenBank are represented. The error bars show the interquartile ranges of the two variables. The size of the point corresponds to the number of partitions (K) used and the points are colored by phylum.

12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30

Figure 6 Correlation coefficients of shell and genomic distances compared to the fraction of shell per genome. Results for the 88 most abundant species in GenBank are represented. The error bars show the interquartile ranges of the shell fraction. The size of the point corresponds to the number of partitions (K) used and the points are colored by phylum.

31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50

Figure 7 Comparison of the persistent between GenBank genomes and MAGs. Results for 78 species are represented. The colors of the hemispheres provide the percentage of common persistent gene families among the total persistent of MAGs (left hemisphere) or GenBank genomes (right hemisphere). The solid, dashed and dotted lines indicate the identity, a fold change of 1.1 and a fold change of 1.2.

51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

bioRxiv preprint doi: <https://doi.org/10.1101/836239>; this version posted November 9, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

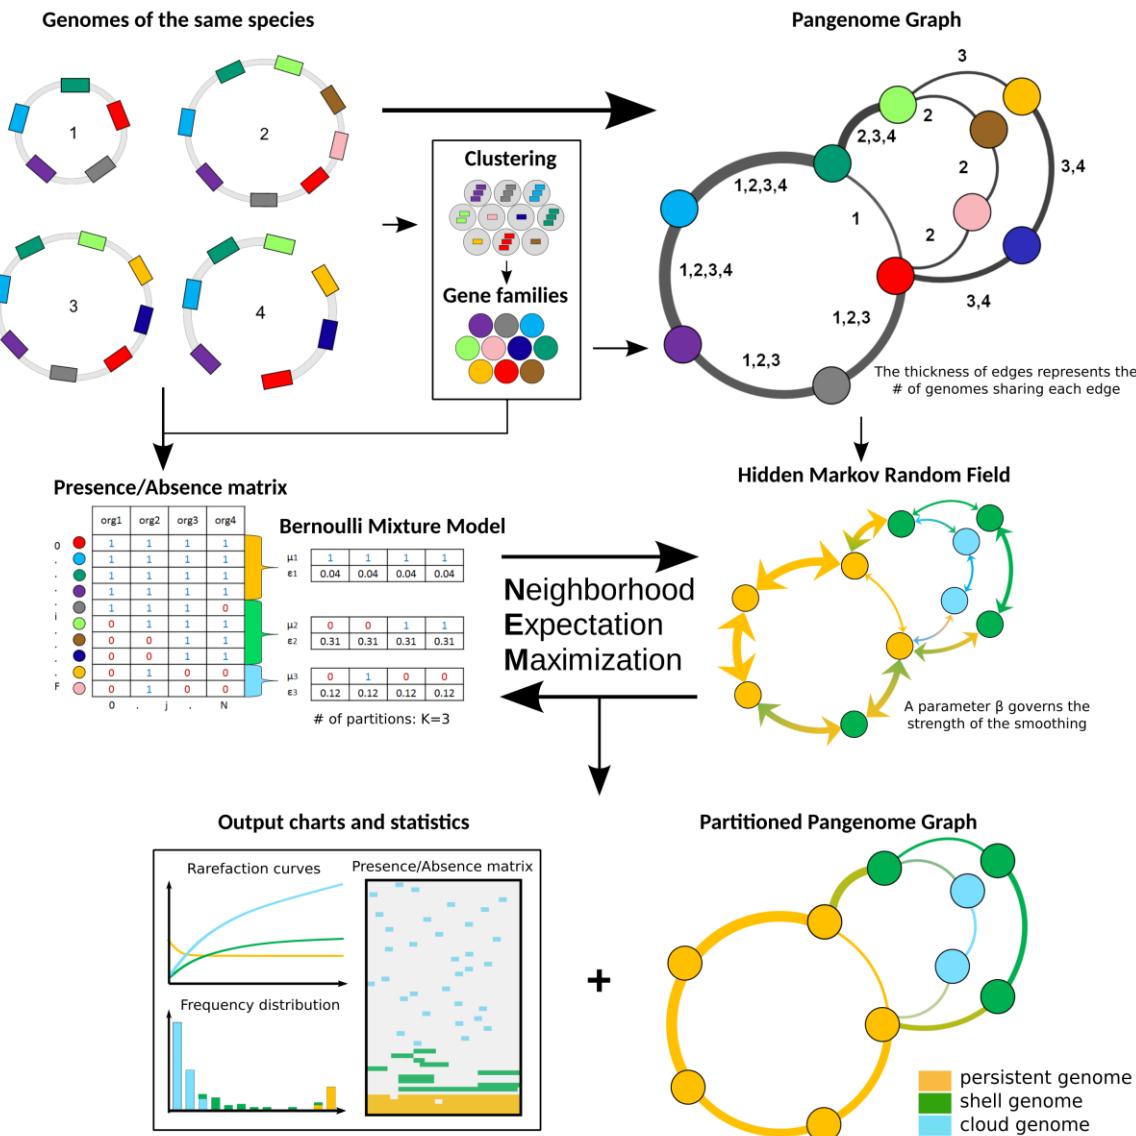
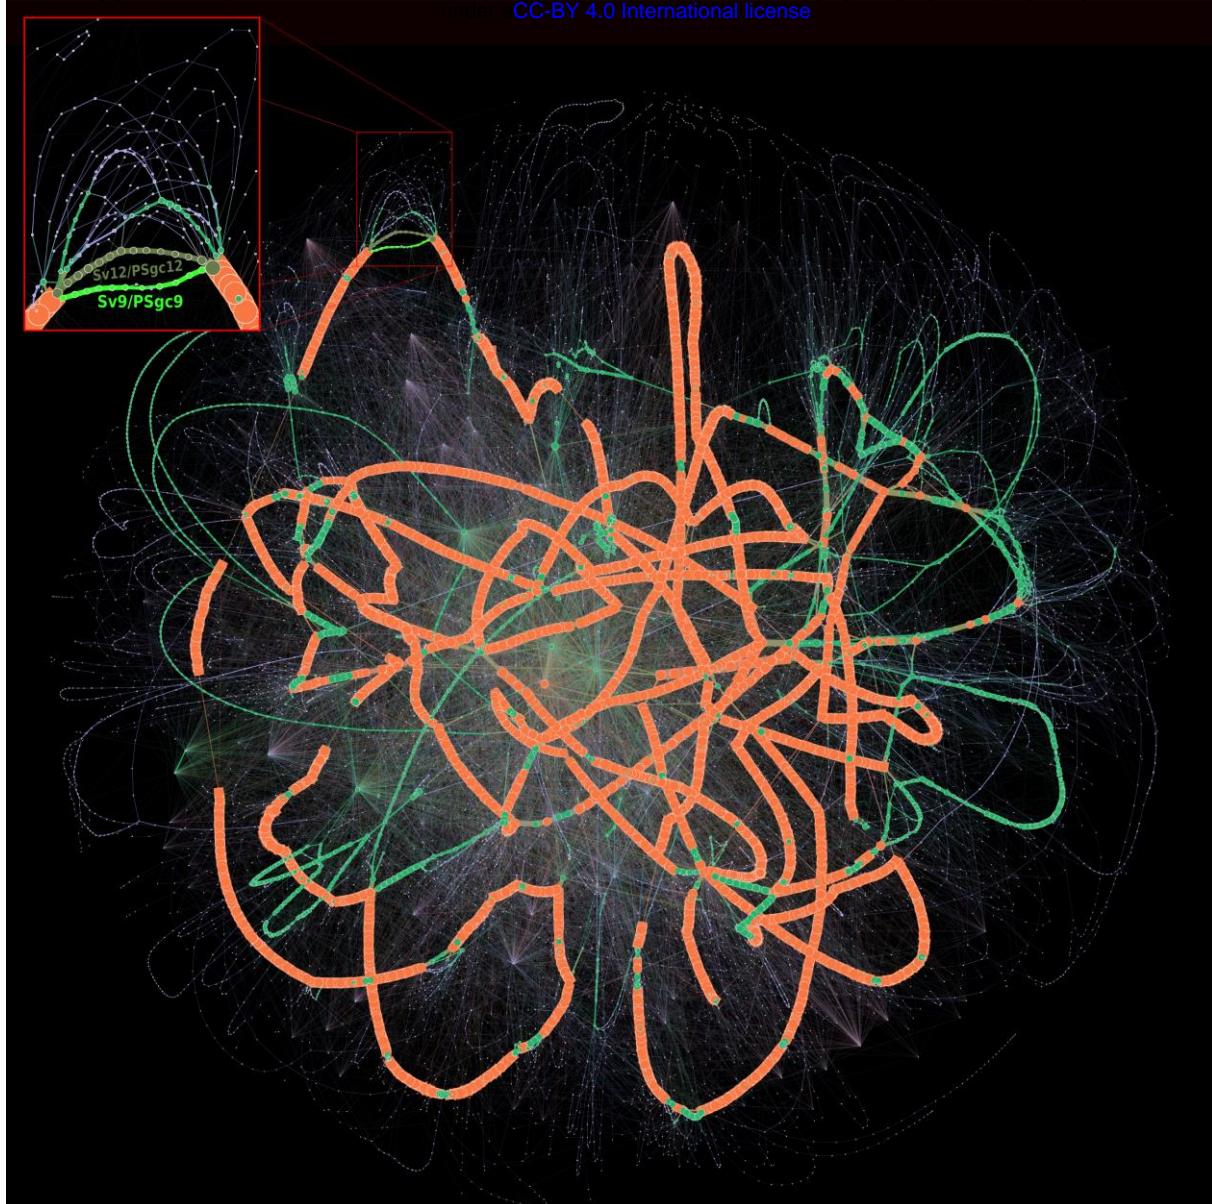


Figure 2

bioRxiv preprint doi: <https://doi.org/10.1101/836239>; this version posted November 9, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available

under a CC-BY 4.0 International license



bioRxiv preprint doi: <https://doi.org/10.1101/836239>; this version posted November 9, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

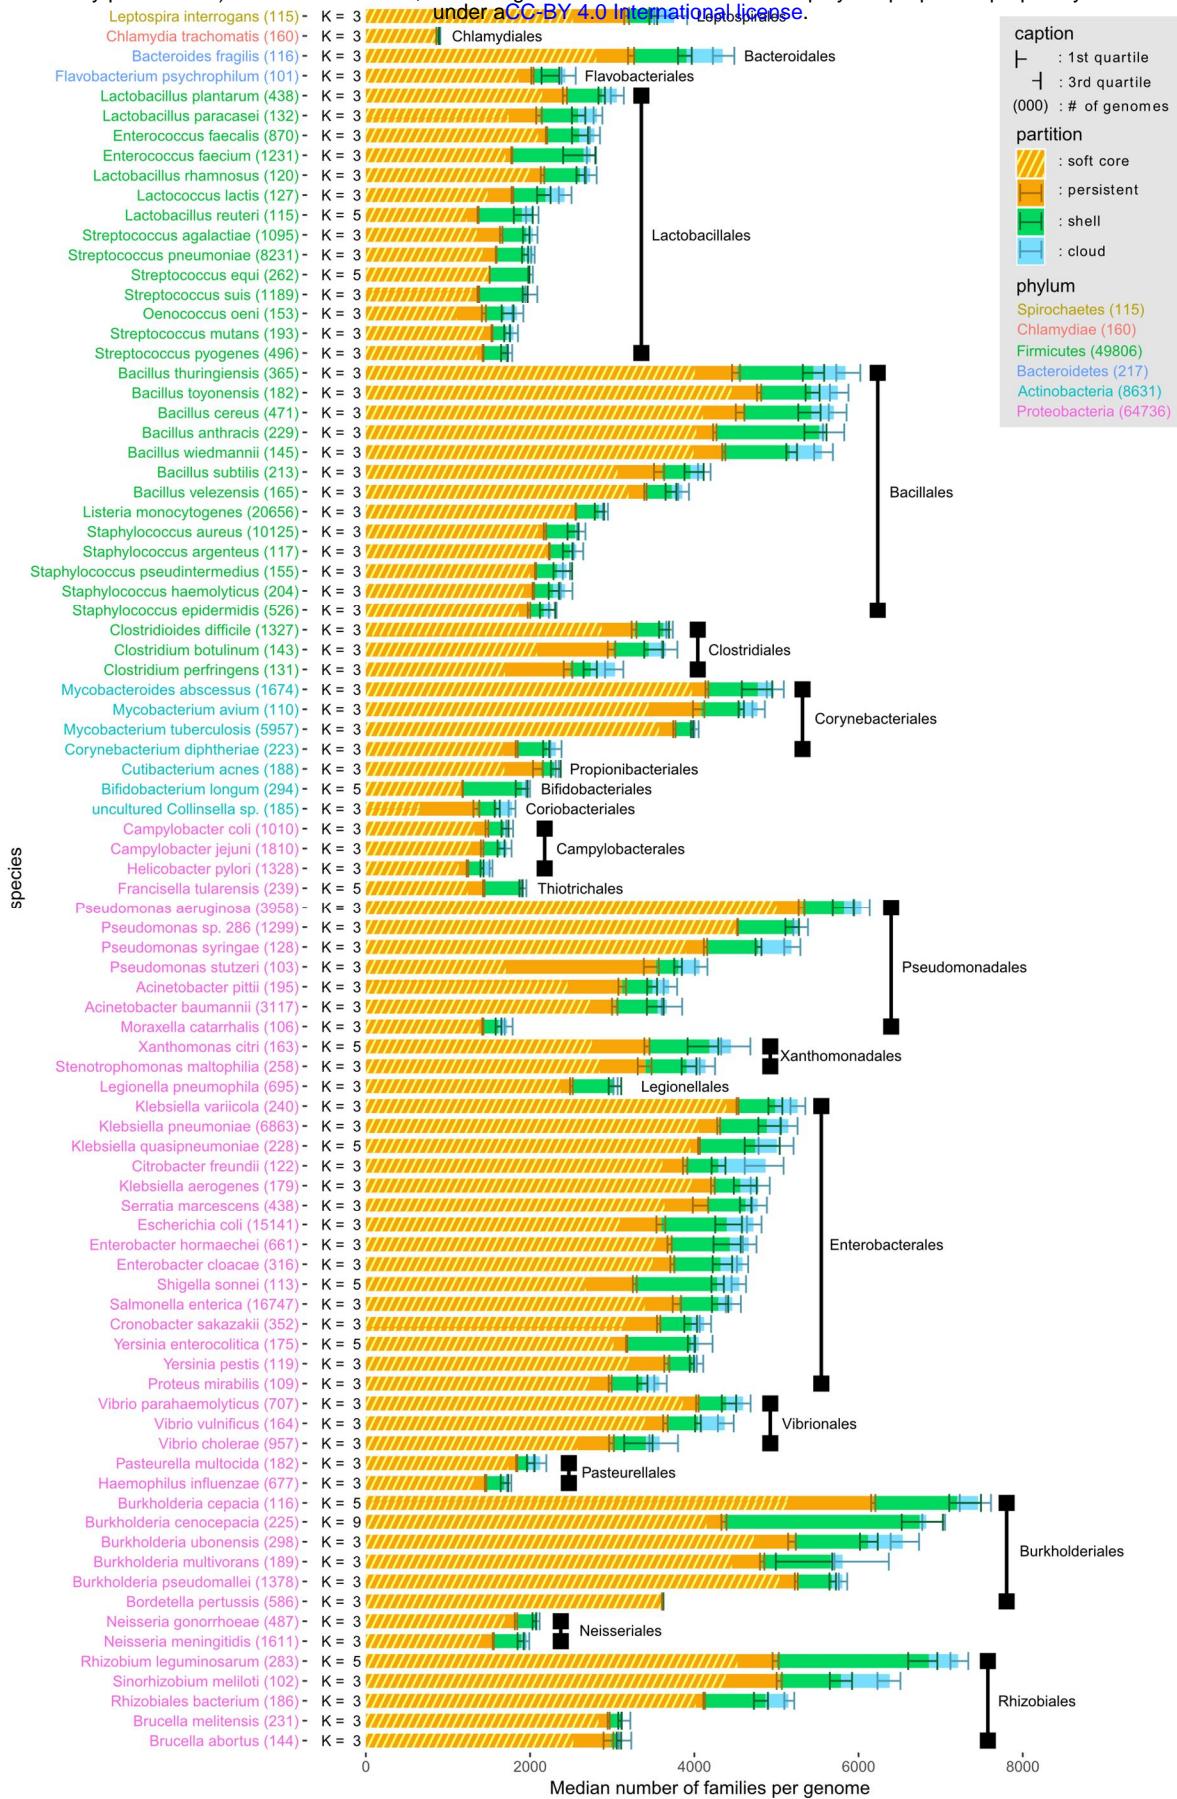


Figure 4

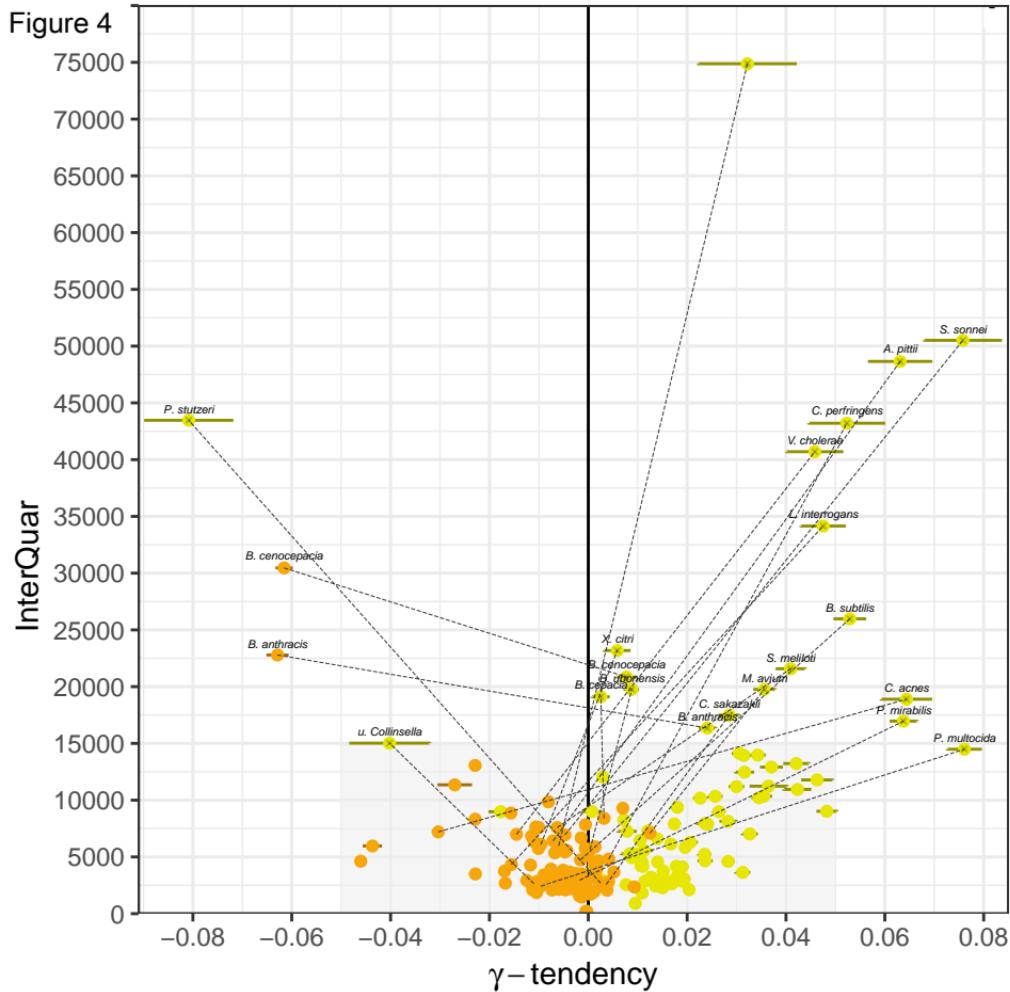


Figure 5

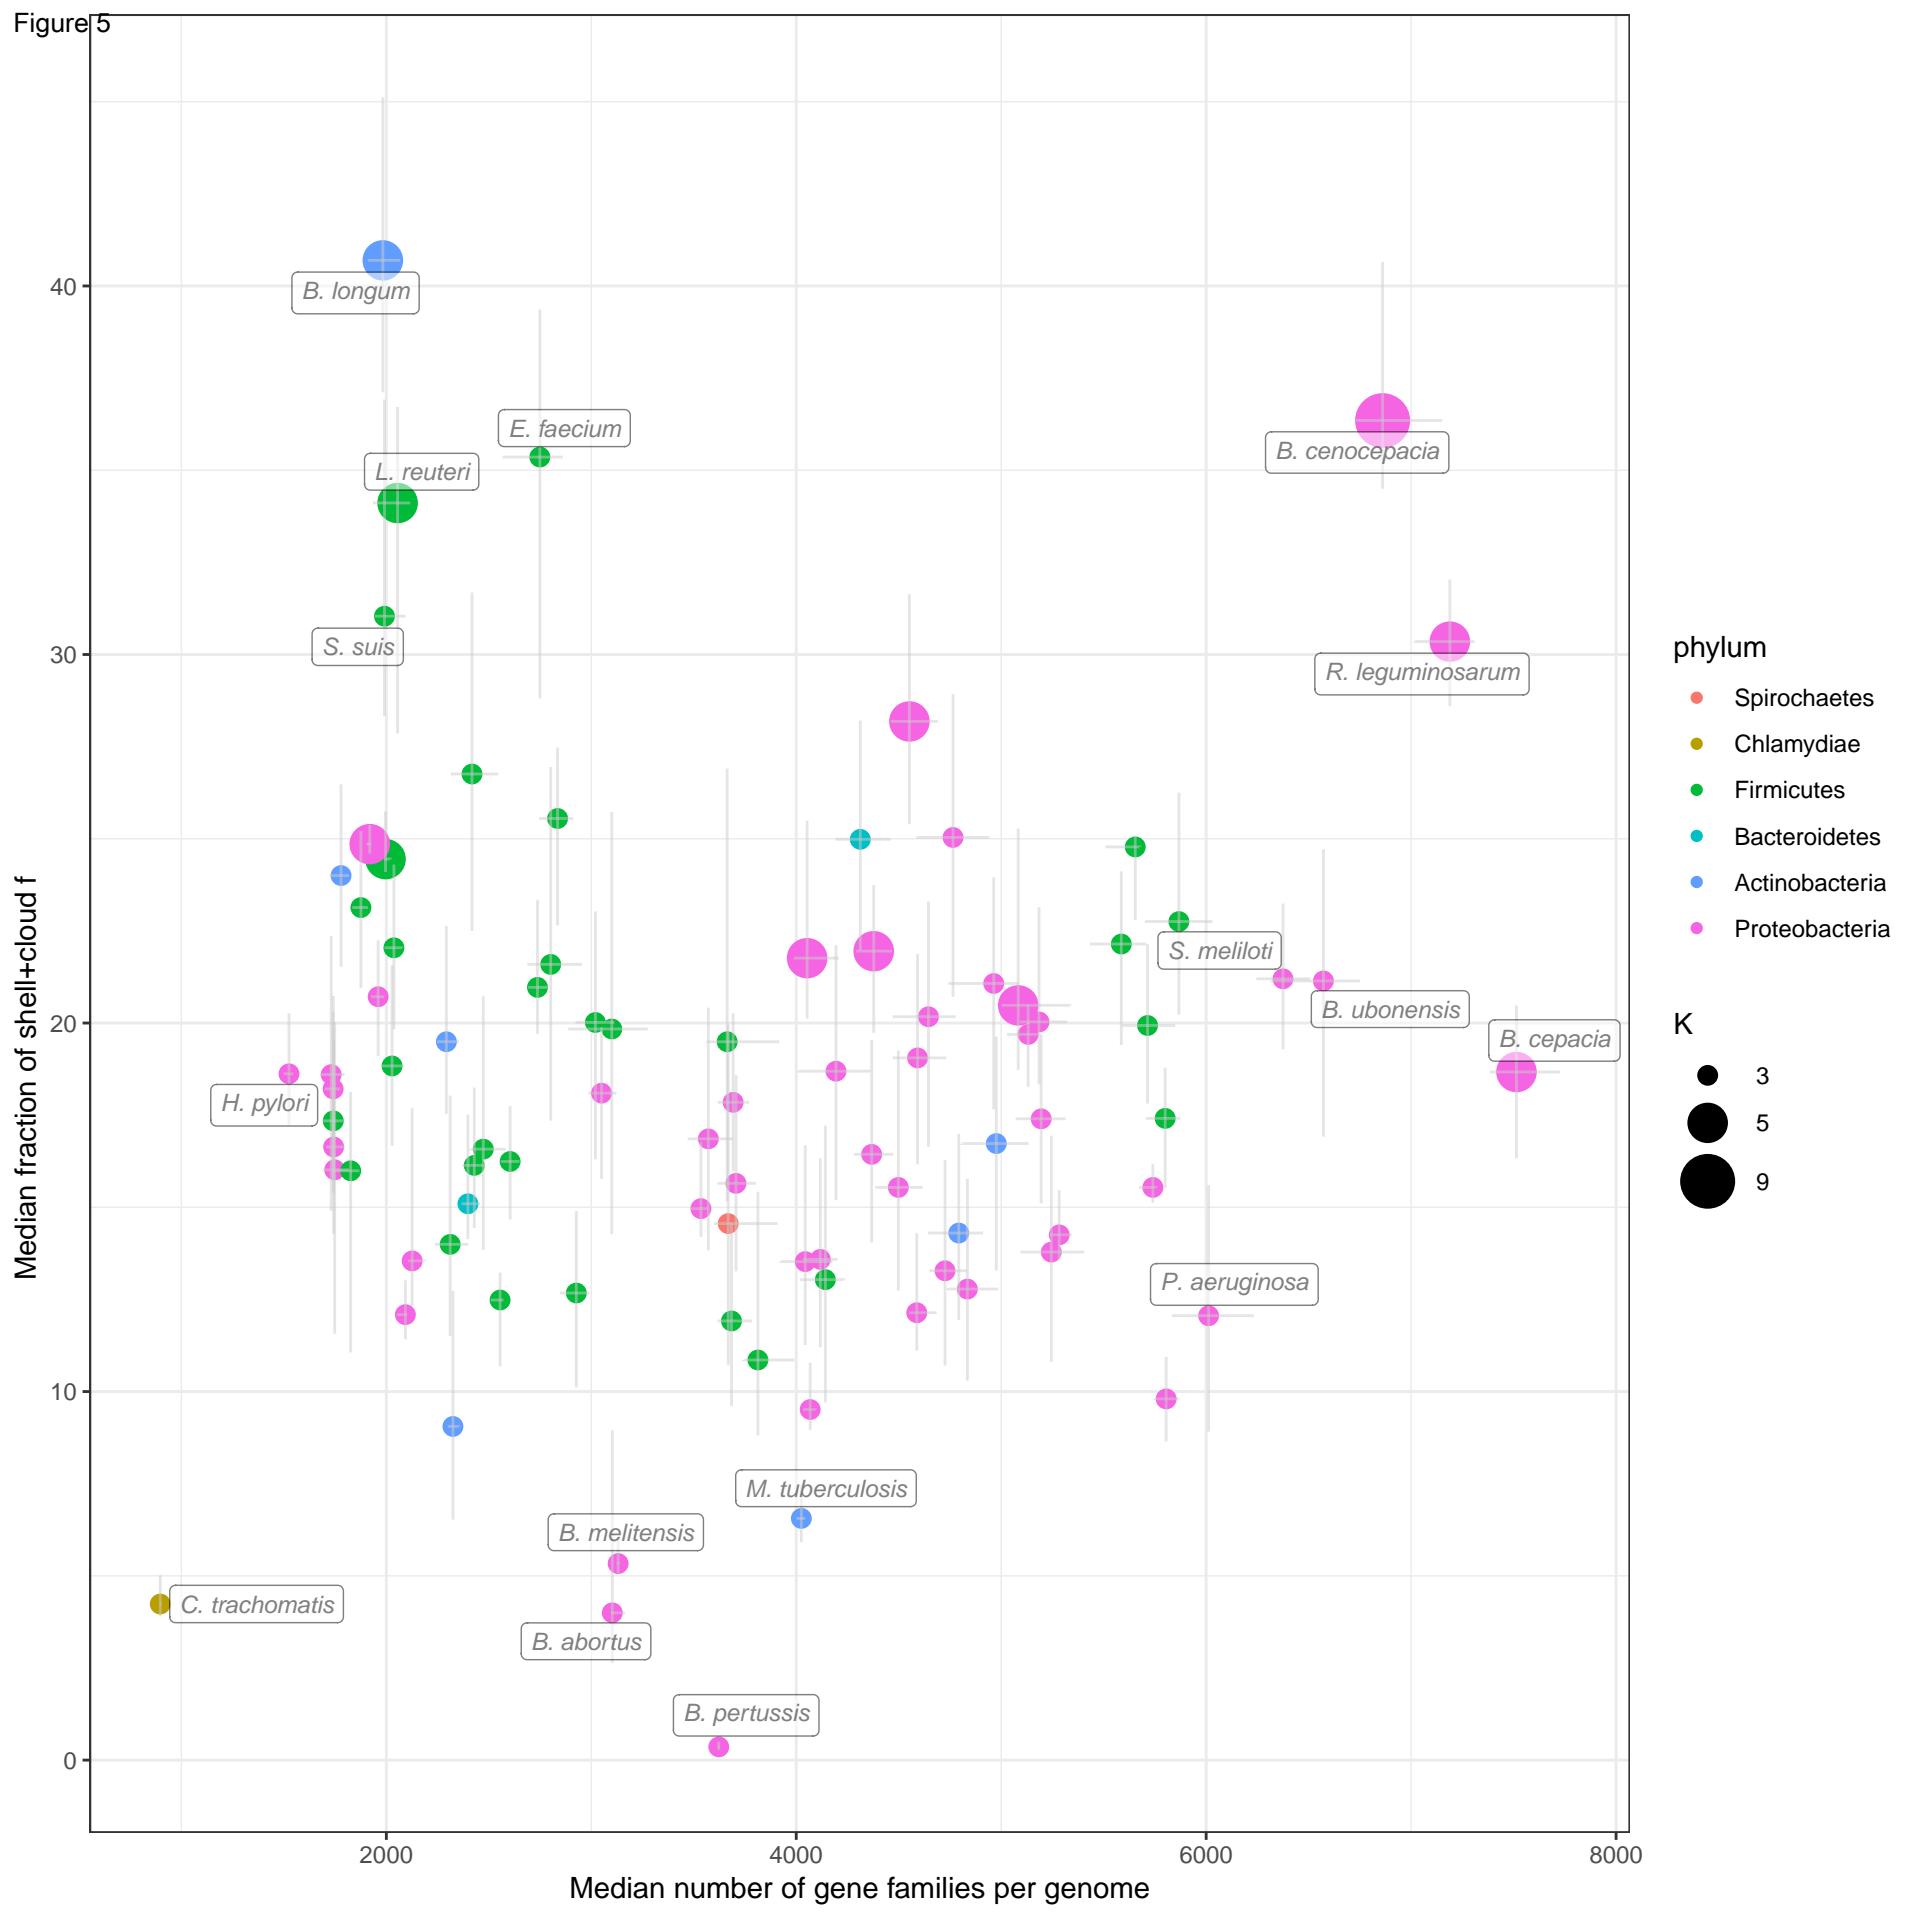


Figure 6

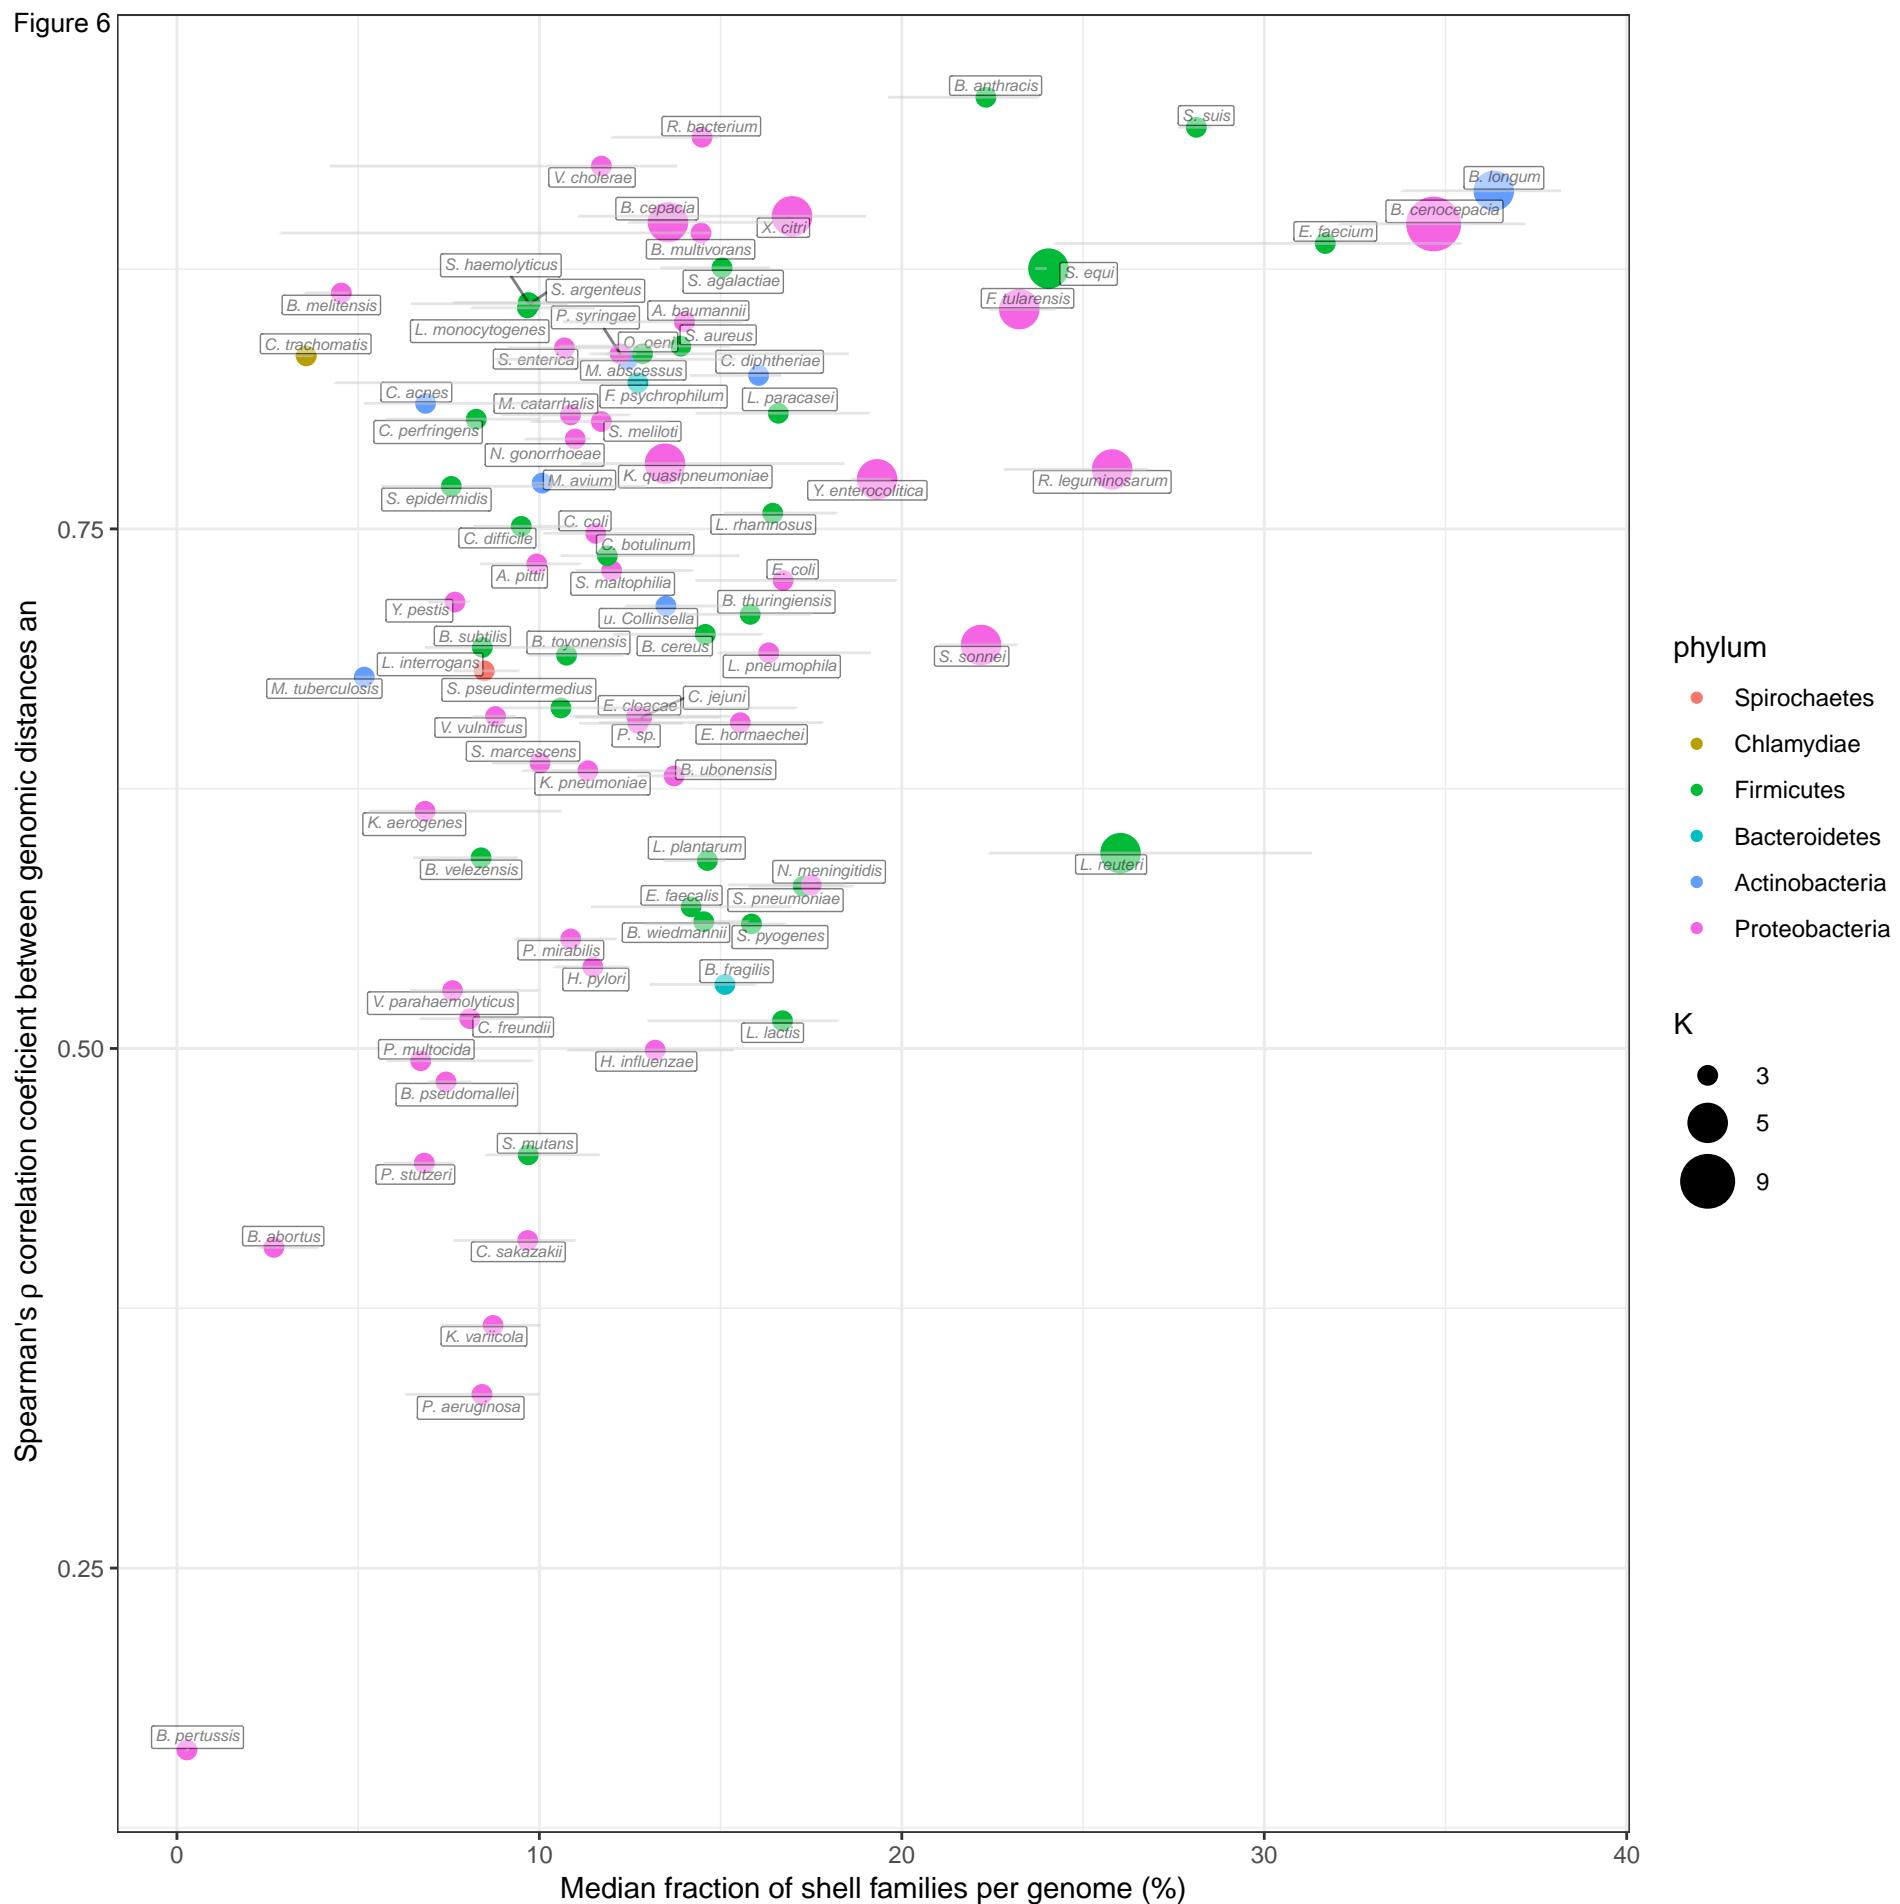


Figure 7

