GeneTonic: an R/Bioconductor package for streamlining the interpretation of RNA-seq data

Federico Marini^{1,2,*}, Annekathrin Ludt¹, Jan Linke^{1,2}, and Konstantin Strauch¹

¹Institute of Medical Biostatistics, Epidemiology and Informatics (IMBEI), University Medical Center of the Johannes Gutenberg University Mainz, Germany

²Center for Thrombosis and Hemostasis (CTH), University Medical Center of the Johannes Gutenberg-University Mainz, Germany.

*To whom correspondence should be addressed (marinif@uni-mainz.de)

May 19, 2021

Abstract

Background: The interpretation of results from transcriptome profiling experiments via RNA sequencing (RNA-seq) can be a complex task, where the essential information is distributed among different tabular and list formats - normalized expression values, results from differential expression analysis, and results from functional enrichment analyses. A number of tools and databases are widely used for the purpose of identification of relevant functional patterns, yet often their contextualization within the data and results at hand is not straightforward, especially if these analytic components are not combined together efficiently.

Results: We developed the GeneTonic software package, which serves as a comprehensive toolkit for streamlining the interpretation of functional enrichment analyses, by fully leveraging the information of expression values in a differential expression context. GeneTonic is implemented in R and Shiny, leveraging packages that enable HTML-based interactive visualizations for executing drilldown tasks seamlessly, viewing the data at a level of increased detail. GeneTonic is integrated with the core classes of existing Bioconductor workflows, and can accept the output of many widely used tools for pathway analysis, making this approach applicable to a wide range of use cases. Users can effectively navigate interlinked components (otherwise available as flat text or spreadsheet tables), bookmark features of interest during the exploration sessions, and obtain at the end a tailored HTML report, thus combining the benefits of both interactivity and reproducibility.

Conclusion: GeneTonic is distributed as an R package in the Bioconductor project (<https://bioconductor.org/packages/GeneTonic/>) under the MIT license. Offering both bird's-eye views of the components of transcriptome data analysis and the detailed inspection of single genes, individual signatures, and their relationships, GeneTonic aims at simplifying the process of interpretation of complex and compelling RNA-seq datasets for many researchers with different expertise profiles.

Keywords: RNA-Seq; Data interpretation; Functional enrichment analysis; Interactive data analysis; Data visualization; Transcriptomics; R; Bioconductor; Shiny; Reproducible research

Background

In modern life and clinical sciences, RNA-sequencing (RNA-seq) is an essential tool for studying gene expression and its regulation (Van den Berge *et al.*, 2019). High-throughput sequencing technologies generate readouts for a large number of molecular entities simultaneously, posing challenges to proper hypothesis generation and data interpretation (Conesa *et al.*, 2016). Among the typical bioinformatic

workflows, differential expression (DE) analysis is often employed to identify the genes showing evidence for statistically significant changes, thus being candidate effectors for regulation across the sampled experimental conditions (Love *et al.*, 2015).

Most studies where these techniques are being adopted result in a list containing tens to thousands of gene candidates, with their associated effect size and significance level - often reported as log₂ fold change (log₂FC) and adjusted p-values, respectively. Putting these results into biological context by leveraging existing knowledge is essential for facilitating the interpretation of data at a systemic level, and enabling novel discoveries (Chen *et al.*, 2016).

Commonly used knowledge bases for the purpose of functional enrichment analysis include Gene Ontology (GO) (Ashburner *et al.*, 2000; Carbon *et al.*, 2019), KEGG (Kanehisa *et al.*, 2017, 2019), REACTOME (Fabregat *et al.*, 2018), and MSigDB (Liberzon *et al.*, 2011, 2015), where the genes are organized either in simple lists (gene sets, or signatures), or as pathways by accounting for the interactions occurring among the respective members; throughout this manuscript, we will use these terms interchangeably. The analysis at the functional level not only aims to reduce the complexity of high dimensional molecular data (grouping thousands of genes and proteins to just several hundreds of coherent entities), but also increasing the explanatory power of the underlying observed mechanisms (Khatri *et al.*, 2012).

A large variety of computational methods and software have been designed for functional enrichment analysis (Xie *et al.*, 2021), and despite their different implementations, they can still be grouped in three main categories, as identified by Khatri *et al.* (2012): 1. Over-Representation Analysis (ORA), contrasting only the set of DE genes against a background of expressed genes; 2. Functional Class Scoring (FCS), including e.g. Gene Set Enrichment Analysis (GSEA, Subramanian *et al.* (2005)) and its different flavors, incorporating a feature-(gene-)level score/statistic, later aggregated at the pathway level to avoid the choice of a binary threshold; 3. Pathway Topology (PT) based approaches, which utilize the additional information of graph/network structure describing the interactions (Nguyen *et al.*, 2018). The most widely adopted methods in this context are ORA and FCS methods, owing to their ease of applicability, fast runtime, and relevance of resulting gene set rankings, as shown in a recent benchmarking effort (Geistlinger *et al.*, 2020).

Visualization techniques are widely used to make sense of enrichment analysis results, where gene sets might also be highly redundant, thus making the prioritization and interpretation of interesting candidates more challenging (Villaveces *et al.*, 2015; Supek and Škunca, 2017). Numerous tools and applications aim to simplify the interpretation step by adopting a diverse range of methods and visual summaries, and these include BiNGO (Maere *et al.*, 2005), ClueGO (Bindea *et al.*, 2009; Mlecnik *et al.*, 2018), GOrilla (Eden *et al.*, 2009), REVIGO (Supek *et al.*, 2011), GOplot (Walter *et al.*, 2015), AgriGO (Tian *et al.*, 2017), NaviGO (Wei *et al.*, 2017), WebGestalt (Liao *et al.*, 2019), CirGO (Kuznetsova *et al.*, 2019), AEGIS (Zhu *et al.*, 2019), FunSet (Hale *et al.*, 2019), hypeR (Federico and Monti, 2020), KeggExp (Liu *et al.*, 2019), Metascape (Zhou *et al.*, 2019), pathfindR (Ulgen *et al.*, 2019), ShinyGO (Ge *et al.*, 2020), ViSEAGO (Brionne *et al.*, 2019), STRING (Szklarczyk *et al.*, 2019), GSOAP (Tokar *et al.*, 2020), GOMCL (Wang *et al.*, 2020), and netGO (Kim *et al.*, 2020).

Aggregating and summarizing the identified categories is an efficient way to capture and distill the main underlying biological aspects, exploiting visual methods that can efficiently encode the essential information of a table. Among the commonly used visualization methods, many apply different ways of grouping and displaying similar genes or gene sets together, including graph-like representations, clustered heatmaps (either genes by samples, or genes by genesets), or wordclouds. Good visualizations enable discovering underlying trends in the data in an unbiased fashion, and are essential components for the proper communication of results in interdisciplinary projects (Supek and Škunca, 2017; Calura and Martini, 2021).

Datasets and gene set collections increase constantly in their size and complexity, constituting a major barrier for the interpretability of transcriptomic data and their enrichment results, to the point

that a potential bottleneck for omics data is the so-called *tertiary* analysis, opposed to mapping and quantification (*primary* analysis) and statistical testing (*secondary* analysis) (Akhmedov *et al.*, 2020). Efficient platforms that enable advanced workflows for a wide range of users can play a big role in providing the required level of interactivity, while guaranteeing the adherence to gold standard methods and to best practices for reproducible analyses (Sandve *et al.*, 2013; Marini and Binder, 2016; Brito *et al.*, 2020).

The different atomic elements for a typical RNA-seq analysis (expression table, results from differential expression, functional enrichment results) can stem from different pipeline outputs, yet they need to be combined together, e.g. in a report created following the rules of literate programming (Knuth, 1984). By providing accessible summaries with proper data visualization and interpretation methods, in formats that facilitate dynamic shareable outputs, such frameworks can greatly reduce the time to generate novel hypotheses and insight. Often, this task is not straightforward to carry out, as different software solutions or environments might be chosen, resulting in different file formats, thus increasing the difficulty for practitioners to explore all relevant aspects of the data at hand, even if common sets of gene and pathway identifiers are adopted.

A number of solutions have been developed in diverse languages (mostly R, Python, Java) to address the challenges listed above, but no software package provides a comprehensive framework for assisting the proper interpretation of RNA-seq data; interested readers can find a comparative overview of the features of the above mentioned tools in Suppl. Table S1.

Here we present GeneTonic, an R/Bioconductor package aiming to streamline the identification of relevant functional patterns, as well as their contextualization in the data and results at hand, by combining in a seamless way all the pieces of information relevant for a transcriptomic analysis. The GeneTonic package is composed by a Shiny web application, with a variety of standalone functions to perform the analysis both interactively as well as in a programmatic way. GeneTonic requires as input the results generated by each analytic step (quantification, DE testing, functional enrichment), which are usually shared as separate tables or spreadsheets by bioinformaticians and core facility service providers, in formats that are suitable to standardization.

GeneTonic makes it easy to generate visualizations, starting from bird's eye perspective summaries (gene-geneset graphs, enrichment maps, also linked to interactive tables in the web application), as well as getting in depth dedicated summaries for each geneset of interest. User actions enable further insight and deliver additional information (e.g. gene info boxes, geneset summaries, and signature heatmaps), with drilldown tasks activated by simple mouse clicks. While simple operations within the call to the `GeneTonic()` main function makes the result set more interpretable, our package also supports built-in RMarkdown reporting as a foundation for computational reproducibility, to conclude an interactive exploration session (Marini and Binder, 2019; Marini *et al.*, 2020). We carefully designed the user interface, enabling the required tasks in a straightforward way, as a result of an open and continuous dialogue with researchers adopting this tool in its early development. Users can learn-by-doing the functionality of GeneTonic via guided tours, creating a common ground for experimentalists and analysts to explore transcriptomic data at the desired depth and efficiently generate novel insights (Poplawski *et al.*, 2016).

GeneTonic connects together a number of R/Bioconductor packages, implementing the current best practices in RNA-seq data analysis, and facilitates the communication between experts of different disciplines. Harmonizing the output of the many analysis steps, possibly performed also with a variety of tools, GeneTonic is a powerful tool for digesting and enjoying any RNA-seq dataset: the interactivity is a compelling means to empower end users for the exploration of many features of interest, and by providing a report with full code snippets, we support analyses that are reproducible and easily extendable. The GeneTonic package is available at <https://bioconductor.org/packages/GeneTonic/>, and a public instance is available for demonstration purposes at <http://shiny.imbei.uni-mainz.de:3838/GeneTonic>.

Implementation

General design of GeneTonic

The GeneTonic package is written in the R programming language, leveraging many existing packages currently available in the Bioconductor project, which constitute the foundation for a broad spectrum of analytic workflows in computational biology and bioinformatics (Huber *et al.*, 2015; Amezquita *et al.*, 2019), and the Shiny framework for interactivity (Chang *et al.*, 2020). The typical use case for GeneTonic expects researchers to run the web application locally, providing the atomic components of a typical RNA-seq analysis workflow (Fig. 1, top section).

GeneTonic is designed to be used after the main steps of DE and functional enrichment analyses have already been completed. While this might seem a limiting factor, we wanted to acknowledge that a plethora of validated methods for performing functional analyses at the pathway level exist, and similarly, well established statistical methods for DE are available (Geistlinger *et al.*, 2020; Van den Berge *et al.*, 2019). Our focus was rather on providing a standardized interface (via so-called shaker functions) to automatically handle the outputs of the different tools which most users might be familiar with, so that GeneTonic retains a wide applicability with respect to the upstream analysis workflows - this is illustrated in the use cases of Suppl. Files 1 and 2.

The required input components are stored as in the DESeq2 workflow (Love *et al.*, 2014), using classes descending from the versatile SummarizedExperiment container, *de facto* the adopted standard for interoperability in the Bioconductor ecosystem (Huber *et al.*, 2015). Tabular information can be provided as simple data frame objects, either imported from textual output of the different tools, or converted internally by the shaker functions of GeneTonic. We encourage users to adopt stable feature identifiers, such as ENSEMBL or Gencode (Yates *et al.*, 2019; Frankish *et al.*, 2019), and enable the automated conversion to HGNC gene symbols via annotation tables.

Most of the implemented functionality can be accessed by a single call to the main `GeneTonic()` function, with the different visualizations and summaries directly available from the dedicated sections of the web application. These functions are also exported for usage in scripted analyses such as RMarkdown HTML reports, making it easy to automate tasks while still creating interactive widgets that can be explored in depth offline (Fig. 1, bottom section).

The user interface (shown in Fig. 2) is structured with the layout provided by the `bs4Dash` package (Granjon, 2019), which implements Bootstrap 4 over the infrastructure of `shinydashboard` (Chang and Borges Ribeiro, 2018). The main features include a sidebar menu (Fig. 2A) to navigate the different sections of the app, a header and a collapsible control panel to provide widgets which define the general behavior of the main panel, displayed in the dashboard body (Fig. 2C).

To instruct users on how to efficiently leverage the exploration components, we enhance the content provided in the use case vignette with guided tours of the interface (Fig. 2B), implemented via the `rintrojs` package (Ganz, 2016). This learning-by-doing paradigm invites the user to perform the actions that reflect typical usages in each module, and can be seen as a dynamic extension of the static documentation format.

Collapsible and tab-based elements allowed us to build a rich user interface, yet without adding too much visual clutter, which would hamper the usability of the analysis sessions - and by that reduce the ability to extract relevant insight.

All the required elements for running GeneTonic are provided at the beginning of the execution, meaning that the navigation throughout the different modules can take place with the usual iteration cycles that build up a full in-depth exploration. As this process can become time-consuming, we implemented a dedicated bookmarking system to temporarily store the genes and gene sets of interest, either by clicking on the dedicated button or with a keystroke (defaulting to the left control key). A summary for these selected features is automatically rendered in the *Bookmarks* section, where the user can generate a full report on the provided input parameters, focusing on the aspects picked up during

Figure 1: Overview of the GeneTonic workflow. Four elements are required to optimally use the functionality of GeneTonic (top section): **dds**, a DESeqDataSet, with the information related to the expression values; **res_de**, a DESeqResults object, storing the results of the differential expression analysis; **res_enrich**, a data frame with the result from the enrichment analysis; **annotation_obj**, a data frame containing different sets of matched identifiers for the features of interest. These can be combined into a GeneTonicList object (middle section), which can be directly fed to GeneTonic and its functions, enabling (bottom section) interactive exploration and a variety of visualizations on the geneset enrichment results, both exploited to generate an integrated HTML report of the provided data.

the live session (Fig. 1, bottom). It is then easy to reconstruct and reproduce the analytic rationale, and share the rendered outputs with cooperation partners, or simply store them for the purpose of careful documentation.

Typical usage workflow

The typical session with GeneTonic can start once the required inputs are provided to the main function (as illustrated in Fig. 1).

In order to use GeneTonic, the following inputs are required: 1. *dds*, a *DESeqDataSet*, the main component in the *DESeq2* framework, storing the information related to the expression matrix; 2. *res_de*, encoded as *DESeqResults* for containing the results of the differential expression analysis; 3. *res_enrich*, i.e. the result from the enrichment analysis, likely converted through one of the shaker functions for preprocessing (or manually, if feeding this from a tool currently not supported), structured as a data frame with a minimal set of required variables (pathway identifier, description, significance level, and affected genes); 4. *annotation_obj*, the gene annotation data frame, i.e. a table with at least two columns, *gene_id* for a set of unambiguous identifiers (e.g. ENSEMBL ids), and *gene_name*, containing a human-readable set of names, e.g. HGNC-based gene symbols.

Conveniently, a single named list containing these inputs (Fig. 1, middle section) can be provided as an alternative format, with many functions of GeneTonic accepting a *gtl* parameter (standing for "GeneTonicList"). This simplifies the creation of context-dependent serialized objects that can be easily shared by data analysts to experimental collaborators.

In its current version, GeneTonic can directly handle the output of different tools, selected for being among the most commonly used in pathway analysis, including *topGO* (Alexa *et al.*, 2006), *clusterProfiler* (Yu *et al.*, 2012), *DAVID* (Huang *et al.*, 2009), *Enrichr* (Kuleshov *et al.*, 2016), *g:Profiler* (Reimand *et al.*, 2019; Raudvere *et al.*, 2019), and *fgsea* (Korotkevich *et al.*, 2021) - these are showcased in the code included in Suppl. Files 1 and 2. We plan to extend the compatibility of GeneTonic with the output of newly developed tools, or alternatively welcome contributions on the project homepage on GitHub (<https://github.com/federicomarini/GeneTonic>).

All the components of the *GeneTonic()* application can be seamlessly used by leveraging sets of shared gene identifiers across the different input objects. This makes it possible to compute aggregate scores for each gene set, e.g. averaging the $\log_2$ fold change of all the affected members of the gene set, or computing a Z-score based on the standardized sum of the number of genes regulated in either direction. As gene sets cannot take into account the topological information of a pathway, this is a valid surrogate means to summarize the effect between conditions at the functional level, and can be visually encoded in the outputs of the GeneTonic dedicated routines.

The process of data exploration and interpretation is iterative by its own nature. GeneTonic supports this by employing a variety of visual summaries (gene-geneset bipartite graphs, enrichment maps, geneset volcano plots, and more in the dedicated app sections - Fig. 2 and Suppl. Fig. 1). We also offer methods to efficiently extract the most meaningful affected biological themes, e.g. by grouping similar categories and selecting a representative pathway for each subset, in order to simplify the redundancy often found in functional enrichment results. Whenever possible, we provide additional information boxes for genes and genesets (Fig. 2D and E) to facilitate drilldown tasks and better understand the whole data components of the project. A number of automatically generated action buttons link directly to external databases, such as AmiGO (Carbon *et al.*, 2019), NCBI (Agarwala *et al.*, 2017), GeneCards (Stelzer *et al.*, 2016), GTEx (Gamazon *et al.*, 2018), enabling more in-depth analysis of particular genesets or genes, without the need to type all the entries of interest.

While the main way of using the functionality of GeneTonic is probably via its web application, we designed all the underlying functions to be able to handle standard objects and classes adopted by the current Bioconductor workflows, and therefore their output can be also incorporated in information-

Figure 2: Screenshot of the Gene-Geneset panel in the GeneTonic application. The main navigation of the app is performed via the sidebar menu (A), while the options common to the different panels are available in the control bar (toggled with the cogs icon, close to the B circle). The main area of the Gene-Geneset panel (C) contains an interactive bipartite graph connecting genesets to their respective components, color coded according to the type and the expression changes. Upon clicking on any node, a Geneset Box (D) or a Gene Box (E) can be automatically generated for facilitating further exploration. The backbone of the graph above (F) represents a compact summary of the bipartite projections, and a table displays the most highly connected nodes (G), as potential hubs with regulatory roles in multiple genesets.

rich HTML reports and existing scripted analyses without additional effort. Indeed, the report itself created via the `happy_hour()` function is an exemplary RMarkdown document, which users can edit and extend as they see fit. Literate programming approaches were initially conceived by the seminal work of Knuth (Knuth, 1984) and have been currently refined in the `knitr` framework (Xie, 2013) and in the Jupyter notebook system (Rule *et al.*, 2019). These techniques constitute a powerful toolkit to ensure the reproducibility of computational analyses (Sandve *et al.*, 2013; Stodden and Miguez, 2014; Brito *et al.*, 2020). The creation of such an HTML document is also intended as the recommended concluding step of a typical usage session for GeneTonic.

In case additional or bespoke visual representations of the input objects (e.g. MA plot for the DE results, customized heatmaps, ...) should be created, the `iSEE` Bioconductor package (Rue-Albrecht *et al.*, 2018) can be used for this purpose. We provide a specific export function, combining the provided inputs into a `SummarizedExperiment` object that can be readily further explored in an instance of `iSEE`, by properly accessing the `assays`, `colData`, and `rowData` slots.

Results

In this section, we will illustrate the functionality of GeneTonic, showcasing the results for the analysis of a human RNA-seq dataset of macrophage immune stimulation, published in Alasoo *et al.* (2018). The data is made available via the `macrophage` Bioconductor package, which contains the files output from the Salmon quantification (version 0.12.0 - Patro *et al.* (2017)), against the Gencode v29 human reference. Expression values, summarized at the gene level, are available from 6 individual donors, in 4 different conditions. We will focus on the comparison between Interferon gamma treated samples vs naive samples.

A comprehensive report on the processing for this dataset and its usage with GeneTonic is included in Suppl. File 1. Suppl. File 2 showcases the usage of GeneTonic on the findings of the work of Mohebiany *et al.* (2020) (A20-deficiency in mouse microglia cells), describing also an alternate entry point for running GeneTonic with objects from the `edgeR` workflow for differential expression (Lun *et al.*, 2016).

Augmenting functional enrichment results with expression data

The majority of functions in GeneTonic requires only a minimal set of information on the pathways enrichment results, i.e. a gene set identifier, its description, and a measure of significance for the enrichment, often specified as a p-value. Nevertheless, it is often beneficial to combine additional knowledge, if provided by the method used to perform the enrichment test; this might include the number of genes annotated to each pathway and the related subset detected as differentially expressed, but more importantly it can incorporate the full set of expression values in the original dataset.

One way to do so is via the `get_aggrscores()` function, which computes the overall pathway Z-score and an aggregated score (such as the mean or the median), which summarize at the gene set level the effect ($\log_2FC$) of the differentially expressed genes annotated as its members. In detail, the gene set Z-score attempts to determine the “direction” of change, and is computed as $z_{gs} = \frac{(DE_{up} - DE_{down})}{\sqrt{(DE_{up} + DE_{down})}}$, where DE_{up} and DE_{down} are the number of up-regulated and down-regulated DE genes annotated to the geneset gs , respectively (Fig. 3A and B).

Alternatively, a sample-level gene set score can be computed, in an approach similar to the implementation of GSVA (Hänzelmann *et al.*, 2013). First, a variance stabilizing transformation is applied to the expression matrix, returning values that show a higher degree of homoscedasticity, thus more amenable to downstream processing and visualization. For each gene, the values are Z-standardized by subtracting the row-wise mean and dividing by the row-wise standard deviation.

Figure 3: Graphical summaries of enrichment results. **A** A geneset volcano plot (`gs_volcano()`), with the $-\log_{10}$ p-value depicted against the geneset Z-score, with a subset of representative affected functions, generated from the output of the `gs_fuzzyclustering()` function. The size of the points maps the information of the geneset size. **B** An enhanced visual summary for the enrichment results, displaying the contributions of the single genes to each gene set (with the directionality as $\log_2$ FC), created with `enhance_table()`. **C** A heatmap for the matrix of sample-wise geneset Z-scores for the same subset used in the other subpanels, generated with `gs_scoresheat()` on the output of `gs_scores()`. **D** A multidimensional scaling plot for genesets, colored by their Z-score, representing the set similarity. This can be generated with the `gs_mdsc()` function.

Finally, for each pathway, we take the subset of the Z values corresponding to its members, and its average is computed and returned as pathway score. We define Z_{ij} as the Z-score for gene i in sample j as $Z_{ij} = \frac{T_{ij} - \bar{T}_i}{sd_i}$, whereas T_{ij} is the correspondent entry in the transformed expression values matrix, $\bar{T}_i$ and sd_i are the mean and standard deviation for the gene i , respectively. The entry GS_{kj} for pathway k in sample j is thus defined as $GS_{kj} = \frac{1}{|P_k|} \sum_{i \in P_k} Z_{ij}$, where P_k is the set of DE genes for pathway k (Fig. 3C).

This extra information about the status of activation/repression for each pathway can be efficiently encoded as aesthetic elements in plots (e.g. the color of a node in a graph, with the geneset Z-score), or directly displayed as a heatmap of the pathway score matrix to compare the activity among the samples.

Exploring the interplay of pathways and genes, interactively

The relationships among pathways and their member genes, or just between different pathways, can quickly become hard to manage when using simple textual or tabular formats. This can be due to the growing size of existing annotations, whereas the increase in detail can also lead to an increase in redundancy, thus making the task of extracting the key biological messages harder.

A number of visualization techniques have been adopted in the last years to simplify this basic yet essential operation (Supek and Škunca, 2017; Calura and Martini, 2021), and a common way to represent this complex interplay is by using graphs. Unipartite graphs are an efficient way to depict the degree of similarity among genesets, where genesets themselves are the nodes, and edges encode for information such as the degree of similarity/overlap between the two nodes (Merico *et al.*, 2010) - see Suppl. Fig. 1. Bipartite graphs (as in Fig. 2) can be naturally adopted to include both genes and genesets as the main node types, with unweighted edges representing in this case the binary membership status for one gene with respect to one geneset (Pomaznoy *et al.*, 2018).

GeneTonic builds upon these foundations and implements the possibility to interact with the nodes upon hovering with the mouse (or clicking on them). The graph objects are generated dynamically, including the desired number of genesets; by default, the top most significant hits in the enrichment results are selected. Interactivity is provided by the `visNetwork` package (Almende B.V. *et al.*, 2019), that wraps the `vis.js` library bindings, building on the `htmlwidgets` framework. Depending on the type of node selected in the main user interface, an information box is populated (Fig. 2D and E).

If selecting a pathway (displayed as a yellow box), the info box will contain details on the geneset (if detected as a Gene Ontology term), and a signature heatmap is displayed, with the variance stabilizing transformed expression data encoded as color to give a fine-grained view of the behavior for all its set members; this is particularly useful to connect the existing biological background with lists of features where no information on the topology is provided, enabling to detect subgroups of correlated expression patterns. Another useful representation can be obtained by coupling a volcano plot (for representing differential expression) with the annotated labels of the members of a geneset; this is implemented in the `signature_volcano()` function, and displayed in the same info box (Suppl. Fig. 1).

Genes are displayed in the graph as ellipses, colored using a divergent palette to encode for the effect size as $\log_2$ fold change; when a gene is selected, a plot for the corresponding expression values is shown, split by experimental variables, and the DE results for the selected gene, together with automatically generated links to external databases opening up in new tabs, to simplify the subsequent exploration steps.

The content available in the *Gene-Geneset* tab is an excellent starting point to get an overview on the provided data. While navigating the interactive graph, it might occur that the user encounters genes or genesets of particular interest; by simply clicking on the Bookmark button in the header section (or alternatively, pressing the left control key) while the node is selected, these elements are stored throughout the session and collected in the *Bookmarks* panel, where one can generate a dedicated

report on these entities.

GeneTonic enables the extraction of a graph backbone, wrapping the efficient implementation of the backbone package (Domagalski *et al.*, 2021) to highlight the salient edges of the bipartite projections for each type of features included, as a way to summarize information contained in large networks (Fig. 2F and G).

Additional insight can be extracted by drilling down the interactive *Enrichment Map* (Merico *et al.*, 2010; Yu *et al.*, 2012), either by focusing on the selected nodes (checking out signature heatmaps or bookmarking the genesets for inserting them into the report), or also by running a variety of community detection algorithms on the graph object returned by the `enrichment_map()` function (Suppl. Fig. 1C). Together with the community membership information, it is then possible to obtain a more compact summary for the functional enrichment results, where the most representative genesets for each subpartition of the graph are selected and returned in tabular format. This network-based approach can be exploited to detect the handful of overarching themes, which might give a more immediate snapshot than the many, often redundant, categories, commonly returned by pathway enrichment algorithms (Suppl. Fig. 1E-F-G).

Summarizing the enrichment results

GeneTonic provides numerous ways to summarize the enrichment results, often leveraging the effectiveness of visual representations to extract insights. The *Overview* and *GSViz* panels serve this purpose, showcasing different views on the dataset at hand, with the main controls provided in the right sidebar.

The geneset volcano plot (Fig. 3A) displays all genesets from the `res_enrich` object and labels the most relevant (or any subset of interest). We use one of the aggregated scores (geneset Z-score, or average $\log_2$ fold change) to determine the horizontal position in the plot. To avoid clutter, it is also possible to reduce the terms based on an overlap threshold, retaining only the most representative ones, and provide this more compact summary to the following visualization routines.

The enhanced table (Fig. 3B) summarizes the top genesets by displaying the $\log_2$ FC of each set's components along a line (one for each set). On top of the static version, this is provided also as an interactive widget, where tooltips activated with the mouse deliver extra information on each dot, representing a single gene.

The complex relationships among genesets and their behavior across samples are just two aspects one can inspect in depth with the implemented methods. Among these, users can generate a genesets-by-sample heatmap, showing the standardized expression values of the members (via the `gs_scoresheat()` function, Fig. 3C), or alternatively a summary heatmap (with `gs_summary_heat()`, Suppl. File 2), which aims to display the redundancy between different sets, while encoding the values of the expression changes. A multi-dimensional scaling (MDS) plot (Fig. 3D) delivers a 2d visualization of the distance among genesets, based on a similarity measure, e.g. their overlap or other criteria, such as their semantic similarity. In a similar fashion, a dendrogram for genesets enables the possibility to use node color, node size, and branch color to encode relevant features, with the tree structure mirroring the distance matrix based on a similarity measure. GeneTonic simplifies the creation of simple summaries for the enrichment, where the essential columns are encoded as graphical parameters of the points, extendable to the case of comparing the same genesets in more than one scenario (e.g. if it is possible to extract more than one contrast from the expression matrix). Switching to polar coordinates, this can be captured in spider plots for one or more `res_enrich` objects (see Suppl. File 2 for more examples of usage).

These visual summaries constitute appealing alternatives to the commonly reported tabular formats, which often fail to provide an overall view for the affected functional landscape.

Wrapping up the session

The *Bookmarks* panel offers the possibility to review and inspect the shortlisted features of interest, where both genes (on the left side of the interface) and genesets (right side) can be exported to text files.

A more comprehensive report, with dynamically generated content based on the user selections, is compiled when starting the `happy_hour()` function. This is made possible by a template RMarkdown document, included in the *GeneTonic* package, which accesses the input elements and the reactive values for the Shiny components. Notably, this functionality can also be used outside an interactive usage session, specifying as parameters the values for the genes and genesets to focus on. In either case, a full HTML document is rendered, whose content mirrors the structure of the info boxes, and can be later shared or stored as a reproducible artifact for the performed analyses.

Another action button creates the serialized version of a *SummarizedExperiment* object, ready to be provided as the main input to *iSEE* (Rue-Albrecht *et al.*, 2018), for further tailored visualizations, either with standard or custom panels of the web application.

Discussion

Interpreting the results of transcriptomic studies can be a complex task, where differential expression analysis is combined with a higher-level pathway enrichment analysis, in order to robustly define the molecular actors that display expression changes, and also to identify the underlying functional patterns. Geneset functional enrichment has been successfully applied to thousands of works, and for this step many methods and approaches have been developed. These tasks are also often shared with alternative workflows other than DE analysis, whereas the aim is to extract meaningful information from large lists of genes, yet it is still a prohibitive task to combine in a straightforward way all the single results from each step. This can be for example due to disjoint sets of identifiers, different output and file formats, and to the difficulties in extracting knowledge while handling large numbers of redundant genesets. Providing concise and biologically meaningful views of the underlying cellular processes, defined via differential expression, is essential in many applications, and a proper visualization framework plays a fundamental role in transforming the otherwise tedious and error/bias-prone task of navigating large textual tables into a more compelling activity (Merico *et al.*, 2010; Supek and Škunca, 2017).

In this work, we introduced *GeneTonic* as a solution to explore all the components of a transcriptome dataset in a more integrative way, instead of having to process them as separated outputs. *GeneTonic* is focused on the analytic step devoted to the interpretation of data, rather than on the implementation of additional methods for detection of functionally enriched biological processes or pathways. Consequently, *GeneTonic* implements a variety of summary and visual representations, while accommodating the output of many commonly adopted enrichment tools, making efficient use of the Shiny framework to deliver interactivity and enable drilldown operations. These would otherwise need to be laboriously addressed in multiple iterations of scripted analyses, either done by the user itself or in collaboration with an external unit, such as a bioinformatics core facility.

Several software packages and web-based portals exist for providing similar functionality, and a comprehensive overview of their salient features is presented in Suppl. Table S1. Naturally, these tools differ in terms of implementation, range of applicability, ease of use, with many proposals offering embedded versions of enrichment tests. Since we developed *GeneTonic* in the R programming language, where many such testing procedures are natively available, we instead focused on the support and integration of their output formats into a common workflow. This can be easily combined with existing analysis pipelines, making our tool well suit for potential wide adoption. The comparison with other tools is also available online (<https://federicomarini>).

github.io/GeneTonic_supplement), linked to a Google Sheet where the individual characteristics of each tool can be updated, in order to provide guidance for users who might be seeking advice on which solution best fits their needs (accessible at <https://docs.google.com/spreadsheets/d/167XV0w18P0FS1d1dt6owN4C2Esx15FU2QTo4D-wclz0/edit?usp=sharing>).

While currently focused on the output of single ORA and FCS enrichment methods, future developments of GeneTonic will implement functionality for combined and ensemble approaches, such as EnrichmentBrowser (Geistlinger *et al.*, 2016) or EGSEA (Alhamdoosh *et al.*, 2016). Moreover, extending such visualizations and interactive summaries to scenarios where multiple omics layers are collected will be a promising avenue for GeneTonic, given the growing number of such datasets becoming available. Finally, we intend to address more refined similarity measurements among genesets, e.g. accounting for information contained in protein-protein interaction networks databases (Yoon *et al.*, 2019), in order to better capture the functional relatedness of the affected pathways.

As bioinformatics evolves constantly into a highly interdisciplinary field, it will become increasingly important to develop common platforms usable by many profiles with substantial differences in their level of programming skills, and GeneTonic's design guidelines adhere to this principle. As a complement to this didactic effect by making the analyses even more open, transparent, and easy to share, GeneTonic could make it easier for bioinformatics skilled users to better understand the systems under investigation, prompting e.g. the development of further tailored methods, which could be a key in obtaining a deeper knowledge of the experimental scenarios.

Conclusion

The identification of relevant functional patterns for the features identified in the differential expression analysis, accounting for the available expression data, remains one of the common bottlenecks for transcriptome-based workflows. GeneTonic provides a web application and many underlying functions to assemble the pieces together, supporting the exploration both interactively as well as in a programmatic way. Combining together the results for quantification, DE testing, and functional enrichment (either generated autonomously, or obtained from collaborators), GeneTonic assists in the unmet yet increasing need of extracting novel knowledge and insights, which can become daunting especially on larger datasets.

GeneTonic has the potential to become an ideal interface between experimental and computational scientists, with the HTML report built via RMarkdown as a milestone for reproducibility, upon conclusion of an interactive session. GeneTonic can be integrated in a wide spectrum of existing bioinformatic pipelines, as it provides functions to convert and input the results of many pathway enrichment tools. This aligns with the principle of interoperability at the heart of the Bioconductor project, which enables a large number of such workflows.

The experience of enjoying transcriptomic data analysis and exploration can be easily shared with reduced communication burden, with both experimental and computational sides empowered in the tasks of realizing complex summaries and visualizations. This will significantly facilitate and democratize the discovery process, bridging the gaps existing between technical and domain expertise.

Additional information

Acknowledgements

This work has been supported by the computing infrastructure provided by the Core Facility Bioinformatics at the University Medical Center Mainz, used also for deploying the demo instance. The authors thank the members of the Core Facility Bioinformatics at the Institute of Molecular Biology

Mainz, Miguel Andrade (IOME, Johannes Gutenberg University of Mainz), Gerrit Toenges and Arsenij Ustjanzew (IMBEI Mainz), and Francesca Finotello (ICBI, Medical University of Innsbruck) for valuable feedback and suggestions.

Funding

The work of FM is supported by the German Federal Ministry of Education and Research (BMBF 01EO1003).

Availability and requirements

Project name: GeneTonic

Project home page: <https://bioconductor.org/packages/GeneTonic/> (release),
<https://github.com/federicomarini/GeneTonic/> (development version)

Archived version: <https://doi.org/10.5281/zenodo.4742518>, package source as gzipped tar archive of the version reported in this article

Project documentation: rendered at <https://federicomarini.github.io/GeneTonic/>

Operating systems: Linux, Mac OS, Windows

Programming language: R

Other requirements: R-4.0.0 or higher, Bioconductor 3.11 or higher

License: MIT

Any restrictions to use by non-academics: none.

Abbreviations

DE: Differential Expression **FCS:** Functional Class Scoring **FDR:** False Discovery Rate **GO:** Gene Ontology **GSEA:** Gene Set Enrichment Analysis **HGNC:** HUGO (Human Genome Organisation) Gene Nomenclature Committee **log₂FC:** base-2 logarithm of the fold change **MA plot:** M (log ratio) versus A (mean average) plot **MDS:** multi-dimensional scaling **MSigDB:** Molecular Signatures Database **NCBI:** National Center for Biotechnology Information **ORA:** Over-Representation Analysis **PT:** Pathway Topology **RNA-seq:** RNA sequencing

Availability of data and materials

The datasets used in this manuscript and its supplements are available from the following articles:

- The data set on the macrophage immune stimulation is included in PubMed ID: 29379200 (<https://doi.org/10.1038/s41588-018-0046-7>). Dataset deposited at the ENA (ERP020977, project id: PRJEB18997) and accessed from the Bioconductor experiment package macrophage package (<https://bioconductor.org/packages/macrophage/>, version 1.7.2)
- The data set on murine A20-deficient microglia is included in PubMed ID: 32023471 (<https://doi.org/10.1016/j.celrep.2019.12.097>). Dataset deposited at the GEO (GSE123033, project id: PRJNA507355) and accessed from the https://github.com/federicomarini/GeneTonic_supplement/ repository

The GeneTonic package can be downloaded from its Bioconductor page <https://bioconductor.org/packages/GeneTonic/> or the GitHub development page <https://github.com/federicomarini/GeneTonic/>. GeneTonic is also provided as a recipe in Bioconda (<https://anaconda.org/bioconda/bioconductor-genetonic>).

The repository available at https://github.com/federicomarini/GeneTonic_supplement/ contains the code used to generate the supplemental material, and the required input data to replicate the analyses presented in the use cases.

Authors' contributions

Conceptualization: Federico Marini, Konstantin Strauch

Data curation: Federico Marini, Annekathrin Ludt, Jan Linke

Formal analysis: Federico Marini, Annekathrin Ludt

Funding acquisition: Federico Marini, Konstantin Strauch

Methodology: Federico Marini, Annekathrin Ludt

Project administration: Federico Marini

Resources: Federico Marini, Konstantin Strauch

Software: Federico Marini, Annekathrin Ludt, Jan Linke

Supervision: Federico Marini, Konstantin Strauch

Visualization: Federico Marini, Annekathrin Ludt

Writing - original draft: Federico Marini, Konstantin Strauch

Writing - review & editing: Federico Marini, Annekathrin Ludt, Jan Linke, Konstantin Strauch

All authors read and approved the final version of the manuscript.

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Bibliography

- Agarwala, R., Barrett, T., Beck, J., Benson, D. A., Bollin, C., Bolton, E., Bourexis, D., Brister, J. R., Bryant, S. H., Canese, K., Charowhas, C., Clark, K., DiCuccio, M., Dondoshansky, I., Feolo, M., Funk, K., Geer, L. Y., Gorenkov, V., Hlavina, W., Hoepfner, M., Holmes, B., Johnson, M., Khotomlianski, V., Kimchi, A., Kimelman, M., Kitts, P., Klimke, W., Krasnov, S., Kuznetsov, A., Landrum, M. J., Landsman, D., Lee, J. M., Lipman, D. J., Lu, Z., Madden, T. L., Madej, T., Marchler-Bauer, A., Karsch-Mizrachi, I., Murphy, T., Orris, R., Ostell, J., O'Sullivan, C., Palanigobu, V., Panchenko, A. R., Phan, L., Pruitt, K. D., Rodarmer, K., Rubinstein, W., Sayers, E. W., Schneider, V., Schoch, C. L., Schuler, G. D., Sherry, S. T., Sirotkin, K., Siyan, K., Slotta, D., Soboleva, A., Soussov, V., Starchenko, G., Tatusova, T. A., Todorov, K., Trawick, B. W., Vakarov, D., Wang, Y., Ward, M., Wilbur, W. J., Yaschenko, E., and Zbicz, K. (2017). Database Resources of the National Center for Biotechnology Information. *Nucleic Acids Research*, **45**(D1), D12–D17.
- Akhmedov, M., Martinelli, A., Geiger, R., and Kwee, I. (2020). Omics Playground: a comprehensive self-service platform for visualization, analytics and exploration of Big Omics Data. *NAR Genomics and Bioinformatics*, **2**(1), 1–10.
- Alasoo, K., Rodrigues, J., Mukhopadhyay, S., Knights, A. J., Mann, A. L., Kundu, K., Hale, C., Dougan, G., and Gaffney, D. J. (2018). Shared genetic effects on chromatin and gene expression indicate a role for enhancer priming in immune response. *Nature Genetics*, **50**(3), 424–431.
- Alexa, A., Rahnenführer, J., and Lengauer, T. (2006). Improved scoring of functional groups from gene expression data by decorrelating GO graph structure. *Bioinformatics*, **22**(13), 1600–1607.
- Alhamdoosh, M., Ng, M., Wilson, N. J., Sheridan, J. M., Huynh, H., Wilson, M. J., and Ritchie, M. E. (2016). Combining multiple tools outperforms individual methods in gene set enrichment analyses. *Bioinformatics*, **33**(September 2016), btw623.
- Almende B.V., Thieurmél, B., and Robert, T. (2019). *visNetwork: Network Visualization using 'vis.js' Library*. R package version 2.0.9.
- Amezquita, R., Carey, V., Carpp, L., Geistlinger, L., Lun, A., Marini, F., Rue-Albrecht, K., Risso, D., Soneson, C., Waldron, L., Pagès, H., Smith, M., Huber, W., Morgan, M., Gottardo, R., and Hicks, S. (2019). Orchestrating Single-Cell Analysis with Bioconductor. *bioRxiv*, page 590562.

- Ashburner, M., Ball, C. A., Blake, J. A., Botstein, D., Butler, H., Cherry, J. M., Davis, A. P., Dolinski, K., Dwight, S. S., Eppig, J. T., Harris, M. A., Hill, D. P., Issel-Tarver, L., Kasarskis, A., Lewis, S., Matese, J. C., Richardson, J. E., Ringwald, M., Rubin, G. M., and Sherlock, G. (2000). Gene ontology: Tool for the unification of biology. *Nature Genetics*, **25**(1), 25–29.
- Bindea, G., Mlecnik, B., Hackl, H., Charoentong, P., Tosolini, M., Kirilovsky, A., Fridman, W. H., Pagès, F., Trajanoski, Z., and Galon, J. (2009). ClueGO: A Cytoscape plug-in to decipher functionally grouped gene ontology and pathway annotation networks. *Bioinformatics*, **25**(8), 1091–1093.
- Brionne, A., Juanchich, A., and Hennequet-Antier, C. (2019). ViSEAGO: A Bioconductor package for clustering biological functions using Gene Ontology and semantic similarity. *BioData Mining*, **12**(1), 1–13.
- Brito, J. J., Li, J., Moore, J. H., Greene, C. S., Nogoy, N. A., Garmire, L. X., and Mangul, S. (2020). Recommendations to enhance rigor and reproducibility in biomedical research. *GigaScience*, **9**(6), 1–6.
- Calura, E. and Martini, P. (2021). Summarizing RNA-Seq Data or Differentially Expressed Genes Using Gene Set, Network, or Pathway Analysis. In E. Picardi, editor, *RNA Bioinformatics*, volume 2284, chapter 9, pages 147–179. Humana, New York, NY.
- Carbon, S., Douglass, E., Dunn, N., Good, B., Harris, N. L., Lewis, S. E., Mungall, C. J., Basu, S., Chisholm, R. L., Dodson, R. J., Hartline, E., Fey, P., Thomas, P. D., Albou, L. P., Ebert, D., Kesling, M. J., Mi, H., Muruganujan, A., Huang, X., Poudel, S., Mushayahama, T., Hu, J. C., LaBonte, S. A., Siegle, D. A., Antonazzo, G., Attrill, H., Brown, N. H., Fexova, S., Garapati, P., Jones, T. E., Marygold, S. J., Millburn, G. H., Rey, A. J., Trovisco, V., Dos Santos, G., Emmert, D. B., Falls, K., Zhou, P., Goodman, J. L., Strelets, V. B., Thurmond, J., Courtot, M., Osumi, D. S., Parkinson, H., Roncaglia, P., Acencio, M. L., Kuiper, M., Lreid, A., Logie, C., Lovering, R. C., Huntley, R. P., Denny, P., Campbell, N. H., Kramarz, B., Acquaah, V., Ahmad, S. H., Chen, H., Rawson, J. H., Chibucos, M. C., Giglio, M., Nadendla, S., Tauber, R., Duesbury, M. J., Del, N. T., Meldal, B. H., Perfetto, L., Porras, P., Orchard, S., Shrivastava, A., Xie, Z., Chang, H. Y., Finn, R. D., Mitchell, A. L., Rawlings, N. D., Richardson, L., Sangrador-Vegas, A., Blake, J. A., Christie, K. R., Dolan, M. E., Drabkin, H. J., Hill, D. P., Ni, L., Sitnikov, D., Harris, M. A., Oliver, S. G., Rutherford, K., Wood, V., Hayles, J., Bahler, J., Lock, A., Bolton, E. R., De Pons, J., Dwinell, M., Hayman, G. T., Laulederkind, S. J., Shimoyama, M., Tutaj, M., Wang, S. J., D'Eustachio, P., Matthews, L., Balhoff, J. P., Aleksander, S. A., Binkley, G., Dunn, B. L., Cherry, J. M., Engel, S. R., Gondwe, F., Karra, K., MacPherson, K. A., Miyasato, S. R., Nash, R. S., Ng, P. C., Sheppard, T. K., Shrivatsav Vp, A., Simison, M., Skrzypek, M. S., Weng, S., Wong, E. D., Feuermann, M., Gaudet, P., Bakker, E., Berardini, T. Z., Reiser, L., Subramaniam, S., Huala, E., Arighi, C., Auchincloss, A., Axelsen, K., Argoud, G. P., Bateman, A., Bely, B., Blatter, M. C., Boutet, E., Breuza, L., Bridge, A., Britto, R., Bye-A-Jee, H., Casals-Casas, C., Coudert, E., Estreicher, A., Famiglietti, L., Garmiri, P., Georgioud, G., Gos, A., Gruaz-Gumowski, N., Hatton-Ellis, E., Hinz, U., Hulo, C., Ignatchenko, A., Jungo, F., Keller, G., Laiho, K., Lemercier, P., Lieberherr, D., Lussi, Y., Mac-Dougall, A., Magrane, M., Martin, M. J., Masson, P., Natale, D. A., Hyka, N. N., Pedruzzi, I., Pichler, K., Poux, S., Rivoire, C., Rodriguez-Lopez, M., Sawford, T., Speretta, E., Shypitsyna, A., Stutz, A., Sundaram, S., Tognolli, M., Tyagi, N., Warner, K., Zaru, R., Wu, C., Chan, J., Cho, J., Gao, S., Grove, C., Harrison, M. C., Howe, K., Lee, R., Mendel, J., Muller, H. M., Raciti, D., Van Auken, K., Berriman, M., Stein, L., Sternberg, P. W., Howe, D., Toro, S., and Westerfield, M. (2019). The Gene Ontology Resource: 20 years and still GOing strong. *Nucleic Acids Research*, **47**(D1), D330–D338.
- Chang, W. and Borges Ribeiro, B. (2018). *shinydashboard: Create Dashboards with 'Shiny'*. R package version 0.7.1.
- Chang, W., Cheng, J., Allaire, J., Xie, Y., and McPherson, J. (2020). *shiny: Web Application Framework for R*. R package version 1.4.0.2.
- Chen, Y., Lun, A. T. L., and Smyth, G. K. (2016). From reads to genes to pathways: differential expression analysis of RNA-Seq experiments using Rsubread and the edgeR quasi-likelihood pipeline. *F1000Research*, **5**, 1438.
- Conesa, A., Madrigal, P., Tarazona, S., Gomez-Cabrero, D., Cervera, A., McPherson, A., Szcześniak, M. W., Gaffney, D. J., Elo, L. L., Zhang, X., and Mortazavi, A. (2016). A survey of best practices for RNA-seq data analysis. *Genome Biology*, **17**(1), 13.
- Domagalski, R., Neal, Z. P., and Sagan, B. (2021). Backbone: An R package for extracting the backbone of bipartite projections. *PLoS ONE*, **16**(1), e0244363.
- Eden, E., Navon, R., Steinfeld, I., Lipson, D., and Yakhini, Z. (2009). GOrilla: a tool for discovery and visualization of enriched GO terms in ranked gene lists. *BMC Bioinformatics*, **10**(1), 48.
- Fabregat, A., Jupp, S., Matthews, L., Sidiropoulos, K., Gillespie, M., Garapati, P., Haw, R., Jassal, B., Korninger, F., May, B., Milacic, M., Roca, C. D., Rothfels, K., Sevilla, C., Shamovsky, V., Shorser, S., Varusai, T., Viteri, G., Weiser, J., Wu, G., Stein, L., Hermjakob, H., and D'Eustachio, P. (2018). The Reactome Pathway Knowledgebase. *Nucleic Acids Research*, **46**(D1), D649–D655.
- Federico, A. and Monti, S. (2020). hyper: an R package for geneset enrichment workflows. *Bioinformatics*, **36**(4), 1307–1308.
- Frankish, A., Diekhans, M., Ferreira, A. M., Johnson, R., Jungreis, I., Loveland, J., Mudge, J. M., Sisu, C., Wright, J., Armstrong, J., Barnes, I., Berry, A., Bignell, A., Carbonell Sala, S., Chrast, J., Cunningham, F., Di Domenico, T., Donaldson, S., Fiddes, I. T., García Girón, C., Gonzalez, J. M., Grego, T., Hardy, M., Hourlier, T., Hunt, T., Izuogu, O. G., Lagarde, J., Martin, F. J., Martínez, L., Mohanan, S., Muir, P., Navarro, F. C., Parker, A., Pei, B., Pozo, F., Ruffier, M., Schmitt, B. M., Stapleton, E., Suner, M. M., Sycheva, I., Uszczynska-Ratajczak, B., Xu, J., Yates, A., Zerbino, D., Zhang, Y., Aken, B., Choudhary, J. S., Gerstein, M., Guigó, R., Hubbard, T. J., Kellis, M., Paten, B., Reymond, A., Tress, M. L., and Flicek, P. (2019). GENCODE reference annotation for the human and mouse genomes. *Nucleic Acids Research*, **47**(D1), D766–D773.
- Gamazon, E. R., Segrè, A. V., van de Bunt, M., Wen, X., Xi, H. S., Hormozdiari, F., Ongen, H., Konkashbaev, A., Derks, E. M., Aguet, F., Quan, J., Nicolae, D. L., Eskin, E., Kellis, M., Getz, G., McCarthy, M. I., Dermizakis, E. T., Cox, N. J., and Ardlie, K. G. (2018). Using an atlas of gene regulation across 44 human tissues to inform complex disease- and trait-associated variation. *Nature Genetics*, **50**(7), 956–967.
- Ganz, C. (2016). rintrojs: A Wrapper for the Intro.js Library. *The Journal of Open Source Software*, **1**(6), 2016.
- Ge, S. X., Jung, D., and Yao, R. (2020). ShinyGO: a graphical gene-set enrichment tool for animals and plants. *Bioinformatics*, **36**(8), 2628–2629.
- Geistlinger, L., Csaba, G., and Zimmer, R. (2016). Bioconductor's EnrichmentBrowser: Seamless navigation through combined results of set- & network-based enrichment analysis. *BMC Bioinformatics*, **17**(1), 45.
- Geistlinger, L., Csaba, G., Santarelli, M., Ramos, M., Schiffer, L., Turaga, N., Law, C., Davis, S., Carey, V., Morgan, M., Zimmer, R., and Waldron, L. (2020). Toward a gold standard for benchmarking gene set enrichment analysis. *Briefings in Bioinformatics*, **00**(August 2019), 1–12.

- Granjon, D. (2019). *bs4Dash: A 'Bootstrap 4' Version of 'shinydashboard'*. <https://rinterface.github.io/bs4Dash/index.html>, <https://github.com/RinterRface/bs4Dash>.
- Hale, M. L., Thapa, I., and Ghersi, D. (2019). FunSet: An open-source software and web server for performing and displaying Gene Ontology enrichment analysis. *BMC Bioinformatics*, **20**(1), 359.
- Hänzelmann, S., Castelo, R., and Guinney, J. (2013). GSEA: Gene set variation analysis for microarray and RNA-Seq data. *BMC Bioinformatics*, **14**.
- Huang, D. W., Sherman, B. T., and Lempicki, R. A. (2009). Systematic and integrative analysis of large gene lists using DAVID bioinformatics resources. *Nature Protocols*, **4**(1), 44–57.
- Huber, W., Carey, V. J., Gentleman, R., Anders, S., Carlson, M., Carvalho, B. S., Bravo, H. C., Davis, S., Gatto, L., Girke, T., Gottardo, R., Hahne, F., Hansen, K. D., Irizarry, R. a., Lawrence, M., Love, M. I., MacDonald, J., Obenchain, V., Oleś, A. K., Pagès, H., Reyes, A., Shannon, P., Smyth, G. K., Tenenbaum, D., Waldron, L., and Morgan, M. (2015). Orchestrating high-throughput genomic analysis with Bioconductor. *Nature Methods*, **12**(2), 115–121.
- Kanehisa, M., Furumichi, M., Tanabe, M., Sato, Y., and Morishima, K. (2017). KEGG: New perspectives on genomes, pathways, diseases and drugs. *Nucleic Acids Research*, **45**(D1), D353–D361.
- Kanehisa, M., Sato, Y., Furumichi, M., Morishima, K., and Tanabe, M. (2019). New approach for understanding genome variations in KEGG. *Nucleic Acids Research*, **47**(D1), D590–D595.
- Khatri, P., Sirota, M., and Butte, A. J. (2012). Ten years of pathway analysis: Current approaches and outstanding challenges. *PLoS Computational Biology*, **8**(2), e1002375.
- Kim, J., Yoon, S., and Nam, D. (2020). netGO: R-Shiny package for network-integrated pathway enrichment analysis. *Bioinformatics*, pages 1–4.
- Knuth, D. E. (1984). Literate Programming. *The Computer Journal*, **27**(2), 97–111.
- Korotkevich, G., Sukhov, V., Budin, N., Shpak, B., Artyomov, M. N., and Sergushichev, A. (2021). Fast gene set enrichment analysis. *bioRxiv*.
- Kuleshov, M. V., Jones, M. R., Rouillard, A. D., Fernandez, N. F., Duan, Q., Wang, Z., Koplev, S., Jenkins, S. L., Jagodnik, K. M., Lachmann, A., McDermott, M. G., Monteiro, C. D., Gundersen, G. W., and Ma'ayan, A. (2016). Enrichr: a comprehensive gene set enrichment analysis web server 2016 update. *Nucleic Acids Research*, **44**(W1), W90–W97.
- Kuznetsova, I., Lugmayr, A., Siira, S. J., Rackham, O., and Filipovska, A. (2019). CirGO: an alternative circular way of visualising gene ontology terms. *BMC Bioinformatics*, **20**(1), 84.
- Liao, Y., Wang, J., Jaehnig, E. J., Shi, Z., and Zhang, B. (2019). WebGestalt 2019: gene set analysis toolkit with revamped UIs and APIs. *Nucleic Acids Research*, **47**(W1), W199–W205.
- Liberzon, A., Subramanian, A., Pinchback, R., Thorvaldsdottir, H., Tamayo, P., and Mesirov, J. P. (2011). Molecular signatures database (MSigDB) 3.0. *Bioinformatics*, **27**(12), 1739–1740.
- Liberzon, A., Birger, C., Thorvaldsdóttir, H., Ghandi, M., Mesirov, J. P., and Tamayo, P. (2015). The Molecular Signatures Database Hallmark Gene Set Collection. *Cell Systems*, **1**(6), 417–425.
- Liu, X., Han, M., Zhao, C., Chang, C., Zhu, Y., Ge, C., Yin, R., Zhan, Y., Li, C., Yu, M., He, F., and Yang, X. (2019). KeggExp: A web server for visual integration of KEGG pathways and expression profile data. *Bioinformatics*, **35**(8), 1430–1432.
- Love, M. I., Huber, W., and Anders, S. (2014). Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biology*, **15**(12), 550.
- Love, M. I., Anders, S., Kim, V., and Huber, W. (2015). RNA-Seq workflow: gene-level exploratory analysis and differential expression. *F1000Research*, **4**, 1070.
- Lun, A. T. L., Chen, Y., and Smyth, G. K. (2016). It's DE-licious: A Recipe for Differential Expression Analyses of RNA-seq Experiments Using Quasi-Likelihood Methods in edgeR. In E. Mathé and S. Davis, editors, *Statistical Genomics*, number April, chapter 19, pages 391–416. Humana Press, New York, NY.
- Maere, S., Heymans, K., and Kuiper, M. (2005). BiNGO: A Cytoscape plugin to assess overrepresentation of Gene Ontology categories in Biological Networks. *Bioinformatics*, **21**(16), 3448–3449.
- Marini, F. and Binder, H. (2016). Development of Applications for Interactive and Reproducible Research: a Case Study. *Genomics and Computational Biology*, **3**(1), 39.
- Marini, F. and Binder, H. (2019). pcaExplorer: an R/Bioconductor package for interacting with RNA-seq principal components. *BMC Bioinformatics*, **20**(1), 331.
- Marini, F., Linke, J., and Binder, H. (2020). ideal: an R/Bioconductor package for interactive differential expression analysis. *BMC Bioinformatics*, **21**(1), 565.
- Merico, D., Isserlin, R., Stueker, O., Emili, A., and Bader, G. D. (2010). Enrichment Map: A Network-Based Method for Gene-Set Enrichment Visualization and Interpretation. *PLoS ONE*, **5**(11), e13984.
- Mlecnik, B., Galon, J., and Bindea, G. (2018). Comprehensive functional analysis of large lists of genes and proteins. *Journal of Proteomics*, **171**, 2–10.
- Mohebiany, A. N., Ramphal, N. S., Karram, K., Di Liberto, G., Novkovic, T., Klein, M., Marini, F., Kreutzfeldt, M., Härtner, F., Lacher, S. M., Bopp, T., Mittmann, T., Merkler, D., and Waisman, A. (2020). Microglial A20 Protects the Brain from CD8 T-Cell-Mediated Immunopathology. *Cell Reports*, **30**(5), 1585–1597.e6.
- Nguyen, T., Mitrea, C., and Draghici, S. (2018). Network-Based Approaches for Pathway Level Analysis. *Current Protocols in Bioinformatics*, **61**(1), 8.25.1–8.25.24.

- Patro, R., Duggal, G., Love, M. I., Irizarry, R. A., and Kingsford, C. (2017). Salmon provides fast and bias-aware quantification of transcript expression. *Nature Methods*.
- Pomaznoy, M., Ha, B., and Peters, B. (2018). GOnet: A tool for interactive Gene Ontology analysis. *BMC Bioinformatics*, **19**(1), 1–8.
- Poplawski, A., Marini, F., Hess, M., Zeller, T., Mazur, J., and Binder, H. (2016). Systematically evaluating interfaces for RNA-seq analysis from a life scientist perspective. *Briefings in Bioinformatics*, **17**(2), 213–223.
- Raudvere, U., Kolberg, L., Kuzmin, I., Arak, T., Adler, P., Peterson, H., and Vilo, J. (2019). g:Profiler: a web server for functional enrichment analysis and conversions of gene lists (2019 update). *Nucleic Acids Research*, **47**(W1), W191–W198.
- Reimand, J., Isserlin, R., Voisin, V., Kucera, M., Tannus-Lopes, C., Rostamianfar, A., Wadi, L., Meyer, M., Wong, J., Xu, C., Merico, D., and Bader, G. D. (2019). Pathway enrichment analysis and visualization of omics data using g:Profiler, GSEA, Cytoscape and EnrichmentMap. *Nature Protocols*, **14**(2), 482–517.
- Rue-Albrecht, K., Marini, F., Soneson, C., and Lun, A. T. (2018). iSEE: Interactive SummarizedExperiment Explorer. *F1000Research*, **7**(0), 741.
- Rule, A., Birmingham, A., Zuniga, C., Altintas, I., Huang, S. C., Knight, R., Moshiri, N., Nguyen, M. H., Rosenthal, S. B., Pérez, F., and Rose, P. W. (2019). Ten simple rules for writing and sharing computational analyses in Jupyter Notebooks. *PLoS Computational Biology*, **15**(7), 1–8.
- Sandve, G. K., Nekrutenko, A., Taylor, J., and Hovig, E. (2013). Ten Simple Rules for Reproducible Computational Research. *PLoS Computational Biology*, **9**(10), e1003285.
- Stelzer, G., Rosen, N., Plaschkes, I., Zimmerman, S., Twik, M., Fishilevich, S., Stein, T. I., Nudel, R., Lieder, I., Mazor, Y., Kaplan, S., Dahary, D., Warshawsky, D., Guan-Golan, Y., Kohn, A., Rappaport, N., Safran, M., and Lancet, D. (2016). The GeneCards Suite: From Gene Data Mining to Disease Genome Sequence Analyses. *Current Protocols in Bioinformatics*, **54**(1), 1.30.1–1.30.33.
- Stodden, V. and Miguez, S. (2014). Best Practices for Computational Science: Software Infrastructure and Environments for Reproducible and Extensible Research. *Journal of Open Research Software*, **2**(1), e21.
- Subramanian, A., Tamayo, P., Mootha, V. K., Mukherjee, S., Ebert, B. L., Gillette, M. A., Paulovich, A., Pomeroy, S. L., Golub, T. R., Lander, E. S., and Mesirov, J. P. (2005). Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences*, **102**(43), 15545–15550.
- Supek, F. and Škunca, N. (2017). Visualizing GO Annotations. In *The Gene Ontology Handbook*, volume 1446, pages 207–220. Humana Press, New York, NY.
- Supek, F., Bošnjak, M., Škunca, N., and Šmuc, T. (2011). REVIGO Summarizes and Visualizes Long Lists of Gene Ontology Terms. *PLoS ONE*, **6**(7), e21800.
- Szklarczyk, D., Gable, A. L., Lyon, D., Junge, A., Wyder, S., Huerta-Cepas, J., Simonovic, M., Doncheva, N. T., Morris, J. H., Bork, P., Jensen, L. J., and von Mering, C. (2019). STRING v11: protein–protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets. *Nucleic Acids Research*, **47**(D1), D607–D613.
- Tian, T., Liu, Y., Yan, H., You, Q., Yi, X., Du, Z., Xu, W., and Su, Z. (2017). AgriGO v2.0: A GO analysis toolkit for the agricultural community, 2017 update. *Nucleic Acids Research*, **45**(W1), W122–W129.
- Tokar, T., Pastrello, C., and Jurisica, I. (2020). GSOAP: A tool for visualisation of gene set over-representation analysis. *Bioinformatics*, pages 3–5.
- Ulgen, E., Ozisik, O., and Sezerman, O. U. (2019). pathfindR: An R Package for Comprehensive Identification of Enriched Pathways in Omics Data Through Active Subnetworks. *Frontiers in Genetics*, **10**(SEP), 1–33.
- Van den Berge, K., Hembach, K. M., Soneson, C., Tiberi, S., Clement, L., Love, M. I., Patro, R., and Robinson, M. D. (2019). RNA Sequencing Data: Hitchhiker’s Guide to Expression Analysis. *Annual Review of Biomedical Data Science*, **2**(1), 139–173.
- Villaveces, J. M., Koti, P., and Habermann, B. H. (2015). Tools for visualization and analysis of molecular networks, pathways, and -omics data. *Advances and Applications in Bioinformatics and Chemistry*, **8**(1), 11–22.
- Walter, W., Sánchez-Cabo, F., and Ricote, M. (2015). GOplot: An R package for visually combining expression data with functional analysis. *Bioinformatics*, **31**(17), 2912–2914.
- Wang, G., Oh, D.-h., and Dassanayake, M. (2020). GOMCL: a toolkit to cluster, evaluate, and extract non-redundant associations of Gene Ontology-based functions. *BMC Bioinformatics*, **21**(1), 139.
- Wei, Q., Khan, I. K., Ding, Z., Yermen, S., and Kihara, D. (2017). NaviGO: Interactive tool for visualization and functional similarity and coherence analysis with gene ontology. *BMC Bioinformatics*, **18**(1), 177.
- Xie, C., Jauhari, S., and Mora, A. (2021). Popularity and performance of bioinformatics software: the case of gene set analysis. *BMC Bioinformatics*, **22**(1), 191.
- Xie, Y. (2013). *Dynamic Documents with R and knitr*. Chapman & Hall/CRC, New York.
- Yates, A. D., Achuthan, P., Akanni, W., Allen, J., Allen, J., Alvarez-Jarreta, J., Amode, M. R., Armean, I. M., Azov, A. G., Bennett, R., Bhai, J., Billis, K., Boddu, S., Marugán, J. C., Cummins, C., Davidson, C., Dodiya, K., Fatima, R., Gall, A., Giron, C. G., Gil, L., Grego, T., Haggerty, L., Haskell, E., Hourlier, T., Izuogu, O. G., Janacek, S. H., Juettemann, T., Kay, M., Lavidas, I., Le, T., Lemos, D., Martinez, J. G., Maurel, T., McDowall, M., McMahon, A., Mohanan, S., Moore, B., Nuhn, M., Oheh, D. N., Parker, A., Parton, A., Patricio, M., Sakthivel, M. P., Abdul Salam, A. I., Schmitt, B. M., Schuilenburg, H., Sheppard, D., Sycheva, M., Szuba, M., Taylor, K., Thormann, A., Threadgold, G., Vullo, A., Walts, B., Winterbottom, A., Zadissa, A., Chakiachvili, M., Flint, B., Frankish, A., Hunt, S. E., Ilesley, G., Kostadima, M., Langridge, N., Loveland, J. E., Martin, F. J., Morales, J., Mudge, J. M., Muffato, M., Perry, E., Ruffier, M., Trevanion, S. J., Cunningham, F., Howe, K. L., Zerbino, D. R., and Flicek, P. (2019). Ensembl 2020. *Nucleic Acids Research*, **48**(D1), D682–D688.
- Yoon, S., Kim, J., Kim, S.-K., Baik, B., Chi, S.-m., Kim, S.-y., and Nam, D. (2019). GScluster: network-weighted gene-set clustering analysis. *BMC Genomics*, **20**(1), 352.

-
- Yu, G., Wang, L.-G., Han, Y., and He, Q.-Y. (2012). clusterProfiler: an R Package for Comparing Biological Themes Among Gene Clusters. *OMICS: A Journal of Integrative Biology*, **16**(5), 284–287.
- Zhou, Y., Zhou, B., Pache, L., Chang, M., Khodabakhshi, A. H., Tanaseichuk, O., Benner, C., and Chanda, S. K. (2019). Metascape provides a biologist-oriented resource for the analysis of systems-level datasets. *Nature Communications*, **10**(1), 1523.
- Zhu, J., Zhao, Q., Katsevich, E., and Sabatti, C. (2019). Exploratory Gene Ontology Analysis with Interactive Visualization. *Scientific Reports*, **9**(1), 1–9.

Supplemental Figure 1: **Screenshot of the *Enrichment Map* panel in the GeneTonic application.** The sidebar menu (A) controls the main navigation in the app, and a common set of options is toggled with the cogs icon (B). The main area of the *Enrichment Map* panel (C) contains an interactive graph for the enrichment map of the genesets, connected according to their similarity, and color coded according to the specified geneset property (here, the Z-score). Upon clicking on any geneset, a Geneset Box (D) is displayed for further exploration (e.g. to show a volcano plot with the geneset members labelled). The geneset distillery (E) enables the exploration of meta-genesets, derived by computing clusters on the graph object underlying the enrichment map. From the tabular representation, it is possible to visualize meta-genesets as heatmaps (F), or display a modal popup containing the enrichment map where the cluster assignments of the genesets are shown (G).