

GobyWeb: Simplified Management and Analysis of Gene Expression and DNA Methylation Sequencing Data

Kevin C. Dorff¹, Nyasha Chambwe^{2,3}, Zachary Zeno¹, Manuele Simi¹, Rita Shaknovich⁴, Fabien Campagne^{1,2*}

1 The HRH Prince Alwaleed Bin Talal Bin Abdulaziz Alsaud Institute for Computational Biomedicine, The Weill Cornell Medical College, New York, New York, United States of America, **2** Department of Physiology and Biophysics, The Weill Cornell Medical College, New York, New York, United States of America, **3** Tri-Institutional Training Program in Computational Biology and Medicine, The Weill Cornell Medical College, New York, New York, United States of America, **4** Department of Pathology and Department of Medicine; The Weill Cornell Medical College, New York, New York, United States of America

Abstract

We present GobyWeb, a web-based system that facilitates the management and analysis of high-throughput sequencing (HTS) projects. The software provides integrated support for a broad set of HTS analyses and offers a simple plugin extension mechanism. Analyses currently supported include quantification of gene expression for messenger and small RNA sequencing, estimation of DNA methylation (i.e., reduced bisulfite sequencing and whole genome methyl-seq), or the detection of pathogens in sequenced data. In contrast to previous analysis pipelines developed for analysis of HTS data, GobyWeb requires significantly less storage space, runs analyses efficiently on a parallel grid, scales gracefully to process tens or hundreds of multi-gigabyte samples, yet can be used effectively by researchers who are comfortable using a web browser. We conducted performance evaluations of the software and found it to either outperform or have similar performance to analysis programs developed for specialized analyses of HTS data. We found that most biologists who took a one-hour GobyWeb training session were readily able to analyze RNA-Seq data with state of the art analysis tools. GobyWeb can be obtained at <http://gobyweb.campagnelab.org> and is freely available for non-commercial use. GobyWeb plugins are distributed in source code and licensed under the open source LGPL3 license to facilitate code inspection, reuse and independent extensions <http://github.com/CampagneLaboratory/gobyweb2-plugins>.

Citation: Dorff KC, Chambwe N, Zeno Z, Simi M, Shaknovich R, et al. (2013) GobyWeb: Simplified Management and Analysis of Gene Expression and DNA Methylation Sequencing Data. PLoS ONE 8(7): e69666. doi:10.1371/journal.pone.0069666

Editor: Nicholas James Provart, University of Toronto, Canada

Received: March 5, 2013; **Accepted:** June 11, 2013; **Published:** July 23, 2013

Copyright: © 2013 Dorff et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: This work was supported by grant UL1 RR024996 (NIH/NCRR) of the Clinical and Translation Science Center at Weill Cornell Medical College, by grant LLS 6304-11 from the Lymphoma and Leukemia Society Translational Research Program, and by R01 MH086883 (NIH/NIMH). Support was also provided by the Tri-Institutional Training Program in Computational Biology and Medicine. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: fac2003@campagnelab.org.

Introduction

High-Throughput sequencing (HTS) instruments have been used to develop a variety of cost-effective assays. Each of these assays leverage the ability of second generation sequencing instruments to output millions of short sequence reads in a few days. It is now not uncommon to generate about three billion 100 base pair long sequence reads per week with one HiSeq 2000 instrument (many core facilities have several similar instruments). Such throughput makes it possible to multiplex assays, which has contributed to reducing the cost of assaying each single sample. Reductions in sequencing costs are making it possible for research groups to produce datasets with tens to hundreds of biological or clinical samples.

With increasing sequencing throughput, the management and analysis of large datasets produced with HTS assays have become a significant challenge for most research groups. Indeed, HTS data analysis is now recognized as a bottleneck of most research studies.

While many programs have been developed to process HTS data on the command line, only a few integrated systems have been developed that can help investigators process large amounts of data with a simple user interface. Existing systems with a user

interface are often restricted to analysis of a single type of data (e.g., see [1,2]), which forces users to work with different tools to analyze gene expression data or DNA methylation data, for instance. Systems that provide both a user interface and support multiple types of data have been offered commercially, but these systems often operate as black boxes and cannot be inspected in detail or extended.

To address these problems, we developed GobyWeb as a web application that can help users with no programming or command line experience analyze HTS datasets efficiently. GobyWeb takes advantage of compute grids to parallelize applications and dramatically accelerate computations for large datasets. This new tool provides intuitive and consistent analysis workflows that make it possible to track data and results for large projects. This report describes the user interface we have designed for GobyWeb, the types of analyses currently supported by the software, and the computational requirements for local installation. We present examples of analyses that can be conducted with the system. A plugin mechanism is used to implement all types of analysis and makes it possible to customize or extend an installed instance of GobyWeb for future or custom analysis needs. Importantly, creating new plugins requires shell-scripting experience, but does

not necessitate a strong parallel computing experience. In contrast to commercial systems, GobyWeb plugins are distributed in source code, in order to promote code inspection, reuse, modifications or extensions. We compare GobyWeb to several analysis software and systems previously described in the peer-review literature and demonstrate substantial advantages in storage requirement, computational performance and ease of use.

Results

Software Overview

We designed GobyWeb with the following main goals:

- Provide an intuitive user interface that biologists with limited bioinformatics experience can use effectively to analyze their datasets.
- Offer direct download of intermediary and final analysis results in well-defined formats to allow bioinformaticians to perform visualization or custom analyses.
- Support validated analyses for gene expression and DNA methylation.
- Provide mechanisms to track data. The system offers tags for each data element that can be recorded and used at a later time to retrieve data from the web interface. For instance, GobyWeb tags are listed in this manuscript following a description of an analysis and can be used to locate analyses in the GobyWeb demonstration system (<http://gobyweb-demo.apps.campagnelab.org/>).
- Facilitate information sharing among members of a team of investigators.
- Offer efficient analyses that can process large datasets on small compute grids in a few hours.
- Support common types of analyses with scalable analysis workflows.
- Support future extensions via the definition of plugins for new alignment or analysis methods.
- Minimize the cost of data storage. GobyWeb takes advantage of the Goby file formats [3], that can store alignment data in about one tenth the size required with the BAM format [4].

User Interface

The GobyWeb user interface consists of the menus shown in Figure 1. The application menus are organized along three categories: Browse, Actions and Account. Figure 2 presents an overview of the data flows through the system and summarizes the types of interactions end-users have with the deployed system. We describe the options offered by these menus and data flows in File S1 description and Figures S1, S2, S3, as well as in video tutorials online (see <http://gobyweb.campagnelab.org>).

Importantly, this user interface is accessible from most modern web browsers. Upload of large data files must be done over a fast network connection at the start of an analysis. However, after uploads are completed, the system design makes it possible to perform analysis over connections typical of residential or mobile Internet access.

User Experience

Earlier versions of GobyWeb have been deployed and made available to a local user base for the last two years at our institution. We found that a one-hour training session was sufficient for most users to start using GobyWeb to analyze their own RNA-Seq datasets. A second one-hour session covered

advanced analysis for projects involving DNA methylation or DNA-Seq and visualization. Material described in these training sessions is also available online as videos on the project web site. Our experience with this user base suggests that new users who take the introductory one-hour training, and subsequently register and upload data to the system, are able to obtain differential expression results for an RNA-Seq project within one week, with no or minimal interactions with the system's administrators. Most users have no command line or R experience, yet are able to routinely use STAR, GSNAP, DESeq or EdgeR with GobyWeb to analyze their RNA-Seq data [5–8].

System Architecture

Figure 3 presents a high-level overview of the architecture of the GobyWeb software. Briefly, GobyWeb consists of a web front-end, a persistent data store, and a compute grid. The system uses production quality infrastructure components that are widely available in many academic institutions. The system can be installed within an institutional firewall, on an intranet or as an Internet facing application. Each user is required to obtain registration credentials with the system. User authentication enables rich data management capabilities and makes it possible to audit data and CPU usage on a per-user basis.

Compute Grid

We configured a GobyWeb instance with a compute grid of three nodes. Each node contains 4 Intel Xeon X5660 processors at 2.80 GH and offers 24 effective threads and 48 GB of memory (a configuration that was purchased for less than \$9,000 per node three years ago, assuming a three year equipment life, the cost of each node is 34.22 cents per hour, excluding cost of electricity). This compute grid is referred to in the following as the *small benchmark grid* and can run a total of 72 effective threads in parallel. It is difficult to ascertain the exact speed of the nodes used by published benchmarks performed in a cloud environment. For instance, in the Myrna benchmark evaluation [1], EC2 Extra Large High CPU Instances were used, but these 'instances' are virtual images and have been deployed on different hardware over the years. We are unable to provide a direct performance comparison with these tools because many of the cloud-based software cannot also be deployed on local servers. Nevertheless, for illustration purposes, we will compare the performance of the small benchmark grid to the 80 cores cluster mentioned in the Myrna benchmark (small cluster) [1].

RNA-Seq Data Analysis

As an illustration of the capabilities of the GobyWeb software, we uploaded 72 human mRNA-Seq samples, previously used to benchmark Myrna [1]. Reads were trimmed to 35 bp and concatenated across the Yale and Argonne sites to closely replicate the benchmark conditions described by Langmead et al [1]. This dataset consists of approximately 1.1 billion 35 bp reads. Upload of these reads to GobyWeb took 30 minutes, compared to 1 h 15 minutes as reported previously with Myrna [1]. We aligned reads to the genome with GobyWeb, using the BWA and GSNAP alignment plugins. The Myrna benchmark used the Bowtie aligner, which is often faster than BWA. Despite this difference, alignments with BWA and GobyWeb completed in 2 h 22 m, when Myrna reported 2 h 56 m (10 worker configuration). Detailed benchmark times are presented in Table 1, with data for Myrna obtained from [1]. Since we have performed our benchmark on different hardware than used in the Myrna benchmark, Table 1 is not an exact comparison but suggests that GobyWeb is competitive when aligning reads on small clusters.

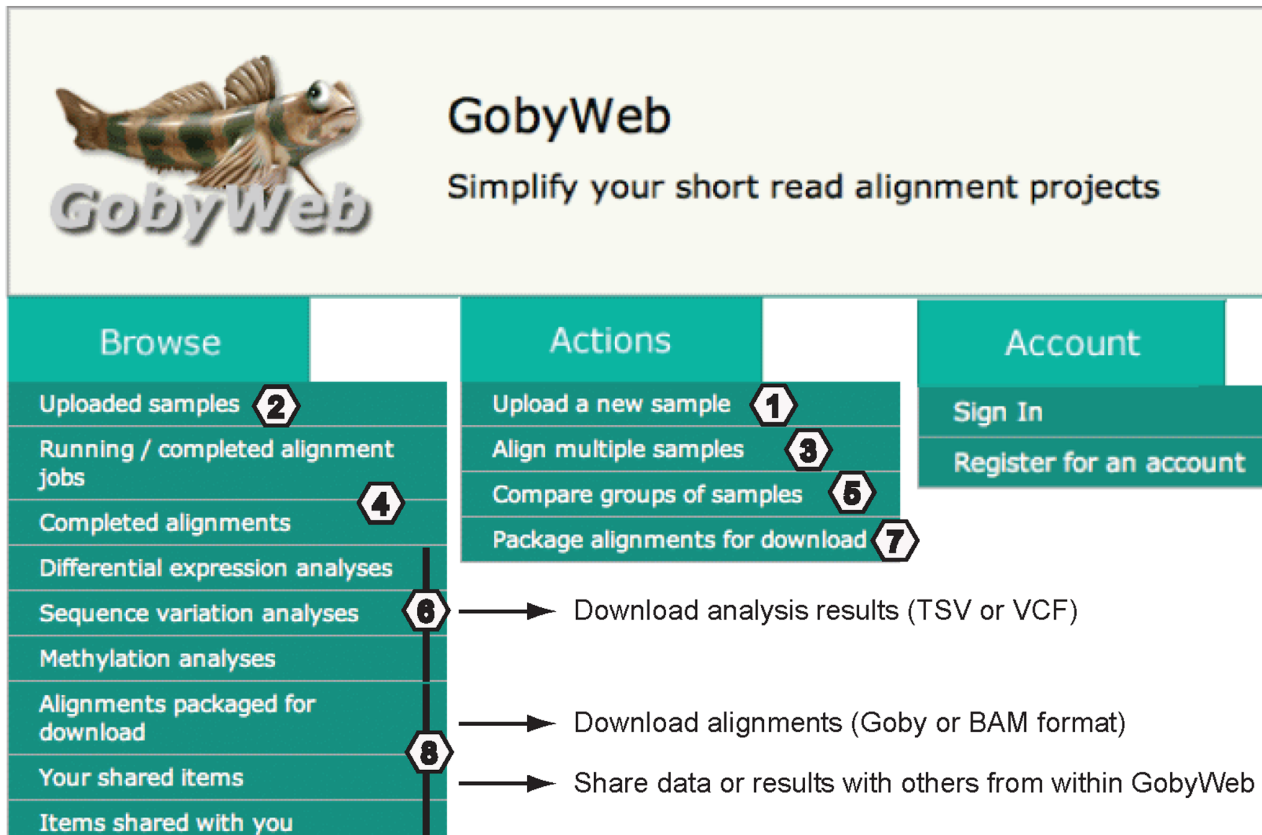


Figure 1. GobyWeb user interface menus. Increasing numbers indicate the typical order in which a user would navigate the interface, from data upload (1) to download or sharing results with others (8). See Supporting information description for a detailed description of each step. doi:10.1371/journal.pone.0069666.g001

Because both BWA and Bowtie are unable to align RNA-Seq reads through splice junctions, we also aligned this dataset with

GSNAP [6]. Spliced alignment with GSNAP was much slower, requiring 25 h 20 m to align the 1.1 billion reads, but as expected

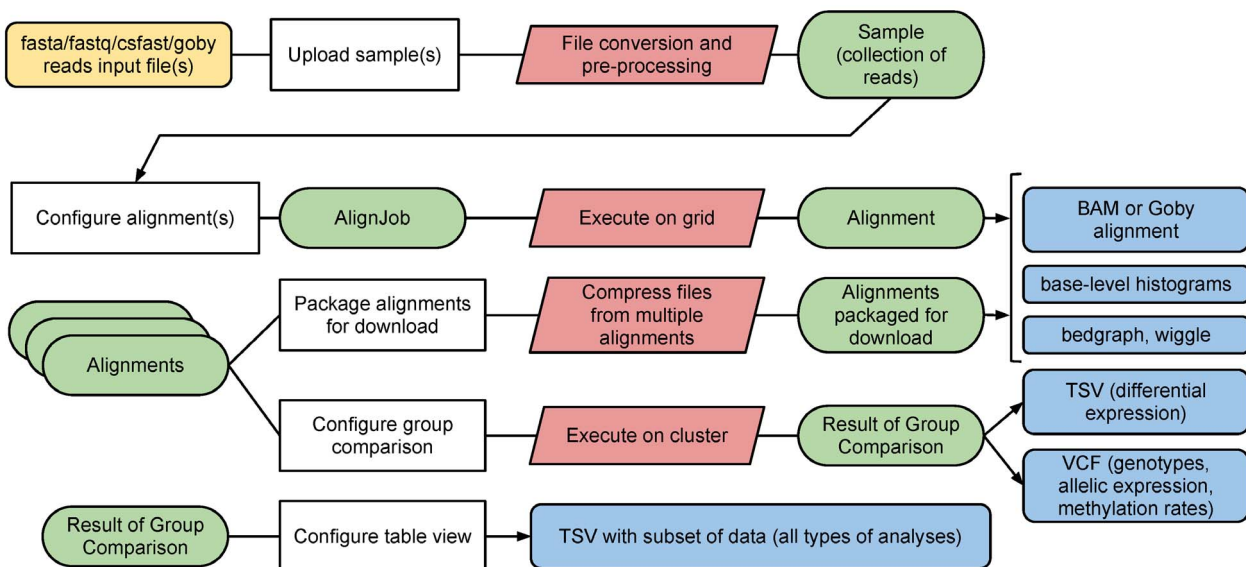


Figure 2. Overview of flows of data in the GobyWeb system. A typical project starts with upload of data files (in yellow, top left). Tasks that run on the compute grid are shown in red. Items of data represented in GobyWeb are shown in green and have dedicated web user interface views. Most views offer the option of downloading result files (in blue) in formats compatible with third-party software. doi:10.1371/journal.pone.0069666.g002

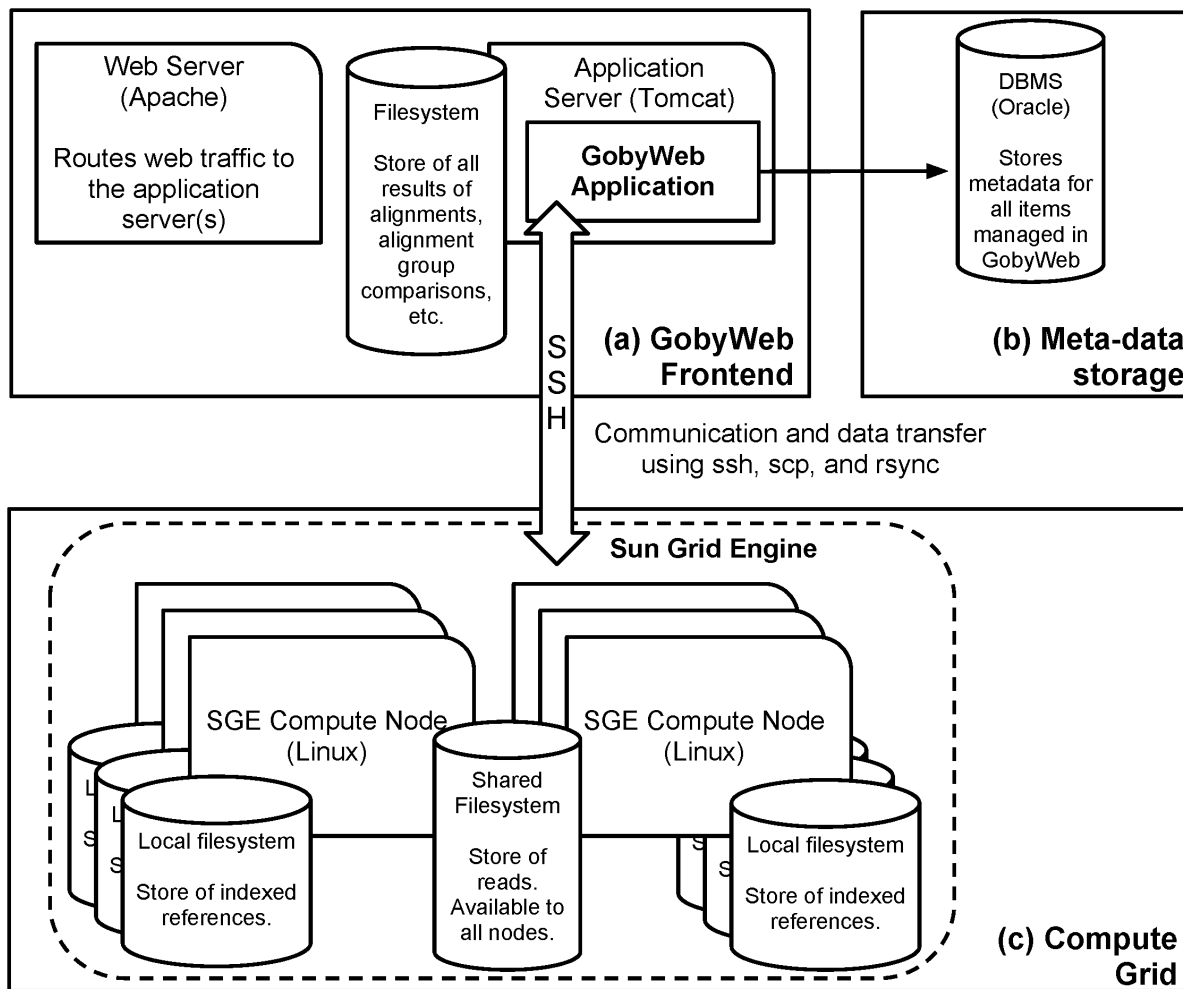


Figure 3. Overview of the system architecture. An installation of GobyWeb relies on three pieces of infrastructure. (a) The web front-end is deployed as a Java web application on one or more application server(s). Several servers can be used to scale the application up under heavy usage. (b) Meta-data about samples, alignments, analyses and users are stored persistently in a Database Management System (DBMS). (c) A compute grid is used to process large datasets efficiently. All datasets (reads, alignments, processed results) are stored as large files on local disks directly attached to each compute node, and the web application servers, as well as in a shared network file system. The software automatically performs data transfers between the shared file system and local storage disks and optimizes these transfers to maximize the overall analysis throughput of the system. The system relies on production quality software components (Apache web server, Tomcat application server, Oracle/JDBC DBMS, and Sun Grid Engine, Linux and Network File System) that are already available and used in many academic institutions.
doi:10.1371/journal.pone.0069666.g003

mapped many reads to the genome through splice junctions that BWA (and Bowtie) are unable to map.

The STAR aligner [5] is also available as a plugin to GobyWeb. STAR can perform spliced alignments, but in our configuration can only align reads longer than 50 bp (because shorter reads are now uncommon). We used GobyWeb to align about 43 million reads with both GSNAP and STAR (publicly available dataset GEO GSM424349). Table 2 presents the duration of these alignments and indicates that STAR (tag: EBGNHJW) is about 4.8 times faster than GSNAP (tag: HJMAOVP) on this dataset, while providing comparable spliced alignments. To illustrate this later point, we visualized alignments generated with GobyWeb in the Goby format [3] with the Integrative Genomics Viewer (IGV, [9]). This comparison is shown in Figure 4.

Differential expression tests for genes can be conducted with GobyWeb using either DESeq [7] or EdgeR [8,10]. A Goby differential expression plugin also makes it possible to estimate RPKM values and their logarithm for genes in individual samples

and estimate fisher exact test statistics and non-moderated Student *t* tests (with adjustment for multiple testing with the Benjamini Hochberg method). While the Fisher statistic is not recommended for comparison of samples with biological variation [8], the RPKM values in individual samples are useful to create correlation plots. To assess the performance of different expression analysis for genes, we split the 72 samples in two groups (randomly, following [1]) and calculate differential expression with each method supported by GobyWeb. Table 2 summarizes these benchmarks. When Myrna performed the analysis in 80 minutes (10 node cluster), GobyWeb completed an equivalent analysis in 21 minutes (small benchmark grid).

Pathogen Detection

Biological samples can be contaminated by viral or microorganisms other than the organism under study. When left undetected, such contaminations can bias the conclusion of a study [11]. Detecting pathogen contamination in clinical samples

Table 1. Gene Expression Analysis performance.

	GobyWeb			Myrna	
	System	BWA	GSNAP	System	Bowtie
Number of nodes in benchmark system	3			10	
Number of cores in benchmark system	72			80	
Upload time	15 m			1 h15 m	
Alignment time		2 h22 m	25 h20 m		2 h56 m
Differential expression time	21 m			80 m	
Total time	36 m	142 m	1520 m	155 m	176 m
Analysis wall clock time (including upload, alignment and differential expression tests)		225 m	1603 m		331 m
Compute Cost per node		\$0.34	\$0.34		
Approximate cost for computation		\$3.88	\$27.43	\$44.00	
Compute Infrastructure costs (excluding storage, network and web)	\$27,000			\$0	

Wall clock times for analysis of 72 Pickrell et al 36 bp RNA-Seq samples.
doi:10.1371/journal.pone.0069666.t001

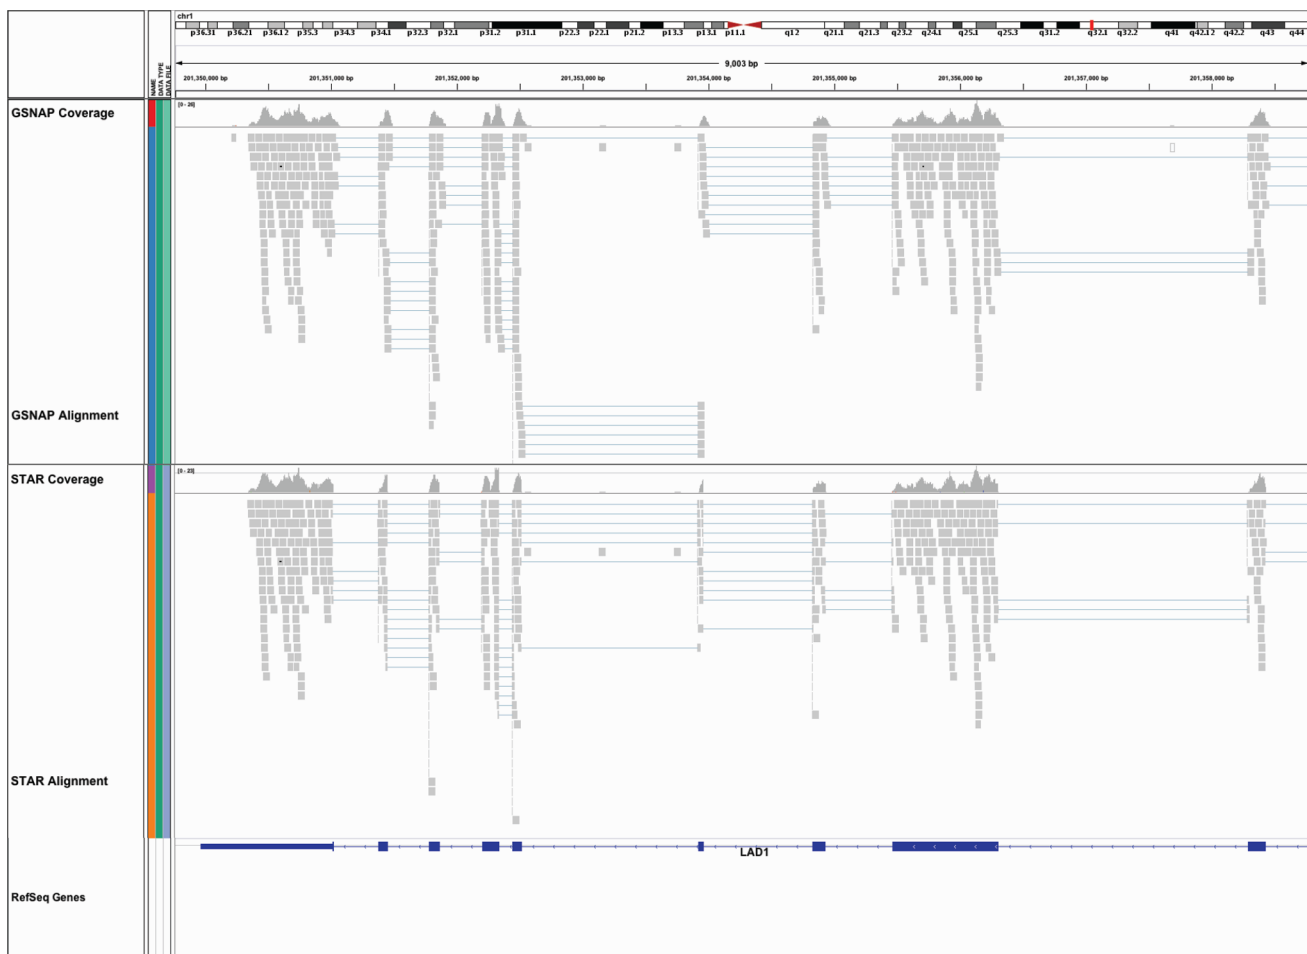


Figure 4. Visualizing spliced RNA-Seq alignments done with GobyWeb and the GSNAP or STAR aligners. This figure was constructed with the Integrative Genomics Viewer (IGV), which directly supports alignments in Goby format. Alignments in the Goby format are substantially smaller than in BAM format, and can be directly downloaded from GobyWeb for interactive visualization with IGV. The plot provides a visual comparison of spliced alignments generated with the GobyWeb GSNAP and STAR plugins over the LAD1 gene (human).
doi:10.1371/journal.pone.0069666.g004

Table 2. Performance of spliced alignments with GSNAP and STAR.

Aligner	GSNAP	STAR
Wall-clock time for alignment and statistics collection	373 m	78 m

Alignments were performed with the GobyWeb and the GSNAP or STAR alignment plugin. One 50 bp single end RNA-Seq sample with about 43 million reads.

doi:10.1371/journal.pone.0069666.t002

is also of great interest [12]. GobyWeb offers a plugin to detect pathogen contamination in samples. Briefly, alignments are processed to extract reads that did not align to the reference genome. Such reads are assembled and mapped to viral, bacterial or fungal transcriptomes. Results are summarized as a table of species matched by each sample or group of samples, contigs for assembled reads, and table of detailed contig mapping information. The process makes it possible to detect contamination by viral, bacterial or fungal organisms in a variety of samples. While several command line tools have been developed to detect pathogens in sequencing data (e.g., [11–13]), the pathogen detection plugin is tightly integrated with GobyWeb and makes it possible to routinely screen samples for pathogen contamination. To measure performance, we analyzed the 72 RNA-Seq Pickrell samples with the pathogen detection plugin (searching viral genomes, using the Minia assembler and stripping Illumina adapters from the reads prior to assembly). The analysis completed in 53 minutes on the small benchmark grid and detected two viruses and one phage (see Table 3). Detected viruses include the *Human herpesvirus 4*/Epstein Barr Virus (EBV) with more than 10 contigs per sample, and the *Macacine herpesvirus 4*. Enterobacteria phage phiX174, sometimes used as spike-in for quality control on the Illumina platform was also detected in two samples. Detecting EBV in HapMap samples is expected because the HapMap cell lines were produced by transforming B lymphocytes with the EBV [14]. Detection of the *Macacine herpesvirus 4* (MH4) is likely to be artifactual, since MH4 is a virus of the same genus as the EBV, and is detected in each sample with less than 10 contigs.

DNA Methylation

GobyWeb supports the analysis of bisulfite-converted reads. DNA samples that have been processed with an experimental protocol such as RRBS, ERBS or methyl-Seq make it possible to estimate methylation at specific cytosine bases in biological samples. DNA methylation analyses start with aligning the reads to a reference genome while allowing for the type of mismatches

introduced during bisulfite conversion. To this end, GobyWeb offers a choice of alignment tools: GSNAP [6], Last [15], or Bismark aligner [16]. Each of these tools is implemented as a plugin, and it is easy to add support for new methods.

To measure the performance of bisulfite alignment with these tools, we uploaded 6 Reduced Representation Bisulfite Sequencing (RRBS) samples, the Dnmt samples [17], to GobyWeb. Each sample contained between 30 and 37 million reads. The six samples (201 million 36 bp reads) uploaded in 23 minutes. Aligning these reads against the mouse MM9 genome required 14 h and 3 minutes with GSNAP, 4 h 13 m with Bismark and 2 h 33 minutes with LAST (see Table 4).

Analysis of DNA methylation data often requires estimating methylation rates at observed sites across the genome, and performing tests of differential methylation. GobyWeb offers two plugins to help with these analyses. The first plugin estimates methylation rates for each observed cytosine (base-level analysis). We simulated bisulfite conversion to compare methylation rates obtained with the Goby plugin to rates estimated from files produced with Bismark and found comparable agreement for both methods (Figure S4). The second plugin estimates average methylation rates over a set of pre-defined annotations (e.g., CpG islands, promoter regions, gene body regions). Both these plugins perform tests of differential methylation when groups of samples are defined by the end-user. The plugins support up to 10 different groups and an arbitrary number of comparisons between pairs of groups (statistics of differential methylation are reported for each site for each comparison defined by the user). Results can be viewed and downloaded with web-based table views. GobyWeb table views are fully interactive, support filters on multiple columns, and scale gracefully to support results with hundreds of millions of rows (Figure 5). In addition to table view, both methylation plugins produce files suitable for visualization with IGV. The region-based methylation plugin produces files in the IGV format, while the base-level plugin produces VCF files that support viewing base-level methylation estimates across sets of samples. Figure 6 presents an example of visualization produced with these two plugins for the Dnmt samples.

Taken together, these capabilities are substantial improvements over software tools previously published.

Discussion

Command Line Tools

Development of HTS technology has spurred the development of specific computational approaches to process the data, such as alignment programs (e.g., BWA, Bowtie, or GSNAP) or approaches for calling differentially expressed genes (e.g., DeSeq, EdgeR). However, most tools were designed for users comfortable

Table 3. Pathogen detection performance.

Viral organism detected	N	Comments
Human herpesvirus 4 type 1 (EBV)	72	The HapMap lymphocyte samples were transformed with EBV to yield individual lymphoblastoid cell lines. Detecting EBV in these samples is therefore expected.
Human herpesvirus 4 (EBV)	72	
Macacine herpesvirus 4	23	Likely mis-detected because of close homology with the EBV virus (less than 10 viral contigs per sample are detected in a subset of samples).
Enterobacteria phage phi X 174	2	Likely spike-in with Illumina PhiX phage DNA.

Pathogen detection took 53 m for the 72 Pickrell et al RNA-Seq samples. N: Number of samples where GobyWeb identified at least one viral contig from the specified organism. See tag: DNOAOZI.

doi:10.1371/journal.pone.0069666.t003

Table 4. DNA methylation analyses.

	GobyWeb				Analysis of individual cytosines	Analysis of annotated regions
	System	GSNAP	Bismark	Last		
Upload	78 m					
Alignment		14 h03 m	4 h13 m	2 h33 m		
Differential Methylation Analysis					17 m	5 m
Complete base level analysis, wall clock time		860 m	975 m	170 m		

We analyzed 6 RRBS samples organized in two groups to detect differentially methylated regions and bases.
doi:10.1371/journal.pone.0069666.t004

with the UNIX/Linux command line. This category of users rarely includes the biologists who generate the datasets. This fact contributes to creating an analysis bottleneck since bioinformaticians are needed even for the most routine data analyses. Beyond this mismatch, command line tools do not fully address the kind of practical problems that investigators encounter when their studies require the collection and analysis of tens of samples. This is especially true when each sample requires several Gigabytes of storage. For these projects, even experienced command line users can benefit from intuitive user interfaces that help with routine analyses, and can improve data organization and analysis reproducibility. In our experience, most researchers need data management capability to help with large HTS analysis projects. GobyWeb is an integrated analysis system that provides strong data management capabilities with a convenient user interface.

Core Facility Pipelines

Many bioinformatics groups have integrated command line tools into internal pipelines to facilitate data processing. Because pipelines are often implemented as scripts, the same limitations discussed for command line tools apply, and these pipelines are typically not exposed to biologists who generated the data, limiting results communicated to biologists to pre-determined sets of reports. In contrast to users of core facilities that maintain in-house pipelines, users of GobyWeb can query their own datasets directly, freeing time for bioinformatician to evaluate new methods, develop or install new plugins and generally focus on more interesting problems than running the same analysis ten times. The software also supports collaborative patterns where a set of users (e.g., members of a bioinformatics lab or core facility personnel) runs standard alignments and analyses for the type of data, and end-users browse and query the results in various ways, or run additional analyses trying different algorithms or parameters.

Gene Expression and DNA Methylation

The plugins distributed with GobyWeb provide state of the art methods for analysis of gene expression and DNA methylation data. STAR and GSNAP alignments can perform spliced alignments efficiently, DESeq or EdgeR statistics are available to call differentially expressed genes or splice sites with differential usage across groups of samples. Bisulfite converted reads can be mapped with Bismark, or the Last aligner, and analyzed across groups to yield differential methylation statistics at single cytosines or annotated regions. GobyWeb generates file formats that can be directly visualized in IGV to follow up on findings of differential methylation and integrate these observations with other annotations or data. Together these features provide an integrated

analysis system to study gene expression and DNA methylation. We anticipate that methods developed as R scripts such as MethylKit or.


Cloud Computing

Several systems have been developed to run analysis of HTS data on compute clouds. Some of these systems provide capabilities to process collections of samples. However, these systems are often limited to one type of data and/or require users to transfer data beyond their institutional firewall. Example of such systems include Myrna [1], which focuses on RNA-Seq data, MethylKit, which focuses on base-level methylation data [18], SIMPLEX [2], which focuses on exome data, or Clovr, for bacterial genome assembly [19]. Systems that support a single type of data require users to work with multiple user interfaces for projects that require the integrated analysis of different assays. GobyWeb improves upon these systems by offering one convenient interface and numerous types of analyses. Our experience suggests that users can learn to use GobyWeb effectively in two, one hour, training sessions. Much of the material covered in these sessions is also offered as training videos online [20]. Cloud-based systems also often lack strong data management capabilities to help users work with many samples. A drawback of GobyWeb compared to cloud-based systems is that it currently cannot be easily deployed to a commercial cloud environment and is limited to a local grid. This is a drawback because cloud-based systems can procure on-demand compute capacity for periods when project activity spikes. We chose to focus on local grid deployment for the initial release of GobyWeb because the proximity of the analysis grid to sequencers deployed at an institution has significant performance advantages as the data volume of typical projects continues to grow. Costs of a local grid are also typically lower than cloud solutions when compute needs are sustained, and access to a server room and system administration team are available [21]. Because of its focus on internal grids, GobyWeb can be deployed in intranets with no Internet connectivity when data confidentiality is a strong requirement.

Commercial Systems

Several proprietary systems provide data management features and are commercially available (e.g., GeneSpring NGS, Avadis NGS, Partek Genomics Suite). Beyond significant costs, commercially available systems are closed source, most rely on programs that were neither described nor evaluated in peer-reviewed publications, and many are not readily extensible to support new types of analyses. GobyWeb is offered free of charge for academic institutions, integrates many state of the art academic tools, and provides a mechanism to describe analyses as plugins, which are

Browse Results of Methylation analysis

This action allows you to browse the results of a Methylation analysis. You can sort by column values, filter the data (see the input boxes under each column), determine the subset of columns to display (see the control panel at the bottom of the page). You can export the current view of the data by clicking the "Download filtered data as TSV" button. Filtering Tips: 

Analysis of INDUCED / CONTROL for tag RQLDONK

Filtered list of elements
Download filtered data as TSV

CHROM	POS	REF	ALT	INFO[Context]	INFO[Delta_MR[INDUCED/CONTROL]]	INFO[empiricalP[IN
2	122175241	G	A	CpG	58	0.0013048379
2	122187773	C	T	CpG	71	0.0062861047
2	123915726	C	T	CpG	24	0.008558251
2	123915761	C	T	CpG	38	9.036674E-4
2	123915962	G	A	CpG	45	0.007330666
2	123915992	C	T	CpG	92	0.0015074909
2	123915993	G	A	CpG	54	0.0028309398
2	123916010	C	T	CpG	85	0.002119586
2	123916016	C	T	CpG	68	0.0032031597
2	123916017	G	A	CpG	86	0.0010401481
2	123916033	G	A,C	CpG	57	0.0060527106
2	123916044	G	A,C	CpG	86	5.3971895E-4
2	123916052	G	A	CpG	81	0.0010401481
2	123916381	C	T	CpG	99	0.0062861047
2	123916426	G	A	CpG	88	0.0016894651
2	123916438	G	A	CpG	67	0.004766484
2	123916441	G	A	CpG	72	0.0035836513
2	123916446	G	A	CpG	70	0.0016894651
2	123916458	G	A	CpG	75	0.0018755752
2	123916750	G	A	CpG	49	0.0030268491
2	124758349	G	A	CpG	60	0.0062861047

Showing 1 to 21 of 8,019 entries (filtered from 363,251 total entries)

All Columns

Filter:

- my_primary_key
- CHROM
- POS
- ID
- REF
- ALT
- QUAL
- FILTER
- INFO[BIOART_COORDS]
- INFO[Strand]
- INFO[Context]
- INFO[LOD[INDUCED/CONTROL]]
- INFO[LOD_SE[INDUCED/CONTROL]]
- INFO[LOD_Z[INDUCED/CONTROL]]

Showing 120 of 120

Show/Hide Columns

Visible Columns

Filter:

- CHROM
- POS
- REF
- ALT
- INFO[Context]
- INFO[Delta_MR[INDUCED/CONTROL]]
- INFO[empiricalP[INDUCED/CONTROL]]
- INFO[DMR[INDUCED/CONTROL]]
- INFO[MR[INDUCED]]
- INFO[MR[CONTROL]]
- INFO[#C_Group[INDUCED]]
- INFO[#Cm_Group[INDUCED]]
- INFO[#C_Group[CONTROL]]
- INFO[#Cm_Group[CONTROL]]

Showing 20 of 20

Figure 5. Scalable Table Views. GobyWeb offers web-based table views that scale to support tables of results with hundred of millions of rows. Users can subset the table to keep specific columns, as well as rows that match complex filters on column values. This mechanism makes it possible for end-users to work with very large tables and download only interesting subsets of the data, even over slow Internet connections. In this snapshot, the table viewer displays results from a base-level methylation analysis (tag = RQLDONK). The panel "Filtered list of elements" displays the current view of the table. The panel at the bottom makes it possible for end-users to select which subset of columns they need to visualize/download. The filters help users identify columns by keyword. Text boxes under each column are used to enter filtering criteria on the specific column. doi:10.1371/journal.pone.0069666.g005

distributed with source code to facilitate peer-review and future extensions.

Parallel Computing

We have designed the GobyWeb system from the ground up for parallel computing. A number of paradigms have been proposed to deploy HTS data analysis on parallel systems. Most command line tools developed by the bioinformatics community are either

sequential or limited to node parallelism, where the program runs parts of the work on parallel threads on the same machine. This type of parallelism requires adding additional core or processors inside a single node to scale to larger datasets. In contrast, grid or cluster parallel computing paradigms can split workloads and coordinate a large set of nodes to achieve parallel speed-ups. GobyWeb takes advantage of the grid paradigm and makes it possible to reduce computing time by adding nodes to a compute

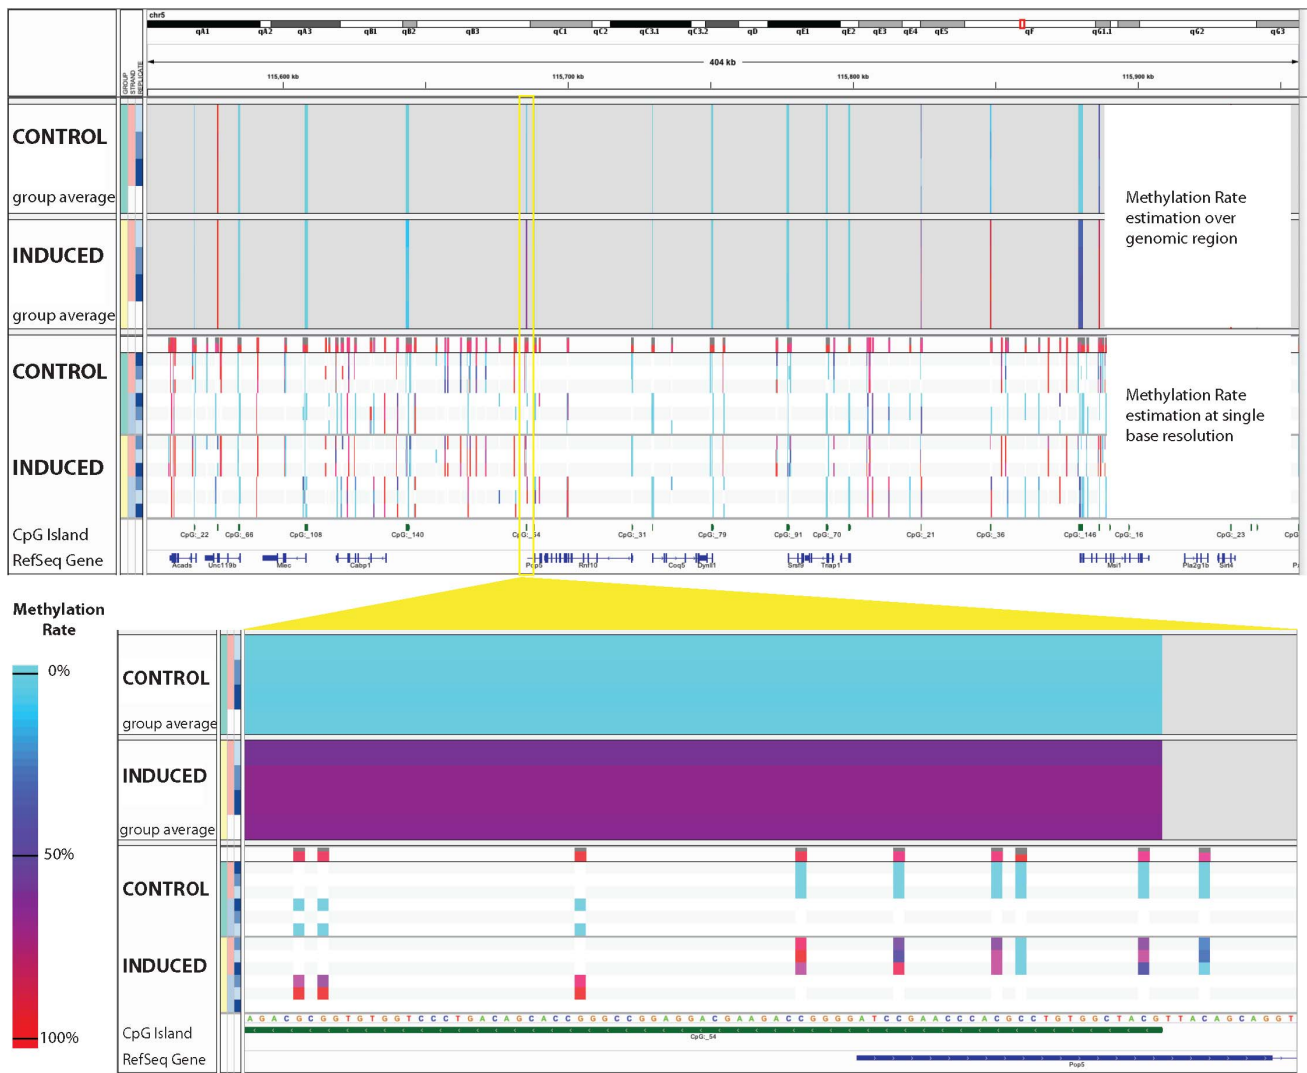


Figure 6. DNA methylation data analyzed with GobyWeb and visualized with IGV. GobyWeb produces data files in formats directly supported by the Integrative Genomics Viewer (IGV). This figure presents the results of methylation analysis over regions and individual bases for the Dnmt public datasets [17]. The bottom insert shows a smaller region with more details of the methylation rate at individual bases. Three rows per strand are shown, corresponding to 3 control and 3 induced samples. Integration with IGV makes it possible to visualize DNA methylation rates alongside other types of annotations or data types supported by IGV. The genomic region shown was selected among the regions that show one of the smaller p-values when comparing the control and induced group (empirical p-value, GobyWeb). doi:10.1371/journal.pone.0069666.g006

grid. Several programming methods have been proposed to take advantage of collections of compute nodes. The MapReduce approach, described in [22] and applied to some bioinformatics problem efficiently co-locates data with compute resources. While very efficient, the approach requires most programs to be rewritten to fit the MapReduce requirements. In contrast to MapReduce, the GobyWeb plugin system can integrate and run in parallel a variety of software without major redesign and reimplementation of their algorithms. To achieve this, we require that the software is able to run on specific parts of very large files and to combine part results into a complete result set. This requirement is often much easier to meet than the typical algorithm redesign and re-implementation required for a MapReduce solution. A current limitation of the parallelization paradigm used in GobyWeb is that scripts must have access to a shared file system.

Workflow Systems

Workflow systems are used widely to interactively construct custom pipelines to integrate and/or query a variety of datasets with different tools. Taverna [23] and Galaxy [24] are well known workflow systems developed to support bioinformatics applications [23,24]. Taverna relies on web services to perform computations. It is not well adapted to process large, multi-gigabyte, HTS datasets and can run into serious performance issues when the volume of data exceeds the capability of the network connectivity between the Taverna application and the services. In contrast, Galaxy can be installed locally to store large datasets and process them on a local compute grid. Galaxy provides a tool-box of several key HTS tools, which end users can apply to analyze their datasets. Because Galaxy is a general workflow system, it is inherently more flexible than GobyWeb, making it possible for end users to assemble custom analysis pipelines. In contrast GobyWeb is more rigid: it only supports a limited set of predefined analyses

(plugins currently have strict sets of inputs and outputs). However, we have designed this limited set to cover a wide range of common HTS analysis, and have been able to optimize these analyses for very large HTS datasets. A key advantage of GobyWeb is that all data are stored in compressed binary formats [3], which are several orders of magnitude smaller than the uncompressed text files used internally by Galaxy. For instance, alignment files stored by GobyWeb are about 10 times smaller than BAM files, with BAM files about three times smaller than SAM text files or equivalent tab delimited files. GobyWeb is therefore expected to require about 30 times less storage space for analysis than Galaxy. We note that GobyWeb currently uses a more rigid plugin system than Galaxy. Rigid systems can be more rigorously tested, which is critical in some domains, such as clinical sequencing, while more flexible workflow systems make it possible to experiment with new analyses quickly. Rigid and flexible plugin systems can be seen as complementary. For instance, tab delimited data exported from GobyWeb after alignment and group comparison can be further analyzed in Galaxy with custom, project specific pipelines. Table 5 provides a tabular comparison of a few systems available to analyze HTS data.

Software Licensing

GobyWeb can be freely downloaded for non-commercial academic use. Detailed installation instructions are provided on the web site. GobyWeb plugins are distributed under the LGPL3 license and can be obtained from the GitHub repository linked from the main project web site.

Methods

Web Application

GobyWeb is implemented as a web application written with the Grails framework (<http://grails.org/>). We have deployed the application in a Tomcat application server. Object persistence was implemented with Hibernate (<http://www.hibernate.org/>) and an Oracle backend. HTS read files are multi-gigabyte files that frequently exceed the 2 GB size limit that HTTP file uploads can support. Large file uploads are supported in GobyWeb with the Jumploader applet (<http://jumploader.com/>).

User Interface Design

The application was developed iteratively and the user interfaces were designed as a team, then implemented and continually adjusted in response to user feedback. Initial releases were limited to the Campagne laboratory and direct collaborators, but a version of GobyWeb has been released to a wider audience of more than 70 investigators (faculty, post-doctoral fellows and students from the Weill Medical College of Cornell University, Memorial Sloan Kettering Cancer institute and Hospital for Special Surgery) since January 2010.

Grid Computing

Embarrassingly parallel analyses are split into chunks and scheduled as array job on an Sun/Oracle Grid Engine (OGE, <http://www.oracle.com/us/products/tools/oracle-grid-engine-075549.html>). OGE supports arbitrary job dependencies and this feature is used extensively to interleave user jobs and increase overall job throughput on the compute grid. Results of computations for independent chunks are merged by a post-process OGE job that is configured to start when all parts of the array job have finished. Importantly, the post-process job is configured to start even if some parts of the array job fail. This is useful when working with aligners or other software that

sometimes can reproducibly fail when run on specific data items (in such cases, GobyWeb produces an output for the data that could be processed, and provide visual indication of failure for the chunks of the parallel job that failed to compute). We use the Goby framework (<http://goby.campagnelab.org>, [3]) to split read and alignment files and combine results efficiently.

Computational Analyses

Logic for computational analyses is implemented in Bash scripts. These shell scripts automate data transfers between the web server and the compute cluster as required by end-user analyses. The scripts schedule long running and parallel jobs on the Oracle Grid Engine instance with appropriate dependencies for each type of analysis. Scripts are also responsible for informing the web application of changes in the analysis status (job submission, start, intermediate steps, failure or completion). GobyWeb bash scripts are organized with a plugin architecture that makes it possible for developers and administrators to extend GobyWeb with custom alignment or analysis methods.

Plugin Architecture

GobyWeb currently supports three types of plugins: resource, aligner and alignment analysis. Plugins are organized in directories that contain a plugin definition file (`config.xml`), a plugin script (`script.sh`) and optional plugin files (named as described in the plugin definition file). Briefly, resource plugins provide access to data files or executables that other plugins can share access to, such as samtools, vcf-tools or executables for alignment programs. Aligner plugins make it possible to extend GobyWeb to perform alignments with new alignment tools. Aligner plugins can generate either BAM or Goby alignments and run each sample file either sequentially, or in parallel. Alignment analysis plugins make it possible to extend GobyWeb with methods for analysis a set of alignment results. Source code for GobyWeb plugins is distributed at GitHub (<http://github.com/CampagneLaboratory/gobyweb2-plugins>). Developing new plugins consists in creating an XML plugin configuration file and providing a bash script with predefined shell functions that interface between GobyWeb and the analysis or alignment tool. The `config.xml` files make it possible for plugin developers to define parameters of the plugins, which get injected into the web user interface at run-time and displayed to end-users in a user friendly way. Step by step instructions for creating new plugins are provided in the Software Development Kit section of the GobyWeb web site.

Plugins and Parallelization

Both aligner and alignment analysis plugins can be written to run either in a node-parallel mode, or in a grid-parallel mode. The node parallel plugins can take advantage of thread parallelism on a single node. They are simple to implement because it is only necessary to define one process function per plugin. However, node parallelism limits the maximum data processing throughput that can be obtained when a grid with several nodes is available. Node parallel plugins are also only efficient if the programs they use can take advantage of thread parallelism. In order to support programs that run sequentially, and increase data processing throughput by using more nodes of a compute grid, the plugin system also supports grid-parallel plugins. Grid-parallel plugins need to implement four shell functions, which are used to (i) determine how to split a file into chunks, (ii) count the number of chunks generated in the first step, (iii) process a chunk to produce a result, and (iv) combine many chunk results to produce a complete result.

Table 5. Comparison of GobyWeb with other NGS Analysis Systems.

System	Efficient Use of Storage capacity	HTS data assays supported	Flexibility	Automatically installs dependencies:		HTS data/General purpose system	Parallel analyses on compute grid
				software	data		
Anduril [32]	++	Varied, depends on available workflows	+++++(5)	Y	Some	General purpose	N
Galaxy [24]	+(1)	RNASeq pipelines provided	+++ (6)	Y	Y	General purpose	Limited (11)
GenePattern [33]	++ (2)	Varied, depends on available workflows	+++	Some (8)	N (9)	General purpose	N (12)
GobyWeb (this report)	+++ (3)	RNA-Seq, DNA methylation, DNA-Seq	++ (7)	Y	Y	HTS data	Y (13)
MyRna [1]	++ (4)	RNA-Seq	+	Y	N	HTS data	Y (14)
Taverna [23]	++	Varied, depends on available workflows	+++ (6)	N/A	Y (10)	General purpose	Y (15)

Notes: (1) Supports text files.

(2) Supports any file format.

(3) Supports Goby compressed file formats.

(4) Supports BAM.

(5) Provides a scripting language specialized for pipeline development.

(6) Pipelines can be constructed by connecting components graphically.

(7) Pipelines are optimized for analysis of large HTS datasets.

(8) For example, installs Perl on Windows, but not on other platforms.

(9) Data often needs to be installed on execution server.

(10) Relies on remote web services.

(11) Very few file formats support parallelization.

(12) Runs analyses on a single server.

(13) All analyses are fully parallel.

(14) Deploys to Amazon AWS for Hadoop.

(15) Some extensions support execution on grid.

doi:10.1371/journal.pone.0069666.t005

Status Monitoring

Status updates are communicated to the web application by posting messages to pre-defined URLs. Scripts and plugins accomplish this by calling a command line tool with status information and the tag of the analysis that the script is running. The web application stores status updates in the Hibernate persistence store (e.g., oracle database). Upon receiving a completed status update, the web application checks that result files have been created and are valid. When this is verified, the information about the job is saved to the Hibernate persistence store.

Read Storage

Upon sample upload, fasta, fastq or csfasta read files are converted to Goby compact-reads format with the fasta-to-compact tool [3]. This transformation is performed to allow efficient parallelization of the alignment jobs (see Grid computing and Read alignments sections). Reads are transferred from the web application server to the shared file system visible to the compute cluster (copies are performed with scp). A number of pre-processing steps are executed to validate input files, obtain read length and base quality statistics, as well as associate weights to reads (such as heptamer weights required by the Hansen et al approach [25]). Read conversion and processing steps are executed on the OGE grid and proceeds in parallel for each sample file uploaded.

Read Alignments

For plugin aligners that support parallel processing, alignments are split as OGE array jobs with n tasks, and a post-process job (see Grid computing section). The number of tasks is determined by dividing the read file size by the chunk size (in practice, we use `CHUNK_SIZE = 50,000,000`, for approximately 50 MB chunks). When an array job component starts executing on the cluster, OGE sets the `SGE_TASK_ID` variable to the index of the task in the array job. We determine the start and end offsets within the reads file that contains the reads that this task should align as follows:

`START_POSITION = (SGE_TASK_ID - 1) * CHUNK_SIZE`. Similarly, `END_POSITION = START_POSITION + CHUNK_SIZE - 1`. We then run the aligner plugin 'plugin_align' function with these `START_POSITION` and `END_POSITION` values, since these options restrict the alignment to only the reads contained within these file byte offsets. When the array tasks finish, Goby alignment results are concatenated with the concatenate-alignment Goby tool, calling the tool recursively to combine at most 100 pieces at a time and therefore limiting the number of files open at any given time. This process makes it possible to scale alignment concatenation to alignments with billions of reads. To support aligners that do not provide native Goby support, we use the goby compact-to-fasta tool with `-s START_POSITION` and `-e END_POSITION` to extract Fasta or Fastq format for the reads of the part. Aligners that do not produce Goby alignment format require the implementation of conversion scripts that convert the aligner output to Goby format. An example of this method is provided for the last/lastag aligners. For plugin aligners whose definition file indicates that they do not support parallel processing, we run the alignment of each sample as an independent OGE job with exclusive access to a node, and call the plugin_align function of the plugin with the entire reads file. Such aligner plugins are expected to leverage multi-core architectures on each node.

GSNAP Plugins

The GSNAP aligner is integrated with GobyWeb in the `GSNAP_GOBY` and `GSNAP_BAM` plugins. The former generates alignments in the Goby format, while the latter yields alignments in the BAM format. Both plugins require a recent version of GSNAP ($> = 2011.11.17$) compiled with Goby support. The GSNAP aligner, or the STAR aligner, is the recommended choice when aligning RNA-Seq samples to a genome because they can perform spliced alignments efficiently (i.e., map reads that span one or several exon-exon junctions). GSNAP also supports mapping reads from samples converted with sodium bisulfite (e.g., as in the Methyl-Seq and RRBS protocols). When aligning bisulfite converted samples, the `GSNAP_GOBY` plugin aggressively trims reads to remove adapter sequences (these sequences have to be provided by Goby administrators and can be obtained from the vendor of the sequencing instrument and kits). Adapter trimming is performed with the Goby trim mode and only trims adapters if the trimmed part of the read would exceed four base pairs. Trimmed reads are then aligned with GSNAP in methylation mode with parameters that disable indel matches and terminal trimming (`-mode cmet -m 1 -i 100-terminal-trimming = 100`).

BWA Plugins

The BWA aligner is integrated in the `BWA_GOBY` and `BWA_BAM` plugins. These plugins require the version of BWA patched to support Goby file format (available from <http://campagnelab.org/software/goby/bwa/>). `BWA_GOBY` is grid-parallel and produces Goby alignment files, while `BWA_BAM` is node parallel and produces BAM alignment files. BWA plugins are recommended when aligning DNA-Seq samples.

Last Plugin

The `LAST_GOBY` plugin runs a recent version of the Last aligner [15] (currently version 230). This aligner is recommended when the reads are from an organism for which no reference genome is available and when one needs to align to a reference that is a distant homolog from the organism of interest. The Last aligner is also recommended when aligning small RNA sequences.

Last Bisulfite Plugin

The `LAST_BISULFITE` plugin runs the Last aligner [15] with matrices suitable to align bisulfite converted reads. The plugin searches both forward and reverse strands of a specially indexed reference sequence, and combines the results following the recipe described at <http://last.cbrc.jp/doc/bisulfite.txt>.

RNA-Seq

GobyWeb currently offers five plugins to analyze RNA-Seq results. These plugins calculate counts in parallel, combine results from parallel splits, normalize counts and estimate statistics of differential expression. `DIFF_EXP_GOBY` estimates counts over annotations with the Goby alignment-to-annotation-counts. This plugin outputs counts, RPKM and $\log_2(\text{RPKM})$ for each alignment included in the analysis. This plugin also estimates Student t-test statistics for RPKM values between two groups and fisher-exact test on the raw counts. `DIFF_EXP_DESEQ` estimates counts over annotations with Goby, but uses the R package DESeq to estimate statistics of differential expression. The third plugin `DIFF_EXP_EDGE_R` integrates the EdgeR package with GobyWeb (count estimation is performed with Goby). The fourth plugin, `SEQ_VAR_GOBY` has an output format that estimates allelic differential expression. In RNA-Seq data, this plugin

estimates the significance that the reference allele expression is different between the groups under study. The plugin is implemented with the Goby discover-sequence-variant mode – format allelic-frequencies. The fifth plugin, `SPLICING_DIFF_EXP`, counts the number of reads spanning splice junctions to determine alternative splicing usage (counts spanning junctions are determined with Goby and statistics of differential splicing usage are determined either with DESeq or EdgeR).

RNA-Seq Normalization

GobyWeb supports four normalization methods: hexamer bias removal [25], RPKM/FPKM normalization, upper-quartile normalization [26] and the TMM procedure implemented in EdgeR [8,10]. The method of Hansen et al [25] was implemented to make it possible to remove random hexamer priming biases. Heptamer weights are associated to each read during the post-upload process are used as described previously ([25]) to estimate bias adjusted counts. RPKM normalization is conducted with the Goby alignment-to-annotation-counts and calculated as $r = [(c+1)/(L/1000.0)/(N/1.106)]$, where c is the count, or number of read fragment that overlap with a given annotation, L is the length of the annotation and N a normalization factor. For FPKM/RPKM normalization, N is taken to be the total number of read fragments aligned in a sample. The upper-quartile normalization method is implemented with the same formula, but using the 75 percentile of annotation counts as the value of N , as described previously by Dudoit and colleagues [26].

Pathogen Detection

To determine the presence of pathogen in sequenced samples, this plugin (`CONTAMINANT_EXTRACT`) proceeds in three steps. (1) Reads that do not map the reference sequence (unmapped reads) are retrieved from a set of alignments and reads files. (2) Unmapped reads are optionally trimmed from adapters. (3) Unmapped reads are assembled into contigs. GobyWeb can use either the Trinity assembler or the Minia assembler [27]. Minia is the default choice and recommended for performance and memory usage. (4) Contigs are used to search a large and diverse database of (a) viral, (b) bacterial, or (c) fungal organisms. In the current version of GobyWeb, viral, bacterial and fungal RNA sequences are used and obtained from RefSeq [28]. Contigs that match in these databases are considered annotated if they match over more than 150 bp and have an E-value less than $1e-6$. The species matched by annotated contigs in these diverse databases are recorded and associated to the sample that contributed the unmapped reads. Contig sequences are provided in Fasta format.

DNA-Seq

GobyWeb offers a few approaches to analyze DNA-Seq data. A samtools plugin (`SEQ_VAR_SAMTOOLS`) makes it possible to call genotypes with samtools mpileup, or to estimate allelic association tests for alignments in the BAM format. The `SEQVAR_GOBY` plugin supports calling genotypes as well as performing allelic association tests for alignments in the Goby format. The `SEQVAR_GOBY` plugin is grid-parallel and optimized for comparisons involving large numbers (>50) of samples across groups. This plugin should be considered experimental until we publish the results of large-scale validation tests. The output is generated for both plugins in the Variant Calling Format 4.1, which can be viewed in IGV (VCF).

Detection of Somatic Variations

GobyWeb offers two plugins to detect somatic variants when both germline and somatic DNA are analyzed (e.g., to compare germline DNA to tumor DNA to identify somatic variations in the tumor). The first plugin uses Mutect [29]. The second plugin uses the Goby 2.2+ somatic variation mode (<http://goby.campagnelab.org>). Both plugins process samples in parallel. The Goby somatic variation plugin should be considered experimental until we publish the results of large-scale validation tests.

Methyl-Seq

GobyWeb estimates methylation rates and calls sites of differential methylation across groups of samples. These analyses are implemented in the `SEQ_VAR_GOBY_METHYLATION` plugin. This plugin uses the Goby discover-sequence-variants mode with the methylation output format. This mode implements statistics of differential methylation with a fisher exact-test at individual genomic positions where at least one cytosine is observed. The use of fisher exact test statistics to call differential methylation has been described recently [30]. P-values are adjusted for multiple testing with the Benjamini-Hochberg method across all sites tested in the genome. Methylation rates and statistics are written to VCF format. This plugin also support an empirical p-value estimation (empirical-p) methods that takes into account biological variability within groups and will be described elsewhere.

Variant Calling Format

VCF files are annotated with VCF-annotate to map sites of variations to gene, RefSNP rs id identifiers and variation predicted effect, when possible. Annotated VCF files are then sorted in genomic position (vcf-sort) and indexed with tabix [31]. This process makes it possible to load VCF files produced with GobyWeb directly in the Integrated Genome Viewer (IGV). We have extended IGV to display VCF files that encode methylation rates, as produced by GobyWeb and Goby. VCF files produced for DNA-Seq datasets with GobyWeb can also be visualized with IGV.

Other Plugins

We frequently add new GobyWeb plugins or improve existing ones. The definitive source of information about plugins is the GitHub GobyWeb plugin repository. Plugin configuration files offer a version number for each plugin that is displayed on the user interface and makes it possible to track changes to plugin software over time.

Supporting Information

Figure S1 Uploading reads into GobyWeb to create a new Sample. Read files can be uploaded in a variety of file formats. When the checkbox “Create Multiple Samples” is not selected, individual files are concatenated to yield a single independent biological sample. When the box is not checked, multiple samples are created and associated with the meta-data described on the form.

(TIFF)

Figure S2 Consistent alignment of multiple samples. GobyWeb supports selecting an arbitrary number of samples for alignment. Configuration of the alignments is entered once through the user interface and applied consistently across all the jobs that will be started.

(TIF)

Figure S3 Visual status for alignment running on compute grid. The figure shows the visual status for an alignment in progress against a large sample (30 GB compressed reads were split into more than 600 chunks and were scheduled for alignment). GobyWeb aligns and sorts each chunk, then concatenates the sorted alignments pieces to yield a completely sorted alignment. Alignments are post-processed to derive base level histograms as well as statistics such as number of aligned reads and number of sequence variations at each cycle. (TIF)

Figure S4 Comparison between estimates of methylation rates produced with Bismark and Last/Goby. GobyWeb can align bisulfite converted reads with either the Bismark or the Last aligner. Furthermore, alignments of bisulfite-converted reads can be processed to estimate methylation rates with either Goby or a simple script that post-processes the Bismark result files. Here, (A) we simulated reads from a uniform distribution of methylation rates over a 5 MB region of the human genome, at 50X or 250X average coverage and compare the estimate of methylation with the methylation estimate

produced by each analysis method. We find (B) that both methods yield comparable agreement with true methylation rates and correlate well with each other when average coverage >50X (data simulated for a target of 50X coverage includes regions of the genome where actual coverage is lower than 50X, these sites tend to have larger disagreement with true methylation).

(TIFF)

File S1 Supplemental description and Figures S1-S4. (PDF)

Acknowledgments

We thank the Tri-Institutional Training Program in Computational Biology and Medicine for support to NC. We thank Dr. Ji-Eun Oh for suggestions for improving the GobyWeb user interface.

Author Contributions

Conceived and designed the experiments: FC. Performed the experiments: KCD NC ZZ MS FC. Analyzed the data: KCD NC FC. Contributed reagents/materials/analysis tools: RS. Wrote the paper: FC.

References

- Langmead B, Hansen KD, Leek JT (2010) Cloud-scale RNA-sequencing differential expression analysis with Myrna. *Genome Biol* 11: R83.
- Fischer M, Snajder R, Pabinger S, Dander A, Schossig A, et al. (2012) SIMPLEX: cloud-enabled pipeline for the comprehensive analysis of exome sequencing data. *PLoS One* 7: e41948.
- Campagne F, Dorff K, Chambwe N, Robinson JT, Mesirov JP, et al. (2012) Compression of structured high-throughput sequencing data. Preprint at arXiv.org.
- Li H, Handsaker B, Wysoker A, Fennell T, Ruan J, et al. (2009) The Sequence Alignment/Map format and SAMtools. *Bioinformatics* 25: 2078–2079.
- Dobin A, Davis CA, Schlesinger F, Drenkow J, Zaleski C, et al. (2012) STAR: ultrafast universal RNA-seq aligner. *Bioinformatics*.
- Wu TD, Nacu S (2010) Fast and SNP-tolerant detection of complex variants and splicing in short reads. *Bioinformatics* 26: 873–881.
- Anders S (2010) Analysing RNA-Seq data with the DESeq package. *Molecular biology*: 1–17.
- Robinson MD, McCarthy DJ, Smyth GK (2010) edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics* 26: 139–140.
- Robinson JT, Thorvaldsdottir H, Winckler W, Guttman M, Lander ES, et al. (2011) Integrative genomics viewer. *Nat Biotechnol* 29: 24–26.
- Robinson M (2009) edgeR: Methods for differential expression in digital gene expression datasets. *Bioconductor*: 1–7.
- Lin Z, Puetter A, Coco J, Xu G, Strong MJ, et al. (2012) Detection of murine leukemia virus in the Epstein-Barr virus-positive human B-cell line JY, using a computational RNA-Seq-based exogenous agent detection pipeline, PARSEs. *J Virol* 86: 2970–2977.
- Wylie KM, Mihindukulasuriya KA, Sodergren E, Weinstock GM, Storch GA (2012) Sequence analysis of the human virome in febrile and afebrile children. *PLoS One* 7: e27735.
- Yozwiak NL, Skewes-Cox P, Stenglein MD, Balmaseda A, Harris E, et al. (2012) Virus identification in unknown tropical febrile illness cases using deep sequencing. *PLoS Negl Trop Dis* 6: e1485.
- Hapmap Consortium (2003) The International HapMap Project. *Nature* 426: 789–796.
- Kielbasa SM, Wan R, Sato K, Horton P, Frith MC (2011) Adaptive seeds tame genomic sequence comparison. *Genome Res* 21: 487–493.
- Krueger F, Andrews SR (2011) Bismark: a flexible aligner and methylation caller for Bisulfite-Seq applications. *Bioinformatics* 27: 1571–1572.
- Steine EJ, Ehrlich M, Bell GW, Raj A, Reddy S, et al. (2011) Genes methylated by DNA methyltransferase 3b are similar in mouse intestine and human colon cancer. *J Clin Invest* 121: 1748–1752.
- Akalin A, Kormaksson M, Li S, Garrett-Bakelman FE, Figueroa ME, et al. (2012) methylKit: a comprehensive R package for the analysis of genome-wide DNA methylation profiles. *Genome Biol* 13: R87.
- Angiuoli SV, Matalka M, Gussman A, Galens K, Vangala M, et al. (2011) CloVR: a virtual machine for automated and portable sequence analysis from the desktop using cloud computing. *BMC Bioinformatics* 12: 356.
- Dorff KC, Chambwe N, Zeno Z, Simi M, Shakhovich R, et al. (2000–2013) GobyWeb software: <http://gobyweb.campagnelab.org>. Available: <http://gobyweb.campagnelab.org>. Accessed.
- Dudley JT, Pouliot Y, Chen R, Morgan AA, Butte AJ (2010) Translational bioinformatics in the cloud: an affordable alternative. *Genome Med* 2: 51.
- Dean J, Ghemawat S. MapReduce: Simplified Data Processing on Large Clusters; 2004.
- Hull D, Wolstencroft K, Stevens R, Goble C, Pocock MR, et al. (2006) Taverna: a tool for building and running workflows of services. *Nucleic Acids Res* 34: W729–732.
- Goecks J, Nekrutenko A, Taylor J (2010) Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences. *Genome Biol* 11: R86.
- Hansen KD, Brenner SE, Dudoit S (2010) Biases in Illumina transcriptome sequencing caused by random hexamer priming. *Nucleic Acids Res* 38: e131.
- Bullard JH, Purdom E, Hansen KD, Dudoit S (2010) Evaluation of statistical methods for normalization and differential expression in mRNA-Seq experiments. *BMC Bioinformatics* 11: 94.
- Chikhi R, Rizk G. Space-efficient and exact de Bruijn graph representation based on a Bloom filter; 2012.
- Wheeler DL, Church DM, Lash AE, Leipe DD, Madden TL, et al. (2001) Database resources of the National Center for Biotechnology Information. *Nucleic Acids Res* 29: 11–16.
- Cibulskis K, Lawrence MS, Carter SL, Sivachenko A, Jaffe D, et al. (2013) Sensitive detection of somatic point mutations in impure and heterogeneous cancer samples. *Nat Biotechnol* 31: 213–219.
- Li M, Chen WD, Papadopoulos N, Goodman SN, Bjerregaard NC, et al. (2009) Sensitive digital quantification of DNA methylation in clinical samples. *Nat Biotechnol* 27: 858–863.
- Li H (2011) Tabix: fast retrieval of sequence features from generic TAB-delimited files. *Bioinformatics* 27: 718–719.
- Ovaska K, Laakso M, Haapa-Paananen S, Louhimo R, Chen P, et al. (2010) Large-scale data integration framework provides a comprehensive view on glioblastoma multiforme. *Genome Med* 2: 65.
- Reich M, Liefeld T, Gould J, Lerner J, Tamayo P, et al. (2006) GenePattern 2.0. *Nat Genet* 38: 500–501.