

SOFTWARE

Open Access



LAPIS is a fast web API for massive open virus sequencing data

Chaoran Chen^{1,2*} , Alexander Taepper^{1,3} , Fabian Engelniederhammer⁴ , Jonas Kellerer⁴ ,
Cornelius Roemer^{2,5}  and Tanja Stadler^{1,2*} 

*Correspondence:
chaoran.chen@bsse.ethz.ch;
tanja.stadler@bsse.ethz.ch

¹ Department of Biosystems
Science and Engineering, ETH
Zürich, Basel, Switzerland

² Swiss Institute of Bioinformatics,
Basel, Switzerland

³ School of Computation,
Information and Technology -
Informatics, TU Munich, Munich,
Germany

⁴ TNG Technology Consulting
GmbH, Unterföhring, Germany

⁵ Biozentrum, University of Basel,
Basel, Switzerland

Abstract

Background: Recent epidemic outbreaks such as the SARS-CoV-2 pandemic and the mpox outbreak in 2022 have demonstrated the value of genomic sequencing data for tracking the origin and spread of pathogens. Laboratories around the globe generated new sequences at unprecedented speed and volume and bioinformaticians developed new tools and dashboards to analyze this wealth of data. However, a major challenge that remains is the lack of simple and efficient approaches for accessing and processing sequencing data.

Results: The Lightweight API for Sequences (LAPIS) facilitates rapid retrieval and analysis of genomic sequencing data through a REST API. It supports complex mutation- and metadata-based queries and can perform aggregation operations on massive datasets. LAPIS is optimized for typical questions relevant to genomic epidemiology. Using a newly-developed in-memory database engine, it has a high speed and throughput: between 25 January and 4 February 2023, the SARS-CoV-2 instance of LAPIS, which contains 14.5 million sequences, processed over 20 million requests with a mean response time of 411 ms and a median response time of 1 ms. LAPIS is the core engine behind our dashboards on genspectrum.org and we currently maintain public LAPIS instances for SARS-CoV-2 and mpox.

Conclusions: Powered by an optimized database engine and available through a web API, LAPIS enhances the accessibility of genomic sequencing data. It is designed to serve as a common backend for dashboards and analyses with the potential to be integrated into common database platforms such as GenBank.

Background

Pathogen genomic sequencing data are a key public health resource for responding to epidemic outbreaks. During the early stages of an outbreak, genomic sequencing data are essential for understanding the origin, evolution, and extent of spread of the pathogen [15, 16]. At later stages, sequencing data are the primary early indicator of evolutionary and epidemiological changes, as demonstrated repeatedly with SARS-CoV-2 variants [9, 19, 20]. Rapid analysis of sequencing data is therefore a crucial component for evidence-based public health responses. Although a lot of infrastructure for generating and analyzing genomic



© The Author(s) 2023. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

sequencing data in real-time was established during the SARS-CoV-2 pandemic, major challenges remain [4, 7, 14, 15].

The unprecedented scale of SARS-CoV-2 sequence generation, coupled with enormous popular interest in these data, highlights a need for user-friendly tools for analyzing massive sequence data sets. One such category of tools is web dashboards. Once set up, these can be used by a wide audience without requiring programming and data science knowledge. Examples of popular dashboards that digest massive SARS-CoV-2 data sets include the CDC's COVID Data Tracker [6], CoVariants [13], Outbreak.info [12], and our own CoV-Spectrum dashboard [8]. Another category of tools that facilitate quick, ad-hoc analyses are "notebooks" like Jupyter Notebooks and R Markdown scripts. Notebooks are useful to data scientists with programming knowledge to quickly perform their own statistical analyses and generate their own plots. Combined, dashboards and notebooks allow different users to access different visualizations and focus on different aspects of the data. In this way, everyone from experts like scientists and public health agencies to the general public can benefit from sequence data.

Many of these tools for sequence data analysis require common operations on sequence data like filtering, stratification, and aggregation. For instance, filtering for sequences with certain mutations and calculating the relative frequency of mutations are commonly performed operations for genome sequencing data. Although these operations are simple in principle, the gigantic size of modern genome sequence data sets makes them non-trivial. Over 14 million SARS-CoV-2 sequences are available, and up to hundreds of thousands of new sequences are added weekly. General-purpose database systems such as PostgreSQL are not optimized for genomic sequence analysis on this scale.

Our resource LAPIS (Lightweight API for Sequences) is designed to perform common data operations on millions of genomic sequences within milliseconds, facilitating interactive data exploration. Using a self-written in-memory database engine, it is optimized for filtering and aggregating large genomic sequencing data sets. Accessible through a web API (application programming interface), we believe that LAPIS can serve as a common backend for many dashboards and analyses (e.g., through notebooks). This would relieve scientists and dashboard builders from the costly but boring task of developing their own databases and implementing common basic operations. Instead, they would be free to focus on analysis and visualization tasks. Furthermore, LAPIS streamlines the direct download of cleaned and pre-processed data including aligned and unaligned sequences.

In contrast to data repositories like GenBank [3], LAPIS is not a broad database but a targeted data service. While GenBank contains sequences from more than 400,000 species and aims to provide a general and stable data source, LAPIS supports features specific to an outbreak species like lineage/clade annotation and filtering by mutations from a reference genome. In this way, LAPIS aims to support answering current research and public health questions about emerging pathogen threats.

Results

Functionalities

LAPIS implements many of the same functionalities as GenBank and additionally supports novel download, filter, and aggregation functionalities to support outbreak analysis (Table 1).

Table 1 Feature comparisons between GenBank and LAPIS

Feature	GenBank	LAPIS
Download metadata	✓	✓
Download unaligned sequences	✓	✓
Download aligned sequences	×	✓
Download protein amino acid sequences	×	✓
Download mutations	×	✓
Filter by basic metadata (country, date, etc.)	✓	✓
Filter by lineage/clade	✓	✓
Filter by mutations	×	✓
Perform aggregation	×	✓

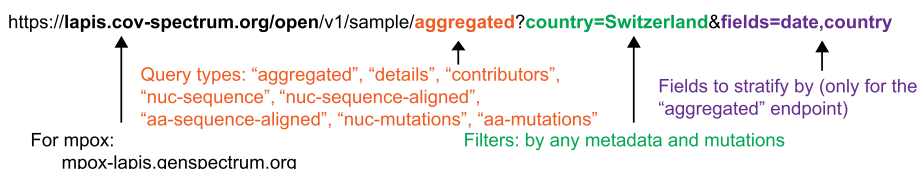


Fig. 1 Components of a query link

The simplest way to use LAPIS is to encode a query in a URL prefixed with a particular LAPIS endpoint. Each LAPIS endpoint supports a different type of query and returns a different type of data (e.g., aggregated data, sequence data, mutations, etc.). Figure 1 illustrates a URL query structure. In the following sections, we explain the different parts of a query in more detail.

Aggregation and stratification

LAPIS implements two types of endpoints: endpoints that provide aggregated data and endpoints that provide per-sample data. We describe the first type in this section and the second type in the next.

The aggregated endpoint counts the number of samples that fulfill user-defined filters in a query. If the fields parameter is not set, it returns the total number of samples. By setting fields, we can stratify the data. E.g., /aggregated?fields=pangoLineage,country will return the number of samples per Pango lineage and country. The fields parameter accepts all metadata and lineage-defining fields but not mutations or insertions.

To calculate the distribution of mutations and insertions, LAPIS offers the endpoints nuc-mutations, aa-mutations, nuc-insertions, and aa-insertions. They return the number of occurrences of mutations in a set of samples and their proportions. When calculating the proportions, the unknown or ambiguous bases are excluded. For example, if there are 10 sequences, 3 sequences have a mutation from A to G at position 5, 3 sequences have the reference base A, and 4 sequences have an N (i.e., unknown) at position 5, the proportion of the mutation A5G is $\frac{3}{6} = 0.5$ (and not $\frac{3}{10} = 0.3$).

Data download

LAPIS can also be used to obtain non-aggregated data. The `details` endpoint returns the metadata and supports an optional `fields` parameter that can be used to limit the desired metadata fields. The `nuc-sequence` and `nuc-sequence-aligned` endpoints return the original and aligned nucleotide sequences, respectively. Finally, the `aa-sequence-aligned/gene` endpoint (e.g., `aa-sequence-aligned/S` for the SARS-CoV-2 Spike protein) returns the aligned amino acid sequences.

Filters and advanced variant queries

By default (i.e., without specifying additional parameters), a query is evaluated on the whole set of sequences. To query a subset of sequences, a wide range of filters is available. It includes filtering by metadata, lineage names, and mutations. As shown in Fig. 1, filters can be set by adding request parameters to the end of the URL. If multiple filters are set, the samples that fulfill all of them will be selected.

For ordinal data like dates, there are two available filters: one with a `From`-suffix for the lower bound and one with a `To`-suffix for the upper bound. E.g., `dateFrom=2023-01-01&dateTo=2023-01-31` will filter for samples from January 2023.

LAPIS additionally supports two different ways to specify a variant. The simple approach is similar to the metadata filters and can be used to filter samples that fulfill all of a list of conditions. Possible parameters for the SARS-CoV-2 instance include `pangoLineage`, `aaMutations`, `nucMutations`, `aaInsertions`, etc. For the mutations/insertions, it is possible to use a comma-separated list. An example of a simple variant filter would be `pangoLineage=XBB.1* &aaMutations=S:E484R,S:K417T`. The `*` behind XBB.1 means all sub-lineages of the Pango lineage XBB.1 will also be included in the query. Insertion queries may contain wildcards, for instance, `ins:1000:AAT?`. This filters for all sequences with an insertion that starts with AAT between positions 1000 and 1001.

The second approach is using advanced variant queries. Advanced variant queries support more than the conjunction of a list of conditions – they also allow Boolean logic and threshold queries. One example is shown in Fig. 2. Examples of real-world, user-defined advanced variant queries can be found in the CoV-Spectrum Collections¹ where users can define and monitor sets of variants specified by advanced variant queries. In particular, the threshold queries have proven highly valuable. For example, they have been recently used to group sequences that share the same number of mutations in the receptor binding domain (RBD) [5].

“Maybe” queries

Advanced variant queries further support “maybe” queries. These queries find sequences that *might* have a certain mutation instead of definitely having a certain mutation. By default, when filtering for a mutation, LAPIS returns sequences for which the mutation is confirmed. E.g., the query `A5G` selects sequences with a G at position 5. This is a conservative way of filtering. In practice, we don’t know the values at every position of every

¹ <https://cov-spectrum.org/collections>.

(B.1.1.529* | (S:K417N & !S:954H)) & [2-of: C241T, C4321T, C9534T]

① ②⑤ ③ ⑥⑦ ③ ⑥ ⑧ ④ ④ ④

- ① Pango lineage
- ② Including sublineages (e.g., BA.1, BA.2, ...)
- ③ Amino acid mutation
- ④ Nucleotide mutation
- ⑤ Logical disjunction / OR
- ⑥ Logical conjunction / AND
- ⑦ Logical negation / NOT
- ⑧ Threshold query: at least 2 of the following expressions must be fulfilled

Fig. 2 Components of an advanced variant query

sequence: for some sequences, we might have a N (=unknown/everything is possible) or another ambiguity code that includes G such as K (=G or T) at position 5. For those samples, it is possible that, in reality, they do have the mutation A5G. This implies that the aggregated endpoint usually² provides the lower-bound number of samples when we filter for mutations. “maybe” queries allow us to obtain the corresponding upper bound. For the A5G example, sequences with a N, X, R, S, V, D and B at position 5 will also be included. Maybe queries are part of the advanced variant queries. For example, we can query `maybe(5 G) & maybe(6T)`. In fact, we can write arbitrary variant query expressions in a `maybe()` clause. Equivalent to the previous example, we can write `maybe(5 G & 6T)`. A more complex example would be `maybe((S:10K & !S:11 H) & [2-of: 100A, 101T, 102 G])`.

While the previous examples appear simple and intuitive, it is not always straightforward to determine the semantics of a maybe query. Let us consider the nucleotide sequence ATGCNT. It has one unknown at position 5. The sequence would neither match the query 5A nor 5C but it would match `maybe(5A)` and it would also match `maybe(5C)`. What’s about `maybe(5A) & maybe(5C)`? From a Boolean logic perspective, if we consider `maybe(5A)` and `maybe(5C)` to be true, then their conjunction must be true as well. On the other hand, a sequence cannot have two different bases at the same position; thus, shouldn’t `maybe(5A) & maybe(5C)` be a contradiction and unconditionally false? LAPIS would evaluate `maybe(5A) & maybe(5C)` for the aforementioned sequence to be true. The main reason we decided on this semantic is that it is possible to evaluate it efficiently.³

Performance

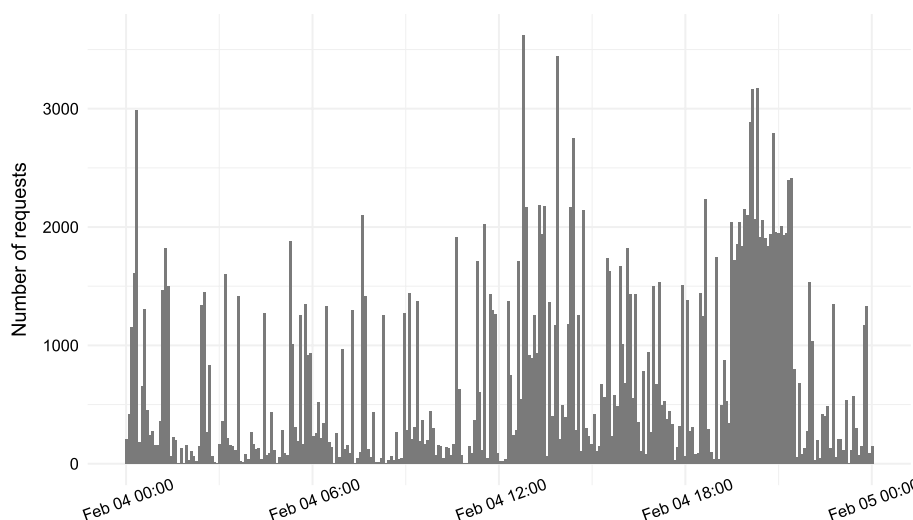
LAPIS is computationally efficient. It has proven capable of reliably processing millions of requests per day with most response times within a few hundred milliseconds as the backend to our CoV-Spectrum dashboard.

² There is an exception for queries that contain a negation.

³ Although the alternative might appear more intuitive in some cases, we believe that it cannot be efficiently computed. We have not proven it formally but we think that the complexity of the alternative semantics grows exponentially with the length of the query.

Table 2 Empirical data on the usage and performance of the endpoints

Endpoint	Number of requests	Cache hit	Response time (mean/median, in ms)
aa-insertions	34,212 (0.17%)	14.19%	654/254
aa-mutations	101,346 (0.49%)	20.86%	1050/267
aggregated	20,302,815 (98.85%)	72.66%	404/1
nuc-insertions	34,221 (0.17%)	14.20%	639/252
nuc-mutations	66,731 (0.32%)	24.32%	1203/266

**Fig. 3** Number of requests within a day. Each bar represents one minute. In total, there were 208249 requests

We currently run the LAPIS instance for SARS-CoV-2 data from GISAID on an AWS `r5.8xlarge` server (256 GB RAM, 32 vCPUs).⁴ Between 25 January⁵ and 4 February 2023, it processed over 20 million requests with a mean response time of 411 ms and a median response time of 1 ms (Table 2). This low median response time was possible because 72% of all responses had been cached (“section [Caching](#)”), which greatly reduces response time (Fig. 4). Altogether, 83% of requests to the SARS-CoV-2 instance of LAPIS were processed within 500 ms.

LAPIS often has to process many requests in parallel. It is quite common to have very few requests in one minute and over a thousand in the next (Fig. 3). The CoV-Spectrum collections are a major reason for that. In the user-defined collections, users can see information about many variants simultaneously. When a collection page is opened, the web application sends one request per variant to the server at the same time, and some collections (e.g., collection 24⁶) have hundreds of variants. When we consider only requests that were executed when the server had less than 100 parallel requests

⁴ <https://aws.amazon.com/ec2/instance-types/>.

⁵ We only started logging the response times in the afternoon of 24 January 2023.

⁶ <https://cov-spectrum.org/collections/24>.

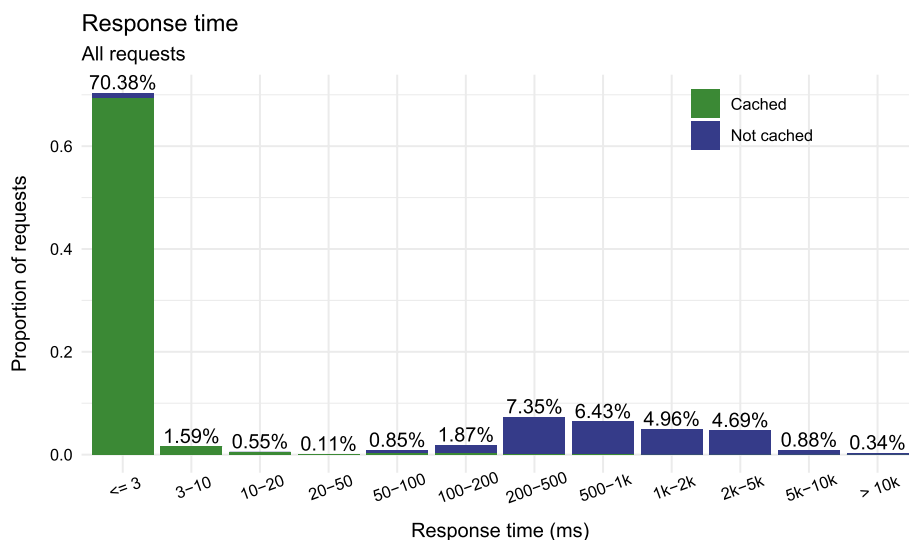


Fig. 4 Proportion of requests for the different response time bins, stratified by cache status

(that is the case for 79% of the requests), 97% of the requests were processed within 500 ms (Fig. 5).

In summary, the computational efficiency of LAPIS makes it suitable as a back-end for other tools and websites, including responsive and interactive dashboards and workflows. LAPIS achieves computational efficiency through a newly-developed data processing engine (see “section [Data query engine](#)”) that is optimized for genomic data. It can perform common operations like searching for nucleotide mutations and amino acid changes in millions of sequences and hundreds of gigabytes of data within tens to hundreds of milliseconds.

Discussion

The unique filtering, aggregation, and download functionalities supported by LAPIS, coupled with high computational efficiency, make LAPIS a key resource for the real-time analysis of genomic sequencing data from ongoing outbreaks. LAPIS is currently available for all openly accessible SARS-CoV-2⁷ and mpox⁸ sequencing data on GenBank [3]. We also maintain a private SARS-CoV-2 instance with sequencing data from GISAID [11], which serves as the backend for our CoV-Spectrum dashboard.⁹

LAPIS’ SARS-CoV-2 instances highlight the value of this approach as dataset size grows. As of January 2023, more than 14,500,000 SARS-CoV-2 sequences are available on GISAID, reaching a size of over 400 GB. LAPIS is capable of querying this entire dataset efficiently, supporting an interactive user experience on our CoV-Spectrum dashboard. CoV-Spectrum mainly presents aggregated data: it visualizes temporal, geographic, and mutational distributions of variants through a large variety of charts, tables, and maps. It solely uses LAPIS for retrieving genomic data and thanks to the flexibility

⁷ <https://lapis.cov-spectrum.org/open>.

⁸ <https://mpox-lapis.genspectrum.org>.

⁹ <https://cov-spectrum.org>.

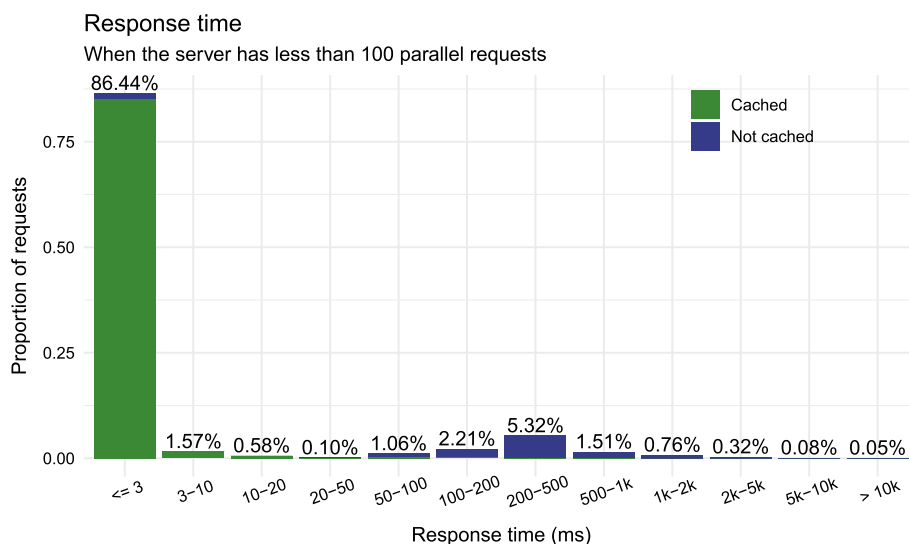


Fig. 5 Same as Fig. 4 but only for requests executed when the server has less than 100 parallel requests

of LAPIS, it was possible to develop new features in CoV-Spectrum without the need of extending or adapting LAPIS.

With LAPIS' mpox instance, we demonstrated the adaptability of the API approach. At the start of the mpox outbreak in 2022, within a few days of the release of the first sequence, we set up a LAPIS instance to support rapid sharing and easy access to open genomic data. It was accompanied by the MpoxSpectrum dashboard¹⁰ which, in addition to providing overview plots, enabled users to look up samples, download pre-processed metadata and aligned sequences, and open them in the Nextclade tool. To use the Nextclade integration feature, users can select sequences of interest on the Mpox-Spectrum dashboard and Nextclade will download the sequences from LAPIS for quality analysis. Further, just four hours after we publicized LAPIS for mpox on Twitter, Taxonomium announced the launch of a mpox service using LAPIS data as data source [17, 18].

These successes highlight that LAPIS fills a necessary role in addressing common challenges for accessing and analyzing genomic sequencing data. As demonstrated with mpox, LAPIS is easily extendable to other organisms. While supporting a new pathogen currently requires changes to the code base, we are actively working to generalize the LAPIS code to enable users to deploy instances with their own data and for other pathogens, possibly containing additional private metadata, via a configuration file. This will allow independent groups to run LAPIS instances for different use cases, akin to how Nextstrain publishes phylogenetic analyses for a limited number of pathogens but also provides the same analysis tools as an open-source resource for researchers to set up their own analyses. We hope to increase the incentive for data sharing in the public domain with this open-source philosophy: with the support of the API, researchers can directly analyze their own shared data within the global genomic context.

¹⁰ <https://mpox.genspectrum.org>.

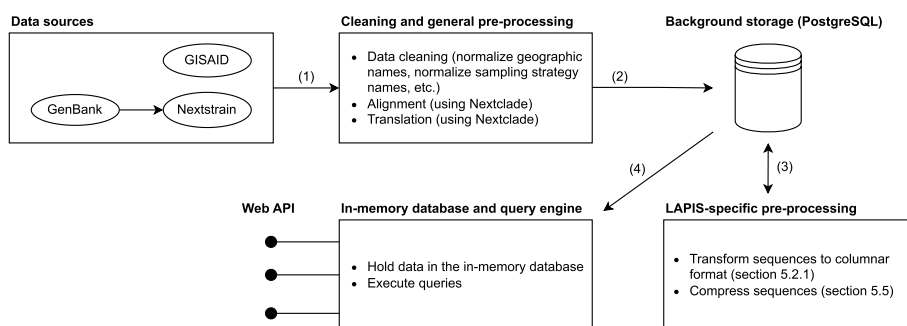


Fig. 6 Data pre-processing workflow

Going forward, we see great potential for database platforms such as GenBank to directly integrate APIs with functionalities like LAPIS’s into their framework. This avoids the necessity of hosting data in a second database and allows researchers to benefit from functionality provided by an API such as LAPIS for many different organisms. On the research side, this requires developing techniques for efficiently querying even larger genomic data sets. The current implementation of LAPIS is capable of supporting up to around 20 to 30 million sequences of length 30kbp. We are working on better algorithms to push this boundary.

Conclusions

In summary, we introduce an in-memory database engine for genomic sequencing data which can be accessed through an API. This framework facilitates the analysis of millions of sequences in real time, meaning users can interactively query and filter sequencing data. In particular, our framework supports the analysis of open genomic sequencing data and enables researchers and authorities to rapidly analyze the evolution and epidemiology of pathogens for evidence-based public health response.

Methods

Data pre-processing

For the three LAPIS instances we currently maintain, we download the raw data from GISAID (SARS-CoV-2) or Nextstrain which retrieved it from GenBank (SARS-CoV-2 and mpox). The raw data contain the genomic (consensus) sequences and corresponding metadata. We pre-process the data in two steps. During the first step, we clean up the metadata, align the sequences to a reference genome, and translate the nucleotide sequences to protein amino acid sequences. For the alignment and translation, we use Nextclade [2] but other tools are equally applicable. The first step is not specific to LAPIS and can be replaced by alternative pipelines that produce an alignment and protein sequences. During the second step, we perform LAPIS-specific transformations and generate compressed columnar sequences (“section [Column-wise storage](#)” and “[Compression](#)”). The pre-processed data are loaded into the in-memory database (“section [Data query engine](#)”) and exposed through a REST API. Figure 6 illustrates the workflow.

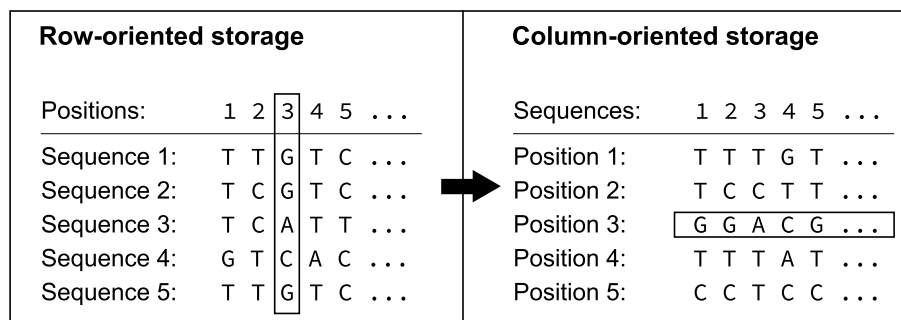


Fig. 7 Transformation of sequences to the columnar format: The row-oriented storage maintains one string per sequence; in contrast, the column-oriented storage keeps one string per position

We store the pre-processed data – both after the first and after the second step – in a PostgreSQL database. Hereby, the PostgreSQL database only serves as a background storage and can be easily replaced by the file system or a different database system. It is not crucial to the performance of LAPIS outside of the pre-processing pipeline because the in-memory database is used for evaluating the queries.

Data query engine

We developed a novel data query engine for our public web API that is tailored to support real-time, interactive genomic surveillance and genomic epidemiology. Specifically, it is designed to support high numbers of requests and fast query processing of genomic sequencing data. Our internal SARS-CoV-2 LAPIS instance based on GISAID data currently receives hundreds of thousands of requests per day, mostly from users of CoV-Spectrum. At the same time, it must support interactive and exploratory analyses where the user is able to switch quickly between different variants, countries, and time periods by responding to most requests within tens to hundreds of milliseconds. Existing database systems are not sufficient for this task.

Column-wise storage

Our approach is based on techniques developed for column-oriented database systems [1]. In the pre-processing step, we transform the sequencing data into a columnar format. For each position in the aligned nucleotide sequence or in the aligned amino acid sequence, we construct a string with the characters of all sequences at that position (Fig. 7). The i -th character in the new, columnar sequence corresponds to the sequence with the ID i . To find sequences with a mutation at a given position, we then only need to read a single string and not filter through each sequence. The columnar sequences are easy to compress (“section Compression”), and by compressing them, we can cache them in memory and eliminate any disk and round-trip time to the database.

Filter insertions

The column store as described in the previous subsection can only store the aligned sequences. It has one column for each base of the reference genome but it cannot store insertions which are parts of a sequence that cannot be directly linked to the

positions of the reference. To filter for insertions, LAPIS uses a dedicated insertion store which maintains for each position of the reference genome a mapping of inserted values to sequences with the insertion. E.g., a mapping of AATGGC at position 1000 to {sequence1, sequence2, sequence3} means that there are three sequences that have the insertion AATGGC between position 1000 and 1001.

To evaluate a query such as `ins:1000:AAT?` as described in “section [Filters and advanced variant queries](#)”, LAPIS looks up the insertions at position 1000 and matches them against the requested pattern. For the SARS-CoV-2 data, this approach works well because insertions are rare, short, and not very diverse. For genomic data with many long and diverse insertions, this method is not very efficient.

Sequence downloads

To download whole sequences, LAPIS first filters the sequences with the in-memory query engine, fetches compressed sequences (“section [Compression](#)”) from the background storage (“section [Data pre-processing](#)”), and decompresses them. If a large set of sequences should be downloaded, it fetches the sequences in small batches and streams them to the user to ensure a low memory footprint.

Discussion

The data engine was first deployed when there were around one million genomic sequences for SARS-CoV-2, and it still performs well for 15 million sequences today. It is a significant improvement to using common relational database systems which are not optimized for genomic sequencing data. The current algorithm is simple and easy to implement. However, it is also rather naive and not using state-of-the-art database engineering techniques. We are working on an improved version with reduced response times and higher throughput and look forward to sharing our results in the near future.

Data versions

To allow the user to use consistent data, each response of LAPIS contains the version of the data. The user can then check if the data versions of multiple requests are the same, and reload if that is not the case. The data version is provided in the HTTP response header `LAPIS-Data-Version`. For JSON responses, the data version is further given in the `dataVersion` field.

For example, this is relevant to compute the proportion of a variant in the sequencing data. For the calculation, we would fetch the number of sequences of the variant and the total number of sequences; that means that two API calls are required. In this case, data could be updated between the two calls which would lead to wrong results because the nominator and denominator to calculate the proportion are incompatible. Comparing the data versions of the two requests would prevent an error.

Caching

To minimize the response time for common requests, LAPIS caches the results of previously evaluated queries in a Redis database. Caching is usually a difficult task

due to the complexity to determine when a cache entry is stale. In the case of LAPIS, however, we have the advantage that we do not have a continuous stream of small data changes but perform rare (e.g., once a day) but big updates. This allows us to distinguish different versions of the data (“section [Data versions](#)”).

Each cached result is linked to a data version. If the user defines a data version in a request, and the result generated from the data of the specified version is in the cache, it can be returned immediately. If the user does not define a data version, LAPIS will check if the result for the most recent data version is cached. Figures 4 and 5 and Table 2 show the proportions of cache hits.

Compression

We compress the genome sequences before inserting them into the database. We use Zstd (level 3) [10] which gives us a good balance between compression ratio and speed. For the compression of the whole nucleotide and amino acid sequences, we use the respective reference sequence as the pre-defined dictionary to improve the compression ratio. For the columnar sequences, a pre-defined dictionary is not needed for a good compression ratio as it is intrinsically easy to compress. We achieve a compression ratio of 94% for the unaligned sequences, 99.3% for the aligned sequences, and 96% for the sequences stored in the column-oriented format.

Abbreviations

API	Application programming interface
CDC	Centers for Disease Control and Prevention
LAPIS	Lightweight API for Sequences

Acknowledgements

We acknowledge all labs for the timely submission of their SARS-CoV-2 and mpox sequencing data to the ISNDC (GenBank, ENA, and DDBJ) and GISAID. We acknowledge the Nextstrain team for pre-processing the GenBank data. These efforts are core for LAPIS to share sequencing data from outbreaks as they are generated. We thank Sarah Nadeau for valuable discussions and feedback on the manuscript. We thank François Bienvenu for valuable discussions on the “maybe” queries.

Author contributions

CC: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Visualization, Supervision, Writing—original draft. AT: Methodology, Software, Writing—review & editing. FE: Software, Writing—review & editing. JK: Software, Writing—review & editing. CR: Conceptualization, Writing—review & editing. TS: Conceptualization, Funding acquisition, Resources, Supervision, Writing—original draft. All authors read and approved the final manuscript.

Funding

Open access funding provided by Swiss Federal Institute of Technology Zurich. The authors acknowledge funding through the Swiss National Science Foundation (31CA30_196267), the Federal Office of Public Health, Switzerland (Dienstleistungsvertrag “Analysen zur Überwachung von SARS-CoV-2: Varianten-Monitoring und now-casting Hospitalisationen mit Zusatzoption Ad-hoc Analysen”), and ETH Zürich.

Availability of data and materials

The code is released under the GPL-3.0 license at <https://github.com/GenSpectrum/LAPIS>. The current study did not generate new genomic sequencing data. Part of the analyzed data is publicly available in the ISNDC (GenBank/ENA/DDBJ) repositories (<https://www.insdc.org/>). The remaining part of the analyzed SARS-CoV-2 data is available in GISAID (<https://gisaid.org/>) to which users with a GISAID account may have access.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Received: 1 August 2022 Accepted: 23 May 2023

Published online: 05 June 2023

References

- Abadi D. The design and implementation of modern column-oriented database systems. *Found Trends® Databases*. 2012;5(3):197–280.
- Aksamentov I, Roemer C, Hodcroft E, Neher R. Nextclade: clade assignment, mutation calling and quality control for viral genomes. *J Open Source Softw*. 2021;6(67):3773.
- Benson DA, Cavanaugh M, Clark K, Karsch-Mizrachi I, Ostell J, Pruitt KD, Sayers EW. GenBank. *Nucleic Acids Res*. 2017;46(D1):D41–7.
- Black A, MacCannell DR, Sibley TR, Bedford T. Ten recommendations for supporting open pathogen genomic analysis in public health. *Nat Med*. 2020;26(6):832–41.
- Callaway E. COVID 'variant soup' is making winter surges hard to predict. *Nature*. 2022;611(7935):213–4.
- Centers for Disease Control and Prevention. Cdc covid data tracker (2023).
- Chen C, Nadeau S, Topolsky I, Beerenwinkel N, Stadler T. Advancing genomic epidemiology by addressing the bioinformatics bottleneck: challenges, design principles, and a swiss example. *Epidemics*. 2022;39: 100576.
- Chen C, Nadeau S, Yared M, Voinov P, Xie N, Roemer C, Stadler T. CoV-spectrum: analysis of globally shared SARS-CoV-2 data to identify and characterize new variants. *Bioinformatics*. 2021;38(6):1735–7.
- Chen C, Nadeau SA, Topolsky I, Manceau M, Huisman JS, Jablonski KP, Fuhrmann L, Dreifuss D, Jahn K, Beckmann C, Redondo M, Noppen C, Risch L, Risch M, Wohlwend N, Kas S, Bodmer T, Roloff T, Stange M, Egli A, Eckerle I, Kaiser L, Denes R, Feldkamp M, Nissen I, Santacroce N, Burcklen E, Aquino C, de Gouvea AC, Moccia MD, Grüter S, Sykes T, Opitz L, White G, Neff L, Popovic D, Patrignani A, Tracy J, Schlapbach R, Dermizakis ET, Harshman K, Xenarios I, Pegeot H, Cerutti L, Penet D, Blin A, Elies M, Althaus CL, Beisel C, Beerenwinkel N, Ackermann M, Stadler T. Quantification of the spread of SARS-CoV-2 variant B.1.1.7 in Switzerland. *Epidemics*. 2021;37: 100480.
- Collet Y. Rfc 8878: Zstandard compression and the 'application/zstd' media type (2021).
- Elbe S, Buckland-Merrett G. Data, disease and diplomacy: GISAID's innovative contribution to global health. *Global Chall*. 2017;1(1):33–46.
- Gangavarapu K, Latif AA, Mullen JL, Alkuzweny M, Hufbauer E, Tsueng G, Haag E, Zeller M, Aceves CM, Zaiets K, Cano M, Zhou X, Qian Z, Sattler R, Matteson NL, Levy JI, Lee RTC, Freitas L, Maurer-Stroh S, GISAID Core and Curation Team, Suchard MA, Wu C, Su AI, Andersen KG, Hughes LD. Outbreak.info genomic reports: scalable and dynamic surveillance of SARS-CoV-2 variants and mutations. *Nat Methods*. 2023.
- Hodcroft EB. Covariants: Sars-cov-2 mutations and variants of interest (2021).
- Hodcroft EB, Maio ND, Lanfear R, MacCannell DR, Minh BQ, Schmidt HA, Stamatakis A, Goldman N, Dessimoz C. Want to track pandemic variants faster? Fix the bioinformatics bottleneck. *Nature*. 2021;591(7848):30–3.
- Knyazev S, Chhugani K, Sarwal V, Ayyala R, Singh H, Karthikeyan S, Deshpande D, Baykal PI, Comarova Z, Lu A, Porozov Y, Vasylyeva TI, Wertheim JO, Tierney BT, Chiu CY, Sun R, Wu A, Abedalthagafi MS, Pak VM, Nagaraj SH, Smith AL, Skums P, Pasaniuc B, Komissarov A, Mason CE, Bortz E, Lemey P, Kondrashov F, Beerenwinkel N, Lam TT-Y, Wu NC, Zelikovsky A, Knight R, Crandall KA, Mangul S. Unlocking capacities of genomics for the COVID-19 response and future pandemics. *Nat Methods*. 2022.
- Li J, Lai S, Gao GF, Shi W. The emergence, genomic diversity and global spread of SARS-CoV-2. *Nature*. 2021;600(7889):408–18.
- Sanderson T. Taxonium, a web-based tool for exploring large phylogenetic trees. *eLife*. 2022;11:e82392.
- Sanderson, T. Tweet (2022).
- Tegally H, Moir M, Everatt J, Giovanetti M, Scheepers C, Wilkinson E, Subramoney K, Moyo S, Amoako DG, Baxter C, Althaus CL, Anyaneji UJ, Kekana D, Viana R, Giandhari J, Lessells RJ, Maponga T, Maruapula D, Choga W, Matshaba M, Mayaphi S, Mbhele N, Mbulawa MB, Msomi N, Naidoo Y, Pillay S, Sanko TJ, San JE, Scott L, Singh L, Magini NA, Smith-Lawrence P, Stevens W, Dor G, Tshiabuila D, Wolter N, Preiser W, Treurnicht FK, Venter M, Davids M, Chiloane G, Mendes A, McIntyre C, O'Toole A, Ruis C, Peacock TP, Roemer C, Williamson C, Pybus OG, Bhiman J, Glass A, Martin DP, Rambaut A, Gaseitsiwe S, von Gottberg A, de Oliveira T. Continued emergence and evolution of omicron in South Africa: new BA.4 and BA.5 lineages. 2022.
- Viana R, Moyo S, Amoako DG, Tegally H, Scheepers C, Althaus CL, Anyaneji UJ, Bester PA, Boni MF, Chand M, Choga WT, Colquhoun R, Davids M, Deforche K, Doolabh D, du Plessis L, Engelbrecht S, Everatt J, Giandhari J, Giovanetti M, Hardie D, Hill V, Hsiao N-Y, Iranzadeh A, Ismail A, Joseph C, Joseph R, Koopile L, Pond SLK, Kraemer MUG, Kuate-Lere L, Laguda-Akingba O, Lesetedi-Mafoko O, Lessells RJ, Lockman S, Lucaci AG, Maharaj A, Mahlangu B, Maponga T, Mahlakuwe K, Makatini Z, Marais G, Maruapula D, Masupu K, Matshaba M, Mayaphi S, Mbhele N, Mbulawa MB, Mendes A, Misana K, Mnguni A, Mohale T, Moir M, Moruisi K, Mosepele M, Motsatsi G, Motswaledi MS, Mphoyakgosi T, Msomi N, Mwangi PN, Naidoo Y, Ntuli N, Nyaga M, Olubayo L, Pillay S, Radibe B, Ramphal Y, Ramphal U, San JE, Scott L, Shapiro R, Singh L, Smith-Lawrence P, Stevens W, Strydom A, Subramoney K, Tebeila N, Tshiabuila D, Tsui J, van Wyk S, Weaver S, Wibmer CK, Wilkinson E, Wolter N, Zarebski AE, Zuze B, Goedhals D, Preiser W, Treurnicht F, Venter M, Williamson C, Pybus OG, Bhiman J, Glass A, Martin DP, Rambaut A, Gaseitsiwe S, von Gottberg A, de Oliveira T. Rapid epidemic expansion of the SARS-CoV-2 omicron variant in Southern Africa. *Nature*. 2022;603(7902):679–86.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.