

Global Benchmark Database

Ashlin Iser   

Karlsruhe Institute of Technology, Germany

Christoph Jabs   

HIIT, University of Helsinki, Finland

Abstract

This paper presents Global Benchmark Database (GBD), a comprehensive suite of tools for provisioning and sustainably maintaining benchmark instances and their metadata. The availability of benchmark metadata is essential for many tasks in empirical research, e.g., for the data-driven compilation of benchmarks, the domain-specific analysis of runtime experiments, or the instance-specific selection of solvers. In this paper, we introduce the data model of GBD as well as its interfaces and provide examples of how to interact with them. We also demonstrate the integration of custom data sources and explain how to extend GBD with additional problem domains, instance formats and feature extractors.

2012 ACM Subject Classification Theory of computation → Logic; Theory of computation → Design and analysis of algorithms; Information systems → Information integration

Keywords and phrases Maintenance and Distribution of Benchmark Instances and their Features

Digital Object Identifier 10.4230/LIPIcs.SAT.2024.18

Supplementary Material *Software (GBD Source Code)*: <https://github.com/Udopia/gbd> [17]

archived at [swh:1:dir:37fb54c7420fd71b6c4930bcf7e78d5129325a6d](https://swh.1.dir:37fb54c7420fd71b6c4930bcf7e78d5129325a6d)

Software: <https://github.com/Udopia/gbdc> [16]

archived at [swh:1:dir:705258b1bd5c3415069fbd156759945a8ad725c0](https://swh.1.dir:705258b1bd5c3415069fbd156759945a8ad725c0)

Software: <https://github.com/Udopia/gbdeval> [15]

archived at [swh:1:dir:159fae7ac54d9f678ee38ed882dae693cd5f14cf](https://swh.1.dir:159fae7ac54d9f678ee38ed882dae693cd5f14cf)

Dataset: <https://github.com/Udopia/gbd-data> [14]

archived at [swh:1:dir:1764eef19455a0bdfdb6389a04a055977338304ff](https://swh.1.dir:1764eef19455a0bdfdb6389a04a055977338304ff)

Funding *Christoph Jabs*: Work financially supported by Research Council of Finland grant 356046.

Acknowledgements Many thanks to all the people who have supported GBD, either by contributing code or data, or by providing infrastructure. Special thanks also go to the early adopters whose successes and feedback have been and continue to be inspiring and motivating.

1 Introduction

The idea to create Global Benchmark Database (GBD) arose from the need to make available and sustainably maintain benchmark instances and their metadata for the problem of propositional satisfiability (SAT). To this end, we specified a simple hash function for *instance identification* to serve as the primary key for SAT instance data sets. A proof of concept was presented and discussed at the 2018 Pragmatics of SAT (POS) workshop, demonstrating the specification of a SAT instance identifier and methods for labeling and querying for instances [19]. From there, GBD has matured into a comprehensive suite of tools for provisioning and sustainably maintaining benchmark instances of various problem domains and their metadata. Supported problem domains include propositional satisfiability (SAT), maximum satisfiability (MaxSAT), and pseudo-Boolean optimization (PBO).

In essence, the purpose of GBD is to act as a conduit between data science and empirical research on NP-hard problem classes by facilitating the seamless integration of benchmark data into existing workflows. GBD provides tools for distributing benchmark instances and



© Ashlin Iser and Christoph Jabs;

licensed under Creative Commons License CC-BY 4.0

27th International Conference on Theory and Applications of Satisfiability Testing (SAT 2024).

Editors: Supratik Chakraborty and Jie-Hong Roland Jiang; Article No. 18; pp. 18:1–18:10

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

feature databases, as well as tools for transforming instances and extracting instance features. It also provides a set of prebuilt feature databases for identifying instance equivalence classes, categories, or labels, as well as analytical instance features. Examples of successful applications of GBD include the sanitization and selection of benchmarks for SAT competitions [13], domain-specific solver evaluations [8], and the analysis of solver portfolios and solver prediction models [2]. Moreover, the authors of the latest award-winning state-of-the-art SAT solvers use GBD in their empirical evaluations. [11, 25]. GBD is open source and available under the MIT license [18].

In this paper, we describe the conceptual design of the GBD data model and present the interfaces for interacting with GBD, as well as the extension capabilities provided by the data model, including concrete examples of their use. We start with a brief description of the GBD data model and query language in Section 2. Section 3 presents the GBD interfaces and gives concrete examples of how to use GBD in practice. Finally, Section 4 explains how to extend GBD with additional problem domains, instance formats, and feature extractors. We conclude with a summary and an outlook on future work in Section 5.

2 Data Model and Queries

Conceptually, GBD consists of an extensible set of contexts $\mathcal{C} = \{C_0, C_1, \dots\}$. Each context represents a problem domain, or more specifically, a particular instance format of a problem domain. A context is defined as a tuple $C = (\mathcal{X}, \text{id})$, where \mathcal{X} is the set of benchmark instances and $\text{id} : \mathcal{X} \rightarrow \text{PK}$ is a context-specific instance identification function that maps each benchmark instance to an instance identifier unique within the context. Contexts are instantiated with data sources $D = \{\mathcal{F}_1, \mathcal{F}_2, \dots\}$ which provide instance features $\mathcal{F}_i \subseteq \text{PK} \times 2^{\Sigma^*}$ assigning values to benchmark instances, where values are represented by strings over an alphabet Σ .¹

Feature Types

In practice, GBD distinguishes between one-to-one and one-to-many features. One-to-one features are defined as those that provide a single value per instance, with a default value associated with them. In contrast, one-to-many features can have zero to any number of values per instance. One-to-one features facilitate the initial setup and subsequent maintenance, as they are initialized to the default value automatically and subsequent calls to `gbd set` overwrite any previous settings. In contrast, the values of one-to-many features accumulate over time and must be explicitly cleared. Example 1 illustrates the practical use of the different feature types in the context of `cnf` instances.

► **Example 1.** Let `cnf` := (\mathcal{X}, id) be the context where \mathcal{X} is the set of SAT benchmark instances in DIMACS CNF format, and let a data source `meta.db` := $\{\text{track}, \text{result}, \dots\}$ be given, which provides the `cnf` instance features `track` and `result`. In the example, the feature `track` indicates the competition tracks in which the instance was used and is modeled as a one-to-many feature, while `result` indicates the satisfiability of the instance and is defined as a one-to-one feature with the default value “unknown”.

¹ The data sources used in the examples of this section are available at <https://benchmark-database.de>.

Context Mappings

For any two contexts $C_i = (\mathcal{X}_i, \text{id}_i)$, $C_j = (\mathcal{X}_j, \text{id}_j)$ with instance identifiers PK_i, PK_j , a data source can provide a special feature $\mathcal{F} \subseteq \text{PK}_i \times \text{PK}_j$ that establishes a relationship between instances in \mathcal{X}_i and \mathcal{X}_j . In GBD, such features are referred to as context mappings. Context mappings can be used to map instances from different problem domains to each other if there is a reduction procedure that transforms instances from one domain to instances in the other domain. Example 2 illustrates the practical use of context mappings to relate instances of the propositional satisfiability problem to instances of the k -independent set problem.

► **Example 2.** Given the context of `cnf` instances of the SAT problem and the context `kis` of graph-instances of the k -Independent Set problem. Any SAT instance can be transformed to a graph and a number k in such a way that the satisfiability of the SAT instance is equivalent to the existence of an independent set of size k in the graph [23]. A context mapping feature can be used to map instances in `cnf` to their corresponding instances in `kis`.

Context mappings are useful not only for relating instances from different problem domains, but also for relating instances from different contexts within the same problem domain. Application scenarios include relating different instance formats, analyzing different encodings of application instances, or distinguishing sanitized or otherwise preprocessed instances from their original counterparts (cf. Use case 9 below).

Queries

GBD provides a query language designed to filter instances of a given context by specifying constraints over the instance features. A GBD query is a simple constraint c , or a compound constraint c_1 and c_2 , or c_1 or c_2 for sub-constraints c_1, c_2 , and parentheses can be used to indicate precedence. A simple constraint is of the form $f \circ e$ with a feature name f , a value e , and an operator $\circ \in \{=, !=, <, >, <=, >=, \text{like}, \text{unlike}\}$ with the usual semantics. For SQL-inspired operators $\circ \in \{\text{like}, \text{unlike}\}$, a constraint can also be of the form $f \circ \%e$, $f \circ e\%$, or $f \circ \%e\%$ to indicate postfix, prefix, or infix constraints, respectively. Value e can also be a term, where a term is either a constant number, or a feature name in parentheses (f') ², or a compound term $(t_1 \circ t_2)$ from arithmetic operators $\circ \in \{+, -, *, /\}$ and terms t_1, t_2 .

► **Example 3.** Given the `cnf` context and the data source `meta.db` providing the instance feature `track`, a query can be of the form `track = main_2023` to filter for instances from the Main track of the 2023 SAT competition.

► **Example 4.** Given the `cnf` context and the data source `meta.db` providing the instance features `track` and `filename`, a query can be of the form `track = anni_2022 and filename like rphp%` to filter for instances from the Anniversary track of the 2022 SAT competition whose filenames begin with the string `rphp`.

► **Example 5.** Given the `cnf` context and the data source `base.db` providing the instance features `variables` and `clauses`, a query can be of the form `variables > (clauses)` to filter for instances with more variables than clauses.

² Parentheses are used here to disambiguate terms and character strings.

3 Applications

This section introduces the `gbd-tools` package, which is available in the Python Package Index (PyPI) and can be installed using the command `pip install gbd-tools`. The package provides the command line tool `gbd`, a web service, and the Python interface class `GBD`. We start with instructions on how to configure `gbd-tools` in Section 3.1, and then present use cases of the command line tool `gbd` in Section 3.2. The web service and Python interface are presented in Sections 3.3 and 3.4, respectively.

3.1 Data Source Configuration

Data sources for instantiating GBD are specified as a list of files, where each file can be either a `sqlite3` database created by GBD, or a `csv` file for importing data from other sources. The requirements for the `csv` files are that they have a header line containing the feature names, and they must provide the `hash` column containing the instance identifiers. Example 6 shows a typical example of integrating two data sources before analyzing runtime experiments.

► **Example 6.** Let GBD be instantiated with the two data sources `meta.db` and `runtimes.csv`. The file `meta.db` provides the feature `family`, indicating the domain of each instance. The file `runtimes.csv` provides the solver runtimes `baseline` and `incumbent`. The instance identifier in the `hash` column facilitates the integration of the two data sources and thus the domain-wise analysis of experimental results.

Data sources can be registered by setting the environment variable `GBD_DB` to a colon-separated list of paths. It is also possible to specify or override the data sources with the option `-d/--db` on each call of the `gbd` command-line tool. If a data source does not exist, GBD offers to create it. Ready-made feature databases are available for download from <https://benchmark-database.de>. The `gbd info` command can be used to display the registered data sources, their names, and the features provided. The name of a data source is automatically generated from its filename by removing the extension, and is needed in some parameters and queries to explicitly refer to a specific data source (cf. Use cases 3 and 7).

Each data source is associated with a context, so a data source can only provide features for exactly one context. Each context is identified by a name, e.g., `opb` for the context of pseudo-Boolean optimization instances. The set of available contexts, including their names and descriptions, can be displayed with the command `gbd info -c`. To bind a data source to a context, its filename is prefixed with the context name to which the data belongs (cf. Use cases 9 and 10). Otherwise, the data source is treated as belonging to the default context `cnf`.

When features from different contexts are queried simultaneously, a context mapping feature is used to create a relational join between the instance identifiers of the different contexts. The naming convention for context mapping features is `to_{cxt}`, where `cxt` is a context name. The `to_{cxt}` feature maps the identifiers of the context of the data source in which it resides to the identifiers of the context named by `cxt`. GBD recognizes context mapping features according to this naming convention and generates the appropriate foreign key relationships to automatically join features from different contexts.

3.2 Command-Line Tool `gbd`

In this section, we present the command line tool `gbd` and its subcommands for initializing GBD with benchmark instances, querying for instances and features, extracting features from instances, and transforming instances. The full set of subcommands can be viewed with

`gbd --help`. In the use cases presented in this section, we assume that the `cnf` data sources `meta.db`³ and `mylocal.db` (to be initialized in the following) are configured as data sources by the environment variable `GBD_DB`, and that `/path/to/instances` contains the benchmark instances from SAT Competition 2023.⁴

3.2.1 Database Initialization

In order to use GBD to organize your own experiments, it is necessary for GBD to know where the locally available benchmark instances are located. Integrating a set of paths to benchmark instances into GBD requires computing the instance identifiers and storing the instance paths in a GBD data source. In this process, GBD creates the reserved features `local` and `filename` to associate the local paths and filenames with the respective identifiers. Note that projecting to an instance identifier automatically eliminates duplicates. The initialization process may require a considerable amount of time, even for a moderate number of instances, as is the case with SAT competition benchmark sets. However, it is only required once for the registration of benchmarks.

Use case 1 illustrates initializing GBD with a number of benchmark instances present on the local volume using the `gbd init local` command. The `-j/--jobs` option is used to set the number of parallel jobs, the `--target` option is used to specify the name of the database in which to create the features, and the parameter after the local subcommand is used to specify the path to the benchmark instances. The `gbd init local` command recursively searches the specified directory for benchmark instances using the file extensions associated with the context of the target database.

■ **Use case 1** Local database initialization.

```
gbd init -j16 --target mylocal local /path/to/instances
```

3.2.2 Queries for Instances and Features

The `gbd get` command is used to filter for specific instances and retrieve their features. A comprehensive set of options can be displayed with `gbd get --help`. Filtering is typically done with GBD queries, and a list of features to be returned is given by the `-r/--resolve` option. Use case 2 illustrates how to filter for cryptographic instances in the SAT Competition 2023 benchmark set and return their `local` paths.

■ **Use case 2** Filtering and feature resolution.

```
gbd get "track=main_2023 and family like crypto%" -r local
```

Feature names from different databases may overlap in some configurations. In this case, GBD picks the feature values from the first database in the list that provides the feature. To override this behavior, feature names can be prefixed with the database name, separated by a colon. Use case 3 shows how to filter by the feature `track` from the database `meta` and select the feature `local` from the database `mylocal`.

■ **Use case 3** Explicit feature database.

```
gbd get "meta:track=main_2023" -r mylocal:local
```

³ <https://benchmark-database.de/getdatabase/meta>

⁴ https://benchmark-database.de/?track=main_2023

To control the handling of multiple values per hash, the `-c/--collapse` option can be used, as shown in Use case 4. The command collapses the values of the feature `local` to a single value by using the minimum value.

■ **Use case 4** Collapse.

```
gbd get "track=main_2023" -r mylocal:local -c min
```

For grouping instances by a feature other than the instance identifier, the option `-g/--group` can be used as illustrated in Use case 5. The command groups the instances by the feature `isohash` and selects a local path for each group using the minimum function.

■ **Use case 5** Grouping.

```
gbd get "track=main_2023" -r local -c min -g isohash
```

3.2.3 Manual Data Acquisition

In addition to subcommands for creating, deleting, renaming, and copying features, the `gbd` command line tool provides a subcommand for manually setting feature values. Use case 6 demonstrates the use of the `gbd set` command to set the value of the `family` feature to `hardware-verification` for all instances whose filename starts with the string `manol`.

■ **Use case 6** Setting feature values.

```
gbd set family=hardware-verification "filename like manol%"
```

3.2.4 Instance Feature Extraction

Feature extractors are accessible through subcommands of `gbd init`, which can be listed with `gbd init --help`. Each feature extractor is bound to a set of contexts and can only be executed on instances from those contexts. Multiple feature extractors are provided by the `gbdc` extension module which are documented at <https://udopia.github.io/gbdc>. Note that ready-made feature databases are available for download from our instance of the GBD web interface (cf. Section 3.3).

Use case 7 illustrates the use of the subcommand `gbd init base` to extract default features denoted as `base` from a set of `cnf` benchmark instances. In the example, the option `-j/--jobs` is used to specify the number of parallel jobs to use for the feature extraction. The option `--target` is used to specify the name of the database in which to create the features. This requires that a database with the name `mybase` is registered as a data source. The parameter after the subcommand `base` is used to specify the query for the instances to extract features from.

■ **Use case 7** Base feature extraction.

```
gbd init -j16 --target mybase base "track=main_2023"
```

It is often necessary to determine which instances are isomorphic to each other. The `isohash` feature is a hash value that over-approximates the class of isomorphic instances and can be used to identify and group isomorphic instances and to subsequently eliminate them from benchmark sets (cf. Use case 5). Use case 8 illustrates how to use the `gbd init isohash` command to compute `isohash` for a set of `cnf` benchmark instances.

■ **Use case 8** Isohash calculation.

```
gbd init -j16 --target meta isohash "track=main_2023"
```

3.2.5 Instance Transformation

GBD also provides instance transformers accessible via subcommands of `gbd transform`. A complete list of available instance transformers can be displayed with `gbd transform --help`. Instance transformers are functions that transform an instance of one context into another instance of a different context. Additionally, an instance transformer automatically creates a context mapping that relates the instances in both contexts, and can store additional features in the database. Instance transformers are implemented as a special type of feature extractor that, in addition to creating instance features, also creates a new benchmark instance in the target context. Several instance transformers are provided by the `gbdc` extension module and are documented in the `gbdc` documentation.

Use case 9 illustrates the use of the `sanitize` transformer for `cnf` instances. The command transforms the instances in the source context `cnf` to instances in the target context `sancnf` and stores their mapping and local path in the database `sancnf_local.db`.

■ **Use case 9** Instance sanitizer.

```
gbd transform --source cnf --target sancnf_local sanitize
```

Note that `sancnf` is an example of a context that exists only to distinguish between sanitized and unsanitized instances and to use the context mapping feature to check if identifiers have changed. In most practical use cases, the sanitized instances will be used to replace the unsanitized instances in the original `cnf` context.

An example of a real transformation between different problem domains is the transformation of `cnf` instances into k -Independent Set (`kis`) instances, as described in Example 2. Use case 10 illustrates the use of the `cnf2kis` transformer to transform the instances in the source context `cnf` into instances in the target context `kis`. The mapping of the instances in both contexts is stored in the database `kis_local.db`.

■ **Use case 10** Transforming CNF to KIS.

```
gbd transform --source cnf --target kis_local cnf2kis
```

3.3 Web Interface

The GBD web interface is implemented as a RESTful web service [7] that provides access to the benchmarks and databases of the environment it runs in, and can be started with the command `gbd serve`. While the GBD command-line tool operates independently of the web service, the web service uses the same GBD API and database configuration as the command-line tool. Our instance of the GBD web interface, accessible via <https://benchmark-database.de>, is configured to run behind an Nginx reverse proxy [24] and is hosted in a Docker container [21]. At the time of writing, we provide access to more than 100 000 benchmark instances from the `cnf`, `wcnf`, and `opb` contexts, with prebuilt feature databases for them.

3.4 Python Interface

The GBD Python API is wrapped in the class `gbd_core.api.GBD` and its documentation can be found at <https://udopia.github.io/gbd>. With the Python API, GBD data sources can be directly integrated into Python scripts and Jupyter notebooks to be used in evaluation and analysis. Use case 11 illustrates the use of the Python API to query for instances and features. Line 2 creates a GBD object that connects to the GBD data sources `base` and `meta`,

which are passed to the constructor as a list of paths. Line 3 creates a list of feature names from the `base` data source and adds the feature family from the `meta` data source. Line 4 queries for the given features of `cnf` instances from the Main track of SAT Competition 2023. The result is returned as a Pandas [20] DataFrame for further analysis.⁵

■ **Use case 11** GBD Python interface.

```
1 from gbd_core.api import GBD
2 with GBD(['gbd/meta.db', 'gbd/base.db']) as gbd:
3     feat = gbd.get_features('base') + ['family']
4     df = gbd.query("track=main_2023", resolve=feat)
```

4 Writing Extensions

The architectural design of GBD allows the integration of new contexts, feature extractors, and instance transformers. The extensibility of the system is achieved through the use of dictionaries, which serve as registries for contexts, feature extractors, and instance transformers. The following section describes the elements required to create such a registry entry. It is important to note that at the time of writing, these registries are hard-coded, which is a limitation in that it requires modification of these dictionaries within the source code of GBD. This issue will be addressed in future versions of GBD, where the registry will be moved to configuration files.

GBD contexts are managed in a dictionary data structure that is initialized in the `gbd_core.contexts` module. The dictionary is indexed with the respective *context name* and contains a pointer to the instance *identification function* and a list of valid instance *filename extensions*. The instance identification function is the function that assigns a unique instance identifier to each benchmark instance, which serves as the primary key in the context-specific feature databases. The list of valid instance filename extensions is used to locate benchmark instances in the local file system during the database initialization process. Once a context is created in the dictionary, context-specific databases can be populated using the context name.

A new feature extractor can be integrated into GBD by registering a feature extractor function in a dictionary data structure that is initialized in the `gbd_init.feature_extractors` module. The dictionary is indexed by the *name* of the feature extractor and contains a *list of the provided features*, which are tuples of feature names and default values. It also contains a pointer to the feature extractor *function* and a list of *contexts* in which it can be used. Registered feature extractors are automatically accessible by their names as subcommands of `gbd init`.

Instance transformers are managed in a dictionary data structure that is initialized in the `gbd_init.instance_transformers` module. An instance transformer function takes an instance in a *source context* and transforms it to a new instance in the *target context*. The dictionary is indexed by the *name* of the instance transformer and, similarly to feature extractors, contains a *list of the provided features*. Additionally, the dictionary contains a pointer to the instance transformer *function itself*, as well as a pointer to a function that generates the *name of the benchmark instance* to be created. Registered instance transformers are automatically accessible by their names as subcommands of `gbd transform`.

⁵ We provide illustrative examples of benchmark data analyses, including portfolio analysis, category prediction, and category-specific ranking, accessible via <https://udopia.github.io/gbdeval/>.

5 Related and Future Work

SatLib was the first public collection of benchmark instances in the SAT problem domain [12], and similar collections have been created for other problem domains such as MaxSAT [1], Quantified Boolean Formulas (QBF) [9], SAT Modulo Theories (SMT) [4], and Mixed Integer Programming (MIP) [10]. SatEx is the first web-based framework for reproducible execution and evaluation of SAT solver experiments [26], followed by the EDACC framework [3]. Instance features were used to predict the fastest solver for an instance [27] and to reduce redundancy in experiments [22]. Aslib is a library of data sets for training and evaluating solver prediction models [5]. Meta-features such as the instance domain have been used to study specialized heuristic configurations [6].

GBD’s data model, which is centered around instance identification functions, enables the data-driven study of NP-hard problem domains in an unprecedentedly sustainable manner. Future work includes the integration of more problem domains, more feature extractors, and more instance transformers. We also plan to improve the configuration capabilities of the data sources, including the ability to easily switch between different configurations, and to improve the extensibility of GBD by making it a matter of changing a configuration file whenever new problem domains, instance formats, feature extractors, and instance transformers are added. We also plan to improve the web interface to make it easier to use and more informative, and to provide a benchmark submission system. Finally, we plan to automate the feature extraction process to instantly provide feature databases for new benchmark instances, and add support for instance generators.

The purpose of coupling and providing instance data is to enable more complex data-driven analyses of the instance spaces of hard algorithmic problems, thus accelerating the growing use of explainable artificial intelligence methods. Future work in this area includes the application of machine-assisted hypothesis generation and testing, and the development of new methods for exploring, analyzing, and explaining algorithmic datasets.

References

- 1 Fahiem Bacchus. MaxSAT Formulas Library, 2006. URL: <http://www.cs.toronto.edu/maxsat-lib/maxsat-instances/>.
- 2 Jakob Bach, Ashlin Iser, and Klemens Böhm. A Comprehensive Study of k-Portfolios of Recent SAT Solvers. In *Intl. Conf. on Theo. and Appl. of Satisf. Test., SAT*, pages 1–18, 2022. doi:10.4230/LIPIcs.SAT.2022.2.
- 3 Adrian Balint, Daniel Gall, Gregor Kapler, and Robert Retz. Experiment Design and Administration for Computer Clusters for SAT Solvers, EDACC. *J. Satisf. Bool. Model. Comput.*, pages 77–82, 2010. doi:10.3233/SAT190078.
- 4 Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The Satisfiability Modulo Theories Library, SMT-LIB, 2016. URL: www.smt-lib.org.
- 5 Bernd Bischl, Pascal Kerschke, Lars Kotthoff, et al. ASlib: A Benchmark Library for Algorithm Selection. *Artif. Intel.*, pages 41–58, 2016. doi:10.1016/j.artint.2016.04.003.
- 6 Jan Elffers, Jesús Giráldez-Cru, Stephan Gocht, Jakob Nordström, and Laurent Simon. Seeking Practical CDCL Insights from Theoretical SAT Benchmarks. In *Intl. Joint Conf. on Artif. Intel., IJCAI*, pages 1300–1308, 2018. doi:10.24963/ijcai.2018/181.
- 7 Roy Thomas Fielding. *REST: Architectural Styles and the Design of Network-based Software Architectures*. PhD thesis, University of California, 2000. URL: <http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm>.
- 8 Nils Froleyks, Marijn Heule, Ashlin Iser, Matti Jarvisalo, and Martin Suda. SAT Competition 2020. *Artif. Intel.*, 2021. doi:10.1016/j.artint.2021.103572.

- 9 E. Giunchiglia, M. Narizzano, L. Pulina, and A. Tacchella. Quantified Boolean Formulas Library, QBFLIB, 2005. URL: www.qbflib.org.
- 10 Ambros Gleixner, Gregor Hendel, Gerald Gamrath, et al. MIPLIB 2017: Data-driven Compilation of the 6th Mixed-Integer Programming Library. *Math. Program. Comput.*, pages 443–490, 2021. doi:10.1007/s12532-020-00194-3.
- 11 Andrew Haberlandt, Harrison Green, and Marijn J. H. Heule. Effective Auxiliary Variables via Structured Reencoding. In *Intl. Conf. on Theo. and Appl. of Satisf. Test., SAT*, pages 1–19, 2023. doi:10.4230/LIPICS.SAT.2023.11.
- 12 Holger Hoos and Thomas Stützle. SATLIB: An Online Resource for Research on SAT. *SAT 2000*, pages 283–292, 2000. URL: <https://api.semanticscholar.org/CorpusID:17486963>.
- 13 Ashlin Iser. Benchmark Compilation for SAT Competition 2023. In *Proc. of SAT Comp. 2023: Solver, Benchmark and Proof Checker Descriptions*, 2023. URL: <http://hdl.handle.net/10138/563824>.
- 14 Ashlin Iser and Christoph Jabs. GBD Data Repository. Dataset, version 1.0., swId: swh:1:dir:1764eef19455a0bfdb6389a04a055977338304ff (visited on 2024-08-02). URL: <https://github.com/Udopia/gbd-data>.
- 15 Ashlin Iser and Christoph Jabs. GBD Evaluation Scripts. Software, version 1.0., swId: swh:1:dir:159fae7ac54d9f678ee38ed882dae693cd5f14cf (visited on 2024-08-02). URL: <https://github.com/Udopia/gbdeval>.
- 16 Ashlin Iser and Christoph Jabs. Global Benchmark Database (Extension Module). Software, version 1.0., swId: swh:1:dir:705258b1bd5c3415069fbd156759945a8ad725c0 (visited on 2024-08-02). URL: <https://github.com/Udopia/gbdc>.
- 17 Ashlin Iser and Christoph Jabs. Udopia/gbd. Software, version 4.8.5., swId: swh:1:dir:37fb54c7420fd71b6c4930bcf7e78d5129325a6d (visited on 2024-08-02). URL: <https://github.com/Udopia/gbd>.
- 18 Ashlin Iser and Christoph Jabs. GBD Tools, 2024. doi:10.5281/zenodo.11093597.
- 19 Ashlin Iser and Carsten Sinz. A Problem Metadata Library for Research in SAT. In *Proceedings of Pragmatics of SAT, POS*, pages 144–152, 2018. doi:10.29007/gdbb.
- 20 Wes McKinney et al. Data Structures for Statistical Computing in Python. In *Python in Science*, pages 51–56, 2010. doi:10.25080/majora-92bf1922-00a.
- 21 Dirk Merkel. Docker: lightweight linux containers for consistent development and deployment. *Linux J.*, 2014.
- 22 Sibylle Möhle and Norbert Manthey. Better Evaluations by Analyzing Benchmark Structure. In *Pragmatics of SAT workshop, POS*, 2016. URL: http://www.pragmaticsofsat.org/2016/reg/POS-16_paper_4.pdf.
- 23 Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- 24 Will Reese. Nginx: the high-performance web server and reverse proxy. *Linux J.*, 2008.
- 25 Dominik Pascal Schreiber. *Scalable SAT Solving and its Application*. PhD thesis, Karlsruhe Institute of Technology, KIT, 2023. doi:10.5445/IR/1000165224.
- 26 Laurent Simon and Philippe Chatalic. SatEx: A Web-based Framework for SAT Experimentation. *Electron. Notes Discrete Math.*, pages 129–149, 2001. doi:10.1016/S1571-0653(04)00318-X.
- 27 Lin Xu, Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. SATzilla: Portfolio-based Algorithm Selection for SAT. *J. Artif. Intel. Res., JAIR*, pages 565–606, 2008. doi:10.5555/1622673.1622687.