

Ethos: A Fast Proof Checker for the Eunoia Logical Framework

Andrew Reynolds¹ , Hans-Jörg Schurr^{1,3,4} , Mallku Soldevila⁵ ,
Haniel Barbosa⁵ , Clark Barrett² , and Cesare Tinelli¹ 

¹ The University of Iowa, Iowa City, USA

² Stanford University, Stanford, USA

³ KU Leuven, Leuven, Belgium

⁴ Vrije Universiteit Brussel, Brussels, Belgium

⁵ Universidade Federal de Minas Gerais, Belo Horizonte, Brazil

Abstract. SMT solvers are used in many safety-critical applications. To provide evidence of the correctness of their answers, some SMT solvers generate externally checkable proof certificates. We present a high-performance checker for SMT proof certificates called *Ethos*. In contrast with other dedicated SMT proof checkers, *Ethos* does not implement a fixed proof calculus. Instead, it allows users to specify their own calculus in the declarative language Eunoia, which extends the familiar SMT-LIB syntax to make that easy and convenient. We give a short overview of Eunoia and then focus on *Ethos* itself. We describe multiple optimization and implementation details which make *Ethos* fast and practical. We also evaluate *Ethos* on proofs generated by *cvc5*, showing that the flexibility of *Ethos* allows us to efficiently check fine-grained proofs, containing no proof holes, over all SMT-LIB logics without floating point arithmetic.

Keywords: SMT · proof checking · logical frameworks

1 Introduction

Satisfiability Modulo Theories (SMT) solvers reason about logical problems expressed in terms of a wide variety of theories (arithmetic, bit-vectors, strings, and so on). They are used as back ends for software verification tools as well as a variety of other safety-critical applications (e.g., [7]). A key concern for these applications is that SMT solvers are large and complex systems. As such, they are at risk of containing bugs that affect their correctness.

A pragmatic but rigorous approach for addressing this concern is for an SMT solver to produce a *proof certificate*, or just *proof* for short, for each its answers. This is typically a textual artifact, expressing a formal proof of the answer’s correctness, which can be checked by an external tool, a *proof checker* [1, 17]. A proof accepted by a proof checker is a strong piece of evidence that the solver’s answer is correct. The reason is that, since proof checking is in general considerably simpler than proof finding, which is ultimately what SMT solvers do, even a proof checker that is not formally verified can strengthen the user’s

trust in the SMT solver’s answer. This approach requires the definition of a format for SMT proofs, which is a serious challenge as different SMT solvers, or even different theory solvers within the same SMT solver, may use vastly different proof systems, or *calculi*, due to differences in the theories they support and in the preprocessing and reasoning methods they implement.

In this paper, we present *Ethos*, a generic proof checker for SMT. Instead of implementing a fixed calculus, *Ethos* supports the *Eunoia* logical framework, which is designed to make specifying proof systems and proofs in them simple and convenient for SMT developers. *Eunoia*’s syntax extends that of Version 2.7 of the SMT-LIB standard [5] with commands for (i) defining theory *signatures*, i.e., the collection of types, constants, and proof rules that constitute a proof calculus, and (ii) expressing proofs themselves. Given a *Eunoia* file that provides both a signature and a proof, *Ethos* checks that the proof correctly applies the rules in the signature. It can also verify that the assumptions used in a proof in fact correspond to the assertions in a specific SMT-LIB problem, thus ensuring that the proof is not just a correct proof but a proof of the correct theorem.

While *Ethos* is not formally verified, it significantly reduces the trusted code base. For example, the source code of the *cvc5* SMT solver [3] consists of over 300K lines of C++, while *Ethos* is less than 10K lines of C++. The signature for *cvc5* proofs is expressed in 6,397 lines of *Eunoia* code. Strong runtime performance was a main goal when implementing *Ethos*, with the aim of reducing the cost, and thus incentivizing the use, of proof checking. To this end, *Ethos* implements multiple optimization techniques, which we discuss below.

Related Work. The main inspiration for *Ethos* comes from two related systems. The first is *Alethe* [15], a proof format and calculus for SMT supported by the Carcara checker [1]. In contrast to *Eunoia*, *Alethe* defines a fixed calculus, described in English in a reference document [4]. The SMT solver *SMTInterpol* [10] uses a custom proof format similar to *Alethe*. The second inspiration is *LFSC*, also a logical framework, and its proof checker [17]. In contrast to *Eunoia*, its syntax does not resemble SMT-LIB and represents proofs as terms. Based on our experience, SMT developers find it unintuitive and difficult to use, which may have hindered its wide adoption.

Both *LFSC* and *Eunoia* are inspired by other logical frameworks, starting with *Edinburgh LF* [8]. Among them, *Dedukti* [2] stands out as explicitly designed for proof exchange. Multiple calculi for SMT solvers and other automated reasoners have been modeled in *Dedukti* [6]. *RARE* is a domain specific language for declaring proofs based on rewriting [11]. *Eunoia* is partially inspired by it and can be seen as a generalization of its functionality to arbitrary proof rules.

Beyond those mentioned so far, there are a number of additional approaches for defining proof formats. *DRAT* [19] is a unified proof format for propositional logic that enables SAT solvers to easily generate proofs for most state-of-the-art solving techniques in a manner that can be efficiently checked. As a result, most contemporary SAT solvers support *DRAT*. Unfortunately, due to the heterogeneous nature of SMT solving, it seems doubtful that a format with a fixed proof calculus could provide similar standardization benefits for SMT. The *eDRAT*

format [9] uses DRAT together with unjustified theory lemmas, which must be checked separately. It does not support proofs for preprocessing.

The TPTP world [18] provides a proof format for automated theorem provers (ATPs). However, it does not prescribe a way to define proof rules, and theorem provers choose their own. A notable consequence of this design choice is that, since TPTP proofs rules can be quite high level, a trusted ATP, as opposed to a simpler checker, is needed to prove individual rule applications correct.

2 Eunoia

In this section, we provide an overview of Eunoia through an example containing a signature and a proof in that signature. The Ethos user manual [13] contains a complete introduction to Eunoia.

Types and Constants. A Eunoia signature starts with a declaration of relevant types and constants (in the HOL sense), as shown in the following. The constants `Bool`, `true`, `false`, and `->` are not declared because they are built-in.

```

1 (declare-const Int Type) (declare-consts <numeral> Int)          theory.eo
2 (declare-const not (-> Bool Bool))
3 (declare-const and (-> Bool Bool Bool) :right-assoc-nil true)
4 (declare-parameterized-const = ((A Type :implicit)) (-> A A Bool))
5 (declare-const BitVec (-> Int Type))
6 (declare-parameterized-const concat ((m Int :implicit) (n Int :implicit))
7   (-> (BitVec n) (BitVec m) (BitVec (eo::add m n))))

```

Eunoia contains several built-in syntactic categories. It is up to the user though to give them a type. In the example above, the `declare-consts` command (line 1) tells Ethos that every (signed) numeral (`0`, `1`, `-1`, ...) has type `Int`. This makes it possible to change the type of numerals depending on the SMT-LIB logic of interest. For example, in the `QF_LIA` logic, numerals have type `Int`, while in `QF_LRA` they have type `Real`. The defined symbols do not carry any semantic information, beyond their type. They are simply constructors of the term language on which the proof rules operate.

The `:right-assoc-nil` annotation (line 3) for the constant `and` is an extension of SMT-LIB's `:right-assoc` annotation. It indicates that `and` can be used as a variadic operator with `true` as a neutral (or *nil*) element.⁶ Like `:right-assoc`, it allows applications of the operator to more than two arguments as syntactic sugar for a nested application. However, it also allows applications of the form `(and a)`, treated as a shorthand for `(and a true)`. This implies that `(and a b c)` is parsed as `(and a (and b (and c true)))` and not `(and a (and b c))`, as would be the case with the `:right-assoc` annotation. One motivation for `:right-assoc-nil` is that, in contrast with the `:right-assoc` annotation, it distinguishes between terms like `(and a b c)` and `(and a (and b c))`, as the latter is parsed as `(and a (and (and b (and c true)) true))`. As we will see, Eunoia has additional features to facilitate working with such operators.

⁶ Eunoia also supports the specular annotations `:left-assoc` and `:left-assoc-nil`.

In Eunoia, it is also possible to define dependent types and parametric constants. In the example above, we first declare the usual polymorphic equality constant with an implicit type argument `A`. Afterwards, we declare a size-indexed type for bit-vectors and the bit-vector concatenation constant. The return type of `concat` (line 6) is defined by a computation performed by the built-in operator `eo::add` for adding two numbers. Eunoia has a large library of such operators. Some, like `eo::add`, are primitive operators that implement fundamental functionality; others are defined as Eunoia programs built from primitive operators.

Proof Rules. After declaring the *theory signature* we can declare a *proof signature*. That consists of proof rules and associated helper programs which can be used to implement side conditions. A file inclusion command allows us to isolate the theory signature in its own file. Proof rules follow the theory signature.

```

1 (include "theory.eo")                                     rules.eo
2 (declare-rule contra ((F Bool)) :premises (F (not F)) :conclusion false)
3 (program pick ((F Bool) (Fs Bool :list) (n Int)) :signature (Bool Int) Bool
4   (((pick      F 0)   F)
5    ((pick (and F Fs) n) (pick Fs (eo::add n -1)))))
6 (declare-rule andE ((Fs Bool) (i Int))
7   :premises (Fs) :args (i) :conclusion (pick Fs i))

```

Every `declare-rule` command takes a rule name (e.g., `contra`, and `andE` in lines 2, 6) and lists parameters used in the declaration. The `:premises` attribute lists patterns for the rule’s premises. To check a proof step using rule `contra`, the proof checker applies pattern matching to two previously derived formulas to find a matching substitution for the parameter `F`. This substitution is then applied to the conclusion, which in general is also a pattern, although not in this case. Rule `andE` illustrates the use of computations in proofs. The rule implements the elimination of variadic applications of `and` using a recursive program, `pick` (lines 3–5), that returns the n -th argument of a conjunctive formula.

Like rules, programs have a name and a list of *local* parameters that serve as pattern variables. They also have a type signature that declares the argument types (`Bool` and `Int` for `pick`) and the return type (`Bool`). As in many functional programming languages, a program is defined by a list of cases that are considered in order. Abstractly, each case is a pair (p, b) , where p lists a term pattern for each program input and b is the case’s body. Each local parameter in b must appear at least once in p (non-linear patterns are supported). A program application is evaluated by a Eunoia checker by matching the application’s arguments against each pattern list. When this succeeds, it applies the matching substitution to the body, evaluates it, and returns the resulting term. If no case matches, the checker raises an error and rejects the entire proof.

The example also demonstrates another feature of Eunoia’s handling of variadic operators: the `:list` annotation. If an operator is annotated with `:list`, its last argument is not handled as syntactic sugar. Without it, the pattern `(and F Fs)` would be parsed as `to (and F (and Fs true))` which, however, would match only with binary conjunctions—of the form `(and t1 t2)`.

Proofs. We can construct a proof with the rules seen so far. With the `reference` command a proof can be linked to an SMT-LIB problem (e.g., `problem.smt2`) to check that all declared constants and assumptions appear in the problem.

```
(reference "problem.smt2") (include "rules.eo")
(declare-const F Bool)
(assume a1 (and F (not F)))
(step s1 (F) :rule andE :premises (a1) :args (0))
(step s2 (not F) :rule andE :premises (a1) :args (1))
(step s3 (false) :rule contra :premises (s1 s2))
```

The proof above has three derivation steps. The `assume` command introduces the assumption `(and F (not F))` and names it `a1`. The next two commands apply rule `andE` to `a1`, deriving formulas `F` and `(not F)`, respectively. The final command concludes `false` from the formulas derived by the previous two steps. The second argument of `step` is its expected conclusion and is optional. If provided, as in this case, it is matched during proof checking against the actual (computed) conclusion, generating an error in case of a mismatch. This is useful for debugging.

3 Proof Checking With Ethos

In this section, we sketch a few details on the implementation of Ethos. We first focus on a description of the general process used by Ethos to check proofs, and then discuss some low-level aspects.

Proof Checking Model. Ethos is designed to be, first and foremost, a *fast* proof checker. This means that it omits some checks that might be *convenient* for the signature programmer but are not *essential*, i.e., their absence cannot lead to Ethos accepting invalid proofs. For example, Ethos does not fully check the types of programs when they are defined. Instead, it checks the monomorphic types of the ground terms that are generated when checking a proof.

For Ethos, checking a proof amounts to evaluating a proof term. Every rule is in fact a program with only one case. For example, the `contra` rule from above is internally translated to the following program.

```
(declare-const Proof Type)
(declare-const pf (-> Bool Proof))
(program contra ((T Bool)) :signature (Proof Proof) Proof
  ((contra (pf T) (pf (not T))) (pf false)))
```

The earlier proof corresponds roughly to the term `(contra (andE A 0) (andE A 1))`, where `A` is the assumption `(and F (not F))`. This term is considered a correct proof if Ethos can evaluate it to `(pf false)` without error. Note that we use the symbols `Proof` and `pf` here for illustration purposes; they are not built-in in Eunoia.

Ethos implements a function `eval(t, σ)`, where t is a Eunoia term and σ is a grounding substitution over the parameters of t . It evaluates t bottom up, replacing each parameter x in it with $\sigma(x)$. When `eval` encounters a term `(c u1 ... un)`, where c is a program with cases $[(p_1, b_1), \dots, (p_m, b_m)]$, it invokes

a helper `match`($p_i, u_1 \dots u_n$) on each pattern p_i of c in order. If this function succeeds, it returns a substitution σ' , and Ethos evaluates `eval`(b_i, σ').

To ensure the conclusion of proof steps are not mistyped, Ethos can compute the type of parameter-free terms. This task is easier than general type checking because for the theories considered by Eunoia, all literals and constants have monomorphic types. Hence, computing types of applications amounts to pattern matching between the domain types and the monomorphic argument types. This approach has the benefit of being easy to implement, and optimize, at the expense of somewhat reduced statically checked guarantees.

Implementation. Ethos is written in C++, and currently has fewer than 10,000 lines of code. Terms are stored using a data structure with hash-consing (each unique AST is allocated exactly once) and reference counting on expressions. Number-like SMT-LIB literals (numerals, rationals, bit-vector constants) are stored in the AST using GMP arbitrary precision datatypes. Most built-in Eunoia operators on these literals translate directly to GMP functions. Ethos also implements the SMT-LIB policy for Unicode strings.

To accelerate proof checking, Ethos caches evaluation results. As mentioned, the evaluator takes as arguments a term to evaluate and a substitution, or *context*, mapping its free parameters to ground terms. The evaluator is implemented as a directed-acyclic graph traversal over terms. Since the evaluation primitives in Eunoia are stateless, Ethos can cache intermediate results in this process. This means that the result of invoking a side condition on a particular set of arguments is only computed once. However, this caching is not global, meaning that invoking the same computation on two independent proof steps will result in computing the result again. Ethos also evaluates parameter-free terms in rule and program bodies during parsing.

Since term construction is often the bottleneck, the core evaluation primitives on lists are optimized at a low level to minimize the number of terms that must be reallocated. For example, `eo::list_erase_all`, which removes all occurrences of a particular element from a list, will reuse the tail of the list after the last occurrence of the element. To ensure the correctness of the optimized implementation, we cross-check these optimizations against a simpler reference implementation of the list operators as ordinary Eunoia programs. The Ethos regression set contains tests that ensure that the built-in evaluation matches the results of these reference implementations.

Finally, Ethos can be extended with plugins. These are modules that register callbacks with the main loop of Ethos. The callbacks notify the plugin when an event occurs, e.g., when a proof step is parsed. The plugin system can be used to implement custom external proof checkers, or converters to other proof formats.

4 Evaluation

The `cvc5` solver supports three proof checking pipelines. It can produce proofs in the Alethe format which can be checked by Carcara. It can produce proofs in a calculus modeled in LFSC and checkable by the LFSC checker. Finally, it can

Table 1. Evaluation on benchmark sets (1) and (2). Proofs are either fully checked (**Verified**), or have holes (**Holey**). For each case, the table shows the sum of the time in seconds to **Generate** a proof and to **Check** it. The last three columns count failures: no proof was generated, checking timed out, some other checking error occurred.

	Verified	Gen.	Check	Holey	Gen.	Check	no Gen.	TO	Other
Alethe	19,502	90,908s	13,869s	867	7,172s	1,192s	23	26	522
LFSC	454	36s	42s	19,275	81,459s	223,989s	23	126	1,062
CPC	20,813	85,868s	93,157s	0	0s	0s	23	43	61
LFSC	705	36s	34s	136,482	229,300s	221,911s	468	460	2,133
CPC	138,912	657,911s	219,267s	0	0s	0s	473	617	247

produce proofs in its native calculus, called CPC, which is modeled in Eunoia and can be checked by Ethos. However, the support for *cvc5* features differs significantly in these three pipelines. The Alethe proof format only supports uninterpreted functions, linear arithmetic, and quantifiers, and the LFSC calculus omits rules for many steps performed by *cvc5*. In those cases, the LFSC proof contains a step marked as *trust*. The Eunoia-based pipeline is the most complete. If *cvc5* is run in *safe* mode, intended for users that require high assurance, the generated CPC proofs contains no *holes* (unjustified steps). An important factor in making this possible was the ease of writing proof rules in Eunoia.

To provide a perspective on the performance of Ethos, we compare the different proof production and checking pipelines of *cvc5*, and show that the Ethos pipeline is competitive with the other two, while also supporting more features (Table 1). We used all non-incremental SMT-LIB benchmarks [12] without floating-points that do not have status *sat* or have an *inferredStatus* of *sat* in the SMT-LIB catalog [14, 16]. To get a collection of benchmarks *cvc5* could solve, we ran it for 60s in *safe mode* with proof generation off.⁷ This gave us a set of 161,188 benchmarks. We split this benchmarks set in two: (1) benchmarks in logics supported by Alethe, and (2) the rest. Set (1) contains 20,940 benchmarks. We then ran *cvc5* for 600s on these sets with proofs turned on, and afterwards each checker on the generated proofs for another 600s. For Eunoia and Alethe, we used fine-grained proof rules for term rewriting steps, which are not supported by LFSC. This is the main reason for the large number of proof holes in LFSC proofs. Support for such rewrite rules is also experimental in Alethe. On commonly solved problems, Ethos is 6.25 times slower than Carcara.

As we gain experience using Ethos in practice, we expect to find more opportunities to improve its performance and usability. We also plan to focus on the support infrastructure around Ethos to make Eunoia easier to use. This includes the development of a standardized unit testing framework for Eunoia signatures, a documentation generator for proof rules, and a language server. One of our goals is to make Eunoia and Ethos robust and flexible enough to be a useful proof checking infrastructure for other SMT solvers and automated theorem provers.

⁷ We used machines with Intel Xeon E5-2620 v4 CPUs and an 8 GiB memory limit.

Acknowledgments. We thank the anonymous reviewers for their feedback on this paper. This work was supported in part by the Stanford Center for Automated Reasoning, a gift from Amazon Web Services, the Defense Advanced Research Projects Agency (DARPA) under contract FA8750-24-2-1001, the European Union (ERC, CertiFOX, 101122653), and Fonds Wetenschappelijk Onderzoek — Vlaanderen (project G070521N). Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of DARPA, the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Andreotti, B., Lachnitt, H., Barbosa, H.: Carcara: An efficient proof checker and elaborator for SMT proofs in the alethe format. In: Sankaranarayanan, S., Sharygina, N. (eds.) *Tools and Algorithms for Construction and Analysis of Systems (TACAS)*, Part I. *Lecture Notes in Computer Science*, vol. 13993, pp. 367–386. Springer (2023). https://doi.org/10.1007/978-3-031-30823-9_19
2. Assaf, A., Burel, G., Cauderlier, R., Delahaye, D., Dowek, G., Dubois, C., Gilbert, F., Halmagrand, P., Hermant, O., Saillard, R.: Dedukti: a logical framework based on the λI -calculus modulo theory. *CoRR* **abs/2311.07185** (2023). <https://doi.org/10.48550/ARXIV.2311.07185>
3. Barbosa, H., Barrett, C.W., Brain, M., Kremer, G., Lachnitt, H., Mann, M., Mohamed, A., Mohamed, M., Niemetz, A., Nötzli, A., Ozdemir, A., Preiner, M., Reynolds, A., Sheng, Y., Tinelli, C., Zohar, Y.: cvc5: A versatile and industrial-strength SMT solver. In: Fisman, D., Rosu, G. (eds.) *TACAS 2022*. LNCS, vol. 13243, pp. 415–442. Springer (2022). https://doi.org/10.1007/978-3-030-99524-9_24
4. Barbosa, H., Fleury, M., Fontaine, P., Schurr, H.J.: The Alethe proof format: An evolving specification and reference (2026), <https://verit.gitlabpages.uliege.be/alethe/specification.pdf>
5. Barrett, C., Fontaine, P., Tinelli, C.: The Satisfiability Modulo Theories Library (SMT-LIB) (2016), <https://www.smt-lib.org>
6. Blanqui, F.: Translating libraries of definitions and theorems between proof systems (invited talk). In: Bertot, Y., Kutsia, T., Norrish, M. (eds.) *15th International Conference on Interactive Theorem Proving, ITP 2024, Tbilisi, Georgia, September 9-14, 2024*. LIPIcs, vol. 309, pp. 2:1–2:1. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2024), <https://doi.org/10.4230/LIPIcs.ITP.2024.2>
7. Bouchet, M., Cook, B., Cutler, B., Druzkina, A., Gacek, A., Hadarean, L., Jhala, R., Marshall, B., Peebles, D., Rungta, N., Schlesinger, C., Stephens, C., Varming, C., Warfield, A.: Block public access: trust safety verification of access control policies. In: Devanbu, P., Cohen, M.B., Zimmermann, T. (eds.) *ESEC/FSE ’20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020*. pp. 281–291. ACM (2020). <https://doi.org/10.1145/3368089.3409728>
8. Harper, R., Honsell, F., Plotkin, G.D.: A framework for defining logics. *J. ACM* **40**(1), 143–184 (1993). <https://doi.org/10.1145/138027.138060>

9. Hitarth, S., Codel, C.R., Lachnitt, H., Dutertre, B.: Extending DRAT to SMT. In: Narodytska, N., Rümmer, P. (eds.) *Formal Methods in Computer-Aided Design, FMCAD 2024*, Prague, Czech Republic, October 15-18, 2024. pp. 1–11. IEEE (2024). https://doi.org/10.34727/2024/ISBN.978-3-85448-065-5_8
10. Hoenicke, J., Schindler, T.: A simple proof format for SMT. In: Déharbe, D., Hyvärinen, A.E.J. (eds.) *Proceedings of the 20th Internal Workshop on Satisfiability Modulo Theories co-located with the 11th International Joint Conference on Automated Reasoning (IJCAR 2022) part of the 8th Federated Logic Conference (FLoC 2022)*, Haifa, Israel, August 11-12, 2022. *CEUR Workshop Proceedings*, vol. 3185, pp. 54–70. CEUR-WS.org (2022), <https://ceur-ws.org/Vol-3185/paper9527.pdf>
11. Nötzli, A., Barbosa, H., Niemetz, A., Preiner, M., Reynolds, A., Barrett, C.W., Tinelli, C.: Reconstructing fine-grained proofs of rewrites using a domain-specific language. In: Griggio, A., Rungta, N. (eds.) *22nd Formal Methods in Computer-Aided Design, FMCAD 2022*, Trento, Italy, October 17-21, 2022. pp. 65–74. IEEE (2022). https://doi.org/10.34727/2022/ISBN.978-3-85448-053-2_12
12. Preiner, M., Schurr, H.J., Barrett, C., Fontaine, P., Niemetz, A., Tinelli, C.: SMT-LIB release 2025 (non-incremental benchmarks) (Aug 2025). <https://doi.org/10.5281/zenodo.16740866>
13. Reynolds, A., Schurr, H.J., Soldevila, M., Tinelli, C.: The Ethos user manual (2026), https://github.com/cvc5/ethos/blob/main/user_manual.md
14. Schurr, H., Bobot, F., Preiner, M., Niemetz, A., Barrett, C.W., Fontaine, P., Tinelli, C.: Exploring the SMT-LIB benchmark library. In: Junges, S., Katz, G. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems - 32nd International Conference, TACAS 2026, Held as Part of the International Joint Conferences on Theory and Practice of Software, ETAPS 2026*, Turin, Italy, April 11-16, 2026, *Proceedings, Part I*. pp. 150–169. *Lecture Notes in Computer Science*, Springer (2026). https://doi.org/10.1007/978-3-032-22752-2_8
15. Schurr, H., Fleury, M., Barbosa, H., Fontaine, P.: Alethe: Towards a generic SMT proof format (extended abstract). In: Keller, C., Fleury, M. (eds.) *Proceedings Seventh Workshop on Proof eXchange for Theorem Proving, PxTP 2021*, Pittsburg, PA, USA, July 11, 2021. *EPTCS*, vol. 336, pp. 49–54 (2021). <https://doi.org/10.4204/EPTCS.336.6>
16. Schurr, H.J., Preiner, M., Niemetz, A., Barrett, C., Fontaine, P., Tinelli, C.: SMT-LIB catalog 2025 (Jul 2025). <https://doi.org/10.5281/zenodo.16290040>
17. Stump, A., Oe, D., Reynolds, A., Hadarean, L., Tinelli, C.: SMT proof checking using a logical framework. *Formal Methods Syst. Des.* **42**(1), 91–118 (2013). <https://doi.org/10.1007/S10703-012-0163-3>
18. Sutcliffe, G.: Stepping stones in the TPTP world. In: Benz Müller, C., Heule, M.J.H., Schmidt, R.A. (eds.) *Automated Reasoning - 12th International Joint Conference, IJCAR 2024*, Nancy, France, July 3-6, 2024, *Proceedings, Part I*. *Lecture Notes in Computer Science*, vol. 14739, pp. 30–50. Springer (2024). https://doi.org/10.1007/978-3-031-63498-7_3
19. Wetzler, N., Heule, M.J.H., Hunt, W.A.: DRAT-trim: Efficient Checking and Trimming Using Expressive Clausal Proofs. In: Sinz, C., Egly, U. (eds.) *Theory and Applications of Satisfiability Testing (SAT)*. *Lecture Notes in Computer Science*, vol. 8561, pp. 422–429. Springer (2014)