The ILTP Problem Library for Intuitionistic Logic

Release v1.1

Thomas Raths*, Jens Otten and Christoph Kreitz
Institut für Informatik, University of Potsdam, Germany.
e-mail: {raths,jeotten,kreitz}@cs.uni-potsdam.de

Abstract. The Intuitionistic Logic Theorem Proving (ILTP) library provides a platform for testing and benchmarking automated theorem proving (ATP) systems for propositional and first-order intuitionistic logic. It includes about 2800 problems in a standardized syntax from 24 problem domains collected from various sources that are appropriate for intuitionistic logic. For each problem intuitionistic status and difficulty rating were obtained by running comprehensive tests of currently available intuitionistic ATP systems on all problems in the library. Thus for the first time the testing and evaluation of intuitionistic ATP systems is put onto a firm basis.

Keywords: ILTP, problem library, benchmarking, experimental evaluation, ATP, intuitionistic logic

1. Introduction

Benchmarking automated theorem proving (ATP) systems using standardized problem sets is a well-established method for measuring their performance. Several such problem libraries have been developed, e.g. the TPTP library for classical first-order logic [18], libraries for (propositional) satisfiability problems, e.g. SATLIB [9], and libraries for termination and induction problems.[1]

The main purpose of such libraries is to put the testing and evaluation of ATP systems onto a firm basis. The practice of using a few specific problems, possibly even tailored towards a particular ATP system, does not result in meaningful data and therefore does not reflect the actual performance of an ATP system. Using a common standardized and comprehensive problem library yields meaningful results and ensures fair system comparisons.

Unfortunately the availability of such libraries for non-classical logics, like intuitionistic or modal first-order logics, is very limited. The aim of the Intuitionistic Logic Theorem Proving (ILTP) library is to close

* The first author's research is sponsored by DARPA under agreement number FA8750-04-2-0216.

[1] For the termination problem library see `http://www.lri.fr/~marche/tpdb/`, for the induction problem libraries see `http://dream.dai.ed.ac.uk/dc/lib.html` and `http://www.cs.nott.ac.uk/~lad/research/challenges/`.

© 2006 *Kluwer Academic Publishers. Printed in the Netherlands.*

this gap for the intuitionistic first-order and propositional logic. The ILTP library fulfils the following main requirements of a benchmark library as described in [18]:

- it is easy to discover and obtain and provides guidelines for its use in evaluating ATP systems; the ILTP library is widely available via the Internet (see Subsection 1.1).
- it is well structured, documented and provides statistics about the library as a whole (see Section 2).
- it is easy to use; the problems are provided in an easy to understand format and conversion tools to other known syntax formats are included (see Subsections 2.3 and 2.4).
- it is large enough for statistically significant testing; the current release v1.1 of the ILTP library contains about 2800 problems.
- it contains problems of varying difficulty; the ILTP library includes problems from simple to unsolved/open problems (see tables in Subsection 2.2).
- it assigns each problem an unambiguous name (see Subsection 2.3) and provides status and difficulty rating for each problem; comprehensive tests of existing intuitionistic ATP systems yielded *intuitionistic* status and rating information (see Subsection 2.2).

This paper will describe the release v1.1 of the ILTP library. It starts by providing information on how to obtain and use the ILTP library, followed by an overview of previous benchmark collections and a description of changes made with respect to release v1.0 of the ILTP library [15]. Section 2 gives detailed information about the contents of the ILTP library, which is divided into a first-order and a propositional part. It describes the domain structure, the presentation form of the problems and presents statistics about intuitionistic status and difficulty rating. Some details of included tools are given as well. The paper concludes in Section 3 with a summary and an outlook on further research.

1.1. Obtaining and Using the ILTP Library

The ILTP library is available at `http://www.iltp.de` . It consists of two parts: a first-order part and a purely propositional part. Both parts of the ILTP library can be downloaded separately. The library is structured into four subdirectories:

`Axioms` - contains the axiom files (only in the first-order part).
`Documents` - contains papers and statistic files.
`Problems` - contains a directory for each domain with problem files.
`TPTP2X` - contains the tptp2X tool and the format files.

The following conditions should be observed when presenting results of ATP systems based on the ILTP library:

- The release number of the ILTP library has to be stated.
- Each problem should be referred to by its unambiguous name.
- No part of the problem may be modified; no reordering of axioms, hypotheses and/or conjectures is allowed. Only the syntax of problems may be changed, e.g., by using the tptp2X tool (see Subsection 2.4).
- The header information of each problem may not be exploited by an ATP system.
- The version of the tested ATP system including all settings must be documented.

It is a good practice to make at least the binaries/executable of an ATP system available whenever performance results or statistics based on the ILTP library are given. This makes the verification and validation of the given performance data possible.

1.2. Previous Problem Collections

For intuitionistic logic several small collections of problems have been published and used for testing existing ATP systems. Sahlin et al. [16] compiled one of the first collections of first-order problems for testing their intuitionistic ATP system ft. The same collection was also used for benchmarking other intuitionistic ATP systems [19, 11]. A second collection of first-order problems was used to test the intuitionistic ATP system JProver [17], which has been integrated into the constructive interactive proof assistants NuPRL [1] and Coq [4].

A collection of propositional problems was compiled by Dyckhoff.[2] It introduces six classes of scalable problems following the methodology of the Logics Workbench [3]. The advantage of this approach is the possibility to study the time complexity behaviour of an ATP system on a specific generic problems as its size increases. But in order to achieve

[2] See `http://www.dcs.st-and.ac.uk/~rd/logic/marks.html`.

more meaningful benchmark results the number of generic problems would have to be increased significantly. Most of the problems in the collection have a rather syntactical nature, often specifically designed with the presence (or absence) of a specific search strategy in mind. To provide a better view of the usefulness of intuitionistic ATP systems on problems arising in practice, like in program synthesis [1], a benchmark collection should cover a broader range of more realistic problems. These kind of problems are typically presented in a first-order logic as already mentioned in Dyckhoff's benchmark collection.

For classical logic the well-known TPTP library [18] provides a large collection of first-order problems — currently more than 8000 — for testing and benchmarking ATP systems for classical logic. Whereas the semantics of classical and intuitionistic logic differs, they share the same syntax. This allows in principle the use of classical benchmark libraries like the TPTP library for benchmarking intuitionistic ATP systems as well.

The TPTP library started as a library of first-order formulas in clausal form [18] and the majority of problems are still in clausal form. All problems in clausal form, i.e. disjunctive or conjunctive normal form, are intuitionistically invalid and therefore useless for intuitionistic reasoning. Furthermore, the conversion of formulas to clausal form does not preserve intuitionistic validity, because it involves intuitionistically invalid laws like $\neg(A \wedge B) \Rightarrow (\neg A \vee \neg B)$ and $\neg\neg A \Rightarrow A$. Adding double negation to classically valid formulas in order to generate intuitionistically valid formulas is of less interest since the resulting problems are just encodings of the classical ones.

Today the TPTP library contains a large number of problems in non-clausal form as well. These problems form a major part of the ILTP library.

1.3. A Short History of the ILTP Library

First experimental evaluations of the intuitionistic ATP systems JProver [17] and ileanTAP[11] on a large set of suitable problems from the TPTP library were conducted in June 2004. Since then the number of problems as well as the number of tested intuitionistic ATP systems was continuously increased. The first public release v1.0 of the ILTP library [15] went online in April 2005. The most recent version is release v1.1 and came out in January 2006. These are the major changes and enhancements that were done compared to release v1.0 of the ILTP library:

- Due to the substantial interest in developing ATP systems for intuitionistic *propositional* logic, the library was split up into a propositional part and a first-order part.

- The number of problems in the library has almost doubled from 1445 problems in the first release to 2754 problems. Three new problem domains have been added.

- The library now includes all non-clausal (FOF) problems of the TPTP library release v3.1.0. The first release of the ILTP library was mainly based on release v2.7.0 of the TPTP library. The format files were adapted to the new TPTP syntax.

- The number of intuitionistic ATP systems considered for determining the intuitionistic status and rating information was increased to eight systems chosen from a total of 14 tested systems. The newest version of all systems was considered. The time limit for status and rating evaluations was increased to 600 seconds.

2. Contents of the ILTP Library

The release v1.1 of ILTP library includes a total of 2363 abstract problems. Some of these problems have alternative presentations or are generic and allow to scale the actual problem size. This results in a total of 2754 problems. Table I shows the overall statistics of the problems in the library.

Table I. Overall statistics of the ILTP library v1.1

Number of problem domains	24	
Number of abstract problems	2363	
Number of generic problems	20	
Total number of problems	2754	(100%)
Number of non-propositional problems	2479	(90%)
Number of propositional problems	275	(10%)
Number of problems with equality	1730	(63%)
Number of pure equality problems	185	(7%)

Table II provides two separate statistics about the first-order and the propositional part of the library. The intuitionistic status (Theorem, Non-Theorem, Unknown, Open) is explained in Subsection 2.2.

Table II. Number of problems in the first-order and the propositional part

Status	Theorem	Non-Theorem	Unknown	Open	$\sum$
First-order part	501	179	457	1413	2550
Propositional part	127	102	46	0	275

The 2754 problems are divided into 24 problem domains. The following subsections provide details about the problem domains, the sources they have been collected from, the intuitionistic status and difficulty rating, the problem naming and presentation norm, and the tools included in the ILTP library.

2.1. The ILTP Domain Structure

The problems of the ILTP library are grouped into 24 problem domains. They have been collected from several sources.

The major source of problems is the TPTP library [18]. Release 3.1.0 of the TPTP library includes 2324 problems in non-clausal form, so-called "first-order form" (FOF).

These problems are classified into 21 domains. Each domain is identified by a three-letter mnemonic, which is also component of the naming scheme (see Subsection 2.3). A more detailed description of these domains can be found in [18]. The domains are as follows: AGT (agents), ALG (general algebra), COM (computing theory), CSR (commonsense reasoning), GEO (geometry), GRA (graph theory), GRP (group theory), HAL (homological algebra), KRS (knowledge representation), LCL (logic calculi), MGT (management), MSC (miscellaneous), NLP (natural language processing), NUM (number theory), PLA (planning), PUZ (puzzles), SET (set theory), SWC (software creation), SWV (software verification), SYN (syntactic), and TOP (topology). Statistics about the problems can be found in Table III and IV.

In addition to the TPTP domains three more domains have been included with a total of 430 problems. These domains are GEJ (constructive geometry), GPJ (non-clausal group theory) and SYJ (intuitionistic syntactic) and contain the following problems.

GEJ – *Constructive Geometry.*
Elementary geometry is formalized constructively according to von Plato [13]. The original axiomatization and a shortened one [6, 7] were used. Additionally some problems were formalized using the reduced axiom set described in [19].

GPJ – *Non-Clausal Group Theory.*
: Group theory examples were taken from [5], e.g. the statement that a group with the identity element e is commutative if $x * x = e$ holds for every element x. All problems were represented in a non-clausal form. The original axiomatization and a modified one [19] were used.

SYJ – *Intuitionistic Syntactic.*
: Problems were taken from existing benchmark collections (see Subsection 1.2) and have no obvious semantic interpretation. The ft collection [16] contains various kind of first-order problems. For some problems the actual abstract or generic problem was added and instances of different sizes included. The JProver collection [17] contains propositional and first-order problems that are classically valid and can be represented in a pure first-order logic. Type information has been removed from all problems. Finally the generic propositional problems from Dyckhoff's benchmark collection were included. For each generic problem the first 20 instances were included for the propositional part of the ILTP library. For the first-order part only three instances were chosen.

The syntax of all three added domains was standardized and adapted to the TPTP syntax format. Each problem file was given a header with useful information (see Subsection 2.3), like intuitionistic status and rating (see Subsection 2.2). Statistics about these problems can be found in Table III and IV.

The set of problems was split into a first-order part and a propositional part. The first-order part of the ILTP library contains 71 of the propositional problems, but does not contain 204 instances from Dyckhoffs problem collection. This keeps the strong emphasis onto first-order problems.

2.2. Intuitionistic Problem Status and Difficulty Rating

In the TPTP library the difficulty of each problem is rated according to the performance of current (classical) state-of-the-art ATP systems. It expresses the ratio of state-of-the-art systems which are *not* able to solve a problem within a given time limit. For example a rating of 0.0 indicates that every state-of-the-art prover can solve the problem, a rating of 0.5 indicates that half of the systems were able to solve it, and a problem with rating 1.0 could not be solved by any state-of-the-art ATP system. This notion of difficulty rating is adapted to the ILTP library. To this end a set of intuitionistic state-of-the-art ATP systems needs to be specified. Comprehensive tests of 14 currently available ATP systems for intuitionistic first-order and propositional logic were performed. For this purpose a test environment was developed for automatically conducting tests and collecting and evaluating the vast amount of data. An eight-processor cluster system was used to simultaneously test eight ATP systems on all 2754 problems in the ILTP library. The time limit allowed to solve a problem was set to 600 seconds. The results are published within the ILTP library (see also [14] for the results of ILTP release v1.0). The following four first-order and four propositional ATP system were selected that solve the highest number of problems: the first-order systems ft (C-version) [16], JProver [17], ileanTAP[11], ileanCoP[12], and the propositional system ft (C-version)[3] [16], LJT [8], PITP [2] and STRIP [10].

Each problem is assigned its intuitionistic status. This status is either Theorem, Non-Theorem, Unknown or Open. Problems with Unknown status have never been solved by an ATP system. A problem has an Open status if its abstract problem has not been solved at all. No theoretical investigations into the intuitionistic validity of the problems in the TPTP library were done. Instead the intuitionistic status of a problem is marked as Theorem or Non-Theorem if any intuitionistic ATP system was able to prove or refute the problem, respectively. All other problems taken from the TPTP library were given the status Unknown or Open as described above.

Table III shows statistics about the intuitionistic status and rating of the problems in the first-order part of the ILTP library. Note that the first-order part also contains some propositional problems (see Subsection 2.1).

The intuitionistic status and rating of all propositional problems was determined by analyzing the experimental results of the four selected state-of-the-art ATP systems for propositional logic. Table IV shows statistics about the intuitionistic status and rating of the problems in the propositional part of the ILTP library.

2.3. Problem Versions, Naming and Presentation

Similar to the TPTP library [18] problems can have alternative presentations resulting in different versions of the original *abstract* problem. These are two main reasons for different versions of a given problem: different axiomatizations and problems that are instances of the same generic (abstract) problem.

The naming of problems was adapted from the TPTP library. Problems are given an unambiguous file name

`DDD.NNN+V[.SSS].p`

[3] The ft system uses a separate proof calculus for propositional logic.

consisting of the three letter mnemonic `DDD` of its domain (see Subsection 2.1), the number `NNN` of the abstract problem, the version number `V`, an optional parameter `SSS` indicating the size of the instance, and an additional letter `p`. For example `SYJ205+1.007.p` is the seventh instance of problem number 205 in the domain SYJ.

All files are presented in Prolog syntax and can easily be modified using Prolog. The tptp2X tool (see Subsection 2.4) can be used to convert the syntax to the input syntax of a specific ATP system. Problem files are split into input clauses and have the syntax of Prolog facts. They are marked with a name and a type, which is either `axiom`, `hypothesis`, `lemma` or `conjecture`. Each problem file includes a header containing information like file name, problem description, status and difficulty rating. This header is marked as a Prolog comment and should not be used by ATP systems to solve a problem. An example file of a problem is given in Figure 1.

```
%--------------------------------------------------------------------------
% File    : GEJ003+1 : ILTP v1.1
% Domain  : Constructive Geometry
% Problem : Theorem 3.3. Uniqueness of constructed points
% Version : [P95] axioms.
% English : If two lines are convergent and there is a point that is
%         : incident with both lines, then this point is equivalent to the
%         : intersection point of these lines.

% Refs    : [P95] J. von Plato. The Axioms of Constructive Geometry. Annals
%           of Pure and Applied Logic 76 (2): 169-200, 1995.
% Source  : [P95]
% Names   :

% Status (intuit.) : Theorem
% Rating (intuit.) : 0.75 v1.1
% Comments :
%--------------------------------------------------------------------------
include('Axioms/GEJ001+1.ax').
include('Axioms/GEJ002+1.ax').
include('Axioms/GEJ003+1.ax').
include('Axioms/GEJ004+1.ax').
%----------------------------
fof(con_name,conjecture,(
( ! [X,Y,Z] : ((con(X,Y) & ( ~apt(Z,X) & ~apt(Z,Y))) => ~dipt(Z,pt(X,Y))))
)).
%--------------------------------------------------------------------------
```

Figure 1. Example of a problem file (`GEJ003+1.p`)

More specifically the header files contains the following fields: `File` contains the problem name and the ILTP release number; `Domain` identifies the problem domain; `Problem` and `English` provide a short and an extended problem description; `Version` describes the specific version; `Refs`, `Source`, `Name` provide information about references to the problem, the original source and established name of the problem, respectively. The `Status (intuit.)` and `Rating (intuit.)` fields specify the *intuitionistic* status and difficulty rating of the problem (see

Subsection 2.2), `Comment` provides additional remarks. The header of problem files taken from the TPTP library additionally includes the fields `Status`, `Rating` and `Syntax`, which specify the *classical* status and difficulty rating, and various syntactic measures of the problem.

2.4. Tools and Prover Database

The TPTP library provides the tptp2X tool for reformatting, transforming and generating TPTP problem files. This tool can be used for all problems in the ILTP library as well. The tptp2X installation files were adapted to work with the ILTP directory structure. So-called format files have been included for all tested intuitionistic ATP systems. They are used together with the tptp2X tool to convert the problems in the ILTP library into the input syntax of the tested ATP systems. The prover database of the library provides information about published intuitionistic ATP systems. For each system some basic information is provided, like author, homepage, short description and references. For each system a test run on two example problems is included. A summary and a detailed list of the performance results on running the system on all problems in the ILTP library are given as well.

3. Conclusion

Like the TPTP library for classical logic, the main motivation for the ILTP library is to put the testing and evaluation of intuitionistic ATP systems onto a firm basis. It is the first systematic attempt to assemble a benchmark library for intuitionistic ATP systems. This will help to ensure that published results reflect the actual performance of an ATP system and make meaningful system evaluations and comparisons possible. We expect that such a library will be fruitful for the development of novel, more efficient calculi and implementations for intuitionistic first-order and propositional logic, which — compared to classical logic — is still in its infancy. The advancement in the field of intuitionistic ATP systems is indicated by the intuitionistic rating of problems, which should gradually decrease as progress is made. Future work includes adding more problems which occur during the practical use of interactive proof assistants like NuPRL [1] or Coq [4]. Extending the library to other non-classical logics like first-order modal logics or fragments of linear logic is under consideration as well.

Like most problem libraries the ILTP library is an ongoing project. We encourage not only its use but also submissions of new problems suitable for evaluating intuitionistic ATP systems. Users are also invited to submit performance data of their ATP system as well as information about their ATP system itself.

References

1. S. Allen, R. Constable, R. Eaton, C. Kreitz L. Lorigo. The NuPRL Open Logical Environment. 17^{th} CADE, LNAI 1831, 170–176, Springer, 2000.
2. A. Avellone, G. Fiorino, U. Moscato. A new $o(nlogn)$-space decision procedure for propositional intuitionistic logic. *Kurt Goedel Society Collegium Logicum*, volume VIII, pp. 17–33, 2004.
3. P. Balsiger, A. Heuerding, S. Schwendimann. Logics workbench 1.0. *7^{th} TABLEAUX Conference*, LNCS 1397, pp. 35–37, Springer, 1998.
4. Y. Bertot, P. Castéran. *Interactive theorem proving and program development*. Texts in Theoretical Computer Science, Springer, 2004.
5. C.-L. Chang, R. Lee. *Symbolic logic and mechanical theorem proving*. Academic Press, 1973.
6. L. Dafa. Replacing the axioms for connecting lines and intersection points by two single axioms. *AAR Newsletter*. 37:1–7, 1997.
7. L. Dafa. A shorter and intuitive axiom to replace the third axiom of compatibility of equality with apartness and convergence. *AAR Newsletter*. 39:1–6, 1998.
8. R. Dyckhoff. Contraction-free sequent calculi for intuitionistic logic. *Journal of Symbolic Logic*, 57:795–807, 1992.
9. H. Hoos, T. Stützle. SATLIB: An online resource for research on SAT. *SAT 2000*, pp. 283–292, IOS Press, 2000. (`http://www.satlib.org`)
10. D. Larchey-Wendling, D. Méry, D. Galmiche. STRIP: Structural sharing for efficient proof-search. *IJCAR-2001*, LNAI 2083, pp. 696–700, Springer, 2001.
11. J. Otten. ileanTAP: An intuitionistic theorem prover. *6^{th} TABLEAUX Conference*, LNAI 1227, pp. 307–312, Springer, 1997.
12. J. Otten. Clausal connection-based theorem proving in intuitionistic first-order logic. *TABLEAUX 2005*, LNAI 3702, pages 245–261, Springer, 2005. (`http://www.leancop.de`)
13. J. von Plato. The axioms of constructive geometry. *Annals of Pure and Applied Logic*, 76(2): 169–200, 1995.
14. T. Raths. Evaluating intuitionistic automated theorem provers. Technical Report, University of Potsdam, 2005.
15. T. Raths, J. Otten, C. Kreitz. The ILTP Library: Benchmarking Automated Theorem Provers for Intuitionistic Logic. *TABLEAUX 2005*, LNAI 3702, pages 333–337 Springer Verlag, 2005.
16. D. Sahlin, T. Franzen, S. Haridi. An intuitionistic predicate logic theorem prover. *Journal of Logic and Computation*, 2:619–656, 1992.
17. S. Schmitt et al. JProver: Integrating connection-based theorem proving into interactive proof assistants. *IJCAR-2001*, LNAI 2083, pp. 421–426, Springer, 2001.
18. G. Sutcliffe, C. Suttner. The TPTP problem library - CNF release v1.2.1. *Journal of Automated Reasoning*, 21: 177–203, 1998. (`http://www.cs.miami.edu/~tptp`)
19. T. Tammet. A resolution theorem prover for intuitionistic logic. *13^{th} CADE*, LNAI 1104, pp. 2–16, Springer, 1996.

Table III. First-order part of ILTP library: status and rating of problems

Domain	Intuitionistic status			Intuitionistic rating				
	Theorem	Non-Theorem	Open/ Unknown	0.00	0.01 –0.33	0.34 –0.66	0.67 –0.99	1.00
AGT	13	0	39	0	0	2	11	39
ALG	9	51	139	0	3	53	4	139
COM	3	0	0	0	0	0	3	0
CSR	1	0	28	0	0	0	1	28
GEO	6	0	71	0	0	0	6	71
GRA	3	0	15	0	0	0	3	15
GRP	1	2	2	0	2	0	1	2
HAL	0	0	9	0	0	0	0	9
KRS	17	31	110	6	8	0	33	111
LCL	1	3	0	0	1	2	1	0
MGT	25	1	52	5	2	1	18	52
MSC	0	0	2	0	0	0	0	2
NLP	11	0	247	3	4	0	4	247
NUM	24	0	60	0	0	1	23	60
PLA	0	0	6	0	0	0	0	6
PUZ	4	0	2	2	0	0	2	2
SET	76	0	248	12	1	1	62	248
SWC	1	0	422	0	1	0	0	422
SWV	1	5	214	1	1	4	0	214
SYN	120	69	177	101	9	10	69	177
TOP	2	0	1	2	0	0	0	1
GEJ	71	0	22	4	4	10	53	22
GPJ	3	0	2	0	1	0	2	2
SYJ	109	17	2	62	21	19	23	3
$\sum$	501	179	1870	198	58	103	319	1872

Table IV. Propositional part of ILTP library: status and rating of problems

Domain	Intuitionistic status			Intuitionistic rating				
	Theorem	Non-Theorem	Open/ Unknown	0.00	0.01 –0.33	0.34 –0.66	0.67 –0.99	1.00
LCL	0	2	0	2	0	0	0	0
SYN	8	13	0	19	1	0	1	0
SYJ	119	87	46	135	14	29	24	50
$\sum$	127	102	46	156	15	29	25	50