

A proof-theoretic perspective on SMT-solving for intuitionistic propositional logic

Camillo Fiorentini¹ and Rajeev Goré² and Stéphane Graham-Lengrand³

¹ Department of Computer Science, University of Milan

² Research School of Computer Science, Australian National University

³ SRI International

Abstract. Claessen and Rósen have recently presented an automated theorem prover, `intuit`, for intuitionistic propositional logic which utilises a SAT-solver. We present a sequent calculus perspective of the theory underpinning `intuit` by showing that it implements a generalisation of the implication-left rule from the sequent calculus LJ_T, also known as G4ip and popularised by Roy Dyckhoff.

1 Introduction

Intuitionistic propositional logic IPL is one of the most important “non-classical” logics due to its constructive reading of implication. There is a long history of automated reasoning techniques for deciding validity of IPL-formulae, but most of them are based on either sequent or tableaux calculi. One of the simplest procedures for IPL is root-first (*a.k.a. backward or goal-directed*) proof search in the LJ_T (a.k.a. G4ip) sequent calculus [2], as it is guaranteed to terminate without implementing loop-checking.

Claessen and Rósen [1] have recently presented an automated theorem prover, `intuit`, for intuitionistic propositional logic, based on a Satisfiability-Modulo-Theories (SMT) approach. Their procedure also terminates without requiring any loop-detection machinery. As of 2015, the `intuit` prover was the best performing IPL prover, at least when evaluated on about 1200 standard benchmarks [1], which include for instance the ILTP library [12].

The SMT approach embraced by `intuit`, organised around the top-level loop of the DPLL(\mathcal{T}) procedure [11], and the proof-theoretic approach based on root-first proof search, appear as radically different methodologies; the potential connections between them was left as an open question. In this paper we reconcile the two approaches, formalising an explicit connection. In particular, we reformulate a variant of `intuit` using (a suitable generalisation of) one of the rules of LJ_T. The procedure builds an explicit proof when the input formula is valid, and builds an explicit Kripke counter-model when the formula is not valid.

2 Syntax and Kripke Semantics of IPL

In this paper, formulae of IPL, denoted by lowercase Greek letters, are built from an infinite set of propositional variables V , the “falsum” constant \perp and

the connectives $\wedge, \vee, \rightarrow$ with negation defined as $\neg\alpha := (\alpha \rightarrow \perp)$. We use $\text{Atm} := V \cup \{\perp\}$ for the set of “atoms”, denoted by lowercase Roman letters.

A rooted Kripke model for IPL is a quadruple $\langle W, \leq, r, \vartheta \rangle$ where W is a non-empty set of “worlds” containing r , and \leq is a reflexive and transitive binary relation over W , and the root world r is minimal wrt \leq , and $\vartheta : W \mapsto 2^{\text{Atm}}$ is a “valuation” mapping each world to a set of propositional variables which obeys the “persistence” condition: $\forall w, v \in W$, if $w \leq v$ and $p \in \vartheta(w)$ then $p \in \vartheta(v)$. Given a Kripke model $\langle W, \leq, r, \vartheta \rangle$, the valuation ϑ can be extended into a “forcing” relation between worlds and formulae as shown below:

$$\begin{aligned} w \Vdash p & \quad \text{iff } p \in \vartheta(w) & w \Vdash \alpha \rightarrow \beta & \text{iff } \forall v \geq w, v \Vdash \alpha \text{ implies } v \Vdash \beta \\ w \Vdash \alpha \wedge \beta & \text{iff } w \Vdash \alpha \text{ and } w \Vdash \beta & w \Vdash \perp & \quad \text{never holds} \\ w \Vdash \alpha \vee \beta & \text{iff } w \Vdash \alpha \text{ or } w \Vdash \beta \end{aligned}$$

A formula α is IPL-valid if, for all Kripke models $\langle W, \leq, r, \vartheta \rangle$, we have $r \Vdash \alpha$. The problem of deciding whether a formula is IPL-valid is known to be PSPACE-complete [13]. For formula set or multiset Γ , we write $w \Vdash \Gamma$ for $\forall \gamma \in \Gamma. w \Vdash \gamma$.

A model M for classical propositional logic (CPL), or “classical model”, is just a set of propositional variables (assigned true). By $M \models \alpha$ we mean that α is true in model M (following the Boolean truth tables). We write $M \models \Gamma$ iff $\forall \gamma \in \Gamma. M \models \gamma$. We write $\Gamma \vdash_{\text{ipl}} \delta$ when the formula $\bigwedge \Gamma \rightarrow \delta$ is IPL-valid and $\Gamma \vdash_{\text{cpl}} \delta$ when it is CPL-valid, that is, when $M \models \gamma$ for all classical models M .

3 The Theorem Prover intuit

The `intuit` theorem prover is an intuitionistic prover built on top of a SAT-solver, following a Satisfiability-Modulo-Theories (SMT) approach. Despite the fact that SMT-solving works primarily in classical logic and with first-order theories, Claessen and Rosén [1] show this approach to be relevant to the problem of deciding IPL-validity. They use a variant of the SMT scheme known as DPLL(\mathcal{T}) [11], where DPLL is the well-known procedure for SAT-solving and \mathcal{T} is here the “theory of intuitionistic implications”. The main loop of DPLL(\mathcal{T}) can be seen as a particular case of Counter-Example Guided Abstraction Refinement (CEGAR), as we describe next, for the particular case of IPL. A formula α whose IPL-validity is to be determined is transformed into a set R of classical “flat clauses”, a set X of intuitionistic “implication clauses”, and an atomic formula q , such that $\vdash_{\text{ipl}} \alpha$ iff $R, X \vdash_{\text{ipl}} q$ (the definitions are in Section 3.1). The *sequent* $R \Rightarrow q$ constitutes an “abstraction” of the input formula α . A SAT-solver tries to find a classical counter-model M for it, in that $M \models R$ but $M \not\models q$. If no such counter-model exists then α is not only CPL-valid but also IPL-valid. Otherwise the SAT-solver returns such a counter-model M , although the existence of M does not necessarily mean that $R, X \not\vdash_{\text{ipl}} q$, as the implication clauses X have so far been ignored. Therefore another procedure “checks” model M , in that it tries to produce a new abstraction $R' \Rightarrow q$ of α that refines $R \Rightarrow q$ (technically, $R \subseteq R'$) and *defeats* model M , meaning that $M \not\models R'$ while still ensuring $\vdash_{\text{ipl}} \alpha$ iff $R', X \vdash_{\text{ipl}} q$. If it fails, then indeed α is not IPL-valid. If

```

1 procedure prove( $R, X, q$ )
2    $s = \text{newSolver}()$ 
3   for  $\varphi \in R$  do addClause( $s, \varphi$ )
4   for  $(a \rightarrow b) \rightarrow c \in X$  do addClause( $s, b \rightarrow c$ )
5   return intuitProve( $s, X, \emptyset, q$ )
6 end
7 procedure intuitProve( $s, X, A, q$ )
8   loop // CEGAR LOOP
9      $\tau_0 \leftarrow \text{satProve}(s, A, q)$ 
10    if  $\tau_0 = \text{Yes}(A')$  then return Yes( $A'$ )
11    else //  $\tau_0 = \text{No}(M)$ , with  $M$  a putative counter-model
12     | if intuitCheck( $s, X, M$ ) then return No( $M$ )
13     | end
14   end
15 end
16 procedure intuitCheck( $s, X, M$ )
17   for  $\iota = (a \rightarrow b) \rightarrow c \in X$  such that  $a \notin M$  and  $b \notin M$  and  $c \notin M$  do
18      $\tau_1 \leftarrow \text{intuitProve}(s, X_\iota, M \cup \{a\}, b)$  //  $X_\iota = X \setminus \{\iota\}$ 
19     if  $\tau_1 = \text{Yes}(A_1)$  then
20       addClause( $s, \bigwedge(A_1 \setminus \{a\}) \rightarrow c$ )
21       return False
22     end
23   end
24   return True
25 end

```

Fig. 1. Main algorithms of `intuit` [1]

it does produce a refinement, then the procedure loops with R' instead of R . Eventually, it either finds a counter-model for α , or exhausts the set of putative counter-models and conclude that α is IPL-valid.

A key element of the approach is that $R \vdash_{\text{cpl}} q$ iff $R \vdash_{\text{ipl}} q$, as the clauses in R are “flat”. A twist of the approach, compared to the standard $\text{DPLL}(\mathcal{T})$ loop, is that the procedure that checks model M has to solve a new IPL-validity problem (for a different R, X, q), so that it recursively calls a new $\text{DPLL}(\mathcal{T})$ loop. In other words `intuit` implements a recursive version of $\text{DPLL}(\mathcal{T})$, although a single SAT-solver is used, incrementally, for all recursive calls.

3.1 `intuit` in detail

Firstly, the formula α is transformed into a formula $\bigwedge \Gamma \rightarrow q$, where q is an atom and Γ is a set of *flat clauses* φ and *implication clauses* ι , where:

$$\begin{aligned}
\varphi &::= \bigwedge A_1 \rightarrow \bigvee A_2 & A_1 \cup A_2 \subseteq \text{Atm} \\
\iota &::= (a \rightarrow b) \rightarrow c & \{a, b, c\} \subseteq \text{Atm} \\
\bigwedge A_1 &\text{ is the conjunction of the atoms in } A_1 \\
\bigvee A_2 &\text{ is the disjunction of the atoms in } A_2 & \text{if } A_1 = \emptyset \text{ then } \varphi = \bigvee A_2
\end{aligned}$$

```

1 procedure intuitProve1( $s, X, A, q$ )
2   loop
3      $\tau_0 \leftarrow \text{satProve}(s, A, q)$ 
4     if  $\tau_0 = \text{Yes}(A')$  then return Yes( $A'$ )
5     else //  $\tau_0 = \text{No}(M)$ 
6       for  $\iota = (a \rightarrow b) \rightarrow c \in X$  such that  $a \notin M$  and  $b \notin M$  and  $c \notin M$  do
7          $\tau_1 \leftarrow \text{intuitProve1}(s, X_\iota, M \cup \{a\}, b)$  //  $X_\iota = X \setminus \{\iota\}$ 
8         if  $\tau_1 = \text{Yes}(A_1)$  then
9           addClause( $s, \bigwedge(A_1 \setminus \{a\}) \rightarrow c$ )
10          go to line 2 (outer loop)
11        end
12      end
13      return No( $M$ )
14    end
15  end
16 end

```

Fig. 2. The function `intuitProve1`

A flat clause where A_1 is empty is simply a disjunction $\bigvee A_2$ of atoms, and simply an atom a when A_2 is the singleton $\{a\}$. Henceforth, we write R, R_1, R' etc. to denote sets of flat clauses; X, X_1, X' etc. for sets of implication clauses; A, A_1, A' etc. for sets of atoms; and $X_\iota, X_{\iota'}$ for the sets $X \setminus \{\iota\}$ and $X \setminus \{\iota'\}$, respectively. A clausification procedure is presented in [1], similar to Tseitin's [15], where clauses are created by naming subformulae with new propositional atoms. Technically, it transforms any IPL formula α into a triple $\text{clausal}(\alpha) = (R, X, q)$ whose cumulative size is linear in the size of α and that is equiprovable to α :

Lemma 1. *For every α with $\text{clausal}(\alpha) = (R, X, q)$, $\vdash_{\text{ipl}} \alpha$ iff $R, X \vdash_{\text{ipl}} q$ [1].*

From now on, we focus on deciding $R, X \vdash_{\text{ipl}} q$. The `intuit` algorithm [1], outlined in Fig. 1 with only slight modifications, consists of three procedures. It exploits a single SAT-solver s that is *incremental*: clauses can be added to s but not removed, and problems can be solved with varying atomic assumptions. So the clauses in s are “global clauses” which must hold at any point of proof-search. Technically, the SAT-solver has the following API, i.e., it supports the following operations, where $R(s)$ denotes the set of clauses stored in s :

`newSolver()`: create a new SAT-solver;
`addClause(s, φ)`: add the flat clause φ to the SAT-solver s ;
`satProve(s, A, q)`: call the SAT-solver s to decide whether $R(s), A \vdash_{\text{cpl}} q$,
 where A is a local set of assumptions and q is an atom.

The call `satProve(s, A, q)` yields one of the following answers:

Yes(A'): thus $A' \subseteq A$ and $R(s), A' \vdash_{\text{cpl}} q$;
 No(M): thus M is a classical model such that $M \models R(s) \cup A$ and $M \not\models q$.

```

1 procedure intuitPR( $s, X, A, q$ )
2    $\tau_0 \leftarrow \text{satProve}(s, A, q)$ 
3   if  $\tau_0 = \text{Yes}(A')$  then return  $\text{Yes}(A')$ 
4   else //  $\tau_0 = \text{No}(M)$ 
5     for  $\iota = (a \rightarrow b) \rightarrow c \in X$  such that  $a \notin M$  and  $b \notin M$  and  $c \notin M$  do
6        $\tau_1 \leftarrow \text{intuitPR}(s, X_\iota, M \cup \{a\}, b)$  //  $X_\iota = X \setminus \{\iota\}$ 
7       if  $\tau_1 = \text{Yes}(A_1)$  then
8         addClause( $s, \bigwedge(A_1 \setminus \{a\}) \rightarrow c$ )
9         return  $\text{intuitPR}(s, X, A, q)$ 
10      end
11    end
12    return  $\text{No}(M)$ 
13  end
14 end

```

Fig. 3. The procedure `intuitPR` (recursive variant of `intuitProve`)

The main function `prove`(R, X, q) of `intuit` yields:

$\text{Yes}(\emptyset)$ if $R, X \vdash_{\text{ipl}} q$;

$\text{No}(M)$ if there is a Kripke model $\mathcal{K} = \langle W, \leq, r, \vartheta \rangle$ such that $\vartheta(r) = M$ and $r \Vdash R \cup X$ and $r \not\Vdash q$; this implies $R, X \not\vdash_{\text{ipl}} q$.

As sketched by Claessen and Rósen [1], if `prove`(R, X, q) returns $\text{No}(M)$, then one can actually build the mentioned model \mathcal{K} by tracking the sets M' returned by `intuitProve`.

To reason about `intuit`, it is convenient to merge the functions `intuitProve` and `intuitCheck` into one recursive function. Firstly, we plug `intuitCheck` into `intuitProve` and obtain the function `intuitProve1` in Fig. 2. Then, we remove the outer loop by replacing the “go to” statement at line 10 with a recursive call; we get the recursive procedure `intuitPR` in Fig. 3. We henceforth consider the `intuit` algorithm as implemented by the main function `prove` in Fig. 1, with function `intuitProve` at line 5 replaced by function `intuitPR` in Fig. 3.

4 Adapting the sequent calculus LJT to clausal forms

The sequent calculus LJT is a variant of Gentzen’s sequent calculus LJ for intuitionistic logic [6, 7] that was discovered many times, as outlined by Roy Dyckhoff [2]. Its rules are given in Fig. 4, where $\Gamma \Rightarrow \alpha$ denotes a sequent whose *antecedent* Γ is a multiset of assumptions, and whose provability in LJT is denoted $\vdash_{\text{LJT}} \Gamma \Rightarrow \alpha$. The key difference from LJ lies in the left introduction rules for implication. In order to introduce an implication $\eta \rightarrow \gamma$ on the left, LJT offers four rules, depending on the form of η , namely: either $\eta = p$, with $p \in V$; or $\eta = \alpha \wedge \beta$; or $\eta = \alpha \vee \beta$; or $\eta = \alpha \rightarrow \beta$ ⁴.

⁴ We follow Troelstra and Schwichtenberg [14] where the calculus is called G4ip; the original ($L \rightarrow \rightarrow$) rule by Dyckhoff [2] has $\Gamma, \beta \rightarrow \gamma \Rightarrow \alpha \rightarrow \beta$ as the left premise.

$$\begin{array}{c}
\frac{}{\Gamma, p \Rightarrow p} \text{Ax} \quad \frac{}{\Gamma, \perp \Rightarrow \delta} L\perp \quad \frac{\Gamma, \alpha, \beta \Rightarrow \delta}{\Gamma, \alpha \wedge \beta \Rightarrow \delta} L\wedge \quad \frac{\Gamma \Rightarrow \alpha \quad \Gamma \Rightarrow \beta}{\Gamma \Rightarrow \alpha \wedge \beta} R\wedge \\
\frac{\Gamma, \alpha \Rightarrow \delta \quad \Gamma, \beta \Rightarrow \delta}{\Gamma, \alpha \vee \beta \Rightarrow \delta} L\vee \quad \frac{\Gamma \Rightarrow \alpha_k}{\Gamma \Rightarrow \alpha_1 \vee \alpha_2} R\vee_k \quad \frac{\Gamma, \alpha \Rightarrow \beta}{\Gamma \Rightarrow \alpha \rightarrow \beta} R\rightarrow \\
\frac{\Gamma, p, \beta \Rightarrow \delta}{\Gamma, p, p \rightarrow \beta \Rightarrow \delta} L0\rightarrow \quad \frac{\Gamma, \alpha \rightarrow (\beta \rightarrow \gamma) \Rightarrow \delta}{\Gamma, (\alpha \wedge \beta) \rightarrow \gamma \Rightarrow \delta} L\wedge\rightarrow \quad \frac{\Gamma, \alpha \rightarrow \gamma, \beta \rightarrow \gamma \Rightarrow \delta}{\Gamma, (\alpha \vee \beta) \rightarrow \gamma \Rightarrow \delta} L\vee\rightarrow \\
\frac{\Gamma, \beta \rightarrow \gamma, \alpha \Rightarrow \beta \quad \Gamma, \gamma \Rightarrow \delta}{\Gamma, (\alpha \rightarrow \beta) \rightarrow \gamma \Rightarrow \delta} L\rightarrow\rightarrow \quad p \in V, k \in \{1, 2\}
\end{array}$$

Fig. 4. The calculus LJT (a.k.a. G4ip)

4.1 Root-first proof search, invertibility and recursivity

The purpose of replacing Gentzen's left-introduction of implication by those four rules is to ensure that root-first proof search terminates. When given a sequent to prove, root-first proof search matches it against the conclusion of one of the rules, and recursively tries to prove each of its premises. In every rule of LJT, the multiset Γ_i, α_i corresponding to the i th premise $\Gamma_i \Rightarrow \alpha_i$ is strictly smaller than the multiset Γ, α corresponding to the conclusion $\Gamma \Rightarrow \alpha$, according to the well-founded multiset ordering based on formula size.⁵ Hence, the recursions of root-first proof search terminate. Note that keeping several copies of an assumption is never useful for proof search, so from now on, the antecedents of sequents will be considered sets. If, following the application of a rule, the recursive call that attempts to prove any one of its premises fails, then proof search attempts to apply another rule or another instance of the rule. Conceptually, a backtrack point was set when the original rule was applied. However, no such backtrack is needed if the rule is *invertible*, in the sense that, whenever the rule's conclusion is provable (in LJT), so is each of its premises. In that case indeed, if any one of the premises is not provable, then neither is the conclusion, so there is no point in trying out another rule. In LJT, all rules are invertible except $R\vee_k$ and $(L\rightarrow\rightarrow)$, and therefore backtrack points need to be set only when applying those two rules. However $(L\rightarrow\rightarrow)$ is *right-invertible* in that, if the conclusion is provable (in LJT), then so is the right premise (while the left premise may or may not be provable). This means that $(L\rightarrow\rightarrow)$ can be seen as a one-premise invertible rule, guarded by a side-condition:

$$(\Gamma, \beta \rightarrow \gamma, \alpha \Rightarrow \beta) \frac{\Gamma, \gamma \Rightarrow \delta}{\Gamma, (\alpha \rightarrow \beta) \rightarrow \gamma \Rightarrow \delta}.$$

When applying this rule, root-first proof search makes a first recursive call that checks that the side-condition holds and, if successful, makes a second one on the

⁵ This is the number of connectives, each conjunction counting for two [2].

premise. Because of invertibility, that second call is *tail-recursive*, as no backtrack point is needed: the output of proof search is the output of the recursive call. The next sections describe how these two recursive calls, for a generalisation of rule $(L \rightarrow \rightarrow)$ satisfying the same invertibility properties, correspond to the two recursive calls of `intuitPR` in Fig. 3 (lines 6 and 9). The (right-)invertibility of that generalised rule means that the recursive call on the (right) premise is tail-recursive and proof-search can thus be implemented by a while loop: namely the $\text{DPLL}(\mathcal{T})$ loop of `intuit`. To see this we specialise LJT to clausal forms.

4.2 LJT specialised to clausal forms

We now consider sequents in clausal form, namely sequents of the form $R, X \Rightarrow q$. The only LJT rule manipulating X is then $(L \rightarrow \rightarrow)$, which becomes:

$$\frac{R, b \rightarrow c, X_\iota, a \Rightarrow b \quad R, X_\iota, c \Rightarrow q}{R, X \Rightarrow q} \quad \begin{array}{l} \iota = (a \rightarrow b) \rightarrow c \text{ in } X \\ X_\iota = X \setminus \{\iota\} \end{array}$$

All other rules concern R and q , or do not apply because the sequent to prove is already in clausal form. Hence, these rules can be replaced by the use of a SAT-solver, remembering that $R \vdash_{\text{cpl}} q$ iff $R \vdash_{\text{ipl}} q$.

To prove the left premiss of $(L \rightarrow \rightarrow)$, atom a is added as an assumption, and will be taken into account by the SAT-solver. As a is not present in the right premise, it must be removed from the assumptions if the same SAT-solver is used for the right premise. Hence, it is useful to refine the notion of sequent in clausal form into sequents of the form $R, X, A \Rightarrow q$, where R is the set of clauses present in the SAT solver, which can only grow bigger, and A is a set of atomic assumptions that can vary from one call to the next, relying on the SAT-solver's API presented in Section 3.1. On such sequents, rule $(L \rightarrow \rightarrow)$ becomes:

$$\frac{R, b \rightarrow c, X_\iota, A, a \Rightarrow b \quad R, X_\iota, A, c \Rightarrow q}{R, X, A \Rightarrow q} \quad \begin{array}{l} \iota = (a \rightarrow b) \rightarrow c \text{ in } X \\ X_\iota = X \setminus \{\iota\} \end{array}$$

We can start relating root-first proof search to `intuit` by relating the application of the above rule to a call to function `intuitPR(s, X, A, q)`, described in Fig. 3, considering $R = R(s)$.

Indeed the left premise of the above rule is very similar to the first recursive call `intuitPR(s, Xι, M ∪ {a}, b)` on line 6 of Fig. 3, except that $b \rightarrow c$ is added to R in the premise, and model M may differ from A : According to its definition on line 4 of Fig. 3, M must satisfy (and therefore contain) all atoms in A , but other atoms could be assigned true in M that are not in A . Note however that $R(s)$ does contain $b \rightarrow c$ if $(a \rightarrow b) \rightarrow c$ is in X : it has been added to s at the beginning of the computation of function `prove` from Fig. 1.

Likewise, the right premise of the above rule is very similar to the second recursive call `intuitPR(s, X, A, q)` on line 9 of Fig. 3, except that the right premise contains the atomic formula c and the recursive call keeps the implication $(a \rightarrow b) \rightarrow c$.

Furthermore, the use of the SAT-solver in procedure `intuitPR` has side-effects: let R^0 denote $R(s)$ at the time when `intuitPR` is called, R^1 denote $R(s)$ at

the time of the first recursive call (line 6), and R^2 denote $R(s)$ at the time of the second recursive call (line 9). We have $R^0 \subseteq R^1 \subseteq R^1 \cup \{\wedge(A_1 \setminus \{a\}) \rightarrow c\} \subseteq R^2$, for a subset A_1 of $M \cup \{a\}$ such that $R^1, X_L, A_1 \vdash_{\text{ip1}} b$. This incremental use of the SAT-solver is not reflected in root-first proof search using rule $(L \rightarrow \rightarrow)$. Hence, rule $(L \rightarrow \rightarrow)$ has to be generalised to account for these differences.

4.3 A generalised version of $(L \rightarrow \rightarrow)$

The first generalisation consists in allowing the addition, in the left premise of $(L \rightarrow \rightarrow)$, of extra atomic assumptions that were not assumptions in the conclusion. By allowing this, the rule can model either the extra atoms that are in M but not in A for the first recursive call of `intuitPR` (line 6 of Fig. 3), or the atoms A_1 that are returned by the call if successful. This ambiguity is systemic to the description of root-first proof search as presented in Section 4.1, which makes a double usage of sequent calculus rules. Consider a sequent calculus rule.

- Firstly, if the conclusion describes the arguments of a proof search, then the premises describe the arguments of the recursive calls; the rule describes the descent into the recursions.
- Secondly, if the recursive calls succeed, then a proof of each premise has been completed, either explicitly or implicitly, and a proof of the conclusion can be constructed; the rule describes the ascent back from the recursions.

It is useful to enhance proof search by considering, for each rule, a variant used for the first purpose and a variant used for the second purpose. Typically for the first purpose, it is convenient to integrate the weakening rule of the sequent calculus (say in LJ) to the axiom rule, and use context-sharing rules, as in Fig. 4. For the second purpose, it is useful to push all weakenings down towards the proof-tree root, using context-splitting rules. This strengthens the proved sequents by pruning the input sequent of all assumptions that were not used in the proof. This was described for instance in [8, 9], which also connects the said pruning to the notion of *conflict analysis* used in SAT and SMT-solving. This is relevant for the connection between LJT and `intuit`, as a call to function `intuitPR(s, X, A, q)`, if succesful, precisely performs this pruning by outputting a subset A' of A that is sufficient for provability. In that spirit, we present the generalisation of $(L \rightarrow \rightarrow)$ in a context-splitting style, emphasising what happens upon the completion of the recursive calls.

The addition of extra atomic assumptions in the left premise of $(L \rightarrow \rightarrow)$ makes the premise *easier* to prove than with the original rule (when seen as a one-premise rule with a side-condition, the rule applies more often). The price to pay for this is that the new assumption that is learnt and that helps proving the right premise, has to be weakened from c to the flat clause $\varphi = \wedge(A_1 \setminus \{a\}) \rightarrow c$, which in `intuitPR` is added to the SAT-solver on line 8 in Fig. 3. Note that clause φ is a consequence of the original problem. Weakening this newly available assumption in turn means that it no longer subsumes the original implication clause ι , which has to stay in the right premise. The resulting rule is rule (ljt) in Fig. 5.

$$\begin{array}{c}
\frac{R \vdash_{\text{cpl}} q}{R, X \Rightarrow q} \text{ (cpl)} \qquad \frac{R_1, X_1 \vdash_{\text{ipl}} \varphi \quad \varphi, R_2, X_2 \Rightarrow q}{R_1, R_2, X_1, X_2 \Rightarrow q} \text{ (cut}_{\text{ipl}}) \\
\hline
\frac{R_1, b \rightarrow c, X, A_1 \Rightarrow b \quad R_2, \bigwedge(A_1 \setminus \{a\}) \rightarrow c, X, (a \rightarrow b) \rightarrow c \Rightarrow q}{R_1, R_2, X, (a \rightarrow b) \rightarrow c \Rightarrow q} \text{ (ljt)}
\end{array}$$

Fig. 5. The calculus LJT_{SAT}

When emulating SAT or SMT-solving, the sequent calculus has to deal with the effectful aspect of these solvers, which learn clauses that are consequences of the input problem, such as clause φ above.⁶ This effect was described in terms of *memoisation* of root-first proof search in [8, 9], and in terms of *cuts* in [4, 9]. Here again we use cuts to model the phenomenon: the added clauses can be deleted at the end of the proof search computation by applying rule $(\text{cut}_{\text{ipl}})$ of Fig. 5.

5 The Calculus LJT_{SAT}

We introduce the calculus LJT_{SAT} (LJT with SAT-solver) for sequents of the form $R, X \Rightarrow q$, whose provability in LJT_{SAT} is denoted $\vdash_{\text{LJT}_{\text{SAT}}} R, X \Rightarrow q$. It consists of the three rules of Fig. 5. Rule (cpl) has the premise judgment $R \vdash_{\text{cpl}} q$ and the sequent $R, X \Rightarrow q$ as conclusion, with X any set of implicational clauses. The rule can be applied if $R \vdash_{\text{cpl}} q$ holds, as checked by a SAT-solver (hence LJT_{SAT}). The rule $(\text{cut}_{\text{ipl}})$ is a cut rule having the judgment $R_1, X_1 \vdash_{\text{ipl}} \varphi$ as left premise. In the proof-search procedure, whenever we apply $(\text{cut}_{\text{ipl}})$, the left-premise is a judgment of the kind $R_0, X_0 \vdash_{\text{ipl}} \varphi$, where the cut formula φ is a clause already stored in the SAT-solver, and we can take for granted that the assertion holds (we do not have to invoke an external prover to check it). The rule (ljt) is a sort of context-splitting generalisation of $(L \rightarrow \rightarrow)$ needed to capture the `intuit` procedure. We point out that the sets R_1 and R_2 may overlap, thus the common part $R_1 \cap R_2$ is kept in both the premises. The formula $\bigwedge(A_1 \setminus \{a\}) \rightarrow c$ in the right premise is needed to guarantee the soundness of the rule, since A_1 is any set of atoms. To get completeness, we have to keep the main formula $(a \rightarrow b) \rightarrow c$ in the right premise; as a side-effect, the termination of proof-search is now trickier to prove. If the sets R_1 and R_2 coincide and $A_1 = \{a\}$ (thus $\bigwedge(A_1 \setminus \{a\}) \rightarrow c$ is the atom c), we get the instance

$$\frac{R, b \rightarrow c, X, a \Rightarrow b \quad R, c, X, (a \rightarrow b) \rightarrow c \Rightarrow q}{R, X, (a \rightarrow b) \rightarrow c \Rightarrow q} \text{ (ljt)}$$

The formula $(a \rightarrow b) \rightarrow c$ in the right premise is now redundant since it is implied by the occurrence of c , thus we recover Dyckhoff's rule $(L \rightarrow \rightarrow)$.

To prove the soundness of LJT_{SAT} , we show that an LJT_{SAT} -derivation can be translated into the calculus LJT. In derivations, a double line marks the application of more than one rule. Firstly, we prove the soundness of rule (cpl) .

⁶ These effectful additions account for the distinction between R^0 , R^1 and R^2 above.

Lemma 2 (Soundness of rule (cpl)). *If $R \vdash_{\text{cpl}} q$ then $\vdash_{\text{LJT}} R \Rightarrow q$.*

Proof. We proceed via contraposition, so suppose $\not\vdash_{\text{LJT}} R \Rightarrow q$. By completeness, there is a Kripke model $\langle W, \leq, r, \vartheta \rangle$ containing a world $w \in W$ such that $w \Vdash \bigwedge R$ and $w \not\Vdash q$. Now consider any flat clause $\varphi = \bigwedge A_1 \rightarrow \bigvee A_2 \in R$. The valuation $\vartheta(w)$ either has $A_2 \cap \vartheta(w) \neq \emptyset$ or $A_2 \cap \vartheta(w) = \emptyset$. If $A_2 \cap \vartheta(w) \neq \emptyset$ then φ is classically true at w . If $A_2 \cap \vartheta(w) = \emptyset$ then reflexivity demands $A_1 \not\subseteq \vartheta(w)$, as otherwise $w \not\Vdash \varphi$, contradicting our assumption. Again, φ is classically true at w . That is, w by itself is a classical model that also makes $\bigwedge R$ true and q false, so $R \not\vdash_{\text{cpl}} q$. By contraposition, if $R \vdash_{\text{cpl}} q$ then $\vdash_{\text{LJT}} R \Rightarrow q$. Notice that this proof only works because R contains flat clauses. \square

A syntactic way of proving Lemma 2 is to consider the proof returned by a (proof-producing) SAT-solver, justifying the unsatisfiability of R, \bar{q} (where \bar{q} is the negation of q) with a resolution proof concluding the empty clause \perp from the flat clauses R , and \bar{q} . Indeed, the resolution rule is perfectly valid in intuitionistic logic if a clause $\bar{a}_1 \vee \dots \vee \bar{a}_n \vee b_1 \vee \dots \vee b_m$ is read as the flat clause $a_1 \wedge \dots \wedge a_n \rightarrow b_1 \vee \dots \vee b_m$, as pointed out by Claessen and Rósen [1]. Removing $q \rightarrow \perp$ from the leaves of the resolution tree leaves q at its root, yielding an intuitionistic proof of $R \vdash_{\text{ipl}} q$. Completeness of LJTT [2] concludes $\vdash_{\text{LJT}} R \Rightarrow q$.

We prove the main lemma for the soundness of LJTT_{SAT}.

Lemma 3. *If $\vdash_{\text{LJT}_{\text{SAT}}} R, X \Rightarrow q$ then $\vdash_{\text{LJT}} R, X \Rightarrow q$.*

Proof. Let \mathcal{D} be an LJTT_{SAT}-derivation of $R, X \Rightarrow q$; we prove the lemma by induction on the depth of \mathcal{D} . If the root rule of \mathcal{D} is (cpl), the assertion follows by Lemma 2. Let us assume that \mathcal{D} is

$$\frac{\mathcal{D}_2 \quad R_1, X_1 \vdash_{\text{ipl}} \varphi \quad \varphi, R_2, X_2 \Rightarrow q}{R_1, R_2, X_1, X_2 \Rightarrow q} (\text{cut}_{\text{ipl}})$$

By the completeness of LJTT [2], there exists an LJTT-derivation \mathcal{E}_1 of $R_1, X_1 \Rightarrow \varphi$. By the induction hypothesis, there exists an LJTT-derivation \mathcal{E}_2 of $\varphi, R_2, X_2 \Rightarrow q$. Since (cut) is admissible in LJTT [2], from \mathcal{E}_1 and \mathcal{E}_2 we get an LJTT-derivation of $R_1, R_2, X_1, X_2 \Rightarrow q$. Otherwise, \mathcal{D} has the form

$$\frac{\mathcal{D}_1 \quad \mathcal{D}_2 \quad \iota = (a \rightarrow b) \rightarrow c}{R_1, b \rightarrow c, X_\iota, A_1 \Rightarrow b \quad R_2, \varphi, X \Rightarrow q} (\text{ljt}) \quad \begin{array}{l} X_\iota = X \setminus \{\iota\} \\ \varphi = \bigwedge (A_1 \setminus \{a\}) \rightarrow c \end{array}$$

Obtaining \mathcal{E}_1 (resp. \mathcal{E}_2) from \mathcal{E}_1 (resp. \mathcal{D}_2) by the induction hypothesis, we get the following LJTT-derivation of $R, X \Rightarrow q$. We use here the fact that in LJTT, weakenings are admissible (so we can assume $a \in A_1$), and so are cuts.

$$\frac{\frac{\mathcal{E}_1 \quad R_1, b \rightarrow c, X_\iota, A_1 \Rightarrow b \quad \overline{R_1, X_\iota, c, A_1 \setminus \{a\} \Rightarrow c} \text{Ax}}{R_1, X, A_1 \setminus \{a\} \Rightarrow c} L \rightarrow \rightarrow \quad \mathcal{E}_2}{\frac{\overline{R_1, X \Rightarrow \varphi} L \wedge, R \rightarrow \quad R_2, \varphi, X \Rightarrow q}{R_1, R_2, X \Rightarrow q} (\text{cut})} \square$$

```

1 procedure LJTSatMain( $R_0, X_0, q_0$ )
2   //  $s, R_0, X_0$  are global parameters
3    $s \leftarrow \text{newSolver}()$ 
4   for  $\varphi \in R_0$  do addClause( $s, \varphi$ )
5   for  $(a \rightarrow b) \rightarrow c \in X_0$  do addClause( $s, b \rightarrow c$ )
6    $\tau \leftarrow \text{LJTSat}(R_0, X_0, \emptyset, q_0, [])$ 
7   if  $\tau = \mathcal{K}$  then return  $\mathcal{K}$ 
8   else //  $\tau = (\mathcal{D}, R, \emptyset)$ , with  $R = \{\varphi_1, \dots, \varphi_n\}$ 
9     return the LJTSAT-derivation
10    
$$\frac{R_0, X_0 \vdash_{\text{iPl}} \varphi_1 \quad \dots \quad R_0, X_0 \vdash_{\text{iPl}} \varphi_n \quad \begin{array}{c} \mathcal{D} \\ R_0, R, X_0 \Rightarrow q_0 \end{array}}{R_0, X_0 \Rightarrow q_0} (\text{cut}_{\text{iPl}}^*)$$

11  end
12 end

```

Fig. 6. The procedure LJTSatMain

By Lemma 3 and the soundness of LJTSAT, we conclude:

Theorem 1 (Soundness of LJTSAT). $\vdash_{\text{LJTSAT}} R, X \Rightarrow q$ implies $R, X \vdash_{\text{iPl}} q$.

6 Proof-search using LJTSAT

We present the proof-search procedure based on the calculus LJTSAT, implemented by the main function LJTSatMain (Fig. 6), which exploits the auxiliary recursive function LJTSat (Fig. 7). They correspond to the functions `prove` and `intuitPR` respectively, enhanced with explicit proof/counter-model-construction. The worlds of the counter-models to be constructed are sequences of implication clauses, with $[]$ denoting the empty sequence and $w::\iota$ denoting the extension of sequence w with clause ι . Function LJTSat takes as its last argument the root world of the counter-model to be constructed, if the input is not IPL-valid. Initially, LJTSatMain calls LJTSat with the empty sequence as root world, as shown in Fig. 6, where $(\text{cut}_{\text{iPl}}^*)$ also denotes a chain of n successive applications of $(\text{cut}_{\text{iPl}})$.

A Kripke model $\mathcal{K} = \langle W, \leq, r, \vartheta \rangle$ is a *counter-model* for a sequent $\sigma = R, X \Rightarrow q$, written $\mathcal{K} \not\models \sigma$, if $r \Vdash R \cup X$ and $r \not\Vdash q$. In our incarnation of `intuit`, the counter-model is obtained by gluing Kripke models, as explained next.

Let Θ be a family $(\mathcal{K}_i)_{i \in I}$ of Kripke models and M a classical model such that, for every $i \in I$, the root r_i of \mathcal{K}_i obeys $r_i \Vdash M$. We write $\text{Mod}(r, \Theta, M)$ for the model $\mathcal{K} = \langle W, \leq, r, \vartheta \rangle$ obtained by gluing all the models in Θ over r , where r is a new world such that $\vartheta(r) = M$. More specifically, if $\Theta = (\langle W_i, \leq_i, r_i, \vartheta_i \rangle)_{i \in I}$, where the sets W_i are pairwise disjoint and none of them contains

```

1 procedure LJTSat( $R, X, A, q, r$ )
2    $\tau_0 \leftarrow \text{satProve}(s, A, q)$ 
3   if  $\tau_0 = \text{Yes}(A')$  then
4     return  $(\mathcal{D}, R(s), A')$  where  $\mathcal{D}$  is
5     
$$\frac{R(s), A' \vdash_{\text{cpl}} q}{R(s), X, A' \Rightarrow q} \text{ (cpl)}$$

6   else //  $\tau_0 = \text{No}(M)$ 
7      $\Theta \leftarrow \emptyset$  // Empty family of Kripke models
8     for  $\iota = (a \rightarrow b) \rightarrow c \in X$  such that  $a \notin M$  and  $b \notin M$  and  $c \notin M$  do
9        $\tau_1 \leftarrow \text{LJTSat}(R \cup \{b \rightarrow c\}, X_\iota, M \cup \{a\}, b, r::\iota)$  // where  $X_\iota = X \setminus \{\iota\}$ 
10      if  $\tau_1 = \mathcal{K}_1$  then  $\Theta \leftarrow \Theta \uplus (\iota \mapsto \mathcal{K}_1)$  // adding  $\mathcal{K}_1$  with index  $\iota$ 
11      else //  $\tau_1 = (\mathcal{D}_1, R_1, A_1)$ 
12         $\tilde{\varphi} \leftarrow \bigwedge (A_1 \setminus \{a\}) \rightarrow c$ 
13        addClause( $s, \tilde{\varphi}$ )
14         $\tau_2 \leftarrow \text{LJTSat}(R \cup \{\tilde{\varphi}\}, X, A, q, r)$ 
15        if  $\tau_2 = \mathcal{K}_2$  then return  $\mathcal{K}_2$ 
16        else //  $\tau_2 = (\mathcal{D}_2, R_2, A_2)$ 
17          return  $(\mathcal{D}, R_1 \cup R_2, A_2)$  where  $\mathcal{D}$  is
18          
$$\frac{\begin{array}{c} \mathcal{D}_1 \\ R, R_1, b \rightarrow c, X_\iota, A_1 \Rightarrow b \end{array} \quad \begin{array}{c} \mathcal{D}_2 \\ R, R_2, \tilde{\varphi}, X, A_2 \Rightarrow q \end{array}}{R, R_1, R_2, X, A_2 \Rightarrow q} \text{ (ljt)}$$

19        end
20      end
21    end
22    return Mod( $r, \Theta, M$ )
23  end
24 end

```

Fig. 7. The procedure LJTSat

r , then $\text{Mod}(r, \Theta, M)$ is the model $\mathcal{K} = \langle W, \leq, r, \vartheta \rangle$ such that

$$W = \{r\} \uplus \bigoplus_{i=1}^n W_i \quad \leq_0 = \{(r, r_1), \dots, (r, r_n)\} \cup \bigcup_{i=1}^n \leq_i \quad \vartheta = \bigcup_{i=1}^n \vartheta_i \cup \{(r, M)\}$$

and \leq is the reflexive-transitive closure of \leq_0 (\uplus denotes disjoint union).

If $\Theta = \emptyset$, then $\mathcal{K} = \langle \{r\}, \{(r, r)\}, r, \vartheta \rangle$ only contains the reflexive world r with $\vartheta(r) = M$. Given a set X and a classical model M , we write X_M for $\{((a \rightarrow b) \rightarrow c) \in X \mid a \notin M, b \notin M, c \notin M\}$.

Lemma 4. *Let $\sigma = R, X \Rightarrow q$ be a sequent, r be a world, M be a classical model and Θ be a family $(\mathcal{K}_\iota)_{\iota \in X_M}$ of Kripke models indexed by X_M such that:*

- (i) $(a \rightarrow b) \rightarrow c \in X$ implies $b \rightarrow c \in R$;
- (ii) $M \models R$ and $M \not\models q$;
- (iii) For every $\iota \in X_M$, we have $\mathcal{K}_\iota \not\models R, X_\iota, M, a \Rightarrow b$.

Then, $\text{Mod}(r, \Theta, M) \not\models \sigma$.

Proof. Let \mathcal{K} be $\text{Mod}(r, \Theta, M) = \langle W, \leq, r, \vartheta \rangle$, whence $r \Vdash M$. We have to show that, at r , all the formulae in $R \cup X$ are forced and q is not forced. By (ii), we immediately get $r \not\Vdash q$. We prove the cases for R and X .

Proof that $r \Vdash R$: Suppose $\varphi \in R$ and assume $\varphi = \bigwedge A_1 \rightarrow \bigvee A_2$ with $A_1 \neq \emptyset$. Let $w \in W$ be any world such that $r \leq w$ and $w \Vdash \bigwedge A_1$; we prove $w \Vdash \bigvee A_2$. If $w = r$, we have $r \Vdash \bigwedge A_1$, which implies $A_1 \subseteq M$. Since $M \models \varphi$, we get $M \models \bigvee A_2$, which implies $r \Vdash \bigvee A_2$. If $w \neq r$, w must be in some W_ι for some $\mathcal{K}_\iota = \langle W_\iota, \leq_\iota, r_\iota, \vartheta_\iota \rangle$ in Θ , with $\iota \in X_M$. Thus $r_\iota \leq w$ in \mathcal{K} . By (iii), $r_\iota \Vdash R$ in \mathcal{K}_ι and $\varphi \in R$, so $r_\iota \Vdash \varphi$ in \mathcal{K}_ι . Hence $r_\iota \Vdash \varphi$ in \mathcal{K} , and the persistence of \Vdash gives $w \Vdash \varphi$ in \mathcal{K} . Since $w \Vdash \bigwedge A_1$, we obtain $w \Vdash \bigvee A_2$. The case $A_1 = \emptyset$ (namely, $\varphi = \bigvee A_2$) is similar.

Proof that $r \Vdash X$: First note that by (i), $r \Vdash b \rightarrow c$ for every $(a \rightarrow b) \rightarrow c \in X$. Choose any $\iota = (a \rightarrow b) \rightarrow c \in X$ and let w be any world such that $r \leq w$ and $w \Vdash a \rightarrow b$; we prove $w \Vdash c$.

If $c \in M$, then $r \Vdash c$ by construction, hence $w \Vdash c$. If $b \in M$ then $r \Vdash b$ by construction, and we already have $r \Vdash b \rightarrow c$, so we get $r \Vdash c$, hence $w \Vdash c$. If $a \in M$ then $r \Vdash a$ and $r \Vdash b \rightarrow c$ and $w \Vdash a \rightarrow b$, giving $w \Vdash c$. The previous three cases are independent, thus $w \Vdash c$ if $a \in M$ or $b \in M$ or $c \in M$.

So suppose $a \notin M$ and $b \notin M$ and $c \notin M$. By (iii), Θ contains a model $\mathcal{K}_\iota = \langle W_\iota, \leq_\iota, r_\iota, \vartheta_\iota \rangle$ such that $\mathcal{K}_\iota \not\models R, X_\iota, M, a \Rightarrow b$. Thus $r_\iota \not\Vdash a \rightarrow b$ in \mathcal{K}_ι , hence $r_\iota \not\Vdash a \rightarrow b$ in \mathcal{K} . By reverse persistence $r \not\Vdash a \rightarrow b$ (in \mathcal{K}), which implies $w \neq r$. There is also a model $\mathcal{K}_{\iota'} = \langle W', \leq', r', \vartheta' \rangle$ of Θ such that $w \in W'$, hence $r \leq r' \leq w$ (in \mathcal{K}). We need to cover the two cases $\iota = \iota'$ and $\iota \neq \iota'$:

1. If $\iota' = \iota$, by (iii) $r_\iota \Vdash a$ in \mathcal{K}_ι , thus $w \Vdash a$ (in \mathcal{K}). Since $r \Vdash b \rightarrow c$, we have $w \Vdash b \rightarrow c$. Then $w \Vdash a \rightarrow b$ gives $w \Vdash c$.
2. If $\iota' \neq \iota$ then $\iota \in X_{\iota'}$, hence $r' \Vdash \iota$ in $\mathcal{K}_{\iota'}$, whence it follows that $w \Vdash \iota$ (in \mathcal{K}). Our initial assumption that $w \Vdash a \rightarrow b$ gives $w \Vdash c$. \square

In Fig. 6 we define the proof-search function `LJTSatMain` such that `LJTSatMain(R_0, X_0, q_0)` returns either an `LJTSAT`-derivation of $\sigma_0 = R_0, X_0 \Rightarrow q_0$ or a counter-model \mathcal{K} for σ_0 . Worlds of \mathcal{K} are sequences of implication clauses, ordered by the prefix order on sequences, and the empty sequence $[]$ is its root world. We set:

$$H_0 = \{ b \rightarrow c \mid (a \rightarrow b) \rightarrow c \in X_0 \} \quad V_0 = \{ p \in V \mid p \text{ occurs in } \sigma_0 \}$$

$$M_{/V_0}(R) = \{ M \subseteq V_0 \mid M \models R \} \quad \text{for any set } R \text{ of flat clauses}$$

The call `LJTSatMain(R_0, X_0, q_0)` defines a SAT-solver s and initializes it by storing all the clauses in $R_0 \cup H_0$; we consider s, R_0 and X_0 as global parameters. It exploits the auxiliary recursive procedure `LJTSat` defined in Fig. 7. A call `LJTSat(R, X, A, q, r)` performed during the computation of the main call `LJTSatMain(R_0, X_0, q_0)` has the following specification.

Input Assumptions (IA):

- $R \subseteq \text{R}(s)$ and $X \subseteq X_0$;
- for every $\varphi \in \text{R}(s)$, we have $R_0, X_0 \vdash_{\text{ipl}} \varphi$;
- r is a sequence of implication clauses.

Output Properties (OP):

- $\text{LJTSat}(R, X, A, q, r)$ yields a triple (\mathcal{D}, R', A') or else a model \mathcal{K} with:
 - $R' \subseteq R(s)$ and $A' \subseteq A$;
 - for every $\varphi \in R(s)$, we have $R_0, X_0 \vdash_{\text{iPl}} \varphi$;
 - \mathcal{D} is an LJT_{SAT} -derivation of $R, R', X, A' \Rightarrow q$;
 - \mathcal{K} has root r and worlds are ordered by the prefix order on sequences;
 - $\mathcal{K} \not\models R, H_0, X, A \Rightarrow q$.

In (IA), $R(s)$ refers to the clauses in the SAT-solver s at the beginning of the computation of $\text{LJTSat}(R, X, A, q, r)$; in (OP), $R(s)$ is the set of clauses in s at the end of the computation. Note that (OP) implies that the call $\text{LJTSat}(R, X, A, q)$ terminates. To prove the correctness of LJTSat , we have to show that, if the assumptions (IA) are matched, then (OP) holds. We need the following property about derivability in IPL.

Lemma 5. $R, X, A \vdash_{\text{iPl}} b$ implies $R, X, (a \rightarrow b) \rightarrow c \vdash_{\text{iPl}} \bigwedge(A \setminus \{a\}) \rightarrow c$.

Proof. Let $A' = A \setminus \{a\}$. If $R, X, A \vdash_{\text{iPl}} b$, then $R, X, A' \vdash_{\text{iPl}} a \rightarrow b$, which implies $R, X, (a \rightarrow b) \rightarrow c, A' \vdash_{\text{iPl}} c$, hence $R, X, (a \rightarrow b) \rightarrow c \vdash_{\text{iPl}} \bigwedge A' \rightarrow c$. \square

To prove correctness, put the following order relation on pairs (R, X) such that R is any set of flat clauses and $X \subseteq X_0$:

$$(R', X') \prec (R, X) \quad \text{iff} \quad (X' \subset X) \text{ or } (X' = X \text{ and } M_{/V_0}(R') \subset M_{/V_0}(R))$$

Since the sets X and $M_{/V_0}(R)$ are finite, the relation \prec is well-founded, hence we can prove correctness of LJTSat (Lemma 6) by induction on \prec .

Lemma 6. *If a call $\text{LJTSat}(R, X, A, q, r)$ satisfies (IA) then (OP) holds.*

Proof. We use the main induction hypothesis (MIH) below and show that the invariant (Inv) holds at any point of the computation described in Fig. 7:

- (MIH): if $(R', X') \prec (R, X)$, then the lemma holds for $\text{LJTSat}(R', X', A', q', r')$;
- (Inv): for every $\varphi \in R(s)$, we have $R_0, X_0 \vdash_{\text{iPl}} \varphi$.

At the start of the computation (Inv) holds by (IA). Let τ_0 be the value computed at line 2. If $\tau_0 = \text{Yes}(A')$, then the triple $(\mathcal{D}, R(s), A')$ is returned at line 4, with \mathcal{D} defined at line 5; by definition of satProve , it holds that $R(s), A' \vdash_{\text{cPl}} q$ and $A' \subseteq A$, hence (OP) holds. Otherwise, since $R \cup H_0 \subseteq R(s)$, we have:

(P0) $\tau_0 = \text{No}(M)$ and $M \models R \cup H_0 \cup A$ and $M \not\models q$.

Without loss of generality, we can assume $M \subseteq V_0$, namely $M \in M_{/V_0}(R)$. If, for every $(a \rightarrow b) \rightarrow c \in X$, the loop condition at line 8 does not hold, then the loop at lines 8–21 is skipped and the model $\text{Mod}(r, \emptyset, M)$ is returned at line 22, which is a counter-model for $R, H_0, A \Rightarrow q$ by Lemma 4; thus (OP) holds. Let us assume that the loop at lines 8–21 is entered. We prove that at every iteration of the loop the following properties hold, where τ_1 and τ_2 are the values computed at lines 9 and 14 respectively, and $\tilde{\varphi} = \bigwedge(A_1 \setminus \{a\}) \rightarrow c$ is defined at line 12:

- (P1) $\tau_1 = (\mathcal{D}_1, R_1, A_1)$ or $\tau_1 = \mathcal{K}_1$ where:
- $R_1 \subseteq R(s)$ and $A_1 \subseteq M \cup \{a\}$;
 - \mathcal{D}_1 is an LJT_{SAT} -derivation of $R, R_1, b \rightarrow c, X_\iota, A_1 \Rightarrow b$;
 - \mathcal{K}_1 has root $r::\iota$ and worlds are ordered by the prefix order on sequences;
 - $\mathcal{K}_1 \not\models R, H_0, X_\iota, M, a \Rightarrow b$.
- (P2) $\tau_2 = (\mathcal{D}_2, R_2, A_2)$ or $\tau_2 = \mathcal{K}_2$ where:
- $R_2 \subseteq R(s)$ and $A_2 \subseteq A$;
 - \mathcal{D}_2 is an LJT_{SAT} -derivation of $R, R_2, \tilde{\varphi}, X, A_2 \Rightarrow q$;
 - \mathcal{K}_2 has root r and worlds are ordered by the prefix order on sequences;
 - $\mathcal{K}_2 \not\models R, \tilde{\varphi}, H_0, X, A \Rightarrow q$.

Let us consider the first iteration of the loop and let $\iota = (a \rightarrow b) \rightarrow c \in X$ be the selected clause (hence, $a \notin M$, $b \notin M$ and $c \notin M$). The call to LJTSat at line 9 satisfies (IA). Since $X_\iota \subset X$, we have $(R \cup \{b \rightarrow c\}, X_\iota) \prec (R, X)$. By (MIH) τ_1 satisfies (OP); this proves (P1). Note that (OP) guarantees that (Inv) holds after the computation of τ_1 . At line 13, $\tilde{\varphi}$ is added to s ; we check that (Inv) is preserved, namely $R_0, X_0 \vdash_{\text{ip1}} \tilde{\varphi}$. By (P1) and Soundness of LJT_{SAT} (Theorem 1), $R, R_1, b \rightarrow c, X_\iota, A_1 \vdash_{\text{ip1}} b$ and, by Lemma 5, we get $R, R_1, b \rightarrow c, X \vdash_{\text{ip1}} \tilde{\varphi}$, hence $R, R_1, b \rightarrow c, X_0 \vdash_{\text{ip1}} \tilde{\varphi}$. Since $R \cup R_1 \cup \{b \rightarrow c\} \subseteq R(s)$, from (Inv) it follows that $R_0, X_0 \vdash_{\text{ip1}} \tilde{\varphi}$. The call to LJTSat at line 14 matches (IA). To apply (MIH), we have to check that:

- (P3) $(R \cup \{\tilde{\varphi}\}, X) \prec (R, X)$.

Clearly, $M_{/V_0}(R \cup \{\tilde{\varphi}\}) \subseteq M_{/V_0}(R)$; to conclude the proof of (P3), we show that the inclusion is strict. Note that $M \in M_{/V_0}(R)$ (see (P0) and the subsequent remark). For a contradiction, assume $M \models \tilde{\varphi}$. Since $A_1 \subseteq M \cup \{a\}$, we get $A_1 \setminus \{a\} \subseteq M$. By definition of $\tilde{\varphi}$, it follows that $M \models c$, namely $c \in M$, a contradiction. Thus, $M \not\models \tilde{\varphi}$, which implies $M \notin M_{/V_0}(R \cup \{\tilde{\varphi}\})$; this proves (P3). We can apply (MIH) to the recursive call at line 14 and we get (P2) and the preservation of (Inv). Let us consider the iteration $k+1$ of the loop ($k \geq 1$). We can repeat the above reasoning to prove that (P1) and (P2) hold at iteration $k+1$; the invariant property (Inv) is crucial to guarantee that the recursive calls at lines 9 and 14 satisfy (IA). We conclude that (P1) and (P2) hold at every iteration of the loop.

Let us assume that, at some iteration of the loop, τ_1 is not a model, namely $\tau_1 = (\mathcal{D}_1, R_1, A_1)$. If τ_2 is a model \mathcal{K}_2 , then \mathcal{K}_2 is returned at line 15 and (OP) follows from (P2). Otherwise, $\tau_2 = (\mathcal{D}_2, R_2, A_2)$ and $(\mathcal{D}, R_1 \cup R_2, A_2)$ is returned at line 17, where \mathcal{D} is the LJT_{SAT} -derivation displayed at line 18; accordingly (OP) holds. Finally, let us assume that, at every iteration, τ_1 is a model. Since X is finite, the loop eventually ends and the model $\text{Mod}(r, \Theta, M)$ is returned at line 22. At that point, Θ has been completed into a family $(\mathcal{K}_\iota)_{\iota \in X_M}$ such that for every ι in X_M , \mathcal{K}_ι has root $r::\iota$ and its worlds are ordered by the prefix order on sequences, by (P1). $\text{Mod}(r, \Theta, M)$ has root r and also uses the prefix order on sequences. So in order to prove (OP), we only have to check that $\text{Mod}(r, \Theta, M)$ is a counter-model for $R, H_0, X, A \Rightarrow q$. Let $R' = R \cup H_0 \cup M$ and $\sigma' = R', X \Rightarrow q$; by (P0) and (P1), it follows that σ', M and Θ satisfy the assumptions of Lemma 4. Since $A \subseteq M$, $\text{Mod}(r, \Theta, M) \not\models \sigma'$ and (OP) holds. \square

By Lemma 6, we get:

Theorem 2 (Correctness of LJTSatMain). $\text{LJTSatMain}(R_0, X_0, q_0)$ returns either an LJT_{SAT} -derivation of $\sigma_0 = R_0, X_0 \Rightarrow q_0$ or a counter-model for σ_0 .

Proof. Consider the call $\text{LJTSat}(R_0, X_0, \emptyset, q_0, [])$ at line 6. When LJTSat is called, $R(s) = R_0 \cup H_0$. Let $b \rightarrow c \in H_0$; since X_0 contains a formula of the kind $(a \rightarrow b) \rightarrow c$ and $(a \rightarrow b) \rightarrow c \vdash_{\text{ipl}} b \rightarrow c$, it follows that $R_0, X_0 \vdash_{\text{ipl}} b \rightarrow c$. Thus, the call to LJTSat satisfies (IA); by Lemma 6, the returned value τ satisfies (OP). If τ is a counter-model \mathcal{K} , then \mathcal{K} is returned at line 7; since $\mathcal{K} \not\models R_0, H_0, X_0 \Rightarrow q$, we get $\mathcal{K} \not\models \sigma_0$. Otherwise $\tau = (\mathcal{D}, R, \emptyset)$, where \mathcal{D} is an LJT_{SAT} -derivation \mathcal{D} of $R_0, R, X_0 \Rightarrow q_0$ and $R_0, X_0 \vdash_{\text{ipl}} \varphi$, for every $\varphi \in R$. Accordingly, the returned derivation, displayed at line 10, is an LJT_{SAT} -derivation of σ_0 . \square

As a consequence, we get:

Theorem 3 (LJT_{SAT} completeness). $R, X \vdash_{\text{ipl}} q$ implies $\vdash_{\text{LJT}_{\text{SAT}}} R, X \Rightarrow q$.

Let us consider the call $\text{LJTSatMain}(R_0, X_0, q_0)$; we show that we can build an LJT-derivation \mathcal{D}_φ of $R_0, X_0 \Rightarrow \varphi$, for every clause φ stored in the SAT-solver s during the computation. Let $\varphi = b \rightarrow c$ be a clause introduced at the beginning of LJTSatMain (line 5 in Fig. 6). Then, $(a \rightarrow b) \rightarrow c \in X_0$ and, setting $X' = X_0 \setminus \{(a \rightarrow b) \rightarrow c\}$, \mathcal{D}_φ is:

$$\frac{\frac{\frac{R_0, X', b, b \rightarrow c, a \Rightarrow b}{R_0, X_0, b \Rightarrow c} \quad \frac{R_0, X', c, b \Rightarrow c}{R_0, X_0 \Rightarrow b \rightarrow c} L \rightarrow \rightarrow}{R \rightarrow}}{R_0, X_0 \Rightarrow b \rightarrow c}$$

Let us consider the clause $\tilde{\varphi}$ added in the loop of LJTSat when the clause $\iota = (a \rightarrow b) \rightarrow c$ is considered (line 13 in Fig. 7). By Point (P1) in the proof of Lemma 6, there exists an LJT_{SAT} -derivation \mathcal{D}_1 of $R, R_1, b \rightarrow c, X_\iota, A_1 \Rightarrow b$. Note that $R \cup R_1 \cup \{b \rightarrow c\}$ has the form $R_0 \cup \{\varphi_1, \dots, \varphi_n\}$, where the clauses $\varphi_1, \dots, \varphi_n$ are in s and $X_\iota \subseteq X_0$. Let \mathcal{D}'_1 be the LJT_{SAT} -derivation:

$$\frac{\mathcal{D}_1 \quad R_0, X_0 \vdash_{\text{ipl}} \varphi_1 \quad \dots \quad R_0, X_0 \vdash_{\text{ipl}} \varphi_n \quad R_0, \varphi_1, \dots, \varphi_n, X_\iota, A_1 \Rightarrow b}{R_0, X_0, A_1 \Rightarrow b} (\text{cut}_{\text{ipl}}^*)$$

By construction, for every judgment $R_0, X_0 \vdash_{\text{ipl}} \varphi'$ occurring in \mathcal{D}'_1 , the clause φ' is in s , thus we can assume that the LJT-derivation $\mathcal{D}_{\varphi'}$ has already been defined. We can turn \mathcal{D}'_1 into an LJT-derivation \mathcal{E}_1 of $R_0, X_0, A_1 \Rightarrow b$.

If $a \notin A_1$, then $\tilde{\varphi} = \bigwedge A_1 \rightarrow c$ and $\mathcal{D}_{\tilde{\varphi}}$ is the LJT-derivation

$$\frac{\mathcal{E}_1 \quad \frac{\frac{b, b \rightarrow c, a \Rightarrow b}{b, (a \rightarrow b) \rightarrow c \Rightarrow c} \text{Ax} \quad \frac{b, c \Rightarrow c}{L \rightarrow \rightarrow} \text{Ax}}{R_0, X_0, A_1 \Rightarrow c} (\text{cut})}{R_0, X_0 \Rightarrow \bigwedge A_1 \rightarrow c} L \wedge, R \rightarrow$$

If $A_1 = \emptyset$ (hence $\tilde{\varphi} = c$) the bottom applications of $L\wedge$ and $R\rightarrow$ are crossed out. Let $a \in A_1$ and $A'_1 = A_1 \setminus \{a\}$. Thus, $\tilde{\varphi} = \bigwedge A' \rightarrow c$ and $\mathcal{D}_{\tilde{\varphi}}$ is the derivation

$$\frac{\frac{\frac{\mathcal{E}_1}{R_0, X_0, A'_1, a \Rightarrow b} R \rightarrow}{R_0, X_0, A'_1 \Rightarrow a \rightarrow b} \quad \frac{\frac{\frac{\frac{b, b \rightarrow c, a \Rightarrow b}{a \rightarrow b, b \rightarrow c, a \Rightarrow b} \text{Ax}}{a \rightarrow b, (a \rightarrow b) \rightarrow c \Rightarrow c} L0 \rightarrow}{a \rightarrow b, (a \rightarrow b) \rightarrow c \Rightarrow c} \text{Ax}}{a \rightarrow b, (a \rightarrow b) \rightarrow c \Rightarrow c} L \rightarrow \rightarrow}{\frac{R_0, X_0, A'_1 \Rightarrow c}{R_0, X_0 \Rightarrow \bigwedge A'_1 \rightarrow c} L\wedge, R \rightarrow} \text{(cut)}$$

If $A' = \emptyset$ (hence $\tilde{\varphi} = c$) the bottom applications of $L\wedge$ and $R\rightarrow$ are skipped.

By the above discussion, we can enhance the procedures `LJTSatMain` and `LJTSat` so that, whenever a flat clause φ is added to the SAT-solver, an LJT-derivation \mathcal{D}_φ of $R_0, X_0 \Rightarrow \varphi$ is stored. Let us assume that `LJTSatMain`(R_0, X_0, q_0) returns an `LJTSAT`-derivation \mathcal{D} of $\sigma_0 = R_0, X_0 \Rightarrow q_0$. Proceeding as in the proof of Lemma 3, we can exploit the derivations \mathcal{D}_φ to translate \mathcal{D} into an LJT-derivation of σ_0 .

7 Discussion, Further Work and Conclusions

The construction of Kripke countermodels from failed proof search in (a variant of) LJT was already explored by Dyckhoff and Pinto [3]. In this paper we have merged proof search and countermodel construction into one procedure based on LJT while benefitting at the same time from the insights from `intuit` regarding incremental SAT-solvers.

The rule ($L\rightarrow\rightarrow$) from Fig. 4 can be interpreted semantically as follows by reading it from conclusion to premises. The antecedent of the conclusion requires the current world w to make $(a \rightarrow b) \rightarrow c$ true, to make all members of Γ true and to make δ false. If $w \Vdash c$ then we have the right premise. Else $w \not\Vdash c$ and therefore, $w \not\Vdash a \rightarrow b$. But that means that there exists a $v \geq w$ such that $v \Vdash \Gamma$ and $v \Vdash a$ and $v \not\Vdash b$, implying that $v \Vdash b \rightarrow c$, which is the left premise.

We have shown that our simpler recursive version of `intuit` can be reconciled with this semantic view if we generalise the rule ($L\rightarrow\rightarrow$) into the rule (`ljt`). By doing so, we can utilise an incremental SAT-solver to implement the right premise of the rule (`ljt`) by “restarting” the SAT-solver with additional flat clauses learned during the process of finding the derivation of the left premise of (`ljt`).

There are many sequent and natural deduction calculi that contain rules which have this “here” or “at some successor” flavour. For example, the LSJ calculus of Ferrari et al [5] and the traditional tableau calculus for linear temporal logic PLTL [16]. Can we extend our insights to such calculi to obtain **incremental** SAT-based decision procedures for these calculi too [10]?

Another direction for future work is the extent to which this approach relies on the clausification of the input formula. Indeed, LJT is able to natively treat any IPL formula. Technically, could the calculus `LJTSAT` be extended to any sequent, not necessarily in clausal form? If so how would the interaction with the SAT solver be organised?

Acknowledgment This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 689176.

References

1. Koen Claessen and Dan Rosén. SAT modulo intuitionistic implications. In *LPAR*, volume 9450 of *Lecture Notes in Computer Science*, pages 622–637. Springer, 2015.
2. Roy Dyckhoff. Contraction-free sequent calculi for intuitionistic logic. *J. Symb. Log.*, 57(3):795–807, 1992.
3. Roy Dyckhoff and L Pinto. Implementation of a loop-free method for construction of counter-models for intuitionistic propositional logic, 1996. University of St Andrews Report CS/96/8.
4. Mahfuza Farooque, Stéphane Graham-Lengrand, and Assia Mahboubi. A bisimulation between DPLL(T) and a proof-search strategy for the focused sequent calculus. In Alberto Momigliano, Brigitte Pientka, and Randy Pollack, editors, *Proceedings of the 2013 International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2013)*. ACM Press, September 2013.
5. Mauro Ferrari, Camillo Fiorentini, and Guido Fiorino. Contraction-free linear depth sequent calculi for intuitionistic propositional logic with the subformula property and minimal depth counter-models. *J. Autom. Reasoning*, 51(2):129–149, 2013.
6. Gerhard Gentzen. Untersuchungen über das logische schließen. i. *Mathematische Zeitschrift*, 39:176–210, 1934.
7. Gerhard Gentzen. Investigations into logical deduction i. In M. Szabo, editor, *The collected papers of Gerhard Gentzen*. North-Holland, Amsterdam, 1969.
8. Stéphane Graham-Lengrand. Psyche: a proof-search engine based on sequent calculus with an LCF-style architecture. In Didier Galmiche and Dominique Larchey-Wendling, editors, *Automated Reasoning with Analytic Tableaux and Related Methods - 22th International Conference, TABLEAUX 2013, Nancy, France, September 16-19, 2013. Proceedings*, volume 8123 of *Lecture Notes in Computer Science*, pages 149–156. Springer, 2013.
9. Stéphane Graham-Lengrand. *Polarities & Focussing: a journey from Realisability to Automated Reasoning*. Habilitation thesis, Université Paris-Sud, 2014.
10. Jianwen Li, Shufang Zhu, Geguang Pu, Lijun Zhang, and Moshe Y. Vardi. Sat-based explicit ltl reasoning and its application to satisfiability checking. *Formal Methods in System Design*, <https://doi.org/10.1007/s10703-018-00326-5>, 2019.
11. Robert Nieuwenhuis, Alberto Oliveras, and Cesare Tinelli. Solving SAT and SAT modulo theories: from an abstract Davis-Putnam-Logemann-Loveland procedure to DPLL(T). *Journal of the ACM*, 53(6):937–977, 2006.
12. Thomas Rath, Jens Otten, and Christoph Kreitz. The ILTP problem library for intuitionistic logic. *Journal of Automated Reasoning*, 38(1):261–271, Apr 2007.
13. Richard Statman. Intuitionistic propositional logic is polynomial-space complete. *Theoretical Computer Science*, 9:67–72, 07 1979.
14. A.S. Troelstra and H. Schwichtenberg. *Basic Proof Theory*, volume 43 of *Cambridge Tracts in Theoretical Computer Science*. Camb. Univ. Press, 2ed edition, 2000.
15. G. S. Tseitin. *On the Complexity of Derivation in Propositional Calculus*, pages 466–483. Springer Berlin Heidelberg, Berlin, Heidelberg, 1983.
16. Pierre Wolper. Temporal logic can be more expressive. *Information and Control*, 56(1/2):72–99, 1983.