

Elementary Elimination of Prenex Cuts in Disjunction-free Intuitionistic Logic*

Matthias Baaz¹ and Christian G. Fermüller²

- 1 Institute of Discrete Mathematics and Geometry
Vienna University of Technology, Austria
baaz@logic.at
- 2 Theory and Logic Group 185.2
Vienna University of Technology, Austria
chrisf@logic.at

Abstract

The size of shortest cut-free proofs of first-order formulas in intuitionistic sequent calculus is known to be non-elementary in the worst case in terms of the size of given sequent proofs with cuts of the same formulas. In contrast to that fact, we provide an elementary bound for the size of cut-free proofs for disjunction-free intuitionistic logic for the case where the cut-formulas of the original proof are prenex. To emphasize the non-triviality of our result, we establish non-elementary lower bounds for classical disjunction-free proofs with prenex cut-formulas and intuitionistic disjunction-free proofs with non-prenex cut-formulas.

1998 ACM Subject Classification F.4.1 Mathematical Logic (Proof Theory)

Keywords and phrases Cut-elimination, sequent calculus, intuitionistic logic

Digital Object Identifier 10.4230/LIPIcs.CSL.2015.94

1 Introduction

The elimination of cuts (*viz.* lemmas) from given sequent calculus proofs has remained in the focus of proof theory ever since Gentzen's seminal paper [4] from 1934/35. It is well known that the worst case complexity of cut-elimination is non-elementary for first-order intuitionistic as well as classical logic [8, 7, 3]. More precisely, there is a sequence of formulas $\langle F_i \rangle_{i \in \omega}$, where F_i has an **LI**-proof of length $\leq f(i)$, for some elementary function $f(i)$, while the shortest cut-free **LI**-proof of F_i is of length $\geq g(i)$ for some non-elementary function $g(i)$. (This means that $g(i)$ grows faster than $2^{\cdot^{\cdot^{\cdot^2^i}}}$, for any stack of 2s that is of constant height). Here **LI** is Gentzen's sequent calculus for first-order intuitionistic logic. The same result holds for the classical sequent calculus **LK**. The result is very robust: it does not matter whether we define the length of a proof as the number of symbols, formulas or just inference steps in it; moreover we may use any of the many known variants of Gentzen's original calculus.

It seems to be difficult to extract tight upper bounds from Gentzen's original cut-elimination procedure for **LI** [4]. Hudelmaier [5] provides a quadruple exponential upper bound for a suitable variant of *propositional LI*. However no elementary upper bound for cut-elimination seems to be known for non-trivial genuine first-order fragments of intuitionistic logics. The purpose of this paper is to show that one can in fact eliminate cuts involving

* This work was partly supported by Austrian Science Fund (FWF) projects P25417 and P26976.



only *prenex* cut-formulas from *disjunction-free* intuitionistic sequent proofs without non-elementary increase in the size of proofs. To obtain that elementary bound a new type of a cut-elimination argument is presented. The result is sharp in at least two respects. As we will show, in the following cases there exist non-elementary lower bounds for cut-elimination: (1) *classical* disjunction-free sequent derivations with prenex cut-formulas, and (2) intuitionistic disjunction-free sequent derivations with *non-prenex cut-formulas*.

The paper is organized as follows. In Section 2 we introduce the sequent calculus $\mathbf{LI}_m^{-\vee}$ for disjunction-free intuitionistic logic and fix corresponding terminology. We then present the overall procedure for eliminating prenex cut-formulas from $\mathbf{LI}_m^{-\vee}$ -derivations in three steps. First, in Section 3, we consider the special case, where all cut-formulas in the given derivation are quantified atomic formulas. We then use this result in Section 4 in a transformation that trades arbitrarily complex prenex cut-formulas for propositional cut-formulas. In Section 5 we show how the remaining propositional cut-formulas can be eliminated from $\mathbf{LI}_m^{-\vee}$ -derivations. In all three cases the depth of the resulting derivation will be elementarily bounded by the depth of the derivation we start with. In Section 6, we show that not only the depth, but also the size of the final cut-free proof is elementarily bounded in terms of the size of the original proof with arbitrary complex prenex cuts. While this follows straightforwardly for languages without function symbol, a further transformation step is involved in the presence of function symbols. In Section 7 we contrast the elementary upper bound for the elimination of prenex cuts with two cases where there exists a non-elementary lower bound for cut-elimination: disjunction-free classical logic with prenex atomic cuts and disjunction-free intuitionistic logic with non-prenex cuts.

2 Preliminaries

We work in a standard first-order language without identity. Terms are built up from constants and variables using function symbols, as usual. We follow Gentzen in syntactically distinguishing free variables a_1, a_2, \dots and bound variables $x_1, x_2, \dots, y_1, \dots$. Atomic formulas – *atoms*, for short – are of the form $P(t_1, \dots, t_n)$, where P is a n -ary predicate symbol and t_1, \dots, t_n are terms. Formulas of (general) intuitionistic logic are built up from atoms using the propositional connectives $\neg, \wedge, \vee, \supset$ and the quantifiers \forall, \exists . If there are no occurrences of \vee we speak of *disjunction-free* intuitionistic logic. The size $|F|$ of a formula F is the number of symbols occurring in it. A formula is *prenex* if it is of the form $Q_1x_1 \dots Q_nx_nA$, where A is propositional, i.e., quantifier-free. If the quantifier-free part A of a prenex formula is an atom we speak of a *prenex atom*.

We consider the following variant of Gentzen's original calculus \mathbf{LI} that we will refer to as $\mathbf{LI}_m^{-\vee}$. There is at most one formula at the right side of the sequent arrow, denoted here as \vdash ; whereas on the left hand side we may have any finite multiset of formulas. In the following rules Γ and Π denote arbitrary finite multisets of formulas, Δ is either a single formula or empty. Multiset-union is denoted not by the comma, as usual. As usual, we write Γ, Π instead of $\Gamma \uplus \Pi$ and Γ, A instead of $\Gamma \uplus \{A\}$, etc, where \uplus is the union operator for multisets.

Axioms:

$A \vdash A$ where A is atomic

Cut Rule:

$$\frac{\Gamma \vdash A \quad A, \Pi \vdash \Delta}{\Gamma, \Pi \vdash \Delta} \text{ (cut)}$$

Structural Rules:

$$\frac{\Gamma, A, A \vdash \Delta}{\Gamma, A \vdash \Delta} (\text{contr}) \quad \frac{\Gamma \vdash \Delta}{A, \Gamma \vdash \Delta} (\text{weak}, l) \quad \frac{\Gamma \vdash}{\Gamma \vdash A} (\text{weak}, r)$$

Logical Rules (Propositional and Quantifier Rules):

$$\frac{\Gamma, A \vdash}{\Gamma \vdash \neg A} (\neg, r) \quad \frac{\Gamma \vdash A}{\Gamma, \neg A \vdash} (\neg, l) \quad \frac{\Gamma, A \vdash B}{\Gamma \vdash A \supset B} (\supset, r) \quad \frac{\Gamma \vdash A \quad B, \Pi \vdash \Delta}{\Gamma, \Pi, A \supset B \vdash \Delta} (\supset, l)$$

$$\frac{\Gamma \vdash A \quad \Pi \vdash B}{\Gamma, \Pi \vdash A \wedge B} (\wedge, r) \quad \frac{\Gamma, A \vdash \Delta}{\Gamma, A \wedge B \vdash \Delta} (\wedge_1, l) \quad \frac{\Gamma, B \vdash \Delta}{\Gamma, A \wedge B \vdash \Delta} (\wedge_2, l)$$

$$\frac{\Gamma \vdash A(a)}{\Gamma \vdash \forall x A(x)} (\forall, r) \quad \frac{\Gamma, A(t) \vdash \Delta}{\Gamma, \forall x A(x) \vdash \Delta} (\forall, l) \quad \frac{\Gamma \vdash A(t)}{\Gamma \vdash \exists x A(x)} (\exists, r) \quad \frac{\Gamma, A(a) \vdash \Delta}{\Gamma, \exists x A(x) \vdash \Delta} (\exists, l)$$

In (\forall, r) and (\exists, l) a denotes an eigenvariable, i.e., a variable that does not occur in Γ . In (\forall, l) and (\exists, r) t denotes an arbitrary term.

Besides the missing rules for disjunction, $\mathbf{LI}_m^{-\vee}$ differs from Gentzen's original intuitionistic sequent calculus \mathbf{LI} in the following respects: (1) $\mathbf{LI}_m^{-\vee}$ is based on multisets instead of Gentzen's sequences of formulas, which allows us to dispense with the exchange rule of \mathbf{LI} . (2) Whereas Gentzen uses additive rules for introducing connectives, the logical rules of $\mathbf{LI}_m^{-\vee}$ are multiplicative, except for the left conjunction rules. The subscript m in $\mathbf{LI}_m^{-\vee}$ is intended as a reminder on this fact. (3) We insist on atomic axioms; i.e., axioms where the exhibited formula is atomic.

An $\mathbf{LI}_m^{-\vee}$ -derivation of an *end-sequent* $\Gamma \vdash \Delta$ is an upward growing tree of sequents, obtained from instantiating the above rules as usual, starting with the root node $\Gamma \vdash \Delta$. If the end-sequent is $\vdash F$ we speak of a *proof of F*.

We need a few further technical notions for investigating $\mathbf{LI}_m^{-\vee}$ -derivations. The exhibited formula occurrence (A) in the left and right upper sequents of the *cut*-rule is called *cut-formula*. The exhibited formula occurrence in the lower sequent of a structural, propositional or quantifier rule is called the *principal formula (occurrence)* of the corresponding inference. The formula occurrences exhibited in the upper sequent are called *immediate ancestors* of the corresponding principal formula in the lower sequent. The formulas in Γ, Π, Δ are called *side formulas*. Each occurrence of a side formula, say F , in the lower sequent of a rule has a unique corresponding occurrence of the same side formula F in (one of) the corresponding upper sequent(s). This upper occurrence of F is called the *immediate predecessor* of the lower occurrence of F .

Given a derivation γ , any occurrence F of a formula in γ spawns an *ancestor tree* $\tau_\gamma(F)$ defined inductively as follows:

- the given occurrence of F is the root of $\tau_\gamma(F)$;
- if G is a node in $\tau_\gamma(F)$, where G is principal formula of an inference in γ , then the immediate ancestor(s) of G in γ are (is) the successor node(s) of G ;
- if G is a node in $\tau_\gamma(F)$, where G is a side formula of an inference in γ , then the immediate predecessor of G in γ is the successor node of G .

The *height* $h(\tau_\gamma(F))$ of the ancestor tree is defined as usual, where $h(\tau_\gamma(F)) = 0$ if $\tau_\gamma(F)$ consists only of the root F . Note that an ancestor tree only branches at nodes that are occurrences of principal formulas of contractions, i.e., applications of (contr) , or else of the rules (\supset, l) or (\wedge, r) . Every leaf node of an ancestor tree occurs either in axiom or at the lower sequent of a weakening, i.e. an application of (weak, l) or (weak, r) .

A derivation π is *regular* if each application of (\forall, r) and (\exists, l) in π is associated with a unique eigenvariable, which is converted into a bound variable by the corresponding inference.

A derivation π is *pruned* if every branch of π contains any sequent at most once and moreover each formula occurs at most three times on the left hand side of any sequent.

The size $|\pi|$ of a derivation π (in $\mathbf{LI}_m^{-\forall}$ or any other sequent calculus) is the sum of the sizes of all formulas occurrences in π . The height $h(\pi)$ of π is the largest number of consecutive inferences in π . (In other words, $h(\pi)$ is the maximal length of a branch of π .)

3 Prenex atomic cut-formulas

As outlined in the Introduction, we present the overall cut-elimination procedure in three stages. First, in this section, we consider the special case, where all cut-formulas in the given $\mathbf{LI}_m^{-\forall}$ -derivation are quantified atoms.

We start with the following simple, but crucial observation.

► **Lemma 1.** *Let γ be a cut-free $\mathbf{LI}_m^{-\forall}$ -derivation of $\Gamma \vdash A$, where A is a quantified atom $\mathbf{Q}\vec{x}P(\vec{t})$. Then the ancestor tree $\tau_\gamma(A)$ is not branching. I.e., $\tau_\gamma(A)$ consists in a unique thread of formula occurrences A_1, \dots, A_n , where A_1 is the indicated occurrence of $A = \mathbf{Q}\vec{x}P(\vec{t})$ and A_{i+1} is either the immediate ancestor or the immediate predecessor of A_i in γ .*

Proof. It suffices to observe that all nodes in $\tau_\gamma(A)$ are quantified or unquantified atomic formulas that occur on the right hand side of a sequent. Therefore, each non-leaf node in $\tau_\gamma(A)$ has a unique successor that is either its immediate predecessor in γ or else the immediate ancestor of it for an inference (\forall, r) or (\exists, r) . ◀

Remember that Gentzen [4] replaced the cut-rule by the mix-rule in order to formulation his argument for the eliminability of cuts. We will use a different generalization of (*cut*), called *multi-cut rule* (cut^+):

$$\frac{\Gamma_1 \vdash A_1 \quad \dots \quad \Gamma_n \vdash A_n \quad A_1, \dots, A_n, \Pi \vdash \Delta}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta} (cut^+)$$

Clearly, for $n = 1$ (cut^+) coincides with the ordinary cut-rule (*cut*). On the other hand, (cut^+) is readily simulated by n applications of (*cut*). In the the following, we will assume that in $\mathbf{LI}_m^{-\forall}$ (*cut*) is replaced by (cut^+).

We will call the rightmost upper sequent of an instance of (cut^+) its *main premise*. Let us call the occurrences of A_1, \dots, A_n in the main premise *lhs cut-formulas* (since they occur on the left hand side of the sequent). By the *lhs-depth* of an instance of the multi-cut rule in a derivation γ we mean the maximal height of the ancestor trees $\max(d(\tau_\gamma(A_1)), \dots, d(\tau_\gamma(A_n)))$, where A_1, \dots, A_n are the lhs cut-formulas of this instance of (cut^+).

► **Theorem 2.** *Let π be an $\mathbf{LI}_m^{-\forall}$ -derivation of $\Gamma \vdash \Delta$, where each cut-formula is a prenex atom. Then there exists an elementary function f , such that the following holds: there exists a cut-free $\mathbf{LI}_m^{-\forall}$ -derivation π_0 of $\Gamma \vdash \Delta$, such that $h(\pi_0) \leq f(h(\pi))$.*

Proof. Throughout the proof we will assume implicitly that π is regular and that regularity is restored at each transformation step by using variable-renamed copies of sub-derivations where needed. We first focus on the elimination of (multi-)cuts and investigate the increase in complexity separately. Therefore we may assume without loss of generality that the last inference of π is the only instance of (cut^+) in π . In contrast to Gentzen's procedure (and its variants) we do not need nested induction in our case, but only induction over the lhs-depth d of the (only) multi-cut.

$d = 0$:

This entails that $n = 1$, since otherwise one of the lhs cut-formulas must have been introduced

already earlier in π (i.e., above the main premise), contradicting the assumption that the ancestor trees of the lhs cut-formulas consist of only those formulas themselves. There are two cases:

- (1): If the main premise is an axiom then the application of the cut-rule is clearly redundant.
 (2): If the cut-formula in the main premise has been introduced by weakening, then π has the form

$$\frac{\begin{array}{c} \vdots \\ \Gamma \vdash A \end{array} \quad \frac{\begin{array}{c} \vdots \\ \Pi \vdash \Delta \\ A, \Pi \vdash \Delta \end{array} \text{ (weak, } l\text{)}}{A, \Pi \vdash \Delta} \text{ (cut}^+\text{)}}{\Gamma, \Pi \vdash \Delta}$$

and therefore the end-sequent $\Gamma, \Pi \vdash \Delta$ can be obtained without cut by continuing the upper right sub-derivation of $\Pi \vdash \Delta$ by iterated weakening to restore Γ .

$d > 0$:

We distinguish the following cases according to the type of the inference that has the main premise of the multi-cut as its lower sequent. We will refer to this inference as the ‘relevant inference’ in the following.

- (1): If the principal formula of the relevant inference is not among the cut-formula the following subcases arise.

- (1.1): The relevant inference is a unary propositional rule. We only present the case for (\wedge_1, l) , since (\wedge_1, r) , (\supset, r) , (\neg, l) , (\neg, r) are treated analogously. π thus has the form

$$\frac{\begin{array}{c} \vdots \\ \Gamma_1 \vdash A_1 \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \quad \frac{\begin{array}{c} \vdots \\ A_1, \dots, A_n, C, \Pi \vdash \Delta \\ A_1, \dots, A_n, C \wedge D, \Pi \vdash \Delta \end{array} \text{ (}\wedge_1, l\text{)}}{A_1, \dots, A_n, C \wedge D, \Pi \vdash \Delta} \text{ (cut}^+\text{)}}{\Gamma_1, \dots, \Gamma_n, C \wedge D, \Pi \vdash \Delta}$$

To decrease d , π is transformed into

$$\frac{\begin{array}{c} \vdots \\ \Gamma_1 \vdash A_1 \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \quad \frac{\begin{array}{c} \vdots \\ A_1, \dots, A_n, C, \Pi \vdash \Delta \\ \Gamma_1, \dots, \Gamma_n, C, \Pi \vdash \Delta \end{array} \text{ (cut}^+\text{)}}{\Gamma_1, \dots, \Gamma_n, C, \Pi \vdash \Delta} \text{ (}\wedge_1, l\text{)}}{\Gamma_1, \dots, \Gamma_n, C \wedge D, \Pi \vdash \Delta}$$

- (1.2): The relevant inference is a binary propositional rule, We present the case for (\wedge, r) ; the case for (\supset, l) is similar. π has the form

$$\frac{\begin{array}{c} \vdots \\ \Gamma_1 \vdash A_1 \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \quad \frac{\begin{array}{c} \vdots \\ A_l^*, \Pi_l \vdash C \quad A_r^*, \Pi_r \vdash D \\ A_1, \dots, A_n, \Pi \vdash C \wedge D \end{array} \text{ (}\wedge, r\text{)}}{A_1, \dots, A_n, \Pi \vdash C \wedge D} \text{ (cut}^+\text{)}}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash C \wedge D}$$

where $\Pi_l \uplus \Pi_r = \Pi$ and $A_l^* \uplus A_r^* = A_1, \dots, A_n$, by which we mean that the multiset Π partitions into the disjoint sub-multisets Π_l and Π_r . Similarly $A_l^* = A_{i_1}, \dots, A_{i_k}$ and $A_r^* = A_{j_1}, \dots, A_{j_m}$, where $\{i_1, \dots, i_k\} \cup \{j_1, \dots, j_m\} = \{1, \dots, n\}$ are disjoint subsets of indices.

Let γ_l denote the derivation

$$\frac{\begin{array}{c} \vdots \\ \Gamma_{i_1} \vdash A_{i_1} \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma_{i_k} \vdash A_{i_k} \end{array} \quad \frac{\begin{array}{c} \vdots \\ A_l^*, \Pi_l \vdash C \\ \Gamma_{i_1}, \dots, \Gamma_{i_k}, \Pi_l \vdash C \end{array} \text{ (cut}^+\text{)}}{A_l^*, \Pi_l \vdash C}$$

and let γ_r denote the derivation

$$\frac{\begin{array}{c} \vdots \\ \Gamma_{j_1} \vdash A_{j_m} \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma_{j_1} \vdash A_{j_m} \end{array} \begin{array}{c} \vdots \\ A_r^*, \Pi_r \vdash C \end{array}}{\Gamma_{j_1}, \dots, \Gamma_{j_m}, \Pi_r \vdash C} \text{ (cut}^+\text{)}$$

Then π is replaced by

$$\frac{\gamma_l}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash C} \frac{\gamma_r}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash C \wedge D} (\wedge, r)$$

- (1.3):** The relevant inference is a quantifier rule. We only present the case for (\forall, r) ; the case for (\exists, l) is analogous. The cases for (\forall, l) and (\exists, r) are similar, but even simpler, since they not involve variable renaming. For (\forall, r) π has the form

$$\frac{\begin{array}{c} \vdots \\ \Gamma_1 \vdash A_1 \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \frac{\begin{array}{c} \vdots \delta \\ A_1, \dots, A_n, \Pi \vdash F(a) \end{array}}{A_1, \dots, A_n, \Pi \vdash \forall x F(x)} (\forall, r)}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \forall x F(x)} \text{ (cut}^+\text{)}$$

To decrease d , π is transformed into

$$\frac{\begin{array}{c} \vdots \\ \Gamma_1 \vdash A_1 \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \frac{\begin{array}{c} \vdots \delta' \\ A_1, \dots, A_n, \Pi \vdash F(b) \end{array}}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash F(b)} (\text{cut}^+)}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \forall x F(x)} (\forall, r)$$

where b is a free variable that does not occur in $\Gamma_1, \dots, \Gamma_n$ and δ' arises from the δ by renaming a to b everywhere in this derivation.

- (1.4):** If the relevant inference is a structural rule, then the application of (cut^+) can straightforwardly be shifted upwards, similarly to the cases above.
- (2):** If the principal formula of the relevant inference is a cut-formula then the following sub-cases arise.
- (2.1):** If the relevant inference is a contraction (contr) operating on one of the cut-formulas, then π has the form

$$\frac{\begin{array}{c} \vdots \delta \\ \Gamma_1 \vdash A_1 \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \frac{\begin{array}{c} \vdots \\ A_1, A_1, \dots, A_n, \Pi \vdash \Delta \end{array}}{A_1, \dots, A_n, \Pi \vdash \Delta} (\text{contr})}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta} \text{ (cut}^+\text{)}$$

To decrease d , π is transformed into

$$\frac{\begin{array}{c} \vdots \delta' \\ \Gamma_1 \vdash A_1 \end{array} \begin{array}{c} \vdots \\ \Gamma_1 \vdash A_1 \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \frac{\begin{array}{c} \vdots \\ A_1, A_1, \dots, A_n, \Pi \vdash \Delta \end{array}}{\Gamma_1, \Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta} (\text{cut}^+)}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta} (\text{contr})^*$$

where $(\text{contr})^*$ denotes a series of contractions and the sub-derivation δ' is obtained from δ by renaming free variables to ensure regularity (if needed).

(2.2): The relevant inference introduces a cut-formula by weakening. Then π has the form

$$\frac{\frac{\begin{array}{c} \vdots \\ \Gamma_1 \vdash A_1 \end{array} \cdots \frac{\begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \quad \frac{A_2, \dots, A_n, \Pi \vdash \Delta}{A_1, A_2, \dots, A_n, \Pi \vdash \Delta} (weak, l)}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta} (cut^+)}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta} (cut^+)$$

To decrease d , π is transformed into

$$\frac{\frac{\begin{array}{c} \vdots \\ \Gamma_2 \vdash A_2 \end{array} \cdots \frac{\begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \quad \frac{A_2, \dots, A_n, \Pi \vdash \Delta}{A_2, \dots, A_n, \Pi \vdash \Delta} (cut^+)}{\Gamma_2, \dots, \Gamma_n, \Pi \vdash \Delta} (weak, l)^*}{\Gamma_1, \Gamma_2, \dots, \Gamma_n, \Pi \vdash \Delta} (weak, l)^*$$

where $(weak, l)^*$ denotes a series of weakenings.

(2.3): The relevant inference introduces a quantifier of a cut-formula. We present the case for (\exists, l) . (The case for (\forall, l) is similar.) Thus π has the form

$$\frac{\frac{\begin{array}{c} \vdots \delta \\ \Gamma_1 \vdash \exists x A(x) \end{array} \cdots \frac{\begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \quad \frac{A(a), A_2, \dots, A_n, \Pi \vdash \Delta}{\exists x A(x), A_2, \dots, A_n, \Pi \vdash \Delta} (\forall, r)}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta} (cut^+)}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta} (cut^+)$$

where we first assume that the derivation δ has the following form:

$$\frac{\frac{\begin{array}{c} \vdots \gamma_0 \\ \Pi_1 \vdash A(t) \end{array}}{\Pi_1 \vdash \exists x A(x)} (\exists, r)}{\Gamma_1 \vdash \exists x A(x)} \gamma^- \cdots \vdots$$

By Lemma 1, the ancestor tree of the occurrence of $\exists x A(x)$ in the end-sequent of δ consists in a single branch σ of formula occurrences. The indicated instance of (\exists, r) denotes the location in σ , where $\exists x A(x)$ has $A(t)$ as its immediate ancestor; γ^- denotes the part of δ that is obtained by removing the sub-derivation γ_0 of $\Pi_1 \vdash \exists x A(x)$ from δ . The derivation replacing π , whereby d is decreased, can now be presented as

$$\frac{\frac{\begin{array}{c} \vdots \gamma_0 \\ \Pi_1 \vdash A(t) \end{array} \cdots \frac{\begin{array}{c} \vdots \\ \Gamma_n \vdash A_n \end{array} \quad \frac{A(t), A_2, \dots, A_n, \Pi \vdash \Delta}{A(t), A_2, \dots, A_n, \Pi \vdash \Delta} (\forall, r)}{\Pi_1, \dots, \Gamma_n, \Pi \vdash \Delta} (cut^+)}{\Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta} \gamma^+ \cdots \vdots$$

where η' is obtained from η by substituting all occurrences of a by t and γ^+ is obtained from γ^- by replacing the occurrences of $\exists x A(x)$ in the ancestor tree σ by Δ and additionally adding $\Gamma_2, \dots, \Gamma_n, \Pi$ at the left hand side of those sequents where Δ has replaced $\exists x A(x)$.

The other subcase arises if the uppermost occurrence of $\exists x A(x)$ in the derivation δ is not introduced by an application of (\exists, r) , as above, but by weakening. In other words γ_0 ends in $\Pi_1 \vdash$. This derivation leads to a cut-free derivation of $\Gamma_1 \vdash$ and therefore also one of $\Gamma_1, \dots, \Gamma_n, \Pi \vdash \Delta$ by iterated weakening.

This concludes the description of steps for shifting and eventually eliminating an uppermost cut. By repeating the argument, we arrive at the desired cut-free proof π_0 of $\Gamma \vdash \Delta$.

As for the complexity bound, we first investigate the increase in the height of the derivation for the elimination of a single cut. To this aim, note that for any sequent in the original derivation π the maximal number of occurrences of formulas at the left hand side is not increased throughout the cut-elimination procedure. This number is clearly exponentially bounded in terms of $h(\pi)$. Next we look at single steps of moving the cut upwards in the derivation, as indicated in the various cases of the inductive proof. No increase in height arises at the first base case ($d = 0$, cut with an axiom). For the other base case ($d = 0$, cut with a cut-formula introduced by $(weak, l)$), the weakenings needed to turn $\Gamma \vdash \Delta$ into $\Gamma, \Pi \vdash \Delta$ in the new derivation without cut may increase its height by at most h . For the cases where $d > 0$, we observe that in the cases (1.1), (1.2), (1.3), and (1.4) the height of the derivation is increased at most by 1. In case (2.1) the height may be increased by the additional applications of $(contr)$; in case (2.2) the height may be increased by the additional applications of $(weak, l)$; in case (2.3) the height may be increased by appending the derivation γ^+ below the cut. In all three cases the increase is again bounded by h . Since the lhs-depth d is decreased at each step and since $d \leq h$, we arrive at a bound h^2 for the height of a derivation where an uppermost cut in π is eliminated. In repeating the argument for the next cut to be eliminated, we have to replace h by h^2 , resulting in the bound $h^{2^2} = h^4$. Continuing in this manner until the last cut is eliminated, we obtain that the height $h(\pi_0)$ of the final derivation π_0 is bounded by $h^{(2^h)}$, which is clearly elementary in $h(\pi)$. ◀

► **Remark.** Clearly tighter complexity bounds could be extracted. However we are only interested in the contrast between an elementary and a non-elementary increase, here.

► **Remark.** Our argument essentially differs in several respects from other cut-elimination proofs, in particular also from Gentzen's original [4]. That unnested induction over the lhs-depth of a (multi-)cut suffices in our case is due to the observation stated as Lemma 1. It only holds in the absence of disjunction.

4 Complex prenex cut-formulas

In Section 3 we have seen that cuts with prenex *atomic* cut-formulas can be eliminated from $\mathbf{LI}_m^{-\forall}$ -derivations without incurring a non-elementary increase in the height of proofs. In this section we show that arbitrary complex prenex cut-formulas can be replaced by atomic ones for the price of introducing propositional cuts, i.e., applications of *cut*, where the cut-formula is quantifier free. The complexity of this transformation is negligible in our context.

► **Theorem 3.** *Let π be an $\mathbf{LI}_m^{-\forall}$ -derivation of $\Gamma \vdash \Delta$, where each cut-formula is prenex. Then there exists an $\mathbf{LI}_m^{-\forall}$ -derivation π_0^p of $\Gamma \vdash \Delta$ that only contains propositional cut-formulas, such that $h(\pi_0^p) \leq f(h(\pi))$ for some elementary function f .*

Proof. We first present the transformation of π into π_0^p in three separate stages and investigate the corresponding increase in height afterwards. (Again, regularity is assumed without further mentioning throughout the proof.)

Stage 1: We first consider the case, where the last inference of π is the only cut in π . The case where the cut-formula is a prenex atom is covered by Theorem 2. Therefore we assume that the cut-formula is of the form $\vec{Q}\vec{x}B$, where B is a compound propositional formula and $\vec{Q}\vec{x}$ denotes a non-empty string $Q_1x_1 \dots Q_nx_n$ of quantifier occurrences, where $Q_i \in \{\forall, \exists\}$ for $1 \leq i \leq n$. We first make the following observation.

► **Fact 1.** *If either the ancestor tree of the left cut-formula in π or the ancestor tree of the right cut-formula in π only contains quantified formulas, then the corresponding cut-formula can be traced back to ancestors that have been introduced by weakening. In that case the cut can be eliminated just like in the analogous cases in the proof of Theorem 2.*

In the remaining (general) case π can be depicted as follows.

$$\frac{\frac{\frac{\vdots}{\Gamma' \vdash \rho B} (Q_n, r)}{\vdots \vdots} \quad \frac{\frac{\frac{\vdots}{\sigma B, \Pi' \vdash \Delta'} (Q_n, l)}{\vdots \vdots \vdots \vdots} \quad \frac{\frac{\Gamma \vdash \vec{Q}\vec{x}B}{\Gamma, \Pi \vdash \Delta} (cut)}{\Gamma, \Pi \vdash \Delta} (cut)}{\Gamma, \Pi \vdash \Delta} (cut)}$$

where ρB denotes a quantifier-free formula occurrence in the ancestor tree of the cut-formula at the left premise of (cut) , where (Q_n, r) introduces the innermost quantifier $Q_n x_n$ of $\vec{Q}\vec{x}B$. Analogously, for σB and (Q_n, l) at the right part of the derivation. ρ and σ are the substitutions that replace the bound variables in B by appropriate free variables or terms, respectively.

Let $A = P_B(a_1, \dots, a_n, b_1, \dots, b_m)$ be an atom, where P_B is a new predicate symbol, a_1, \dots, a_n are free variables corresponding to the bound variables occurring in B , and b_1, \dots, b_m are free variables corresponding to the free variables occurring in B . We then introduce implications that define A as ‘abbreviation’ of B . Accordingly π is transformed into the following derivation π' .

$$\frac{\frac{\frac{\frac{\vdots}{\Pi \vdash \rho B} \quad \frac{\vdots}{\rho A \vdash \rho A}}{\rho B \supset \rho A, \Pi \vdash \rho A} (\supset, l)}{\forall \vec{x}(B \supset A), \Pi \vdash \rho A} (\forall, l)^*}{\forall \vec{x}(B \supset A), \Gamma \vdash \vec{Q}\vec{x}A} (Q_1, r)} \quad \frac{\frac{\frac{\frac{\vdots}{\Pi \vdash \sigma A} \quad \frac{\vdots}{\sigma B \vdash \sigma B}}{\sigma A, \sigma A \supset \sigma B, \Pi' \vdash \Delta'} (\supset, l)}{\sigma A, \forall \vec{x}(A \supset B), \Pi' \vdash \Delta'} (\forall, l)^*}{\vec{Q}\vec{x}A, \forall \vec{x}(A \supset B), \Pi \vdash \Delta} (Q_1, l)} (cut)}{\forall \vec{x}(B \supset A), \forall \vec{x}(A \supset B), \Gamma, \Pi \vdash \Delta} (cut)}$$

This transformation is repeated for every cut-formula that is not already a prenex atomic formula.

Stage 2: The derivation π' obtained from Stage 1 contains only prenex atomic cut-formulas. Therefore we can apply Theorem 2 to obtain a cut-free derivation π'_0 of $\Pi, \Gamma \vdash \Delta$, where Π denotes the ‘defining implications’ introduced in Stage 1.

Stage 3: For every cut-formula $\vec{Q}\vec{x}B$ of the original derivation π , we replace all occurrences of (instances of) the atoms A , introduced in Stage 1, by (corresponding instances of) B . This results in sub-derivations that have the following form.

$$\frac{\frac{\frac{\vdots}{\Psi_l \vdash \rho B} \quad \frac{\vdots}{\rho B, \Psi_r \vdash \Lambda'}}{\rho B \supset \rho B, \Psi_l, \Psi_r \vdash \Lambda'} (\supset, l)}{\vdots \vdots \vdots \vdots} \quad \frac{\frac{\vdots}{\forall \vec{x}(B \supset B), \Psi_l, \Psi_r, \Psi \vdash \Lambda}}{\Psi_l, \Psi_r, \Psi \vdash \Lambda} (cut)}$$

The indicated instance of (\supset, l) can be replaced by an instance of (cut) to obtain

$$\frac{\Psi_l \vdash \rho B \quad \rho B, \Psi_r \vdash \Lambda'}{\Psi_l, \Psi_r \vdash \Lambda'} (cut)$$

$$\Psi_l, \Psi_r, \Psi \vdash \Lambda$$

Note that the replacement of the atom A by the compound formula B actually results in a derivation that is not a proper $\mathbf{LI}_m^{-\forall}$ -derivation, since it will contain leaf nodes of the form $C \vdash C$, where C is a compound propositional formula. However such ‘improper axioms’ can readily be replaced by derivations from atomic axioms. This finally results in the derivation π_0^p of $\Gamma \vdash \Delta$ that trades prenex atomic cuts for propositional cuts.

It remains to investigate the increase in the height of the derivation. In Stage 1, in the case covered by Fact 1 we argue like in the corresponding case in the proof of Theorem 2: a cut-formula introduced by weakening is eliminated for the (possible) price of at most h additional weakenings, where $h = h(\pi) + 1$. In the general case, exhibited above, the transformation in Stage 1 may increase the depth by the additional introductions of universal quantifiers as indicated. There are as many such inferences as there are quantifier occurrences in the corresponding cut-formula. But these quantifier occurrences have been introduced by corresponding instances of quantifier rules already in the original derivation π and therefore the increase in height is again bounded by h . Repeating this for all cut-formulas, we obtain an overall bound of h^2 for Stage 1.

For Stage 2 we obtain an elementary bound from Theorem 2.

In Stage 3 the increase in height arises from the augmented derivations of improper (non-atomic, but propositional) axioms $C \vdash C$. The height of such derivations is bounded by the size of C , which in turn is not greater than $h(\pi)$, since C must already have been introduced from (atomic) axioms in π . (Remember that we have eliminated cut-formulas that trace back to ancestors introduced by weakening.) Moreover there are at most $h(\pi)$ such formulas.

Summing up, we obtain that the height of the final derivation π_0^p is elementarily bounded in the height of the original derivation π . ◀

5 Eliminating propositional cuts

We complete our cut-elimination proof by showing that propositional cuts can always be eliminated from $\mathbf{LI}_m^{-\forall}$ -derivations without incurring a non-elementary increase in proof size.

► **Theorem 4.** *Let π be an $\mathbf{LI}_m^{-\forall}$ -derivation of $\Gamma \vdash \Delta$, where each cut-formula is propositional. Then there exists a cut-free $\mathbf{LI}_m^{-\forall}$ -derivation π_0 of $\Gamma \vdash \Delta$, such that $h(\pi_0) \leq f(h(\pi))$ for some elementary function f .*

Proof. We will not directly argue about the (ordinary) height $h(\pi)$ of a derivation π , but instead consider the variant $\bar{h}(\pi)$, defined like $h(\pi)$, except that applications of the contraction and weakening rules are not counted. In other words $\bar{h}(\pi)$ is the maximal number of applications propositional and quantifier rules occurring in any branch of π . We will also talk of the *depth of a formula occurrence F* in a derivation π . By this we mean the height $\bar{h}(\pi')$ without counting contractions and weakenings in the sub-derivation π' of π that has as its root the sequent containing the indicated formula occurrence F .

Let $CF(\pi)$ denote the set of different propositional formulas that occur as cut-formulas in π , augmented by all their subformulas. Starting with a cut-formula F of maximal size we will stepwise reduce $CF(\pi)$ until it is empty. The following cases arise.

- (1) **F is atomic:** We consider all sub-derivations of π that end in a cut, where the atomic cut-formula F is a deepest occurrence of F in π . Let

$$\frac{\begin{array}{ccc} F \vdash F & \dots & F \vdash F \\ \vdots \delta & & \vdots \gamma \\ \Gamma \vdash F & & F, \Pi \vdash \Delta \end{array}}{\Gamma, \Pi \vdash \Delta} \text{ (cut)}$$

be such a sub-derivation, where the exhibited axioms in γ are those where the occurrence of F at the left side is in the ancestor tree of the exhibited cut-formula F . (Note that, by assumption, neither δ nor γ contains a cut with F .) The cut is eliminated by replacing these axioms with copies of the sub-derivation δ , resulting in

$$\frac{\begin{array}{ccc} \vdots \delta & \dots & \vdots \delta \\ \Gamma \vdash F & & \Gamma \vdash F \\ & \vdots & \\ & \Gamma, \dots, \Gamma, \Pi \vdash \Delta & \end{array}}{\Gamma, \Pi \vdash \Delta} \text{ (contr)*}$$

The above picture is in fact only adequate if all leaf nodes of the mentioned ancestor tree of the cut-formula F occur in axioms. But such ancestors of the cut-formula may also result from applications of (*weak*, *l*). In each such case the introduction of the corresponding occurrence of F is redundant, which in turn might render also applications of the contraction rule that lead to the cut-formula itself redundant. In fact it can happen that there is no axiom at all that contains an ancestor of the cut-formula. In that case, the cut, together with corresponding applications of (*weak*, *l*) and possibly also (*contr*) is redundant as well.

Another processing step that is left implicit in the above picture is the restoration of regularity to ensure that the eigenvariable condition for quantifier rules is preserved in further transformation steps. This is readily achieved by using variable-renamed copies of the sub-derivation δ .

We additionally eliminate other cuts with the cut-formula F occurring γ in the same manner, i.e., by replacing the axioms that contain the leaf nodes of the ancestor tree of the cut-formula F at the right upper sequent of the corresponding cut by δ (and/or eliminating redundant applications of weakening and contraction). While this renders the left upper sequent of the cut – say $\Pi' \vdash F$ – redundant we have to add the missing multiset Π' of formulas by applying additional weakenings at the position, where originally the cut-rule was applied.

As already indicated, the above transformation is to be applied simultaneously to all deepest occurrences of the atomic formula F as cut-formula in π . There might be further occurrences of F as cut-formulas in the resulting derivation π' , entailing $CF(\pi') = CF(\pi)$. However note that each such occurrence of F must be less deep in π , than the deepest occurrences of F as cut-formula in π . Furthermore note that $\bar{h}(\pi') \leq \bar{h}(\pi) + \bar{h}(\delta)$. By iterating the transformation (always applied to all currently deepest occurrence of the cut-formula F) we arrive at a derivation π'' , where $CF(\pi'') = CF(\pi) - \{F\}$ and $\bar{h}(\pi'') \leq \bar{h}(\pi) \cdot \bar{h}(\delta) \leq \bar{h}(\pi)^2$.

(2) *F* is not atomic: We only consider the case, where *F* is of the form $A \supset B$; the cases for negation and conjunction are similar. We depict π as follows:

$$\frac{\frac{\frac{\vdots}{\Gamma_0, A \vdash B} (\supset, r)}{\Gamma_0 \vdash A \supset B} \quad \frac{\frac{\frac{\vdots \gamma_1^A}{\Pi'_1 \vdash A} \quad \frac{\vdots \gamma_1^B}{B, \Pi''_1 \vdash \Delta_1} (\supset, l)}{A \supset B, \Pi_1 \vdash \Delta_1} \quad \dots \quad \frac{\frac{\frac{\vdots \gamma_n^A}{\Pi'_n \vdash A} \quad \frac{\vdots \gamma_n^B}{B, \Pi''_n \vdash \Delta_n} (\supset, l)}{A \supset B, \Pi_n \vdash \Delta_n}}{\Gamma \vdash A \supset B} \quad \frac{\frac{\vdots \gamma \dots}{A \supset B, \Pi \vdash \Delta} (cut)}{\Gamma, \Pi \vdash \Delta} (cut)$$

where $\Pi_i = \Pi'_i \uplus \Pi''_i$ for $i \in \{1, \dots, n\}$. We first assume that no occurrence of $A \supset B$ in the ancestor trees of the two cut-formulas is introduced by weakening. The exhibited occurrences of (\supset, r) and of (\supset, l) indicate all locations in π , where the first occurrence of the cut-formula in the corresponding ancestor tree is introduced. Let δ' denote the derivation ending in $A, \Gamma \vdash B$ that results from δ by eliminating the exhibited occurrences of (\supset, r) . The exhibited cut is eliminated by transforming π into

$$\frac{\frac{\frac{\frac{\vdots \gamma_1^A}{\Pi'_1 \vdash A} \quad \frac{\vdots \delta'}{A, \Gamma \vdash B} (cut)}{\Gamma, \Pi'_1 \vdash B} (\supset, l)}{\Gamma, \Pi_1 \vdash \Delta_1} \quad \frac{\frac{\vdots \gamma_1^B}{B, \Pi''_1 \vdash \Delta_1} (cut)}{\Gamma, \Pi_1 \vdash \Delta_1} \quad \dots \quad \frac{\frac{\frac{\frac{\vdots \gamma_n^A}{\Pi'_n \vdash A} \quad \frac{\vdots \delta'}{A, \Gamma \vdash B} (cut)}{\Gamma, \Pi'_n \vdash B} (\supset, l)}{\Gamma, \Pi_n \vdash \Delta_n} \quad \frac{\frac{\vdots \gamma_n^B}{B, \Pi''_n \vdash \Delta_1} (cut)}{\Gamma, \Pi_n \vdash \Delta_n} (cut)}{\Gamma, \dots, \Gamma, \Pi \vdash \Delta} (\supset, l) \quad \frac{\vdots \gamma \dots}{\Gamma, \Pi \vdash \Delta} (contr)^*$$

where the various copies of the sub-derivation δ' are variable-renamed to ensure the eigenvariable condition. If an occurrence of $A \supset B$ in the ancestor tree of one of the two cut-formulas is introduced by weakening, rather than by an implication rule, then the cut can be eliminated as usual at the possible expense of additional weakening to keep the side formulas of the redundant upper sequents of the cut in the derivation.

Like in case (1) we also eliminate all other cuts with cut-formula $A \supset B$ that occur in γ (i.e., above the right upper sequent of the exhibited deepest cut in π) in an analogous manner to obtain a derivation π' , where the deepest occurrence of $A \supset B$ as cut-formula is smaller than in π . Also like in case (1), we have $\bar{h}(\pi') \leq \bar{h}(\pi) + \bar{h}(\delta)$ and, by iterating the transformation, obtain a derivation π'' , where $CF(\pi'') = CF(\pi) - \{A \supset B\}$ and $\bar{h}(\pi'') \leq \bar{h}(\pi) \cdot \bar{h}(\delta) \leq \bar{h}(\pi)^2 \leq h(\pi)^2$.

We have seen that the elimination of a single formula from $CF(\pi)$ may be achieved at a quadratic expense in terms of $h = h(\pi)$. Eliminating a second cut-formula from $CF(\pi)$ therefore results in a derivation of height $\leq (h^2)^2 = h^4$. By repeating the transformation for all n formulas in $CF(\pi)$ we thus obtain the bound $h^{(2^n)}$ for a cut-free proof π'_0 of the original end-sequent $\Gamma \vdash \Delta$. Since $n \leq 2^h$, we have $\bar{h}(\pi'_0) \leq g(h(\pi))$ for some elementary function g .

Note that in extracting this bound from our cut-elimination procedure we made essential use of the fact that applications of the contraction and weakening rules are not counted in $\bar{h}(\pi'_0)$. To establish an elementary bound in terms of the ordinary height, as stated in the theorem, we finally have to remove redundant copies of sequents as well as redundant formula occurrences from sequents in π'_0 . More precisely, we prune π'_0 to obtain π_0 as follows. (A similar pruning process is already implicit in [4] and described in more detail in [3].)

(1) Suppose π'_0 contains a redundant copy of a sequent $\Psi \vdash \Lambda$. Then we replace

$$\begin{array}{c} \vdots \gamma \\ \Psi \vdash \Lambda \\ \vdots \vdots \vdots \\ \Psi \vdash \Lambda \\ \vdots \vdots \vdots \\ \Gamma \vdash \Delta \end{array}$$

by

$$\begin{array}{c} \vdots \gamma \\ \Psi \vdash \Lambda \\ \vdots \vdots \vdots \\ \Gamma \vdash \Delta \end{array}$$

(2) To ensure that each sequent contains at most three occurrences of the same formula at the left hand side of any sequent, we proceed as follows. Tracing the derivation downwards from the axioms, we apply the contraction rule immediately below any sequent containing two copies of the same formula at the left hand side. (More than one such application of (*contr*) may be needed, since a binary inference rule may result in several pairs of copies of the same formula.) This leaves us with a derivation, where a formula A may disappear altogether from a lower sequent of an inference rule. In case A is needed as the immediate ancestor of a principal formula in a later inference step, we simply add A by weakening, immediately before the corresponding inference.

A sequent in a pruned derivation has at most $3k + 1$ formula occurrences, where k is the number of different formulas that occur in it. Therefore, there are at most $(m + 1)^{3m+1}$ different sequents in the pruned cut-free derivation π_0 of $\Gamma \vdash \Delta$, if m is the number of different subformulas occurring in π'_0 . Clearly, $(m + 1)^{3m+1}$ also limits the height $h(\pi_0)$ of π_0 . To complete the proof of Theorem 4, it remains to check that m is elementarily bounded in terms of $\bar{h}(\pi'_0)$ and $h(\pi)$. To this aim it suffices to observe that every formula occurring in π'_0 that does not appear in $\Gamma \vdash \Delta$ must appear in the ancestor tree of the principal formula of an application of a quantifier rule in π'_0 . Therefore there are less than $2^{\bar{h}(\pi'_0)}$ such formulas. On the other hand, every formula that occurs in $\Gamma \vdash \Delta$, must either occur in an axiom or as the principal formula of an inference in π , which bounds the number of such formulas by $2^{h(\pi)}$. ◀

6 Elementary bounds on the size of cut-free derivations

Let us bring together the results of the previous sections and see how bounds on the height of cut-free derivations translate into bounds on the size (number of symbols) of derivations. This yields the main result of this paper, which can be formulated as the following corollary to Theorems 2, 3, and 4.

► **Corollary 5.** *Let π be an $\mathbf{LI}_m^{-\vee}$ -derivation of $\Gamma \vdash \Delta$, where each cut-formula is a prenex formula. Then there exists a cut-free $\mathbf{LI}_m^{-\vee}$ -derivation π_0 of $\Gamma \vdash \Delta$, such that $|\pi_0| \leq f(|\pi|)$ for some elementary function f .*

Proof. The following three observations suffice.

- (1) The class of elementary functions is closed under composition. Therefore, by applying first Theorem 3 (which in turn uses Theorem 2) and then Theorem 4, we obtain an elementary upper bound on the height $h(\pi_0)$ of the cut-free derivation π_0 in terms of the height $h(\pi)$ and consequently also the size $|\pi|$ of the original derivation π .
- (2) As we have seen at the end of the proof of Theorem 4, π_0 can be assumed to be pruned. But both, the number of sequents as well as the number of formula occurrences in any pruned derivation are elementary bounded by its height.
- (3) If the language does not contain function symbols then theorem follows immediately from (1) and (2), since then the size of (number of symbols in) a derivation is linear in the number of formula occurrences in it.

The case for languages with function symbols is somewhat more involved. We argue that every cut-free proof π_0 can be transformed into one where the size of terms occurring in it is elementary bounded by the size of terms occurring in the end-sequent of π_0 and the number of formula occurrences in π_0 . To this aim, the derivation is processed upwards (i.e., from the end-sequent towards the axioms) as follows. Whenever an inference of type (\forall, r) or (\exists, l) , introducing $\forall xA(x)$ or $\exists xA(x)$, respectively, is encountered, replace the corresponding the eigenvariable a in all ancestors of those occurrences of $\forall xA(x)$ or $\exists xA(x)$ by a new constant. Whenever an inference of type (\forall, l) or (\exists, r) is encountered, a term t in the upper sequents must have been replaced by a bound variable. Replace all corresponding occurrences of t in all ancestors of the principal formula of this inference by a new variable. The resulting tree of sequents γ is not a valid derivation, in general. However by applying everywhere in γ the most general simultaneous unifier of the pairs of atoms at the left and right side of leaf nodes of γ , these leaf nodes are restored to proper axioms. Finally we re-substitute fresh variables for the new constants introduced earlier. Since the applied substitution (unifier) is most general it is guaranteed that the newly introduced constants turn into variables that satisfy the eigenvariable condition. We thus obtain a proper cut-free derivation, where not only the number of formula occurrences, but also their sizes are properly bounded. (We use the fact that the size of terms in a most general unifier is exponentially bounded by size of terms in the unified atom [6].)

This concludes the proof. \blacktriangleleft

7 Non-elementary lower bounds for related cases of cut-elimination

To emphasize the non-triviality of our main result – elementary cut-elimination for $\mathbf{LI}_m^{-\forall}$ – we contrast it with the following two facts.

- (1) There exists a non-elementary lower bound for cut-elimination in *classical* disjunction-free proofs with prenex atomic cuts.
- (2) There exists a non-elementary lower bound for cut-elimination in intuitionistic disjunction-free proofs with *non-prenex* cuts.

To state (1) more precisely, let $\mathbf{LK}_m^{-\forall}$ be the classical sequent calculus with multiplicative rules, but without rules for disjunction. (In fact it does not really matter which version of the sequent calculus we use. The following result is very robust.)

We say that *elimination of certain types of cuts* for a given sequent calculus is *not elementary bounded* if there exists a sequence π_1, π_2, \dots of derivations of sizes $|\pi_1|, |\pi_2|, \dots$, where $|\pi_i| \leq f(i)$ ($i \geq 1$) for some elementary function f , but $|\pi_i^*| \geq g(i)$ for some non-elementary function g , where π_i^* is the shortest cut-free derivation of the end-sequent of π_i .

\blacktriangleright **Theorem 6.** *The elimination of prenex atomic cuts is not elementary bounded for $\mathbf{LK}_m^{-\forall}$, even if the language does not contain function symbols.*

Proof. By Theorem 3.3 of [1] a derivation with arbitrary cut-formulas can be reduced to a derivation with only prenex cut-formulas at a quadratic expense in terms of its size. (The number of cuts may increase in this transformation.) Since we are in classical logic, disjunction-freeness is inessential: $A \vee B$ may be replaced by $\neg A \supset B$ everywhere at linear expense. We may then use the argument of Theorem 3 in Section 4 to obtain a derivation of the same end-sequent, with only prenex atomic cut-formulas. (The argument does not depend on the restricted form of intuitionistic sequents.) By Corollary 5 the increase in size is elementarily bounded. The formula sequence used by Orevkov [7] to obtain a non-elementary lower bound on cut-elimination for the classical sequent calculus \mathbf{LK} does not contain function symbols. In this manner we obtain the required non-elementary lower bound for the size of cut-free derivations with respect to the size of corresponding derivations with only prenex atomic cut-formulas. \blacktriangleleft

Corresponding to statement (2), above, we have the following.

► **Theorem 7.** *The elimination of non-prenex cuts is not elementary bounded for $\mathbf{LI}_m^{\neg\vee}$, even if the language does not contain function symbols.*

Proof. As in the proof of Theorem 6, we may assume without loss of generality that a classical derivation does not contain disjunctions. We translate any such $\mathbf{LK}_m^{\neg\vee}$ -derivation into an $\mathbf{LI}_m^{\neg\vee}$ -derivation using the following inductively defined formula mapping: $A^+ = \neg\neg A$ if A is an atom, $(A \circ B)^+ = \neg\neg A^+ \circ \neg\neg B^+$ for $\circ \in \{\neg, \wedge, \supset\}$, and $(\mathbf{Q}xA)^+ = \neg\neg \mathbf{Q}xA^+$ for $\mathbf{Q} \in \{\exists, \forall\}$. Moreover let $(\neg\neg A)^-$ be $\neg A$ and write A^- instead of $(A^+)^-$. For a multiset of formulas $\Gamma = \{A_1, \dots, A_n\}$, we define $\Gamma^+ = \{A_1^+, \dots, A_n^+\}$ and $\Gamma^- = \{A_1^-, \dots, A_n^-\}$. We translate a given $\mathbf{LK}_m^{\neg\vee}$ -derivation of the end-sequent $A_1, \dots, A_n \vdash B_1, \dots, B_m$ into an $\mathbf{LI}_m^{\neg\vee}$ -derivation of $A_1^+, \dots, A_n^+, B_1^-, \dots, B_m^- \vdash$ by induction on its depth.

Axioms: $A \vdash A$ is replaced by a derivation of $\neg\neg A, \neg A \vdash$.

Structural rules: For applications of (*cut*)

$$\frac{\Pi \vdash \Psi, A \quad A, \Gamma \vdash \Delta}{\Pi, \Gamma \vdash \Psi, \Delta} \text{ (cut)} \quad \text{translates into} \quad \frac{\frac{\Pi^+, \Psi^-, A^- \vdash}{\Pi^+, \Psi^- \vdash A^+} (\neg, r) \quad A^+, \Gamma^+, \Delta^+ \vdash}{\Pi^+, \Gamma^+, \Psi^-, \Delta^- \vdash} \text{ (cut)}$$

The translations for weakenings and contractions are obvious.

Logical rules: We present the translation for (\wedge_1, l) and (\wedge, r) ; the other cases are analogous.

$$\frac{A, \Gamma \vdash \Delta}{A \wedge B, \Gamma \vdash \Delta} (\wedge_1, l) \quad \text{translates into} \quad \frac{\frac{A^+, \Gamma^+, \Delta^- \vdash}{A^+ \wedge B^+, \Gamma^+, \Delta^- \vdash} (\wedge_1, l) \quad \frac{\Gamma^+, \Delta^- \vdash \neg(A^+ \wedge B^+)}{\neg\neg(A^+ \wedge B^+) [= (A \wedge B)^+], \Gamma^+, \Delta^- \vdash} (\neg, l)}{\frac{\Gamma^+, \Psi^-, A^- \vdash}{\Gamma^+, \Psi^- \vdash A^+} (\neg, r) \quad \frac{\Pi^+, \Delta^-, B^- \vdash}{\Pi^+, \Delta^- \vdash B^+} (\neg, r)}{\Gamma^+, \Pi^+, \Psi^-, \Delta^- \vdash A^+ \wedge B^+} (\wedge, r)}{\neg(A^+ \wedge B^+) [= (A \wedge B)^-], \Gamma^+, \Pi^+, \Psi^-, \Delta^- \vdash} (\neg, l)$$

Any sequence of short $\mathbf{LK}_m^{\neg\vee}$ -derivations, where the corresponding shortest cut-free derivations grow non-elementarily, leads to a sequence of short $\mathbf{LI}_m^{\neg\vee}$ -derivations under this translation. (See [8] for an example with weak prenex quantifiers only in the end-sequent, where disjunctions are readily removed). Note that a lower bound for shortest cut-free derivations corresponds, modulo an exponential increase, to the Herbrand complexity, i.e. the size of the shortest Herbrand sequent (see [2]). Since Herbrand sequents are propositionally valid, they remain valid when double negations are removed. \blacktriangleleft

► Remark. Since classical derivations with arbitrary many cuts can be elementarily transformed into derivations with a single cut (see [2]) we could in fact sharpen Theorem 7 accordingly.

8 Conclusion

We have shown that cut-elimination is elementary for disjunction-free intuitionistic logic with prenex cut-formulas. To achieve this result we had to come up with novel techniques. The first step – elimination of prenex atomic cuts – is specific to the case at hand: it only works for disjunction-free intuitionistic logic. The second step – trading complex prenex cut-formulas for atomic prenex and propositional cuts – uses a scheme that can also be applied in other contexts. The final step – elementary elimination of propositional cuts – although presented in a way tailored to \mathbf{LI}_m^\vee , should also be adaptable to other sequent calculi, thus rendering our results of potential significance beyond disjunction-free intuitionistic logic.

References

- 1 Matthias Baaz and Alexander Leitsch. Cut normal forms and proof complexity. *Annals of Pure and Applied Logic*, 97(1):127–177, 1999.
- 2 Matthias Baaz and Alexander Leitsch. *Methods of Cut-elimination*. Springer, 2011.
- 3 Samuel R. Buss. An introduction to proof theory. In *Handbook of Proof Theory*, volume 137 of *Studies in Logic and the Foundations of Mathematics*, pages 1–78. Elsevier, 1998.
- 4 Gerhard Gentzen. Untersuchungen über das logische Schließen. I & II. *Mathematische Zeitschrift*, 39(1):176–210, 405–431, 1935.
- 5 Jörg Hudelmaier. Bounds for cut elimination in intuitionistic propositional logic. *Archive for Mathematical Logic*, 31(5):331–353, 1992.
- 6 Alberto Martelli and Ugo Montanari. An efficient unification algorithm. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 4(2):258–282, 1982.
- 7 V.P. Orevkov. Lower bounds for increasing complexity of derivations after cut elimination. *Journal of Soviet Mathematics*, 20:2337–2350, 1982.
- 8 Richard Statman. Lower bounds on Herbrand’s theorem. *Proc. of the Amer. Math. Soc.*, 75:104–107, 1979.