Some observations on the logical foundations of inductive theorem proving

Stefan Hetzl
Institute of Discrete Mathematics and Geometry
Vienna University of Technology
Wiedner Hauptstraße 8–10
AT-1040 Vienna
Austria

Tin Lok Wong
Kurt Gödel Research Center for Mathematical Logic
University of Vienna
Währinger Straße 25
AT-1090 Vienna
Austria

26 August, 2017

Abstract

In this paper we study the logical foundations of automated inductive theorem proving. To that aim we first develop a theoretical model that is centered around the difficulty of finding induction axioms which are sufficient for proving a goal.

Based on this model, we then analyze the following aspects: the choice of a proof shape, the choice of an induction rule and the language of the induction formula. In particular, using model-theoretic techniques, we clarify the relationship between notions of inductiveness that have been considered in the literature on automated inductive theorem proving.

1 Introduction

Theories of (natural number) arithmetic have been of great interest to mathematical logicians since the beginning of the 20th century. At that time, work was mainly devoted to questions of consistency. Deep connections that link the definability-theoretic aspects of arithmetic and theories of computation were later discovered and fruitfully developed. In the 1970s and 80s, relevant model-theoretic techniques became mature enough for establishing mathematically interesting unprovability results. While the model theory

of arithmetic evolved into a subject of its own, the connections with theories of computation found their way down to the complexity-theoretic level. Nowadays, theories of arithmetic have penetrated almost every branch of mathematical logic, including mathematical philosophy.

Rather independently of this work, the subject of inductive theorem proving developed in computer science. In this tradition, the central aim is to develop algorithms that find proofs by induction and to implement these algorithms efficiently. This subject is characterized by a great variety of different methods (and systems implementing these methods), for example, rippling [6], theory exploration [12], integration into a superposition prover [24, 15, 33], recursion analysis [2, 28, 7], proof by consistency [14], and cyclic proofs [5, 4]. Recently, a benchmark suite for inductive theorem proving has been presented [13].

The aim of our work is to apply methods and results from the former tradition in mathematical logic to the tradition in computer science. The main advantage of this combination is that it is possible to obtain *unprovability* results (by model-theoretic means) where previously in the literature on inductive theorem proving only empirical observations could be made based on the failure of *a specific algorithm* to find a proof.

A first obstacle in realizing such an application is that the above mentioned approaches to inductive theorem proving are quite different. This makes it difficult to provide a common theoretical basis. However, the final result is typically, in one way or another, explicitly or implicitly, a proof of the goal from instances of an induction scheme and basic axioms from a background theory. We take this observation as a guiding principle for the development of a theoretical model of inductive theorem proving in Section 2. This is the main conceptual contribution of this paper.

In terms of technical contributions we analyze the following aspects of methods for inductive theorem proving: (a) the choice of a proof shape, (b) the choice of the induction rule, and (c) the language of the induction formula. Despite the differences between the existing methods for inductive theorem proving, these aspects play a role in most of them. Mathematically, the technical contributions of this paper are: we show that the equivalence proof shape for inductive proofs (see Section 3) is complete but that the uniform proof shape is not. We show that Walther's method for comparing induction axioms [32] is not complete (Section 4.3). We make the (possibly surprising) observation that the weakest induction axiom for proving an induction axiom may not be that induction axiom itself (Proposition 6.1). We establish the strictness of the implications between several frequently used notions of inductiveness (Section 7). We also include the result, due to Kaye, that PA^- proves the least number principle in the language of rings (Theorem 8.2).

This paper is structured as follows: after developing our model of inductive theorem proving in Section 2, we study the completeness of proof

shapes in Section 3. In Section 4 we describe different formulations of induction and study their equivalence in both a general and a quantifier-free context. In Section 5 we establish non-closure properties of PA^--cuts that will be used in the rest of the paper. In Section 6 we investigate two ways of comparing different induction formulas that prove the same theorem. In Section 7, which is central to this paper, we compare different notions of inductiveness. (A formula is called inductive in the sense of a particular induction rule if it satisfies the hypotheses of this rule provably in the base theory.) In Section 8 we study the effect of the choice of the language for the induction formula using an example involving the $<$-relation.

2 A theoretical model of inductive theorem proving

This section is devoted to developing a theoretical model of inductive theorem proving. This is necessary in order to provide a conceptually and formally clear basis for relating mathematical results to algorithms in automated deduction.

2.1 First-order theories of the natural numbers

In this paper we restrict our attention to first-order theories of the natural numbers (as opposed to, e.g., general inductive data types which is a more common choice in inductive theorem proving). This is a pragmatic choice which is motivated by the following reasons: (1) the mathematical study of these theories is very well developed, and (2) the central problems of inductive theorem proving also surface in this restricted setting. To the mathematical logician this may not seem a restriction since coding allows us to describe an arbitrary inductive data type as a set of natural numbers. However, the increase of the syntactic complexity of formulas and proofs based on coding renders this approach unfit for practical applications, where bounded quantifiers are usually as costly as unbounded quantifiers.

Throughout this paper, we work over the base theory PA^-. It is finitely axiomatized, induction-free, and a fragment of Peano arithmetic (PA). Robinson's Q is another commonly used base theory for arithmetic. There are two main reasons for choosing PA^- instead of Q as the base theory in this paper. First, if one adopts Q as the base theory, then even a slight change in the definitions can make a significant difference in the results. This is due to the fact that many simple arithmetic properties, for example,

$$\forall x, y\ (x < y + 1 \leftrightarrow x \leqslant y),$$

are not provable in Q; cf. Lemma 4.2. We do not want such details to distract us. Second, we can extract useful information about notions of inductiveness in Section 7 over PA^-. On the contrary, Robinson's Q is too

weak to prove any non-trivial implication between our notions of inductiveness. More specifically, Proposition 7.1(e) and (f) become false if one changes the base theory to Q. Nevertheless, many other results in this paper, for example, Theorem 3.1, Proposition 3.2 and Lemma 5.1, remain true over Q.

Definition. Denote the language $\{0, 1, +, \times, <\}$ for ordered rings by $\mathscr{L}_{\mathrm{OR}}$. Abusing notation, if $n \in \mathbb{N}$, then we denote the closed $\mathscr{L}_{\mathrm{OR}}$ term

$$(\cdots((0+\underbrace{1)+1)+\cdots+1)}_{n\text{-many 1's}}$$

by n. Let PA^- denote the theory of the non-negative parts of discretely ordered rings. We axiomatize PA^- by the following.

P1. $\forall x\ \forall y\ \forall z\ \big((x+y)+z = x+(y+z)\big)$.

P2. $\forall x\ \forall y\ (x+y = y+x)$.

P3. $\forall x\ \forall y\ \forall z\ \big((x\times y)\times z = x\times(y\times z)\big)$.

P4. $\forall x\ \forall y\ (x\times y = y\times x)$.

P5. $\forall x\ \forall y\ \forall z\ (x\times(y+z) = (x\times y)+(x\times z))$.

P6. $\forall x\ (x+0 = x)$.

P7. $\forall x\ (x\times 0 = 0)$.

P8. $\forall x\ (x\times 1 = x)$.

P9. $\forall x\ \forall y\ \forall z\ (x<y\wedge y<z\to x<z)$.

P10. $\forall x\ \neg x<x$.

P11. $\forall x\ \forall y\ (x<y\vee x=y\vee x>y)$.

P12. $\forall x\ \forall y\ \forall z\ (x<y\to x+z<y+z)$.

P13. $\forall x\ \forall y\ \forall z\ (z\neq 0\wedge x<y\to x\times z<y\times z)$.

P14. $\forall x\ \forall y\ (x<y\leftrightarrow\exists z\ ((x+z)+1=y))$.

P15. $0<1\wedge\forall x\ (x>0\to x\geqslant 1)$.

P16. $\forall x\ (x\geqslant 0)$.

Here $x\leqslant y$ is an abbreviation for $x<y\vee x=y$.

Definition. If $\theta(x, \bar{z})$ is an $\mathscr{L}_{\mathrm{OR}}$ formula, then the *induction axiom for* θ with respect to the variable x, denoted by $\mathrm{I}_x\theta$ or simply $\mathrm{I}\theta$, is the sentence

$$\forall \bar{z}\ \big(\theta(0, \bar{z}) \wedge \forall x\ \big(\theta(x, \bar{z}) \to \theta(x+1, \bar{z})\big) \to \forall x\ \theta(x, \bar{z})\big).$$

Define

$$\begin{aligned}
\mathrm{IOpen} &= \mathrm{PA}^- + \{\mathrm{I}_x\theta : \theta(x, \bar{z}) \text{ is a quantifier-free } \mathscr{L}_{\mathrm{OR}} \text{ formula}\},\\
\mathrm{I}\Sigma_k &= \mathrm{PA}^- + \{\mathrm{I}_x\theta : \theta(x, \bar{z}) \text{ is an } \mathscr{L}_{\mathrm{OR}} \text{ formula of complexity } \Sigma_k\}, \text{ and}\\
\mathrm{PA} &= \mathrm{PA}^- + \{\mathrm{I}_x\theta : \theta(x, \bar{z}) \text{ is an } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}.
\end{aligned}$$

2.2 The necessity of non-analyticity

It has often been observed in the literature on inductive theorem proving that "within first-order theories that include induction, the cut rule cannot be eliminated" [6, p. 55]. This observation will provide a crucial foundation for the justification of our model of inductive theorem proving. Therefore we would like to discuss it here and, in the process, make it more precise and show how to prove it.

Since we want to speak about cut-elimination we need to speak about the sequent calculus. Which variant of the sequent calculus we use is not of importance for the points discussed here; for the sake of precision let us fix it to be the calculus LK of [9]. A sequent is denoted as $\Gamma \longrightarrow \Delta$. For a theory T and a formula φ, $T \vdash \varphi$ if and only if there is a finite set $T_0 \subseteq T$ and an LK-proof of the sequent $T_0 \longrightarrow \varphi$. Gentzen's cut-elimination theorem states that:

Theorem 2.1. If there is an LK-proof of a sequent $\Gamma \longrightarrow \Delta$, then there is a cut-free LK-proof of $\Gamma \longrightarrow \Delta$.

An important feature of cut-free proofs is that they have the subformula property. In the context of first-order logic this means that every formula that occurs in a cut-free proof of the sequent $\Gamma \longrightarrow \Delta$ is an *instance* of a subformula of a formula that occurs in $\Gamma \longrightarrow \Delta$. A proof that has the subformula property is also called *analytic*.

Since the cut-elimination theorem considers arbitrary first-order sequents, it can also be applied to theories containing induction axioms:

Corollary 2.2. If $\mathrm{PA} \vdash \varphi$ then there is a finite $A_0 \subseteq \mathrm{PA}$ and a cut-free LK-proof of the sequent $A_0 \longrightarrow \varphi$.

So we see that *in the sense of the above corollary*, inductive theories do allow cut-elimination. *However*, A_0 may contain induction axioms on induction formulas which are not instances of subformulas of φ, i.e., non-analytic induction formulas.

The *necessity* of non-analytic induction formulas follows, for example, from Gödel's second incompleteness theorem: recall that, by arithmetizing the syntax of formulas and proofs, one can formulate the consistency of a recursive arithmetical theory as an $\mathscr{L}_{\mathrm{OR}}$ sentence. More specifically, for all $k \in \mathbb{N}$ there is a Π_1 sentence $\mathrm{Con}(\mathrm{I}\Sigma_k)$ expressing the consistency of $\mathrm{I}\Sigma_k$, see for example [10]. We then have:

Theorem 2.3. For all $k \in \mathbb{N}$: $\mathrm{PA} \vdash \mathrm{Con}(\mathrm{I}\Sigma_k)$ but $\mathrm{I}\Sigma_k \nvdash \mathrm{Con}(\mathrm{I}\Sigma_k)$.

Note that this result embodies a very strong non-analyticity requirement: given any $k \geqslant 1$, in order to prove $\mathrm{Con}(\mathrm{I}\Sigma_k)$ not only do we need a non-analytic induction formula, but we need one with more than k quantifier alternations even though $\mathrm{Con}(\mathrm{I}\Sigma_k)$ is only a Π_1 sentence.

This theorem entails the necessity of cut in the following sense. First, formulate induction as the inference rule

$$\frac{\Gamma \longrightarrow \Delta, \psi(0) \quad \Gamma, \psi(x) \longrightarrow \Delta, \psi(s(x))}{\Gamma \longrightarrow \Delta, \forall x\, \psi(x)}\ \mathrm{Ind}$$

with the usual side condition and ψ being an arbitrary formula. Observe that $\mathrm{PA} \vdash \varphi$ if and only if there is an $\mathrm{LK} + \mathrm{Ind}$-proof of $\mathrm{PA}^- \longrightarrow \varphi$. Now, in contrast to LK, the calculus $\mathrm{LK} + \mathrm{Ind}$ does not have cut-elimination:

Corollary 2.4. There is a formula φ such that $\mathrm{PA}^- \longrightarrow \varphi$ has an $\mathrm{LK} + \mathrm{Ind}$-proof but no cut-free $\mathrm{LK} + \mathrm{Ind}$-proof.

Proof. Let $\varphi = \mathrm{Con}(\mathrm{I}\Sigma_k)$ for *any* $k \geqslant 3$. Then, by Theorem 2.3, $\mathrm{PA} \vdash \mathrm{Con}(\mathrm{I}\Sigma_k)$ and consequently there is an $\mathrm{LK} + \mathrm{Ind}$-proof of $\mathrm{PA}^- \longrightarrow \mathrm{Con}(\mathrm{I}\Sigma_k)$. On the other hand, suppose there would be a cut-free $\mathrm{LK} + \mathrm{Ind}$-proof of $\mathrm{PA}^- \longrightarrow \mathrm{Con}(\mathrm{I}\Sigma_k)$. Then, due to the subformula property, all formulas, and in particular: all induction formulas, in this proof would be Σ_3 thus contradicting Theorem 2.3. □

These considerations show that the observation formulated at the beginning of this section can be stated more precisely, and without mentioning the cut rule, as: *inductive theorem proving requires the use of non-analytic induction axioms.*

The relationship of the above considerations to inductive theorem proving is subtle since, on a mathematical level, the necessity of *finding* non-analytic formulas can be avoided by changing the setting. There are a number of ways for achieving that: (1) One can move to second- or even higher-order logic, for example by formulating the theorem-proving problem in terms of the second-order induction axiom. Then, in the presence of second-order quantifiers, the meaning of the subformula property, and with it that of analyticity, changes. Indeed, the non-analytic induction axioms of a PA-proof translate to instances (in the sense of second-order logic) of the

second-order induction axiom. Thus a PA-proof translates into a proof which is analytic in the sense of second-order logic. But this is merely a change in terminology, not in substance. (2) One can use coding, for example as in the proof of the finite axiomatizability of $\mathrm{I}\Sigma_k$ in [19, Theorem I.2.52]. Thus one can formulate the theorem-proving problem in $\mathrm{I}\Sigma_k$ (for a k sufficiently high for practical purposes) as a theorem-proving problem in pure first-order logic, thereby eliminating the need of generating non-analytic formulas. However, coding introduces an overhead which, although constant, is so high that it dominates the complexity of all practically relevant instances to such an extent that this approach is not useful in practice. (3) Similar to but simpler than (2), one can just fix a k, sufficiently high for practical purposes, so that we are interested only in finding proofs whose induction axioms contain at most k symbols. Since there are only a finite number of such induction axioms, the theorem-proving problem, again, becomes a pure first-order problem. Just as in (2), although this avoids the need of generating non-analytic formulas, it introduces a syntactic overhead which renders this approach useless in practice (even though coding does not play a role here).

2.3 A computational observation

We have seen in the previous section that, in order to find a proof by induction, it is necessary to find suitable non-analytic formulas. This is in stark contrast to automated theorem proving in pure first-order logic, where a cut-elimination theorem exists, and all formulas required for proving a goal can be obtained by, e.g., resolution and paramodulation.

Therefore the search space in inductive theorem proving has two dimensions: (1) the search for instances of the induction scheme required to prove the goal and (2) the search for a proof in pure first-order logic of the goal from these instances. We base our model on the following computational observation:

Computational Observation. For sentences φ which are input to an inductive theorem prover: if it is feasible to find formulas $\theta_1, \ldots, \theta_n$ such that $\mathrm{PA}^- + \{\mathrm{I}\theta_1, \ldots, \mathrm{I}\theta_n\} \vdash \varphi$, then it is feasible to find a proof in pure first-order logic of φ from $\mathrm{PA}^- + \{\mathrm{I}\theta_1, \ldots, \mathrm{I}\theta_n\}$.

In order words, the search space for inductive theorem proving extends so much more in the first dimension than in the second that we can afford to disregard the second.

This observation, as stated, is rather imprecise and needs some elaboration. What we mean by "feasibility" is the possibility of an implementation which solves the task successfully on contemporary hardware in a reasonable amount of time. This notion is quite standard in computer science and we believe there is little potential for misunderstanding in this respect. When

we say that "φ is input to an inductive theorem prover" we mean that, in particular, $\mathrm{PA} \vdash \varphi$. The recognition of non-theorems as such, while clearly of high practical value, is a different topic which we do not consider in this paper. But there is more to it: the observation does not apply to arbitrary PA-theorems. To see this, note that the undecidability of provability in PA^- entails that there is no computable bound on the size of a proof of a PA^--theorem. So while, with a general PA^--theorem as input, any choice of $\theta_1, \ldots, \theta_n$ would work, it is not feasible to actually find any PA^--proof. On the other hand, we would usually not expect PA^--theorems as input to an inductive theorem prover since they can be proved already without induction. So, by that phrase, in addition to being a PA-theorem, we mean that φ should be realistic as input to an inductive theorem prover. This subset of PA-theorems is naturally fuzzy and we cannot claim to make it precise here. What we have in mind are goals such as those of the TIP library [13]. They typically consist of a universally quantified atomic formula to be proved from some background axioms consisting also of universally quantified atomic formulas. They have size between a few dozen to several hundred symbols and proofs with a symbolic complexity of one or two orders of magnitude that of the goal.

This observation is based on the following grounds. First, automated theorem proving in pure first-order logic is a subject that has undergone continuous progress for decades and has reached a quite mature state. The regular CASC-competition [29] is a testament to that as is the widespread use of first-order theorem provers in external tools, e.g., sledgehammer [25]. Inductive theorem proving, and in particular the generation of non-analytic induction invariants, does not enjoy a comparable level of stability and maturity (yet?). Secondly, and in terms of concrete evidence, we have considered 53 proofs of problems from the TIP-library. These proofs have been manually entered in the GAPT-system [17] by a student of the first author (for another purpose). Of these 53 proofs, GAPT's built-in first-order prover Escargot, which is a quite simple superposition prover, could re-prove 48 based on the induction axioms alone within a timeout of 1 minute per proof using, on average, 3.3 seconds per proof on standard PC hardware. Last but not least, it is difficult to imagine a method that would generate $\theta_1, \ldots, \theta_n$ such that $\mathrm{PA}^- + \{\mathrm{I}\theta_1, \ldots, \mathrm{I}\theta_n\} \vdash \varphi$ without, at least implicitly, also generating a proof of φ from $\mathrm{PA}^- + \{\mathrm{I}\theta_1, \ldots, \mathrm{I}\theta_n\}$. A noteworthy exception to this are lemma speculation heuristics which are employed in inductive theorem provers, see, e.g. [12, 33], but even in this situation an actual proof is eventually generated.

2.4 One induction axiom is enough

As a final step towards our model of inductive theorem proving, we will see in this section that we can restrict our attention to the use of a single

induction axiom.

Definition. An $\mathscr{L}_{\mathrm{OR}}$ formula $\varphi(x)$ is called *inductive* if $\mathrm{PA}^- \vdash \varphi(0)$ and $\mathrm{PA}^- \vdash \forall x\ (\varphi(x) \to \varphi(x+1))$.

If $\varphi(x)$ is an inductive formula, then $\forall x\ \varphi(x)$ is trivially equivalent over PA^- to the induction axiom

$$\varphi(0) \wedge \forall x\ \big(\varphi(x) \to \varphi(x+1)\big) \to \forall x\ \varphi(x).$$

Therefore, sentences of the form $\forall x\ \varphi(x)$, where $\varphi(x)$ is inductive, can be viewed as particular instances of parameter-free induction axioms. Conversely, as Lemma 2.5 below shows, every induction axiom is equivalent over PA^- to an induction axiom of this form. As a result, every induction axiom corresponds to an inductive formula, and the full induction scheme is equivalent to its parameter-free counterpart. The argument is presumably well known, cf. Kaye [23, Exercise 8.3].

Lemma 2.5. Let $\theta(x, \bar{z})$ be an $\mathscr{L}_{\mathrm{OR}}$ formula. Define $\varphi(x)$ to be

$$\forall \bar{z}\ \big(\theta(0,\bar{z}) \wedge \forall y\ \big(\theta(y,\bar{z}) \to \theta(y+1,\bar{z})\big) \to \theta(x,\bar{z})\big).$$

Then $\varphi(x)$ is inductive and $\mathrm{PA}^- \vdash \mathrm{I}_x\theta \leftrightarrow \forall x\ \varphi(x)$.

Proof. Let us first verify that $\varphi(x)$ is inductive. Work over PA^-. We have $\varphi(0)$ trivially. Suppose x_0 is such that $\varphi(x_0)$ holds, and take $\bar{z}$ such that the hypothesis in $\varphi(x_0+1)$ holds, i.e.,

$$\theta(0,\bar{z}) \wedge \forall y\ \big(\theta(y,\bar{z}) \to \theta(y+1,\bar{z})\big).$$

Then $\theta(x_0,\bar{z})$ must be true since $\varphi(x_0)$, and thus $\theta(x_0+1,\bar{z})$ is also true by the second conjunct displayed above. This shows $\varphi(x_0+1)$.

Next, we verify that $\mathrm{PA}^- \vdash \mathrm{I}_x\theta \leftrightarrow \forall x\ \varphi(x)$. Work over PA^- again. Suppose $\forall x\ \varphi(x)$. Take $\bar{z}$ such that

$$\theta(0,\bar{z}) \wedge \forall y\ \big(\theta(y,\bar{z}) \to \theta(y+1,\bar{z})\big).$$

We want $\forall x\ \theta(x,\bar{z})$. So pick any x_0. We know $\varphi(x_0)$ holds by hypothesis. Thus $\theta(x_0,\bar{z})$ by the definition of $\varphi(x)$, as required.

Conversely, assume $\mathrm{I}_x\theta$ holds, i.e.,

$$\forall \bar{z}\ \big(\theta(0,\bar{z}) \wedge \forall y\ \big(\theta(y,\bar{z}) \to \theta(y+1,\bar{z})\big) \to \forall x\ \theta(x,\bar{z})\big).$$

Let x_0 be arbitrary. We want $\varphi(x_0)$. So take any $\bar{z}$ such that $\theta(0,\bar{z}) \wedge \forall y\ \big(\theta(y,\bar{z}) \to \theta(y+1,\bar{z})\big)$. Then our assumption implies $\theta(x_0,\bar{z})$, which is what we want. □

The key idea behind the lemma above is that inductive formulas are, in a sense, closed under definable conjunction. In particular, any two induction axioms are implied by a third over PA^-.

Proposition 2.6 (Gentzen [18])**.** For all $\mathscr{L}_{\mathrm{OR}}$ formulas θ_0, θ_1, there is an inductive formula $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \to \mathrm{I}\theta_0 \wedge \mathrm{I}\theta_1$.

Proof. Apply Lemma 2.5 to find inductive formulas $\psi_0(x)$ and $\psi_1(x)$ such that $\mathrm{PA}^- \vdash \mathrm{I}\theta_i \leftrightarrow \forall x\ \psi_i(x)$ for each $i < 2$. Define $\varphi(x) = \psi_0(x) \wedge \psi_1(x)$. As $\psi_0(x)$ and $\psi_1(x)$ are both inductive, it is easy to see that $\varphi(x)$ is inductive too. Moreover, the sentence $\forall x\ \varphi(x)$ implies $\forall x\ \psi_0(x) \wedge \forall x\ \psi_1(x)$ and thus also $\mathrm{I}\theta_0 \wedge \mathrm{I}\theta_1$ over PA^-. □

The above proof straightforwardly generalizes to an arbitrary number of $\mathscr{L}_{\mathrm{OR}}$ formulas $\theta_1, \ldots, \theta_n$ and so we obtain:

Corollary 2.7. Let σ be an $\mathscr{L}_{\mathrm{OR}}$ sentence. Then $\mathrm{PA} \vdash \sigma$ if and only if there is an inductive formula $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \to \sigma$. □

Since we only use simple syntactic operations (conjunction in Corollary 2.7 and the definition in the statement of Lemma 2.5) to combine many induction axioms into one, the computational observation is preserved even when restricted to a single induction axiom. Our theoretical model for inductive theorem proving is now the following computational problem.

ITP	
Input:	A sentence σ provable in PA
Output:	An inductive formula $\varphi(x)$ s.t. $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \to \sigma$

Mathematically, this just defines a binary relation (between σ and $\varphi(x)$). But of course, implicitly, we take the perspective of wanting to compute such a $\varphi(x)$ from a given σ. We claim that this computational problem is a suitable theoretical model of the practice of inductive theorem proving. This claim rests on the computational observation made in Section 2.3 and sharpened in this section. As described in Section 2.2, the quantifier complexity of $\varphi(x)$ cannot be bounded by that of σ; in particular $\varphi(x)$ is perhaps not a subformula of σ. In the rest of this paper we will study this problem, in particular by relating it to several of its variants.

3 Variations of the proof shape

The ITP problem as defined above induces a natural proof shape: the combination of (i) a PA^--proof of the induction base, (ii) a PA^--proof of the induction step, and (iii) a PA^--proof of $\forall x\ \varphi(x) \to \sigma$. As long as we use

ordinary successor induction, there is no freedom in the first two proof obligations, there is however in the third. In this section we will consider two variants of ITP which are obtained by modifying (iii).

The sharp-eyed reader may have noticed that the inductive formula $\varphi(x)$ in our proof of Proposition 2.6 actually makes $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \leftrightarrow \mathrm{I}\theta_0 \wedge \mathrm{I}\theta_1$. We can use this to obtain

Theorem 3.1. Let σ be an $\mathscr{L}_{\mathrm{OR}}$ sentence. Then $\mathrm{PA} \vdash \sigma$ if and only if there is an inductive formula $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \leftrightarrow \sigma$.

Proof. The "if" direction is clear from the definition of inductive formulas. For the "only if" direction, apply Proposition 2.6 to find an inductive formula $\psi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \psi(x) \to \sigma$. We verify that

$$\varphi(x) \quad = \quad \neg\sigma \to \psi(x)$$

has the properties we want. First, it is clear that $\mathrm{PA}^- \vdash \sigma \to \forall x\ \varphi(x)$. Second, work over PA^-, and suppose $\neg\sigma$. Since $\mathrm{PA}^- \vdash \forall x\ \psi(x) \to \sigma$, this implies $\exists x\ \neg\psi(x)$. If x_0 is such that $\neg\psi(x_0)$, then $\neg\sigma \wedge \neg\psi(x_0)$ and so $\neg\varphi(x_0)$. We can thus conclude $\mathrm{PA}^- \vdash \neg\sigma \to \neg\forall x\ \varphi(x)$. Finally, the formula $\varphi(x)$ is inductive because it is equivalent to either $x = x$ or $\psi(x)$ depending on whether σ holds or not, and both $x = x$ and $\psi(x)$ are inductive. □

This is a particular case of a more general phenomenon. As a normal form theorem in a very broad sense, it is perhaps reminiscent of the Friedman–Goldfarb–Harrington Theorem and its generalizations [30, 21], which assert that over a sufficiently strong base theory, every $\mathscr{L}_{\mathrm{OR}}$ sentence is equivalent to a consistency statement. This result motivates the consideration of the equivalence version of ITP:

$\mathrm{ITP}_{\mathrm{EQ}}$	
Input:	A sentence σ provable in PA
Output:	An inductive formula $\varphi(x)$ s.t. $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \leftrightarrow \sigma$

Theorem 3.1 ensures that, just as ITP, also $\mathrm{ITP}_{\mathrm{EQ}}$ is a total relation in the sense that for every input there is an output as required. However, for a fixed σ the $\varphi(x)$'s permitted in $\mathrm{ITP}_{\mathrm{EQ}}$ are a strongly restricted subset of those permitted in ITP. This leads to a significant reduction of the search space. Typically, restrictions of the search space play a crucial role for automated theorem proving in practice. We are not aware of a technique that would exploit this reduction of ITP to $\mathrm{ITP}_{\mathrm{EQ}}$. In how far this restriction to equivalent formulas is useful in practice therefore remains unclear for the time being.

Another modification of the proof shape of ITP consists of considering a universally quantified σ, i.e., $\sigma = \forall x\ \psi(x)$, and seeking an inductive formula $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \big(\varphi(x) \to \psi(x)\big)$. This is of relevance to

computer science since it corresponds to the treatment of loops in correctness proofs for imperative programs by loop invariants as in the Hoare calculus [1, 3]. As one may expect, this method does not work for all PA-provable formulas.

Proposition 3.2. There is an $\mathscr{L}_{\mathrm{OR}}$ formula $\psi(x)$ such that $\mathrm{PA} \vdash \forall x\ \psi(x)$ but no inductive formula $\varphi(x)$ makes $\mathrm{PA}^- \vdash \forall x\ \big(\varphi(x) \to \psi(x)\big)$.

Proof. Pick an $\mathscr{L}_{\mathrm{OR}}$ sentence σ that is provable in PA but not in PA^-. Consider the $\mathscr{L}_{\mathrm{OR}}$ formula $\psi(x)$, which is defined to be

$$\sigma \vee x \neq 0.$$

Then $\mathrm{PA} \vdash \forall x\ \psi(x)$ because $\mathrm{PA} \vdash \sigma$. Let $M \models \mathrm{PA}^- + \neg\sigma$, which exists since $\mathrm{PA}^- \nvdash \sigma$. Then $M \models \neg\psi(0)$ by the definition of $\psi(x)$. Therefore, for *no* formula $\varphi(x)$ can

$$\mathrm{PA}^- \vdash \forall x\ (\varphi(x) \to \psi(x)) \wedge \varphi(0). \qquad \square$$

As a result, the following uniform version of ITP

$\mathrm{ITP}_{\mathrm{U}}$	
Input:	A sentence $\forall x\ \psi(x)$ provable in PA
Output:	An inductive formula $\varphi(x)$ s.t. $\mathrm{PA}^- \vdash \forall x\ (\varphi(x) \to \psi(x))$

is not a total relation. Nevertheless, some weak form of completeness is still possible if we restrict ourselves to simple enough $\mathscr{L}_{\mathrm{OR}}$ sentences provable in a sufficiently weak fragment of PA, as the following theorem shows.

Definition. An $\mathscr{L}_{\mathrm{OR}}$ formula is *bounded* if all the quantifiers it contains are of the form $\forall x{<}t$ or $\exists x{<}t$, where t is a term in $\mathscr{L}_{\mathrm{OR}}$ that does not involve the variable x. Bounded formulas are also called Δ_0 *formulas*. The theory $\mathrm{I}\Delta_0$ is $\mathrm{PA}^- + \{\mathrm{I}_x\theta : \theta(x, \bar{z})$ is a bounded $\mathscr{L}_{\mathrm{OR}}$ formula$\}$. Fix a bounded formula $y = 2^x$ such that

$$\begin{aligned}\mathrm{I}\Delta_0 \vdash \forall x, y, y'\ (y = 2^x \wedge y' = 2^x \to y = y')\\ \wedge\, 2^0 = 1 \wedge \forall x, y\ (y = 2^x \leftrightarrow 2y = 2^{x+1}).\end{aligned}$$

Let exp be the axiom $\forall x\ \exists y\ (y = 2^x)$.

See Section V.3(c) in Hájek–Pudlák [19], for example, for a construction of the formula $y = 2^x$.

Theorem 3.3 (Wilkie–Paris)**.** The following are equivalent for a bounded formula $\psi(x)$.

(i) $\mathrm{I}\Delta_0 + \exp \vdash \forall x\ \psi(x)$.

(ii) There is an inductive formula $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ (\varphi(x) \to \psi(x))$.

Proof. See Corollary 8.7 in Wilkie–Paris [34] or Theorem V.5.26 in Hájek–Pudlák [19]. $\square$

4 Different forms of induction

In our definition of ITP we have fixed the induction scheme to the ordinary successor induction. This is by no means the only choice. It is well known that the induction scheme has many equivalent formulations. In this section we will introduce those alternative formulations that we treat in this paper, and start to study their relationship, both in the general and in the quantifier-free setting. Our interest in the quantifier-free setting is motivated by the fact that some methods for inductive theorem proving restrict themselves to quantifier-free induction formulas, see, e.g. [2].

4.1 Different induction schemes

Definition. Let $\theta(x, \bar{z})$ be an $\mathscr{L}_{\mathrm{OR}}$ formula. We define the following induction axioms:

- The *$<$-induction axiom* $\mathrm{I}_x^{<}\theta$ is

$$\forall \bar{z}\ \Big(\forall y\ \big(\forall x{<}y\ \theta(x, \bar{z}) \to \theta(y, \bar{z})\big) \to \forall x\ \theta(x, \bar{z})\Big).$$

- Let $n \in \mathbb{N}$. The *$(n+1)$-step induction axiom* $\mathrm{I}_x^{(n+1)\text{-step}}\theta$ is

$$\forall \bar{z}\ \Big(\bigwedge_{k<n+1} \theta(k, \bar{z}) \wedge \forall x\ \big(\theta(x, \bar{z}) \to \theta(x+n+1, \bar{z})\big) \to \forall x\ \theta(x, \bar{z})\Big).$$

- Let $n \in \mathbb{N}$. The *$(n+1)$-induction axiom* $\mathrm{I}_x^{n+1}\theta$ is

$$\forall \bar{z}\ \Big(\bigwedge_{k<n+1} \theta(k, \bar{z}) \wedge \forall x\ \Big(\bigwedge_{k<n+1} \theta(x+k, \bar{z}) \to \theta(x+n+1, \bar{z})\Big) \to \forall x\ \theta(x, \bar{z})\Big).$$

- The *polynomial induction axiom* $\mathrm{I}_x^{\mathrm{p}}\theta$ is

$$\forall \bar{z}\ \Big(\theta(0, \bar{z}) \wedge \forall x\ \big(\theta(x, \bar{z}) \to \theta(2x, \bar{z}) \wedge \theta(2x+1, \bar{z})\big) \to \forall x\ \theta(x, \bar{z})\Big).$$

Each of these schemes naturally induces a notion of inductiveness.

Definition. Let $\varphi(x)$ be an $\mathscr{L}_{\mathrm{OR}}$ formula with precisely one free variable x.

- $\varphi(x)$ is *$<$-inductive* if $\mathrm{PA}^{-} \vdash \forall y\ \big(\forall x{<}y\ \varphi(x) \to \varphi(y)\big)$.
- Let $n \in \mathbb{N}$. We say $\varphi(x)$ is *$(n+1)$-step inductive* if

$$\mathrm{PA}^{-} \vdash \bigwedge_{k<n+1} \varphi(k) \wedge \forall x\ \big(\varphi(x) \to \varphi(x+n+1)\big).$$

- Let $n \in \mathbb{N}$. We say $\varphi(x)$ is $(n+1)$-*inductive* if

$$\mathrm{PA}^- \vdash \bigwedge_{k<n+1} \varphi(k) \wedge \forall x \Big(\bigwedge_{k<n+1} \varphi(x+k) \to \varphi(x+n+1) \Big).$$

- $\varphi(x)$ is *polynomially inductive*, or simply p-*inductive*, if

$$\mathrm{PA}^- \vdash \varphi(0) \wedge \forall x \, \big(\varphi(x) \to \varphi(2x) \wedge \varphi(2x+1) \big).$$

The first of the above formulations, $<$-induction, is also known as strong induction and is commonly used in mathematics. The consideration of $(n+1)$-step induction is motivated by its use in [6]. The k-induction scheme has become popular in computer-aided verification, see, e.g., [26, 16]. Polynomial induction has been introduced by Buss in [8] for the study of weak arithmetical theories and their relation to computational complexity classes. In this paper, we restrict ourselves to the base-2 polynomial induction scheme in order to keep the exposition sufficiently simple; see Remark 7.5 for some extra information about other bases.

Proposition 4.1. The following are equivalent for every $n \in \mathbb{N}$:

(i) $\mathrm{PA} = \mathrm{PA}^- + \{\mathrm{I}_x\theta : \theta(x, \bar{z}) \text{ is an } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}$;

(ii) $\mathrm{PA}^- + \{\mathrm{I}^{\leqslant}_x\theta : \theta(x, \bar{z}) \text{ is an } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}$;

(iii) $\mathrm{PA}^- + \{\mathrm{I}^{(n+1)\text{-step}}_x\theta : \theta(x, \bar{z}) \text{ is an } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}$;

(iv) $\mathrm{PA}^- + \{\mathrm{I}^{n+1}_x\theta : \theta(x, \bar{z}) \text{ is an } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}$;

(v) $\mathrm{PA}^- + \{\mathrm{I}^{\mathrm{p}}_x\theta : \theta(x, \bar{z}) \text{ is an } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}$.

The proof of Proposition 4.1 is a straightforward exercise. So instead of showing it here, we show its quantifier-free counterpart, Theorem 4.3. Everything carries over *except* the arguments for (i) $\Rightarrow$ (ii) and (v) $\Rightarrow$ (i). The former of these implications was shown to remain true by a more substantial argument in Shepherdson [27]. First we gather a few statements easily provable in PA^- that will be useful at several occasions.

Lemma 4.2. PA^- proves

(a) $\forall x \, \forall y \, \forall z \, (x + z < y + z \to x < y)$,

(b) $\forall x \, \forall y \, \forall z \, (x \times z < y \times z \to x < y)$,

(c) $\forall x \, \forall y \, (x < y \to x + 1 \leqslant y)$,

(d) $\forall x \, \big(x \leqslant n \to \bigvee_{k \leqslant n} x = k \big)$ for every $n \in \mathbb{N}$,

(e) $\forall x \, (x < x + 1)$, and

(f) $\forall x\ (x \neq 0 \to \exists y\ (x = y + 1))$.

Proof. For (a) and (b), see the top of page 18 in Kaye [23]. For (c) and (d), see Proposition 2.1 and Lemma 2.6 in Kaye [23] respectively.

For (e), we see that $x + 1 = (x + 0) + 1$ by axiom P6. So we are done by axiom P14.

For (f), work over PA^-. If $x \neq 0$, then axiom P16 implies $x > 0$, and so we get the y we want by P14. □

Theorem 4.3 (mostly Shepherdson)**.** The following are equivalent for all $n \in \mathbb{N}$:

(i) $\mathrm{IOpen} = \mathrm{PA}^- + \{\mathrm{I}_x\theta : \theta(x, \bar{z}) \text{ is a quantifier-free } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}$;

(ii) $\mathrm{PA}^- + \{\mathrm{I}^{\leq}_x\theta : \theta(x, \bar{z}) \text{ is a quantifier-free } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}$;

(iii) $\mathrm{PA}^- + \{\mathrm{I}^{(n+1)\text{-step}}_x\theta : \theta(x, \bar{z}) \text{ is a quantifier-free } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}$;

(iv) $\mathrm{PA}^- + \{\mathrm{I}^{n+1}_x\theta : \theta(x, \bar{z}) \text{ is a quantifier-free } \mathscr{L}_{\mathrm{OR}} \text{ formula}\}$.

Proof. We refer the reader to the original Shepherdson paper [27] for a proof of (i) ⇒ (ii). We do not include a proof here because it would be too distracting for the present paper to set up all the algebraic materials for the argument. Nevertheless, some of the ideas can be found in Section 8.

Consider (ii) ⇒ (iii). Work over (ii). Fix $\bar{z}$ and let $\theta(x, \bar{z})$ be a quantifier-free $\mathscr{L}_{\mathrm{OR}}$ formula such that $\exists x\ \neg\theta(x, \bar{z})$. Use (ii) to find $x_0 = (\min x)(\neg\theta(x, \bar{z}))$. If $x_0 < n{+}1$, then $\bigvee_{k<n+1} \neg\theta(k, \bar{z})$ by Lemma 4.2(d). So suppose $x_0 \geqslant n{+}1$. Then $x_0 > n$ by (the contrapositive of) Lemma 4.2(c). Apply axiom P14 to find w_0 such that $x_0 = n + w_0 + 1 > w_0$. Then $\theta(w_0, \bar{z})$ holds by the minimality of x_0, but $\neg\theta(w_0 + n + 1, \bar{z})$, as required.

The implication (iii) ⇒ (iv) is clear.

Consider (iv) ⇒ (i). Work over (iv). Fix $\bar{z}$ and let $\theta(x, \bar{z})$ be a quantifier-free $\mathscr{L}_{\mathrm{OR}}$ formula such that $\theta(0, \bar{z}) \wedge \forall x\ \big(\theta(x, \bar{z}) \to \theta(x + 1, \bar{z})\big)$. Then

$$\bigwedge_{k<n+1} \theta(k, \bar{z}) \wedge \forall x\ \Big(\bigwedge_{k<n+1} \theta(x, \bar{z}) \to \theta(x + n + 1, \bar{z})\Big),$$

and so we are done by (iv). □

Following the proof of Theorem V.4.6 in Hájek–Pudlák [19], one can prove IOpen from PA^- plus the polynomial induction scheme for formulas of the form $\forall y{<}t\ \eta(x, y, \bar{z})$ where $\eta(x, y, \bar{z})$ is a quantifier-free $\mathscr{L}_{\mathrm{OR}}$ formula and t is an $\mathscr{L}_{\mathrm{OR}}$ term not involving y. Whether the use of this extra bounded quantifier can be eliminated is not clear to us. Before we prove the quantifier-free analogue of (i) ⇒ (v) of Proposition 4.1, we show the following small lemma which will be handy in a couple of places.

Lemma 4.4. IOpen $\vdash \forall x, d\ \exists q, r\ (d \neq 0 \to x = qd + r \wedge r < d)$ but $\mathrm{PA}^- \nvdash \forall x\ \exists y\ (x = 2y \vee x = 2y + 1)$.

Proof. For the provability part, we follow Kaye [22, page 5]. Work over IOpen. Take any x, d with $d \neq 0$. Using the axioms of PA^- and Lemma 4.2(f), one can routinely verify that $0d \leqslant x < (x+1)d$. Apply induction on q for the atomic formula $qd \leqslant x$ to find q which satisfies $qd \leqslant x < (q+1)d = qd + d$. Then setting $r = x - qd < qd + d - qd = d$ gives us what we want.

For the unprovability part, consider the set $\mathbb{Z}[X]^+$ of all elements of the polynomial ring $\mathbb{Z}[X]$ that either are zero or have positive leading coefficients. It is naturally a model of PA^-, as the reader can verify [23, Section 2.1]. Clearly,

$$\mathbb{Z}[X]^+ \models \neg\exists y\ (X = 2y \vee X = 2y + 1). \qquad \square$$

Proposition 4.5. IOpen proves

$$\forall \bar{z}\ \big(\theta(0, \bar{z}) \wedge \forall x\ \big(\theta(x, \bar{z}) \to \theta(2x, \bar{z}) \wedge \theta(2x + 1, \bar{z})\big) \to \forall x\ \theta(x, \bar{z})\big)$$

for every quantifier-free $\mathscr{L}_{\mathrm{OR}}$ formula $\theta(x, \bar{z})$.

Proof. Work over IOpen. Let $\theta(x, \bar{z})$ be a quantifier-free $\mathscr{L}_{\mathrm{OR}}$ formula and $\bar{z}$ be parameters such that $\exists x\ \neg\theta(x, \bar{z})$. By Proposition 4.3(ii), this has a least witness, say x_0. If $x_0 = 0$, then $\neg\theta(0, \bar{z})$. So suppose $x_0 \neq 0$. Using Lemma 4.4, find y_0 such that $x_0 = 2y_0$ or $x_0 = 2y_0 + 1$. Since $x_0 \neq 0$, we know $y_0 < x_0$ by Lemma 4.2(f) and axiom P14. So $\theta(y_0, \bar{z})$ by the minimality of x_0. We are thus done because $\neg\theta(2y_0, \bar{z}) \vee \neg\theta(2y_0 + 1, \bar{z})$. $\square$

4.2 Variations of induction scheme and proof shape

All results of Section 3, with the possible exception of Theorem 3.3, actually remain true when the induction scheme is replaced by another scheme in Proposition 4.1. The proofs are straightforward modifications of what we presented there, and hence are left to the reader. In this section we abstract one part that may be of independent interest. This stems from the observation that, in a sense, Proposition 2.6 is the only property of PA one needs in establishing Theorem 3.1.

Definition. Let $\mathscr{L}$ be a language. Denote by $\mathscr{L}(\mathsf{X})$ the language obtained from $\mathscr{L}$ by adding one new unary predicate symbol X. Let S be a sentence in $\mathscr{L}(\mathsf{X})$. If $\varphi(x, \bar{z})$ is an $\mathscr{L}$ formula, then define $S\varphi$ to be the universal closure of the $\mathscr{L}$ formula obtained from S by replacing each occurrence of $\mathsf{X}(\dots)$ by $\varphi(\dots, \bar{z})$. The *scheme determined by* S, denoted $S\mathscr{L}$, is defined to be $\{S\varphi : \varphi \in \mathscr{L}\}$. The scheme $S\mathscr{L}$ is *mergeable* over an $\mathscr{L}$ theory B if for all $\mathscr{L}$ formulas ψ_0, ψ_1, there is an $\mathscr{L}$ formula φ such that $B + S\varphi \vdash S\psi_0 \wedge S\psi_1$.

Proposition 2.6 demonstrates the mergeability of the successor induction scheme over PA^-. A very similar proof shows that all other schemes in Proposition 4.1 are also mergeable over PA^-. Another example is the comprehension scheme in second-order arithmetic, which is mergeable over rather weak base theories.

Theorem 4.6. Fix a language $\mathscr{L}$. Let $S\mathscr{L}$ be a scheme and B be an $\mathscr{L}$ theory such that $B \vdash S\top$. The following are equivalent.

(i) $S\mathscr{L}$ is mergeable over B.

(ii) For every $\mathscr{L}$ sentence σ provable from $B + S\mathscr{L}$, there exists an $\mathscr{L}$ formula $\varphi(x, \bar{z})$ such that $B \vdash S\varphi \leftrightarrow \sigma$.

Proof. For (i) $\Rightarrow$ (ii), imitate the proof of Theorem 3.1. Conversely, suppose (ii) holds. Pick arbitrary $\mathscr{L}$ formulas ψ_0 and ψ_1. Define $\sigma = S\psi_0 \wedge S\psi_1$. Then $B + S\mathscr{L} \vdash \sigma$ trivially. By (ii), we get an $\mathscr{L}$ formula φ such that $B \vdash S\varphi \leftrightarrow \sigma$ and hence $B + S\varphi \vdash S\psi_0 \wedge S\psi_1$, as required. □

Corollary 4.7. Let $n \in \mathbb{N}$ and let σ be an $\mathscr{L}_{\mathrm{OR}}$ sentence. The following are equivalent.

(i) $\mathrm{PA} \vdash \sigma$.

(ii) There is a $<$-inductive formula $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \leftrightarrow \sigma$.

(iii) There is an $(n+1)$-step inductive formula $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \leftrightarrow \sigma$.

(iv) There is an $(n+1)$-inductive formula $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \leftrightarrow \sigma$.

(v) There is a p-inductive formula $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \leftrightarrow \sigma$. □

4.3 Walther's method for comparing induction schemes

In the context of inductive theorem proving, Walther [32, Section 7] proposed a method of comparing induction axioms. Let us formulate his method as follows. If B is a finite set of natural numbers and S is a finite set of $\mathscr{L}_{\mathrm{OR}}$ terms, then define $\mathrm{PA}(B, S)$ to be the $\mathscr{L}_{\mathrm{OR}}$ theory axiomatized by PA^- and the following axioms for all $\mathscr{L}_{\mathrm{OR}}$ formulas $\theta(x, \bar{z})$:

$$\forall \bar{z}\ \Big(\bigwedge_{k \in B} \theta(k, \bar{z}) \wedge \bigwedge_{t \in S} \forall x\ \big(\theta(x, \bar{z}) \to \theta(t(x, \bar{z}), \bar{z})\big) \to \forall x\ \theta(x, \bar{z}) \Big).$$

In this terminology, the $(n+1)$-step induction scheme in Proposition 4.1(iii) is essentially

$$\mathrm{PA}(\{k \in \mathbb{N} : k < n+1\}, \{x + n + 1\}).$$

(Some choices of B and S may give rise to a scheme $\mathrm{PA}(B,S)$ that is not true in $\mathbb{N}$, but this does not concern us here.) Walther observed that if $B \subseteq B'$ and $S \subseteq S'$, then $\mathrm{PA}(B,S) \vdash \mathrm{PA}(B',S')$.

In this section we make the observation (not stated in [32]) that this method for comparing induction schemes is incomplete in the sense that the converse implication is not true. To see this, take $m, n \in \mathbb{N}$ such that $m < n$. Then

$$\begin{aligned}&\mathrm{PA}(\{k \in \mathbb{N} : k < n+1\}, \{x+m+1\})\\ &\equiv \mathrm{PA}(\{k \in \mathbb{N} : k < m+1\}, \{x+m+1\})\\ &\equiv \mathrm{PA}(\{k \in \mathbb{N} : k < n+1\}, \{x+n+1\})\end{aligned}$$

by Proposition 4.1. Clearly $\{x+m+1\} \not\subseteq \{x+n+1\}$.

Thanks to Theorem 4.3, the observation in the previous paragraph has an analogue in the quantifier-free context. We formulate this as follows. If B is a finite set of natural numbers and S is a finite set of $\mathscr{L}_{\mathrm{OR}}$ terms, then define $\mathrm{IOpen}(B,S)$ to be the $\mathscr{L}_{\mathrm{OR}}$ theory axiomatized by PA^- and the following axioms for all quantifier-free $\mathscr{L}_{\mathrm{OR}}$ formulas $\theta(x,\bar{z})$:

$$\forall \bar{z}\, \Big(\bigwedge\!\!\!\bigwedge_{k \in B} \theta(k,\bar{z}) \wedge \bigwedge\!\!\!\bigwedge_{t \in S} \forall x\, \big(\theta(x,\bar{z}) \to \theta(t(x,\bar{z}),\bar{z})\big) \to \forall x\, \theta(x,\bar{z}) \Big).$$

Clearly, if $B \subseteq B'$ and $S \subseteq S'$, then $\mathrm{IOpen}(B,S) \vdash \mathrm{IOpen}(B',S')$. The converse is not true because whenever $m, n \in \mathbb{N}$ such that $m < n$, we have

$$\begin{aligned}&\mathrm{IOpen}(\{k \in \mathbb{N} : k < n+1\}, \{x+m+1\})\\ &\equiv \mathrm{IOpen}(\{k \in \mathbb{N} : k < m+1\}, \{x+m+1\})\\ &\equiv \mathrm{IOpen}(\{k \in \mathbb{N} : k < n+1\}, \{x+n+1\})\end{aligned}$$

by Theorem 4.3, but $\{x+m+1\} \not\subseteq \{x+n+1\}$.

5 Non-closure properties of cuts

In this short technical section we will establish an auxiliary result on cuts that will be used in Section 6 and Section 7 for comparing solutions to ITP and notions of inductiveness respectively. Note that the meaning of the word "cut" here is different from that in Section 2.2. In the study of weak theories of arithmetic, these cuts are frequently used in interpretations. They also serve as notions of smallness in many arguments. In the following, we borrow some terminology from Visser [31].

Definition. A *cut* is an inductive formula $\varphi(x)$ such that

$$\mathrm{PA}^- \vdash \forall x, y\, \big(x < y \wedge \varphi(y) \to \varphi(x)\big).$$

An *a-cut* is a cut $\varphi(x)$ such that

$$\mathrm{PA}^- \vdash \forall x\ \big(\varphi(x) \to \varphi(x+x)\big).$$

An *am-cut* is an a-cut $\varphi(x)$ such that

$$\mathrm{PA}^- \vdash \forall x\ \big(\varphi(x) \to \varphi(x \times x)\big).$$

Inductive formulas, cuts, a-cuts, and am-cuts are all different notions. This fact seems to be well known, but we can find no proof of this in the literature. So we include one here. The proof assumes familiarity with some parts of the Hájek–Pudlák book [19].

Lemma 5.1. There are $\mathscr{L}_{\mathrm{OR}}$ formulas $\varphi(x)$ and $\delta(x)$ such that

(1) $\varphi(x)$ is an a-cut but not an am-cut;

(2) $\mathrm{PA}^- \vdash \exists! x\ \delta(x) \wedge \forall x\ \big(\delta(x) \to \varphi(x)\big)$; and

(3) $\mathrm{PA}^- \vdash \neg\forall x\ \big(\varphi(x) \to \varphi(x \times x)\big) \to \forall x\ \big(\delta(x) \to \varphi(x) \wedge \neg\varphi(x \times x)\big)$.

Proof. Let $M \models \mathrm{PA}$ which contains a nonstandard element a that is definable by a bounded formula. For instance, if $M \models \neg\mathrm{Con}(\mathrm{PA})$, then we can take a to be the least (code of a) proof of contradiction from PA in M. Fix a bounded formula $\delta_0(x)$ such that $M \models \exists! x\ \delta_0(x) \wedge \delta_0(a)$. Notice

$$I = 2^{\mathbb{N}a} = \{x \in M : x < 2^{na} \text{ for some } n \in \mathbb{N}\}$$

is an initial segment of M, and $2^{ma} \times 2^{na} = 2^{(m+n)a} \in I$ whenever $m, n \in \mathbb{N}$. So I is itself an $\mathscr{L}_{\mathrm{OR}}$ structure. Since $\delta_0(x)$ is a bounded formula, it is straightforward [19, Remark IV.1.18] to check that a is the unique element which satisfies $\delta_0(x)$ in I. We know $I \models \mathrm{I}\Delta_0$ because $\mathrm{I}\Delta_0$ is preserved under taking initial segments [19, Remark IV.1.21]. Thus

$$\log I = \{x \in I : I \models \exists u\ (2^x = u)\}$$

is an initial segment of I that is closed under $+$ because

$$\begin{aligned}\mathrm{I}\Delta_0 \vdash \forall x, y\ \big(\exists u\ (2^x = u) \wedge y \leqslant x \to \exists v \leqslant u\ (2^y = v)\big)\\ \wedge \forall u, v, x, y\ (2^x = u \wedge 2^y = v \to 2^{x+y} = uv),\end{aligned}$$

as verified in Hájek–Pudlák [19, Lemma V.3.8(iii) and page 302]. We know $\log I$ contains a but not a^2 because

$$\log I = \{x \in I : x < na \text{ for some } n \in \mathbb{N}\}$$

and $a \notin \mathbb{N}$. In particular, it is not closed under $\times$. The whole situation in I can be captured by the sentence

$$\begin{aligned}\sigma \quad = \quad & \varphi_0(0) \wedge \forall x\ (\varphi_0(x) \to \varphi_0(x+x)) \wedge \forall x, y\ (x < y \wedge \varphi_0(y) \to \varphi_0(x))\\ & \wedge \exists! x\ \delta_0(x) \wedge \exists x\ (\delta_0(x) \wedge \varphi_0(x) \wedge \neg\varphi_0(x \times x)),\end{aligned}$$

where $\varphi_0(x)$ denotes the formula $\exists u\ (2^x = u)$. Clearly,

$$\begin{array}{rcl} \varphi(x) & = & \sigma \to \varphi_0(x) \quad \text{and} \\ \delta(x) & = & (\sigma \to \delta_0(x)) \wedge (\neg\sigma \to x = 0) \end{array}$$

have the properties we want. □

Corollary 5.2. (a) There is an inductive formula that is not a cut.

(b) There is a cut that is not an a-cut.

(c) There is an a-cut that is not an am-cut.

Proof. Let $\varphi(x)$ and $\delta(x)$ be as given by Lemma 5.1.

(a) Take $\varphi(x) \vee \exists c\ (\delta(c) \wedge x \geqslant c^2)$.

(b) Take $\exists c, z\ (\varphi(z) \wedge \delta(c) \wedge x \leqslant c^2 + z)$.

(c) Take $\varphi(x)$. □

Solovay's method of shortening cuts [19, Theorem III.3.5] provides a general way of producing cuts with various closure properties. However, a technique for producing cuts with specific non-closure properties in large quantities does not seem available at present. In particular, it is not clear whether there is an inverse of Solovay's method.

Question 5.3. Let $\varphi(x)$ be a cut. Assuming $\mathrm{PA}^- \nvdash \forall x\ \varphi(x)$, can one always find a cut $\psi(x)$ such that PA^- proves

$$\forall x\ \big(\varphi(x) \to \psi(x)\big) \quad \text{and} \quad \forall x\ \varphi(x) \leftrightarrow \forall x\ \psi(x)$$

and $\psi(x)$ is *not* an am-cut?

6 Comparing solutions

Given a PA-theorem σ, our formulation of ITP considers the set of solutions, the search space, to be all inductive $\varphi(x)$ such that $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \to \sigma$. In automated theorem proving, restrictions of the search space (usually such that completeness is retained) play an important role. In this short section, based on the results of Section 5, we make a comment on comparisons of two solutions $\varphi_0(x)$ and $\varphi_1(x)$.

There are (at least) two ways of comparing two induction axioms $\forall x\ \varphi_0(x)$ and $\forall x\ \varphi_1(x)$, where $\varphi_0(x)$ and $\varphi_1(x)$ are inductive formulas. The first way is to say $\forall x\ \varphi_0(x)$ is stronger than $\forall x\ \varphi_1(x)$ when $\mathrm{PA}^- \vdash \forall x\ \varphi_0(x) \to \forall x\ \varphi_1(x)$. The second way is to say $\forall x\ \varphi_0(x)$ is stronger than $\forall x\ \varphi_1(x)$ when $\mathrm{PA}^- \vdash \forall x\ (\varphi_0(x) \to \varphi_1(x))$. Clearly, if an induction axiom is stronger than another one in the second sense, then it is also stronger in the first sense. The converse, however, is not true.

Proposition 6.1. There are cuts $\varphi_0(x), \varphi_1(x)$ such that PA^- proves

$$\forall x\ \varphi_0(x) \leftrightarrow \forall x\ \varphi_1(x) \quad \text{and} \quad \forall x\ (\varphi_0(x) \to \varphi_1(x))$$

but $\mathrm{PA}^- \nvdash \forall x\ (\varphi_1(x) \to \varphi_0(x))$.

Proof. Let $\varphi(x)$ be as given by Lemma 5.1. We show that $\varphi_0(x) = \varphi(x \times x)$ and $\varphi_1(x) = \varphi(x)$ have the properties we want.

As $\varphi(x)$ is an a-cut, it is not hard to check that $\varphi_0(x)$ is a cut. Next, notice $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \to \forall x\ \varphi(x \times x)$ trivially. Thus $\mathrm{PA}^- \vdash \forall x\ \varphi_1(x) \to \forall x\ \varphi_0(x)$. Notice also that $\mathrm{PA}^- \vdash \forall x\ \big(\varphi(x \times x) \to \varphi(x)\big)$ because $\varphi(x)$ is a cut. So $\mathrm{PA}^- \vdash \forall x\ (\varphi_0(x) \to \varphi_1(x))$. Finally, we know $\mathrm{PA}^- \nvdash \forall x\ (\varphi_1(x) \to \varphi_0(x))$ since $\varphi(x)$ is not an am-cut. □

We can view Proposition 6.1 from another angle. One may expect that the weakest induction axiom which can prove an induction axiom $\mathrm{I}\varphi_0$ is $\mathrm{I}\varphi_0$ itself. The proposition above says that this is not the case *if* we adopt the second way of comparing induction axioms: $\mathrm{I}\varphi_0$ may be provable from another induction axiom $\mathrm{I}\varphi_1$ where $\varphi_1(x)$ is a strictly weaker than $\varphi_0(x)$ over PA^- as formulas. One may question whether there is always a weakest induction axiom for proving a given theorem of PA in this sense. We do not know the answer.

Question 6.2. Let $\varphi_0(x)$ be a cut. Assuming $\mathrm{PA}^- \nvdash \forall x\ \varphi_0(x)$, can one always find a cut $\varphi_1(x)$ such that PA^- proves

$$\forall x\ \varphi_0(x) \leftrightarrow \forall x\ \varphi_1(x) \quad \text{and} \quad \forall x\ (\varphi_0(x) \to \varphi_1(x))$$

but $\mathrm{PA}^- \nvdash \forall x\ (\varphi_1(x) \to \varphi_0(x))$?

As one can show by imitating our proof of Proposition 6.1, a positive answer to Question 5.3 implies a positive answer to Question 6.2.

7 Comparing notions of inductiveness

The question of how an inductive theorem prover should choose the induction rule to be applied to its current goal has received a great deal of attention in the literature on inductive theorem proving. Pioneering work on this question has been done in the context of the ACL2 system and its predecessor NQTHM through the introduction of the recursion analysis technique [2]. This technique, which suggests an induction rule based on the type of recursion present in the goal, is specific to goals involving primitive recursive functions; see also Bundy et al. [7] and Stevens [28]. However, the choice of the induction rule also plays an important role in other contexts; see the discussion on "induction revision" in the book [6] by Bundy et al., for example.

The interest in this question comes from the tension between the choice of the induction rule and the choice of the induction formula when proving a goal: the more flexibility we have in choosing the induction rule, the less flexibility we need in choosing the induction formula. In the very extreme case, one can fix an induction formula, e.g., the goal, and search for an induction rule with respect to which this formula is inductive. Thus one can dispose of the difficult task of finding a non-analytic induction formula and simply search for a suitable induction rule. In this section we will carry out a comparison of the formulations of induction introduced in Section 4 from this point of view: we fix a formula $\varphi(x)$ and ask with respect to which induction schemes it is inductive. Let X be any notion of inductiveness defined in Section 4. Then we can state the computational problem

ITP^X	
Input:	A sentence σ provable in PA
Output:	An X-inductive formula $\varphi(x)$ s.t. $\mathrm{PA}^- \vdash \forall x\ \varphi(x) \to \sigma$

By Proposition 4.1 all these problems define total relations. By Theorem 4.3, the quantifier-free versions of ITP^X are equivalent for all X except possibly polynomial induction.

In this section, we investigate implications between these notions of inductiveness. As we will see, some notions turn out to be incomparable. For instance, not all 2-step inductive formulas are 3-step inductive, and not all 3-step inductive formulas are 2-step inductive. The following proposition lists all the implications we can establish.

Proposition 7.1. Let $m, n \in \mathbb{N}$.

(a) The following are equivalent for a formula: it is inductive; it is 1-step inductive; and it is 1-inductive.

(b) If $m+1$ divides $n+1$, then all $(m+1)$-step inductive formulas are $(n+1)$-step inductive.

(c) If $m \leqslant n$, then all $(m+1)$-inductive formulas are $(n+1)$-inductive.

(d) All $(n+1)$-step inductive formulas are $(n+1)$-inductive.

(e) If a formula is $(n+1)$-inductive for some $n \in \mathbb{N}$, then it is $<$-inductive.

(f) A formula $\varphi(x)$ is $<$-inductive if and only if $\forall x' \leqslant x\ \varphi(x')$ is inductive.

Proof. Parts (a)–(d) are easy exercises.

(e) Suppose $\varphi(x)$ is $(n+1)$-inductive. Work over PA^-. Let y_0 be such that $\forall x < y_0\ \varphi(x)$. If $y_0 \leqslant n$, then $\varphi(y_0)$ holds by Lemma 4.2(d) because $\varphi(x)$ is $(n+1)$-inductive. So suppose $y_0 > n$. Use axiom P14 to find x_0

such that $y_0 = x_0+n+1$. Then $\bigwedge\!\!\!\bigwedge_{k<n+1} \varphi(x_0+k)$ because $x_0+k < y_0$ for each $k < n+1$ by P14 again. As $\varphi(x)$ is $(n+1)$-inductive and $y_0 = x_0 + n + 1$, this implies $\varphi(y_0)$.

(f) First, assume $\forall x' \leqslant x\ \varphi(x')$ is inductive. Work over PA^-. Let y_0 be such that $\forall x < y_0\ \varphi(x)$. If $y_0 = 0$, then $\varphi(y_0)$ by our assumption. So suppose $y_0 \neq 0$. Apply Lemma 4.2(f) to find y_0' such that $y_0 = y_0' + 1$. Notice $\forall x \leqslant y_0'\ \varphi(x)$ by Lemma 4.2(e). Our assumption then implies $\forall x \leqslant y_0\ \varphi(x)$, which, in particular, tells us that $\varphi(y_0)$.

Conversely, suppose $\varphi(x)$ is $<$-inductive. Work over PA^- again. With the help of P16, P9 and P10, we have $\varphi(0)$ trivially by $<$-inductiveness. So $\forall x' \leqslant 0\ \varphi(x')$. Let x_0 be such that $\forall x' \leqslant x_0\ \varphi(x')$. Then Lemma 4.2(c) and (a) imply $\forall x' < x_0 + 1\ \varphi(x')$ and so $\varphi(x_0+1)$ by $<$-inductiveness. This shows $\forall x' \leqslant x_0 + 1\ \varphi(x')$, as required. □

We now proceed to show non-implications (based on the non-closure properties of cuts established in Section 5). In particular, we show that beyond Proposition 7.1 there is no further implication between "$<$-inductive", the "$(n+1)$-step inductive"s, and the "$(n+1)$-inductive"s.

Proposition 7.2. Let $m, n \in \mathbb{N}$.

(a) If $(m+1)$ does not divide $(n+1)$, then there is an $(m+1)$-step inductive formula that is not $(n+1)$-step inductive.

(b) If $m > n$, then there is an $(m+1)$-inductive formula that is not $(n+1)$-inductive.

(c) If $n \geqslant 1$, then there is an $(n+1)$-inductive formula that is not $(n+1)$-step inductive.

(d) There is a $<$-inductive formula that is not $(n+1)$-inductive.

Proof. Let $\varphi(x)$ and $\delta(x)$ be $\mathscr{L}_{\mathrm{OR}}$ formulas given by Lemma 5.1. Recall $\varphi(x)$ is not an am-cut. Take $c \in M \models \mathrm{PA}^- + \neg\forall x\ (\varphi(x) \to \varphi(x^2)) + \delta(c)$.

(a) Assume all $(m+1)$-step inductive formulas are $(n+1)$-step inductive. Consider the formula

$$\chi(x) \quad = \quad \varphi(x) \vee \exists y\ (x = (m+1)y).$$

We first show that $\chi(x)$ is $(m+1)$-step inductive. By our assumption on m and n, this will imply $\chi(x)$ is $(n+1)$-step inductive. Work over PA^-. We know $\bigwedge\!\!\!\bigwedge_{k<m+1} \chi(k)$ because $\varphi(x)$ is inductive. Take x_0 satisfying $\chi(x_0)$. If y_0 is such that $x_0 = (m+1)y_0$, then $x_0 + m + 1 = (m+1)(y_0+1)$ by P8 and P5, and thus $\chi(x_0+m+1)$. So suppose $\neg\exists y\ (x_0 = (m+1)y)$. Then $\varphi(x_0)$ holds by the definition

of the formula $\chi(x)$. As $\varphi(x)$ is an inductive formula, this implies $\varphi(x_0 + m + 1)$ and hence $\chi(x_0 + m + 1)$.

Now look at our model M. Let $a = (m+1)(n+1)c^2$, so that $M \models \chi(a)$. Notice $M \models \neg\varphi(c^2)$ by condition (3) in Lemma 5.1. So $M \models \neg\varphi(a + n + 1)$ because $\varphi(x)$ is a cut and $a + n + 1 = c^2 + c^2(mn + m + n) + n + 1 > c^2$ by axiom P14. However, the previous paragraph tells us that $M \models \chi(a + n + 1)$. So $M \models \exists y\ (a + n + 1 = (m + 1)y)$ by the definition of $\chi(x)$.

Find $b \in M$ such that $a + n + 1 = (m + 1)b$. By axiom P14, we know $(m + 1)b = a + n + 1 > a = (m + 1)(n + 1)c^2$. So Lemma 4.2(b) implies $b > (n + 1)c^2$. Apply axiom P14 to find $z \in M$ such that $(n + 1)c^2 + z + 1 = b$. Then

$$a+(m+1)(z+1) = (m+1)(n+1)c^2+(m+1)(z+1) = (m+1)b = a+n+1$$

and thus $(m + 1)(z + 1) = n + 1$ by axioms P11 and P12. Now $n+1 = z+m(z+1)+1$. Hence $n+1 > z$ by axiom P14. Lemma 4.2(d) then tells us $z \in \mathbb{N}$. We can thus conclude that $(m+1)$ divides $(n+1)$, which is what we want.

(b) We show that the formula

$$\rho(x) \quad = \quad \varphi(x) \vee \forall c \left(\delta(c) \to \bigvee_{k<n+1} x = c^2 + k\right)$$

is an example we want.

Suppose $m > n$. We show that $\rho(x)$ is m-inductive. Work over PA^-. If $\forall x\ (\varphi(x) \to \varphi(x^2))$, then $\forall x\ (\rho(x) \leftrightarrow \varphi(x))$ by condition (2) in Lemma 5.1, and so we are done by Proposition 7.1(a) and (c). So suppose $\neg\forall x\ (\varphi(x) \to \varphi(x^2))$. Since $\varphi(x)$ is inductive, we know $\bigwedge_{k<m+1} \rho(k)$. Let c, x_0 be such that $\delta(c) \wedge \bigwedge_{k<m+1} \rho(x_0 + k)$. Take $k < n + 1$. We know $c^2 + k + m \geqslant c^2 + m > c^2 + n$ by axiom P12, and $\neg\varphi(c^2 + k + m)$ since $\varphi(x)$ is a cut and $\neg\varphi(c^2)$. These together imply $\neg\rho(c^2 + k + m)$. So $x_0 \neq c^2 + k$ because $\rho(x_0 + m)$ holds. As the choice of $k < n + 1$ is arbitrary, we know $\varphi(x_0)$ by the definition of $\rho(x)$. It follows that $\varphi(x_0+m+1)$ because $\varphi(x)$ is inductive. Hence $\rho(x_0 + m + 1)$.

We now show that $\rho(x)$ is not $(n + 1)$-inductive using our model M. The definition of $\rho(x)$ tells us $M \models \bigwedge_{k<n+1} \rho(c^2 + k)$. On the one hand, notice $M \models \neg\varphi(c^2)$ by condition (3) in Lemma 5.1. Also $c^2 + n + 1 > c^2$ by axiom P14. Hence $\neg\varphi(c^2 + n + 1)$ since $\varphi(x)$ is a cut. On the other hand, notice $c^2 + n + 1 > c^2 + k$ for any $k < n + 1$ by axiom P12. These together say that $M \models \neg\rho(c^2 + n + 1)$.

(c) We claim that the formula

$$\rho_0(x) \quad = \quad \varphi(x) \vee \forall c\ (\delta(c) \to x = c^2)$$

has the desired properties.

To show that $\rho_0(x)$ is $(n+1)$-inductive, let us work over PA^-. As in the previous part, we assume $\neg\forall x\ (\varphi(x) \to \varphi(x^2))$. Since $\varphi(x)$ is inductive, we know $\bigwedge\!\!\bigwedge_{k<n+1} \rho_0(k)$. Let c, x_0 be such that $\delta(c) \wedge \bigwedge\!\!\bigwedge_{k<n+1} \rho_0(x_0 + k)$. As $n \geqslant 1$, we know $\rho_0(x_0) \wedge \rho_0(x_0+1)$ in particular. Notice $\neg\varphi(c^2+1)$ because $\varphi(x)$ is a cut and $\neg\varphi(c^2)$ by condition (3) in Lemma 5.1. Also $c^2+1 \neq c^2$ by Lemma 4.2(e). Thus $\neg\rho_0(c^2+1)$ holds, from which one deduces $x_0 \neq c^2$. This implies $\varphi(x_0)$ by the definition of $\rho_0(x)$. Since $\varphi(x)$ is inductive, we conclude that $\varphi(x_0 + n + 1)$.

To show that $\rho_0(x)$ is not $(n+1)$-step inductive, we use M as in the proof of (b). Notice $M \models \rho_0(c^2)$ by the definition of $\rho_0(x)$. On the one hand, we know $M \models \neg\varphi(c^2)$ by condition (3) in Lemma 5.1. Also $c^2 + n + 1 > c^2$ by axiom P14. Hence $\neg\varphi(c^2 + n + 1)$ since $\varphi(x)$ is a cut. On the other hand, we know $c^2 + n + 1 > c^2$ by P14. These together say that $M \models \neg\rho_0(c^2 + n + 1)$.

(d) It suffices to prove that the formula $\rho(x)$ defined in the proof of (b) is $<$-inductive.

Work over PA^-. Let c, y_0 be such that $\delta(c)$ and $\forall x{<}y_0\ \rho(x)$. We will show $\rho(y_0)$. If $y_0 = 0$, then $\varphi(y_0)$ because $\varphi(x)$ is inductive, and so $\rho(y_0)$ holds. Therefore, assume $y_0 \neq 0$. Apply Lemma 4.2(f) to find x_0 such that $y_0 = x_0 + 1$. Notice $y_0 > x_0$ by Lemma 4.2(e). So, by the choice of y_0, we know $\rho(x_0)$.

Consider the case when $\neg\forall x\ (\varphi(x) \to \varphi(x^2))$. Then $\neg\varphi(c^2)$ by condition (3) in Lemma 5.1. This implies $c^2 \neq 0$ because $\varphi(x)$ is inductive. Use Lemma 4.2(f) to find w such that $c^2 = w + 1$. Then $\neg\varphi(w)$ since $\varphi(x)$ is inductive and $\neg\varphi(w + 1)$ holds. Also, Lemma 4.2(e) implies $w < c^2 \leqslant c^2 + k$ for each $k < n + 1$. We thus know $\neg\rho(w)$ by the definition of $\rho(x)$. Our assumption on y_0 then implies $x_0 < x_0 + 1 = y_0 \leqslant w < w+1 = c^2$. Thus $x_0 \neq c^2 + k$ for any $k < n+1$ by axiom P9, and so $\varphi(x_0)$ must hold.

If $\forall x\ (\varphi(x) \to \varphi(x^2))$, then $\bigwedge\!\!\bigwedge_{k<n+1} \varphi(c^2+k)$ by conditions (1) and (2) in Lemma 5.1. So $\varphi(x_0)$ in either case because $\rho(x_0)$ holds. Since $\varphi(x)$ is inductive, we deduce $\varphi(x_0 + 1)$ and hence $\rho(y_0)$, as required. □

Buss's notion of polynomial induction does not fit well into the picture.

Proposition 7.3. (a) There is a formula that is $(n+1)$-step inductive for every $n \in \mathbb{N}$ but is not p-inductive.

(b) For every $n \in \mathbb{N}$, there is a p-inductive formula that is not $(n+1)$-inductive.

Proof. (a) By Corollary 5.2(b), there is a cut that is not an a-cut. Such a cut is $(n+1)$-step inductive for every $n \in \mathbb{N}$, but it is not p-inductive.

(b) Let $\varphi(x), \delta(x)$ be $\mathscr{L}_{\mathrm{OR}}$ formulas given by Lemma 5.1. Define $\chi(x)$ to be the formula

$$\varphi(x) \vee \exists c\ \exists s\ \exists \ell \begin{pmatrix} \delta(c) \wedge \bigvee\!\!\!\bigvee_{k<n+1} (s)_0 = c^2 + k \wedge (s)_\ell = x \\ \wedge\, \forall i{<}\ell\ \big((s)_{i+1} = 2(s)_i \vee (s)_{i+1} = 2(s)_i + 1\big) \\ \wedge\, \forall i, j{\leqslant}\ell\ \big(i < j \to (s)_i < (s)_j\big) \end{pmatrix}.$$

Here $(s)_i = x$ is the $\mathscr{L}_{\mathrm{OR}}$ formula expressing "the ith element in the sequence coded by s is x" over PA^- given in Jeřábek [20]. It is then straightforward to see that $\chi(x)$ is p-inductive. However, the formula $\chi(x)$ is not $(n+1)$-inductive because $\mathrm{PA}^- \vdash \exists c\ \big(\delta(c) \wedge \bigwedge\!\!\!\bigwedge_{k<n+1} \chi(c^2+k)\big)$ and $\mathrm{PA}^- \nvdash \exists c\ (\delta(c) \wedge \chi(c^2+n+1))$. □

We do not know whether all p-inductive formulas are $<$-inductive. What we have is only a translation of this question into a model-theoretic language.

Proposition 7.4. The following are equivalent.

(i) There is a p-inductive formula that is not $<$-inductive.

(ii) In some model of PA^-, there is a parameter-free definable element that is neither even nor odd.

Proof. First, suppose (ii) fails. Let $\varphi(x)$ be any $\mathscr{L}_{\mathrm{OR}}$ formula that is not $<$-inductive. Find $c \in M \models \mathrm{PA}^- + \forall x{<}c\ \varphi(x) \wedge \neg\varphi(c)$. If $c = 0$, then $\mathrm{PA}^- \nvdash \varphi(0)$ and we are done. So suppose $c \neq 0$. Notice $c = (\min x)(\neg\varphi(x))$. It is, therefore, a parameter-free definable element of M. The failure of (ii) then gives us $d \in M$ such that $c = 2d$ or $c = 2d+1$. By Lemma 4.2(f) and axiom P14, we know $d < c$. Hence $M \models \varphi(d)$ by the minimality of c. This shows $\varphi(x)$ is not p-inductive.

Conversely, suppose (ii) holds. Let $c \in M \models \mathrm{PA}^-$ in which c is a non-even non-odd parameter-free definable element. Find an $\mathscr{L}_{\mathrm{OR}}$ formula $\delta(x)$ that defines c in M. Define $\varphi_0(x)$ to be

$$\exists! y\ \delta(y) \to \exists y\ (\delta(y) \wedge x < y),$$

and let $\varphi(x)$ be

$$\exists s\ \exists \ell\ \big(\varphi_0((s)_0) \wedge (s)_\ell = x \wedge \forall i{<}\ell\ \big((s)_{i+1} = 2(s)_i \vee (s)_{i+1} = 2(s)_i + 1\big)\big),$$

where $(s)_i = x$ is the $\mathscr{L}_{\mathrm{OR}}$ formula expressing "the ith element in the sequence coded by s is x" in our proof of Proposition 7.3. Then $\varphi(x)$ is p-inductive by construction. Notice $M \models \forall x{<}c\ \varphi_0(x)$. So $M \models \forall x{<}c\ \varphi(x)$ because $\mathrm{PA}^- \vdash \forall x\ (\varphi_0(x) \to \varphi(x))$. However, every element $x \in M$ satisfying $\varphi(x)$ must lie strictly below c or be either even or odd. Hence $M \models \neg\varphi(c)$ by our choice of c. We thus know $\varphi(x)$ is not $<$-inductive. □

Remark 7.5. All results in this paper about base-2 polynomial induction generalize to other bases. Let us say an $\mathscr{L}_{\mathrm{OR}}$ formula $\varphi(x)$ is $(n+2)$-p-*inductive*, where $n \in \mathbb{N}$, if

$$\mathrm{PA}^- \vdash \varphi(0) \wedge \forall x\ \Big(\varphi(x) \to \bigwedge\!\!\!\bigwedge_{k<n+2} \varphi\big((n+2)x + k\big)\Big).$$

In this terminology, the p-inductive formulas are precisely the 2-p-inductive formulas. Let $m, n \in \mathbb{N}$. Then $n+2$ being a power of $m+2$ is a necessary and sufficient condition for every $(m+2)$-p-inductive formula to be $(n+2)$-p-inductive. We omit the proof here.

8 Removing the order from the language

Another aspect of our formulation of ITP is that we search for an inductive formula in the same language as the input sentence. In this section we consider a restriction of the language of the inductive formula. Since most of our arguments about IOpen depend on the presence (or, in fact, the quantifier-free definability) of the order $<$ in $\mathscr{L}_{\mathrm{OR}}$, it is natural to ask whether IOpen really becomes strictly weaker when the induction scheme is restricted to formulas in which $<$ does not appear. It turns out that the induction scheme for quantifier-free formulas in this restricted language is already provable in the base theory PA^-. In fact, as pointed out by Richard Kaye in a conversation in April 2016, the theory PA^- even proves the least number principle (also known as the scheme of strong induction, cf. Proposition 4.1(ii) and Theorem 4.3(ii)) for quantifier-free formulas in the restricted language. With Kaye's permission, we include his proof here. Many ideas in the proof were already present in Shepherdson's contribution [27] to Theorem 4.3.

Definition. Denote the language $\{0, 1, +, \times\}$ for rings by $\mathscr{L}_{\mathrm{R}}$. Define

$$\mathrm{IOpen}(\mathscr{L}_{\mathrm{R}}) = \mathrm{PA}^- + \{\mathrm{I}_x\theta : \theta(x, \bar{z}) \text{ is a quantifier-free } \mathscr{L}_{\mathrm{R}} \text{ formula}\}.$$

Let $\mathrm{LOpen}(\mathscr{L}_{\mathrm{R}})$ be the theory axiomatized by PA^- and the scheme

$$\forall \bar{z}\ \big(\forall y\ \big(\forall x{<}y\ \theta(x, \bar{z}) \to \theta(y, \bar{z})\big) \to \forall x\ \theta(x, \bar{z})\big),$$

where θ ranges over quantifier-free $\mathscr{L}_{\mathrm{R}}$ formulas.

The following elementary algebraic fact plays a key role in Kaye's proof.

Lemma 8.1. Let F be a field. For every polynomial $p(x, \bar{a})$ in a single variable x with coefficients $\bar{a} \in F$, the set $\{x \in F : p(x, \bar{a}) = 0\}$ is either finite or equal to F. □

Theorem 8.2 (Kaye)**.** $\mathrm{PA}^- \vdash \mathrm{LOpen}(\mathscr{L}_\mathrm{R})$.

Proof. Let $M \models \mathrm{PA}^-$. We will show $M \models \mathrm{LOpen}(\mathscr{L}_\mathrm{R})$. Fix $\bar{a} \in M$. Consider a quantifier-free $\mathscr{L}_\mathrm{R}$ formula $\theta(x, \bar{z})$. Put $\neg\theta(x, \bar{z})$ in disjunctive normal form

$$\bigvee_{i \leqslant m} \bigwedge_{j \leqslant n} \big(p_{ij}(x, \bar{z}) = 0 \wedge q_{ij}(x, \bar{z}) \neq 0\big),$$

where the p_{ij}'s and the q_{ij}'s are polynomials over $\mathbb{Z}$. It suffices to show that for each $i \leqslant m$, if some $x \in M$ satisfies $\bigwedge_{j \leqslant n} \big(p_{ij}(x, \bar{a}) = 0 \wedge q_{ij}(x, \bar{a}) \neq 0\big)$, then there is a least such x. So for notational convenience, let us assume $m = 0$, so that $\neg\theta(x, \bar{z})$ becomes

$$\bigwedge_{j \leqslant n} \big(p_{0j}(x, \bar{z}) = 0 \wedge q_{0j}(x, \bar{z}) \neq 0\big).$$

Since M is the non-negative part of a discretely ordered ring, we can further simplify the formula $\neg\theta(x, \bar{z})$ to

$$p(x, \bar{z}) = 0 \wedge q(x, \bar{z}) \neq 0$$

by setting $p(x, \bar{z}) = \sum_{j \leqslant n} (p_{0j}(x, \bar{z}))^2$ and $q(x, \bar{z}) = \prod_{j \leqslant n} q_{0j}(x, \bar{z})$.

Assume $M \models \forall y \, \big(\forall x{<}y \; \theta(x, \bar{a}) \to \theta(y, \bar{a})\big)$. We will construct a strictly decreasing sequence $(c_\ell)_{\ell \in \mathbb{N}}$ of elements of M by recursion as follows. If $M \models \forall x \; \theta(x, \bar{a})$, then we are done already. So suppose not, and take an arbitrary $c_0 \in M \models \neg\theta(c_0, \bar{a})$. Next, suppose c_ℓ is found such that $M \models \neg\theta(c_\ell, \bar{a})$. Using our assumption, pick any $c_{\ell+1} < c_\ell$ in M with $M \models \neg\theta(c_{\ell+1}, \bar{a})$ and carry on.

The result of this construction is an infinite sequence $c_0 > c_1 > c_2 > \cdots$ of elements of M such that $p(c_\ell, \bar{a}) = 0$ and $q(c_\ell, \bar{a}) \neq 0$ for all $\ell \in \mathbb{N}$. Lemma 8.1 then implies $p(x, \bar{a})$ is the zero polynomial. Hence, actually $M \models \forall x \, \big(\neg\theta(x, \bar{a}) \leftrightarrow q(x, \bar{a}) \neq 0\big)$. By an external induction using our assumption in the previous paragraph, one sees that $q(k, \bar{a}) = 0$ for all $k \in \mathbb{N}$. Applying Lemma 8.1 again, we conclude $M \models \forall x \; q(x, \bar{a}) = 0$, as required. □

The obvious argument [19, Lemma I.2.8] shows $\mathrm{LOpen}(\mathscr{L}_\mathrm{R}) \vdash \mathrm{IOpen}(\mathscr{L}_\mathrm{R})$. Therefore, Lemma 4.4 and Theorem 8.2 tell us that $\mathrm{IOpen}(\mathscr{L}_\mathrm{R})$ is strictly weaker than IOpen. In terms of computational problems we obtain that

ITPOpen($\mathscr{L}_\mathrm{R}$)	
Input:	A sentence σ provable in PA
Output:	A quantifier-free $\mathscr{L}_\mathrm{R}$ formula $\theta(x, \bar{z})$ s.t. $\mathrm{PA}^- + \mathrm{I}_x\theta \vdash \sigma$

is defined on strictly fewer inputs than

ITPOpen	
Input:	A sentence σ provable in PA
Output:	A quantifier-free $\mathscr{L}_{\mathrm{OR}}$ formula $\theta(x, \bar{z})$ s.t. $\mathrm{PA}^- + \mathrm{I}_x\theta \vdash \sigma$

To recover the strength of IOpen, it suffices to add back, roughly speaking, a single inequality. The proof can be found in Shepherdson's paper [27], as Leszek Kołodziejczyk pointed out.

Theorem 8.3 (Shepherdson)**.** The following are equivalent over PA^-:

(i) IOpen; and

(ii) the scheme

$$\forall x\, \forall \bar{z}\, \big(\theta(0, \bar{z}) \wedge \forall x'{<}x\, \big(\theta(x', \bar{z}) \to \theta(x' + 1, \bar{z})\big) \to \theta(x, \bar{z})\big),$$

where $\theta(x, \bar{z})$ ranges over formulas of the form $p(x, \bar{z}) < q(x, \bar{z})$ in which $p(x, \bar{z})$ and $q(x, \bar{z})$ are $\mathscr{L}_{\mathrm{R}}$ terms. □

Question 8.4. Does $\mathrm{PA}^- + \{\mathrm{I}_x\theta : \theta(x, \bar{z})$ is an atomic $\mathscr{L}_{\mathrm{OR}}$ formula$\}$ prove IOpen?

9 Conclusion

What do we learn from these results for the automation of inductive theorem proving?

In comparison with successor induction and the uniform proof shape $\mathrm{ITP}_{\mathrm{U}}$, both the use of $<$-induction and that of the non-uniform proof shape ITP come at the price of introducing a universal quantifier to the goals. However, both also come with a significant pay-off: there are strictly more $<$-inductive formulas than $(n+1)$-inductive formulas for any $n \in \mathbb{N}$. Similarly, the uniform proof shape $\mathrm{ITP}_{\mathrm{U}}$ is not complete while the non-uniform proof shape ITP is complete (with respect to PA). Therefore, assuming the theorem proving environment is able to deal with universal quantifiers in the assumptions of a goal efficiently, these results clearly indicate to prefer $<$-induction over $(n+1)$-induction for any $n \in \mathbb{N}$ and the ITP proof shape over the $\mathrm{ITP}_{\mathrm{U}}$ proof shape.

As Theorem 3.1 shows, the non-uniform proof shape ITP can be restricted to the non-uniform equivalence proof shape $\mathrm{ITP}_{\mathrm{EQ}}$ while preserving completeness, thus strongly reducing the search space. Whether this can be exploited for practical applications in inductive theorem proving is unclear to the authors. What would be needed is a practically reasonable procedure for generating universal formulas which are PA^--equivalent to a given sentence. To the best of the authors' knowledge, no such procedure has been

devised for use in inductive theorem proving yet. On a more theoretical note, Theorem 3.1 shows that the essential difficulty of inductive theorem proving does *not* lie in finding an inductive formula which is *stronger* than the goal (in the sense that it implies the goal over PA^-); it is sufficient to find an inductive formula as strong as the goal. Instead, the essential difficulty lies in the non-analyticity of the inductive formula as discussed in Section 2.2.

The results of Section 8 illustrate the importance of the choice of the language for the inductive formulas: allowing a single predicate symbol defined using a single quantifier can increase the strength of the theory considerably.

From a broader perspective, we believe that this paper demonstrates the potential of the use of methods and results from mathematical logic, and in particular, from theories of arithmetic, in inductive theorem proving. In our opinion, a good illustration of this potential is provided by Proposition 7.2. The construction of theorem-specific induction rules has a long history in inductive theorem proving [2, 7, 28, 6]. This approach is typically justified empirically by demonstrating, in the context of a certain algorithm or implementation, that the option to switch to another induction rule allows to solve more goals than before (by that particular algorithm or by a certain implementation within a certain timeout). Proposition 7.2 provides a much more solid foundation to this approach: if there are strictly more formulas which are inductive in sense X than in sense Y an algorithm for inductive theorem proving can get trapped in a situation where the induction formula currently under consideration allows a proof using rule X but does not allow a proof using rule Y, *regardless* of which algorithm or implementation is used, since there simply is no proof.

All of the results in this paper are restricted to arithmetical theories in the strict sense, i.e., theories about the natural numbers. In inductive theorem proving, one typically works in a broader syntactic setting of many-sorted first-order logic with inductive data types such as lists, trees, etc. Since many results about theories of arithmetic depend on results about numbers, rings, etc., it is not straightforward to extend them to this more general syntactic setting. Coding is not an option since it introduces an amount of syntactic complication that is not realistic in inductive theorem proving. However, for the continuation of the line of work in this paper we consider such an extension of the model theory of arithmetic to be of high importance. One prototypical problem along this line is to classify (some of) the benchmark examples from the TIP suite [13] according to the weakest induction schemes (e.g., atomic, quantifier-free) required to prove them.

We believe, as argued in Section 2, that the non-analyticity of induction axioms is a key aspect of inductive theorem proving, so we have considered the set of inductive formulas as the search space in our model. In the case of theorem proving in pure first-order logic, the search space is quite well understood theoretically and consequently we have very powerful tools at our

disposal for navigating in it such as (most general) unification, subsumption, reduction orderings, (resolution) refinements, etc. At present, we do not seem to have a similarly complete theoretical understanding of the set of inductive formulas and how they relate to a given goal. We believe that further work in this direction is crucial for advancing the state of the art in inductive theorem proving.

10 Acknowledgements

The authors would like to thank Zofia Adamowicz, Alan Bundy, Richard Kaye, and Leszek Kołodziejczyk for illuminating discussions around the topics of this paper. The authors would also like to thank the anonymous reviewers whose insightful feedback has led to a considerable improvement of this paper. The first author is supported by the Vienna Science Fund (WWTF) project no. VRG12-004. The second author was financially supported by Austrian Science Fund (FWF) project P24654-N25 while this research was carried out.

References

[1] Krzysztof R. Apt, Frank S. de Boer, and Ernst-Rüdiger Olderog. *Verification of Sequential and Concurrent Programs*. Texts in Computer Science. Springer-Verlag, Dordrecht, third, extended edition, 2009. Reprinted with corrections, 2010.

[2] Robert S. Boyer and J. Strother Moore. *A Computational Logic*. ACM monograph series. Academic Press, New York, 1979.

[3] Aaron R. Bradley and Zohar Manna. *The Calculus of Computation: Decision Procedures with Applications to Verification*. Springer-Verlag, Berlin, 2007.

[4] James Brotherston, Nikos Gorogiannis, and Rasmus L. Petersen. A generic cyclic theorem prover. In Ranjit Jhala and Atsushi Igarashi, editors, *Programming Languages and Systems*, volume 7705 of *Lecture Notes in Computer Science*, pages 350–367, Berlin, 2012. Springer-Verlag.

[5] James Brotherston and Alex Simpson. Sequent calculi for induction and infinite descent. *Journal of Logic and Computation*, 21(6):1177–1216, December 2011.

[6] Alan Bundy, David Basin, Dieter Hutter, and Andrew Ireland. *Rippling: Meta-Level Guidance for Mathematical Reasoning*, volume 56 of

Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 2005.

[7] Alan Bundy, Frank van Harmelen, Jane Hesketh, Alan Smaill, and Andrew Stevens. A rational reconstruction and extension of recursion analysis. In N. S. Sridharan, editor, *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, volume 1, pages 359–365, San Francisco, 1989. Morgan Kaufmann Publishers.

[8] Samuel R. Buss. *Bounded Arithmetic*, volume 3 of *Studies in Proof Theory, Lecture Notes*. Bibliopolis, Naples, 1986. Revision of the 1985 Ph.D. thesis.

[9] Samuel R. Buss. An Introduction to Proof Theory. In *Handbook of Proof Theory* [11], pages 2–78.

[10] Samuel R. Buss. First-Order Proof Theory of Arithmetic. In *Handbook of Proof Theory* [11], pages 79–147.

[11] Samuel R. Buss, editor. *Handbook of Proof Theory*, volume 137 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, Amsterdam, 1998.

[12] Koen Claessen, Moa Johansson, Dan Rosén, and Nicholas Smallbone. Automating inductive proofs using theory exploration. In Maria Paola Bonacina, editor, *Automated Deduction — CADE-24*, volume 7898 of *Lecture Notes in Computer Science*, pages 392–406, Berlin, 2013. Springer-Verlag.

[13] Koen Claessen, Moa Johansson, Dan Rosén, and Nicholas Smallbone. TIP: Tons of inductive problems. In Manfred Kerber, Jacques Carette, Cezary Kaliszyk, Florian Rabe, and Volker Sorge, editors, *Intelligent Computer Mathematics*, volume 9150 of *Lecture Notes in Computer Science*, pages 333–337, Cham, 2015. Springer-Verlag.

[14] Hubert Comon. Inductionless induction. In Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume I, chapter 14, pages 913–962. Elsevier, Amsterdam, 2001.

[15] Simon Cruanes. *Extending Superposition with Integer Arithmetic, Structural Induction, and Beyond*. PhD thesis, École Polytechnique, Palaiseau, France, 2015.

[16] Alastair F. Donaldson, Daniel Kroening, and Philipp Rümmer. Automatic analysis of DMA races using model checking and k-induction. *Formal Methods in System Design*, 39(1):83–113, 2011.

[17] Gabriel Ebner, Stefan Hetzl, Giselle Reis, Martin Riener, Simon Wolfsteiner, and Sebastian Zivota. System description: GAPT 2.0. In Nicola Olivetti and Ashish Tiwari, editors, *8th International Joint Conference on Automated Reasoning (IJCAR)*, volume 9706 of *Lecture Notes in Computer Science*, pages 293–301. Springer, 2016.

[18] Gerhard Gentzen. Zusammenfassung von mehreren vollständigen Induktionen zu einer einzigen. *Archiv für mathematische Logik und Grundlagenforschung*, 2(1):1–3, January 1954.

[19] Petr Hájek and Pavel Pudlák. *Metamathematics of First-Order Arithmetic*. Perspectives in Mathematical Logic. Springer-Verlag, Berlin, 1993.

[20] Emil Jeřábek. Sequence coding without induction. *Mathematical Logic Quarterly*, 58(3):244–248, May 2012.

[21] Joost J. Joosten. Turing jumps through provability. In Arnold Beckmann, Victor Mitrana, and Mariya Soskova, editors, *Evolving Computability*, volume 9136 of *Lecture Notes in Computer Science*, pages 216–225, Cham, 2015. Springer-Verlag.

[22] Richard Kaye. Diophantine induction. *Annals of Pure and Applied Logic*, 46(1):1–40, January 1990.

[23] Richard Kaye. *Models of Peano Arithmetic*, volume 15 of *Oxford Logic Guides*. Clarendon Press, Oxford, 1991.

[24] Abdelkader Kersani and Nicolas Peltier. Combining superposition and induction: A practical realization. In Pascal Fontaine, Christophe Ringeissen, and Renate A. Schmidt, editors, *Frontiers of Combining Systems*, volume 8152 of *Lecture Notes in Computer Science*, pages 7–22, Berlin, 2013. Springer-Verlag.

[25] Lawrence C. Paulson and Jasmin Christian Blanchette. Three years of experience with Sledgehammer, a Practical Link Between Automatic and Interactive Theorem Provers. In Geoff Sutcliffe, Stephan Schulz, and Eugenia Ternovska, editors, *The 8th International Workshop on the Implementation of Logics (IWIL), 2010*, volume 2 of *EPiC Series in Computing*, pages 1–11. EasyChair, 2012.

[26] Mary Sheeran, Satnam Singh, and Gunnar Stålmarck. Checking Safety Properties Using Induction and a SAT-Solver. In Warren A. Hunt Jr. and Steven D. Johnson, editors, *Third International Conference on Formal Methods in Computer-Aided Design (FMCAD)*, volume 1954 of *Lecture Notes in Computer Science*, pages 108–125. Springer, 2000.

[27] John C. Shepherdson. A non-standard model for a free variable fragment of number theory. *Bulletin de l'Académie Polonaise des Sciences. Série des Sciences Mathématiques, Astronomiques et Physiques*, XII(2):79–86, 1964.

[28] Andrew Stevens. A rational reconstruction of Boyer and Moore's technique for constructing induction formulas. In Yves Kodratoff, editor, *ECAI 88*, pages 565–570, London, 1988. Pitman Publishing.

[29] Geoff Sutcliffe. The CADE ATP System Competition - CASC. *AI Magazine*, 37(2):99–101, 2016.

[30] Albert Visser. Faith & falsity. *Annals of Pure and Applied Logic*, 131(1–3):103–131, January 2005.

[31] Albert Visser. Peano Corto and Peano Basso: A study of local induction in the context of weak theories. *Mathematical Logic Quarterly*, 60(1–2):92–117, February 2014.

[32] Christoph Walther. Computing induction axioms. In Andrei Voronkov, editor, *Logic Programming and Automated Reasoning*, volume 624 of *Lecture Notes in Computer Science*, pages 381–392, Berlin, 1992. Springer-Verlag.

[33] Daniel Wand and Christoph Weidenbach. Automatic induction inside superposition. Available at `http://people.mpi-inf.mpg.de/~dwand/datasup/d.pdf`.

[34] Alex J. Wilkie and Jeff B. Paris. On the scheme of induction for bounded arithmetic formulas. *Annals of Pure and Applied Logic*, 35:261–302, 1987.