Pre-logical Relations

Furio Honsell[1,2] and Donald Sannella[1]

[1] Laboratory for Foundations of Computer Science, University of Edinburgh
[2] Dipartimento di Matematica e Informatica, Università di Udine

Abstract. We study a weakening of the notion of logical relations, called *pre-logical relations*, that has many of the features that make logical relations so useful but having further algebraic properties including composability. The basic idea is simply to require the reverse implication in the definition of logical relations to hold only for lambda-expressible functions. Pre-logical relations are the minimal weakening of logical relations that gives composability for extensional structures and simultaneously the most liberal definition that gives the Basic Lemma. The use of pre-logical relations in place of logical relations gives an improved version of Mitchell's representation independence theorem which characterizes observational equivalence for all signatures rather than just for first-order signatures. Pre-logical relations can be used in place of logical relations to give an account of data refinement where the fact that pre-logical relations compose explains why stepwise refinement is sound.

Keywords: Logical relations, typed lambda calculus, semantics, representation independence, data refinement.

1 Introduction

Logical relations are structure-preserving relations between models of typed lambda calculus.

Definition 1.1. *Let $\mathcal{A}$ and $\mathcal{B}$ be Σ-applicative structures. A* logical relation $\mathcal{R}$ over $\mathcal{A}$ and $\mathcal{B}$ *is a family of relations $\{R^\sigma \subseteq [\![\sigma]\!]^\mathcal{A} \times [\![\sigma]\!]^\mathcal{B}\}_{\sigma \in Types(B)}$ such that:*

- $R^{\sigma\to\tau}(f,g)$ *iff* $\forall a \in [\![\sigma]\!]^\mathcal{A}.\forall b \in [\![\sigma]\!]^\mathcal{B}.R^\sigma(a,b) \Rightarrow R^\tau(App_\mathcal{A}\ f\ a, App_\mathcal{B}\ g\ b)$.
- $R^\sigma([\![c]\!]^\mathcal{A}, [\![c]\!]^\mathcal{B})$ *for every term constant $c : \sigma$ in Σ.*

Logical relations are used extensively in the study of typed lambda calculus and have applications outside lambda calculus, for example to abstract interpretation [Abr90] and data refinement [Ten94]. A good reference for logical relations is [Mit96]. An important but more difficult reference is [Sta85].

The Basic Lemma is the key to many of the applications of logical relations. It says that any logical relation over $\mathcal{A}$ and $\mathcal{B}$ relates the interpretation of each lambda term in $\mathcal{A}$ to its interpretation in $\mathcal{B}$.

Lemma 1.2 (Basic Lemma). *Let $\mathcal{R}$ be a logical relation over Henkin models $\mathcal{A}$ and $\mathcal{B}$. Then for all Γ-environments $\eta_\mathcal{A}, \eta_\mathcal{B}$ such that $R^\Gamma(\eta_\mathcal{A}, \eta_\mathcal{B})$ and every term $\Gamma \rhd M : \sigma$, $R^\sigma([\![\Gamma \rhd M : \sigma]\!]^\mathcal{A}_{\eta_\mathcal{A}}, [\![\Gamma \rhd M : \sigma]\!]^\mathcal{B}_{\eta_\mathcal{B}})$.* □

($R^\Gamma(\eta_\mathcal{A}, \eta_\mathcal{B})$ refers to the obvious extension of $\mathcal{R}$ to environments, see page 5 below.)

As structure-preserving relations, logical relations resemble familiar algebraic concepts like homomorphisms and congruence relations but they lack

some of the convenient properties of such concepts. In particular, the composition of two logical relations is not in general a logical relation. This calls into question their application to data refinement at least, where one would expect composition to provide an account of stepwise refinement.

We propose a weakening of the notion of logical relations called *pre-logical relations* (Sect. 3) that has many of the features that make logical relations so useful — in particular, the Basic Lemma still holds for pre-logical relations (Lemma 4.1) – but having further algebraic properties including composability (Prop. 5.5). The basic idea is simply to require the reverse implication in the definition of logical relations to hold only for lambda-expressible functions. Pre-logical relations turns out to be the minimal weakening of logical relations that gives composability for extensional structures (Corollary 6.2) and simultaneously the most liberal definition that gives the Basic Lemma. The use of pre-logical relations in place of logical relations gives an improved version of Mitchell's representation independence theorem (Corollaries 7.5 and 7.6 to Theorem 7.4) which characterizes observational equivalence for all signatures rather than just for first-order signatures. Pre-logical relations can be used in place of logical relations in Tennent's account of data refinement in [Ten94] and the fact that pre-logical relations compose explains why stepwise refinement is sound.

Many applications of logical relations follow a standard pattern where the result comes directly from the Basic Lemma once an appropriate logical relation has been defined. Some results in the literature follow similar lines in the sense that a type-indexed family of relations is defined by induction on types and a proof like that of the Basic Lemma is part of the construction, but the family of relations defined is not logical. Examples can be found in Plotkin's and Jung and Tiuryn's lambda-definability results using I-relations [Plo80] and Kripke logical relations with varying arity [JT93] respectively, and Gandy's proof of strong normalization using hereditarily strict monotonic functionals [Gan80]. In each of these cases, the family of relations involved turns out to be a pre-logical relation (Examples 3.8 and 3.9) which allows the common pattern to be extended to these cases as well. Since pre-logical relations are more general than logical relations and variants like I-relations, they provide a framework within which these different classes can be compared. Here we begin by studying and comparing their closure properties (Prop. 5.6) with special attention to closure under composition.

The definition of pre-logical relations is not new. In [Sch87], Schoett uses a first-order version of algebraic relations which he calls *correspondences*, and he conjectures (p. 281) that for Henkin models, what we have called pre-logical relations (formulated as in Prop. 3.3) would be closed under composition and yield the Basic Lemma. In [Mit90], Mitchell makes the same suggestion, referring to Schoett and also crediting Abramsky and Plotkin, but as an assertion rather than a conjecture. The idea is not developed any further. An independent but apparently equivalent definition of pre-logical relations over cartesian closed categories is given in [PPS98] where they are called *lax logical relations*. It is shown that these compose and that the Basic Lemma holds, and an axiomatic account is provided. Earlier, a closely related notion called *L-relations*

was defined in [KOPTT97] and shown to compose. There appears to be no previous work on pre-logical relations that goes beyond observing that they compose and that the Basic Lemma holds. Another difference to [PPS98] and [KOPTT97] is that our treatment is elementary rather than categorical, and covers also combinatory logics.

2 Syntax and Semantics

We begin with $\lambda^{\to}$, the simply-typed lambda calculus having $\to$ as the only type constructor. Other type constructors will be considered in Sect. 9. We follow the terminology in [Mit96] for the most part, with slightly different notation.

Definition 2.1. *The set $Types(B)$ of* types *over a set B of* base types *(or* type constants*) is given by the grammar $\sigma ::= b \mid \sigma \to \sigma$ where b ranges over B. A* signature Σ *consists of a set B of type constants and a collection C of typed* term constants $c : \sigma$.

Let $\Sigma = \langle B, C\rangle$ be a signature. We assume familiarity with the usual notions of *context* $\Gamma = x_1{:}\sigma_1, \ldots, x_n{:}\sigma_n$ and *Σ-term* M of type σ over a context Γ, written $\Gamma \rhd M : \sigma$, with the meta-variable t reserved for lambda-free Σ-terms. If Γ is empty then we write simply $M : \sigma$. Free variables $FV(M)$ and capture-avoiding substitution $[N/x]M$ are as usual.

Definition 2.2. *A Σ-*applicative structure $\mathcal{A}$ *consists of:*

- *a* carrier *set $[\![\sigma]\!]^{\mathcal{A}}$ for each $\sigma \in Types(B)$;*
- *a function $App_{\mathcal{A}}^{\sigma,\tau} : [\![\sigma \to \tau]\!]^{\mathcal{A}} \to [\![\sigma]\!]^{\mathcal{A}} \to [\![\tau]\!]^{\mathcal{A}}$ for each $\sigma, \tau \in Types(B)$;*
- *an element $[\![c]\!]^{\mathcal{A}} \in [\![\sigma]\!]^{\mathcal{A}}$ for each constant $c : \sigma$ in Σ.*

We drop the subscripts and superscripts when they are determined by the context, and we sometimes abbreviate $App_{\mathcal{A}}^{\sigma,\tau}\, f\, x$ as $f\, x$. Two elements $f, g \in [\![\sigma \to \tau]\!]^{\mathcal{A}}$ are said to be extensionally equal *if $App_{\mathcal{A}}^{\sigma,\tau}\ f\ x = App_{\mathcal{A}}^{\sigma,\tau}\ g\ x$ for every $x \in [\![\sigma]\!]^{\mathcal{A}}$. A Σ-applicative structure is* extensional *when extensional equality coincides with identity.*

*A Σ-*combinatory algebra *is a Σ-applicative structure $\mathcal{A}$ that has elements $K_{\mathcal{A}}^{\sigma,\tau} \in [\![\sigma \to (\tau \to \sigma)]\!]^{\mathcal{A}}$ and $S_{\mathcal{A}}^{\rho,\sigma,\tau} \in [\![(\rho \to \sigma \to \tau) \to (\rho \to \sigma) \to \rho \to \tau]\!]^{\mathcal{A}}$ for each $\rho, \sigma, \tau \in Types(B)$ satisfying the equations $K_{\mathcal{A}}^{\sigma,\tau}\, x\, y = x$ and $S_{\mathcal{A}}^{\rho,\sigma,\tau}\, x\, y\, z = (x\, z)(y\, z)$.*

An extensional combinatory algebra is called a Henkin model. *An applicative structure $\mathcal{A}$ is a* full type hierarchy *when $[\![\sigma \to \tau]\!]^{\mathcal{A}} = [\![\sigma]\!]^{\mathcal{A}} \to [\![\tau]\!]^{\mathcal{A}}$ (the full set-theoretic function space) for every $\sigma, \tau \in Types(B)$ and then it is obviously a Henkin model.*

In a combinatory algebra, we can extend the definition of lambda-free Σ-terms by allowing them to contain S and K; we call these *combinatory Σ-terms.*

A *Γ-environment* $\eta_{\mathcal{A}}$ assigns elements of an applicative structure $\mathcal{A}$ to variables in Γ, with $\eta_{\mathcal{A}}(x) \in [\![\sigma]\!]^{\mathcal{A}}$ for $x : \sigma$ in Γ. A lambda-free Σ-term $\Gamma \rhd t : \sigma$ is interpreted in a Σ-applicative structure $\mathcal{A}$ under a Γ-environment $\eta_{\mathcal{A}}$ in the obvious way, written $[\![\Gamma \rhd t : \sigma]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}$, and this extends immediately to an interpretation of combinatory Σ-terms in combinatory algebras by interpreting K and S as $K_{\mathcal{A}}$ and $S_{\mathcal{A}}$. If t is closed then we write simply $[\![t : \sigma]\!]^{\mathcal{A}}$.

There are various ways of interpreting terms containing lambda abstraction in a combinatory algebra by "compiling" them to combinatory terms so that outermost β holds (see Prop. 2.4 below for what we mean by "outermost β"). In Henkin models, all these compilations yield the same result.

An axiomatic approach to interpreting lambda abstraction requires an applicative structure equipped with an interpretation function that satisfies certain minimal requirements — cf. the notion of *acceptable meaning function* in [Mit96].

Definition 2.3. *A* lambda Σ-applicative structure *consists of a Σ-applicative structure $\mathcal{A}$ together with a function $[\![\cdot]\!]^{\mathcal{A}}$ that maps any term $\Gamma \rhd M : \sigma$ and Γ-environment $\eta_{\mathcal{A}}$ over $\mathcal{A}$ to an element of $[\![\sigma]\!]^{\mathcal{A}}$, such that:*

- $[\![\Gamma \rhd x : \sigma]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}} = \eta_{\mathcal{A}}(x)$
- $[\![\Gamma \rhd c : \sigma]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}} = [\![c]\!]^{\mathcal{A}}$
- $[\![\Gamma \rhd M\ N : \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}} = App_{\mathcal{A}}\,[\![\Gamma \rhd M : \sigma \to \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}\,[\![\Gamma \rhd N : \sigma]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}$
- $[\![\Gamma \rhd \lambda x{:}\sigma.M : \sigma \to \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}} = [\![\Gamma \rhd \lambda y{:}\sigma.[y/x]M : \sigma \to \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}$ *provided* $y \notin FV(M)$
- $[\![\Gamma \rhd M : \sigma]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}} = [\![\Gamma \rhd M : \sigma]\!]^{\mathcal{A}}_{\eta'_{\mathcal{A}}}$ *provided $\eta'_{\mathcal{A}}$ is a Γ-environment such that $\eta_{\mathcal{A}}(x) = \eta'_{\mathcal{A}}(x)$ for all $x \in FV(M)$*
- $[\![\Gamma, x{:}\sigma \rhd M : \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}} = [\![\Gamma \rhd M : \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}$ *for $x \notin \Gamma$*
- $[\![\Gamma, x{:}\sigma \rhd M : \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}[x \mapsto [\![\Gamma \rhd N:\sigma]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}]} = [\![\Gamma \rhd [N/x]M : \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}$

The relationship between lambda applicative structures and combinatory algebras is as follows.

Proposition 2.4. *A lambda applicative structure $\mathcal{A}$ such that $App_{\mathcal{A}}[\![\Gamma \rhd \lambda x{:}\sigma.M : \sigma \to \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}\ a = [\![\Gamma, x{:}\sigma \rhd M : \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}[x \mapsto a]}$ amounts to a combinatory algebra, and vice versa.*

Proof. $\Leftarrow$*: We define $[\![\cdot]\!]^{\mathcal{A}}$ via the standard compilation of lambda terms using K and S to combinatory terms.* $\Rightarrow$*: $K^{\sigma,\tau}_{\mathcal{A}}$ and $S^{\rho,\sigma,\tau}_{\mathcal{A}}$ are the interpretations of the usual lambda terms.* □

The proof of this proposition shows that the interpretation of lambda terms in combinatory algebras via compilation to combinatory terms satisfies the axioms in Def. 2.3 and the additional property in the proposition. Therefore when viewing a combinatory algebra as a lambda applicative structure, this is the interpretation function we have in mind.

3 Algebraic and Pre-logical Relations

We propose a weakening of the definition of logical relations which is closed under composition and which has most of the attractive properties of logical relations. First we change the two-way implication in the condition on functions to a one-way implication which requires preservation of the relation under application.

Definition 3.1. *Let $\mathcal{A}$ and $\mathcal{B}$ be Σ-applicative structures. An* algebraic relation $\mathcal{R}$ over $\mathcal{A}$ and $\mathcal{B}$ *is a family of relations $\{R^{\sigma} \subseteq [\![\sigma]\!]^{\mathcal{A}} \times [\![\sigma]\!]^{\mathcal{B}}\}_{\sigma \in Types(B)}$ such that:*

– *If* $R^{\sigma\to\tau}(f,g)$ *then* $\forall a \in [\![\sigma]\!]^{\mathcal{A}}.\forall b \in [\![\sigma]\!]^{\mathcal{B}}.R^{\sigma}(a,b) \Rightarrow R^{\tau}(App_{\mathcal{A}}\, f\, a, App_{\mathcal{B}}\, g\, b)$.
– $R^{\sigma}([\![c]\!]^{\mathcal{A}}, [\![c]\!]^{\mathcal{B}})$ *for every term constant* $c : \sigma$ *in* Σ.

In lambda applicative structures, we additionally require the relation to preserve lambda abstraction in a sense that is analogous to the definition of *admissible relation* in [Mit96]. First, we extend a family of relations $\mathcal{R} = \{R^{\sigma} \subseteq [\![\sigma]\!]^{\mathcal{A}} \times [\![\sigma]\!]^{\mathcal{B}}\}_{\sigma\in Types(B)}$ to a relation on Γ-environments: $R^{\Gamma}(\eta_{\mathcal{A}}, \eta_{\mathcal{B}})$ if $R^{\sigma}(\eta_{\mathcal{A}}(x), \eta_{\mathcal{B}}(x))$ for every $x{:}\sigma$ in Γ.

Definition 3.2. *Let* $\mathcal{A}$ *and* $\mathcal{B}$ *be lambda* Σ*-applicative structures. A* pre-logical relation over $\mathcal{A}$ and $\mathcal{B}$ *is an algebraic relation* $\mathcal{R}$ *such that given* Γ*-environments* $\eta_{\mathcal{A}}$ *and* $\eta_{\mathcal{B}}$ *such that* $R^{\Gamma}(\eta_{\mathcal{A}}, \eta_{\mathcal{B}})$*, and a term* $\Gamma, x : \sigma \rhd M : \tau$*, if* $R^{\sigma}(a,b)$ *implies* $R^{\tau}([\![\Gamma, x : \sigma \rhd M : \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}[x\mapsto a]}, [\![\Gamma, x : \sigma \rhd M : \tau]\!]^{\mathcal{B}}_{\eta_{\mathcal{B}}[x\mapsto b]})$ *for all* $a \in [\![\sigma]\!]^{\mathcal{A}}$ *and* $b \in [\![\sigma]\!]^{\mathcal{B}}$*, then* $R^{\sigma\to\tau}([\![\Gamma \rhd \lambda x{:}\sigma.M : \sigma \to \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}, [\![\Gamma \rhd \lambda x{:}\sigma.M : \sigma \to \tau]\!]^{\mathcal{B}}_{\eta_{\mathcal{A}}})$.

This amounts to defining pre-logical relations as simply the class of relations that make the Basic Lemma hold, as we shall see in Lemma 4.1 below. (Indeed, since the Basic Lemma for pre-logical relations is an equivalence rather than a one-way implication, an alternative at this point would be to take the conclusion of the Basic Lemma itself as the definition of pre-logical relations.)

A simpler and therefore more appealing definition is obtained if we consider combinatory algebras, where the requirement above boils down to preservation of S and K:

Proposition 3.3. *Let* $\mathcal{A}$ *and* $\mathcal{B}$ *be* Σ*-combinatory algebras. An algebraic relation* $\mathcal{R}$ *over* $\mathcal{A}$ *and* $\mathcal{B}$ *is pre-logical iff* $R(S^{\rho,\sigma,\tau}_{\mathcal{A}}, S^{\rho,\sigma,\tau}_{\mathcal{B}})$ *and* $R(K^{\sigma,\tau}_{\mathcal{A}}, K^{\sigma,\tau}_{\mathcal{B}})$ *for all* $\rho, \sigma, \tau \in Types(B)$.

Proof. Directly from the definitions. □

If we incorporate S and K into the signature Σ, then pre-logical relations are just algebraic relations on combinatory algebras. One way of understanding the definition of pre-logical relations is that the reverse implication in the definition of logical relations is required to hold only for lambda-expressible functions. For combinatory algebras these are exactly the functions that are denotable by combinatory terms, and thus this requirement is captured by requiring preservation of S and K.

The use of the combinators S and K in the above proposition is in some sense arbitrary: the same result would be achieved by taking any other combinatory basis and changing the definition of combinatory algebra and the interpretation function accordingly. It would be straightforward to modify the definitions to accommodate other variants of lambda calculus, for instance λ_I (where in $\lambda x{:}\sigma.M$, the term M is required to contain x) for which a combinatory basis is B, C, I, S, or linear lambda calculi. For languages that include recursion, such as PCF, one would add a Y combinator.

As usual, the binary case of algebraic resp. pre-logical relations over $\mathcal{A}, \mathcal{B}$ is derived from the unary case of *algebraic* resp. *pre-logical predicates* for the product structure $\mathcal{A}\times\mathcal{B}$. We omit the obvious definitions. The properties of pre-logical relations carry over to pre-logical predicates. For most results about pre-logical relations below there is a corresponding result about algebraic relations over applicative structures. We omit these for lack of space.

The fact that pre-logicality is strictly weaker than logicality is demonstrated by the following examples which also provide a number of general methods for defining pre-logical relations.

Example 3.4. Consider the signature Σ containing the type constant nat and term constants $0 : nat$ and $succ : nat \to nat$ and let $\mathcal{A}$ be the full type hierarchy over $\mathbb{N}$ where 0 and $succ$ have their usual interpretations. The predicate

$$P^\sigma(v) \Leftrightarrow v \text{ is the value of a closed } \Sigma\text{-term}$$

is a pre-logical predicate over $\mathcal{A}$. (This is easy to see since $\mathcal{A}$ is a combinatory algebra and so the analogue of Prop. 3.3 for pre-logical predicates applies.) It is not a logical predicate: any function $f \in [\![nat \to nat]\!]^{\mathcal{A}}$, including functions that are not lambda definable, takes values in P to values in P and so must itself be in P. □

Example 3.5. The identity relation on a lambda applicative structure is a pre-logical relation but it is logical iff the structure is extensional. □

Example 3.6. A *Σ-homomorphism* between lambda Σ-applicative structures $\mathcal{A}$ and $\mathcal{B}$ is a type-indexed family of functions $\{h^\sigma : [\![\sigma]\!]^{\mathcal{A}} \to [\![\sigma]\!]^{\mathcal{B}}\}_{\sigma \in Types(B)}$ such that for any constant $c : \sigma$ in Σ, $h^\sigma([\![c]\!]^{\mathcal{A}}) = [\![c]\!]^{\mathcal{B}}$, $h^\tau(App_{\mathcal{A}}^{\sigma,\tau}\ f\ a) = App_{\mathcal{B}}^{\sigma,\tau}\, h^{\sigma\to\tau}(f)\ h^\sigma(a)$ and $h^{\sigma\to\tau}([\![\Gamma \rhd \lambda x{:}\sigma.M : \sigma \to \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}) = [\![\Gamma \rhd \lambda x{:}\sigma.M : \sigma \to \tau]\!]^{\mathcal{B}}_{h(\eta_{\mathcal{A}})}$ where $h(\eta_{\mathcal{A}}) = \{x \mapsto h^\sigma(\eta_{\mathcal{A}}(x))\}$ for all $x{:}\sigma$ in Γ. Any Σ-homomorphism is a pre-logical relation. In particular, interpretation of terms in a lambda applicative structure with respect to an environment, viewed as a relation from the lambda applicative structure of terms, is a pre-logical relation but is not in general a logical relation. □

Example 3.7. Let $\mathcal{A}$ and $\mathcal{B}$ be lambda applicative structures and define $R^\sigma \subseteq [\![\sigma]\!]^{\mathcal{A}} \times [\![\sigma]\!]^{\mathcal{B}}$ by $R^\sigma(a, b)$ for $a \in [\![\sigma]\!]^{\mathcal{A}}, b \in [\![\sigma]\!]^{\mathcal{B}}$ iff there is a closed term $M : \sigma$ such that $[\![M : \sigma]\!]^{\mathcal{A}} = a$ and $[\![M : \sigma]\!]^{\mathcal{B}} = b$. This is a pre-logical relation but it is not in general a logical relation. Generalizing: the inverse of any pre-logical relation is obviously pre-logical and according to Prop. 5.5 below the composition of any two pre-logical relations is pre-logical. Then observe that the above relation is just the composition of closed term interpretation in $\mathcal{B}$ (which is pre-logical according to Example 3.6) and the inverse of closed term interpretation in $\mathcal{A}$. □

Example 3.8. Plotkin's *I-relations* [Plo80] give rise to pre-logical relations. The family of relations on the full type hierarchy consisting of the tuples which are in a given I-relation at a given world (alternatively, at all worlds) is a pre-logical relation which is not in general a logical relation. Similarly for Jung and Tiuryn's *Kripke logical relations with varying arity* [JT93]: the family of relations consisting of the w-tuples which are in a given Kripke logical relation with varying arity at world w is a pre-logical relation which is not in general a logical relation. Although both kinds of relations were originally defined over full type hierarchies, note that it makes perfect sense to consider them over arbitrary Henkin models. □

Example 3.9. Let $\mathcal{A}$ be an applicative structure. Given order relations R^b on $[\![b]\!]^{\mathcal{A}}$ for each base type b, we can define the *hereditarily monotonic functionals* as the equivalence classes of those elements of $\mathcal{A}$ which are self-related with respect to the following inductively defined family of relations on $\mathcal{A} \times \mathcal{A}$:

$$R^{\sigma\to\tau}(f,g) \text{ iff } \forall a \in [\![\sigma]\!]^{\mathcal{A}}.\forall b \in [\![\sigma]\!]^{\mathcal{B}}.R^{\sigma}(a,b) \Rightarrow (R^{\tau}(App_{\mathcal{A}}\ f\ a, App_{\mathcal{B}}\ g\ b) \wedge R^{\tau}(App_{\mathcal{A}}\ f\ a, App_{\mathcal{B}}\ f\ b) \wedge R^{\tau}(App_{\mathcal{A}}\ g\ a, App_{\mathcal{B}}\ g\ b))$$

(This defines simultaneously at each type both the class of functionals we are interested in and the order relation itself.) Notice that this method defines a pre-logical relation but not a logical relation.

Gandy's *hereditarily strict monotonic functionals* [Gan80] can be defined using the above technique with just a small modification of the clause for functionals.

$$R^{\sigma\to\tau}(f,g) \text{ iff } \forall a \in [\![\sigma]\!]^{\mathcal{A}}.\forall b \in [\![\sigma]\!]^{\mathcal{B}}. R^{\sigma}(a,b) \Rightarrow (f \neq g \Rightarrow (R^{\tau} \setminus \Delta^{\tau})(App_{\mathcal{A}}\ f\ a, App_{\mathcal{B}}\ g\ a) \wedge (a \neq b \Rightarrow ((R^{\tau} \setminus \Delta^{\tau})(App_{\mathcal{A}}\ f\ a, App_{\mathcal{B}}\ f\ b) \wedge (R^{\tau} \setminus \Delta^{\tau})(App_{\mathcal{A}}\ g\ a, App_{\mathcal{B}}\ g\ b))))$$

Again we have only a pre-logical relation (with respect to the language of λ_I) and not a logical relation. We can define the continuous functionals, as used in models of PCF, using a similar pattern. □

4 The Basic Lemma

We will now consider the extension of the Basic Lemma to pre-logical relations. In contrast to Lemma 1.2, we get a two-way implication which says that the requirements on pre-logical relations are exactly strong enough to ensure that the Basic Lemma holds. The reverse implication fails for logical relations as Example 3.4 shows (for logical predicates).

Lemma 4.1 (Basic Lemma for pre-logical relations). *Let $\mathcal{R} = \{R^{\sigma} \subseteq [\![\sigma]\!]^{\mathcal{A}} \times [\![\sigma]\!]^{\mathcal{B}}\}_{\sigma \in Types(B)}$ be a family of relations over lambda Σ-applicative structures $\mathcal{A}$ and $\mathcal{B}$. Then $\mathcal{R}$ is a pre-logical relation iff for all Γ-environments $\eta_{\mathcal{A}}, \eta_{\mathcal{B}}$ such that $R^{\Gamma}(\eta_{\mathcal{A}}, \eta_{\mathcal{B}})$ and every Σ-term $\Gamma \rhd M : \sigma$, $R^{\sigma}([\![\Gamma \rhd M : \sigma]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}, [\![\Gamma \rhd M : \sigma]\!]^{\mathcal{B}}_{\eta_{\mathcal{B}}})$.*

Proof. $\Rightarrow$*: The proof is by induction on the structure of M. For variables, we use the assumption that $R^{\Gamma}(\eta_{\mathcal{A}}, \eta_{\mathcal{B}})$. For constants, we use the fact that pre-logical relations are required to respect constants. For an application $\Gamma \rhd M\ N : \tau$, we use the inductive hypothesis for M and N and the fact that pre-logical relations are closed under application. As for $\Gamma \rhd \lambda x{:}\sigma.M : \sigma \to \tau$, by the induction hypothesis we know that for all $\Gamma, x{:}\sigma$-environments η'_A, η'_B such that $R^{\Gamma, x:\sigma}(\eta'_{\mathcal{A}}, \eta'_{\mathcal{B}})$ we have $R^{\sigma}([\![\Gamma, x{:}\sigma \rhd M : \sigma]\!]^{\mathcal{A}}_{\eta'_{\mathcal{A}}}, [\![\Gamma, x{:}\sigma \rhd M : \sigma]\!]^{\mathcal{B}}_{\eta'_{\mathcal{B}}})$. If $R^{\sigma}(a,b)$ then applying this to $\eta'_{\mathcal{A}} = \eta_{\mathcal{A}}[x \mapsto a]$ and $\eta'_{\mathcal{B}} = \eta_{\mathcal{B}}[x \mapsto b]$ and taking the condition in the definition of pre-logical relations gives the desired result.*

$\Leftarrow$*: The first condition of algebraic relations follows by taking M to be $x\ y$ and $\eta_{\mathcal{A}} = \{x \mapsto f, y \mapsto a\}$ $\eta_{\mathcal{B}} = \{x \mapsto g, y \mapsto b\}$ and the second condition*

for a term constant c follows by taking M to be c. The additional condition for pre-logical relations holds a fortiori. □

The "only if" direction of this result is the analogue in our setting of the general version of the Basic Lemma in [Mit96], but where $\mathcal{R}$ is only required to be pre-logical.

If one applies the Basic Lemma for pre-logical relations to Henkin models, the "only if" part of the result is exactly the usual formulation (Lemma 1.2 above), except that $\mathcal{R}$ is only required to be pre-logical.

Corollary 4.2. *Let* $\mathcal{R} = \{R^\sigma \subseteq [\![\sigma]\!]^{\mathcal{A}} \times [\![\sigma]\!]^{\mathcal{B}}\}_{\sigma \in Types(B)}$ *be a family of relations over Henkin models* $\mathcal{A}$ *and* $\mathcal{B}$*. Then* $\mathcal{R}$ *is a pre-logical relation iff for all* Γ*-environments* $\eta_{\mathcal{A}}, \eta_{\mathcal{B}}$ *such that* $R^\Gamma(\eta_{\mathcal{A}}, \eta_{\mathcal{B}})$ *and every* Σ*-term* $\Gamma \rhd M : \sigma$*,* $R^\sigma([\![\Gamma \rhd M : \sigma]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}, [\![\Gamma \rhd M : \sigma]\!]^{\mathcal{B}}_{\eta_{\mathcal{B}}})$*.* □

5 Properties of Pre-logical Relations

A logical relation on lambda applicative structures is pre-logical provided it is admissible in the following sense.

Definition 5.1 ([Mit96]). *A logical relation* $\mathcal{R}$ *on lambda applicative structures* $\mathcal{A}$ *and* $\mathcal{B}$ *is* admissible *if given* Γ*-environments* $\eta_{\mathcal{A}}$ *and* $\eta_{\mathcal{B}}$ *such that* $R^\Gamma(\eta_{\mathcal{A}}, \eta_{\mathcal{B}})$*, and terms* $\Gamma, x{:}\sigma \rhd M, N : \tau$*,*

$$\forall a \in [\![\sigma]\!]^{\mathcal{A}}, b \in [\![\sigma]\!]^{\mathcal{B}}. R^\sigma(a,b) \supset R^\tau([\![\Gamma, x{:}\sigma \rhd M : \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}[x \mapsto a]}, [\![\Gamma, x{:}\sigma \rhd N : \tau]\!]^{\mathcal{B}}_{\eta_{\mathcal{B}}[x \mapsto b]})$$

implies

$$\forall a \in [\![\sigma]\!]^{\mathcal{A}}, b \in [\![\sigma]\!]^{\mathcal{B}}. R^\sigma(a,b) \supset R^\tau(App_{\mathcal{A}}\, [\![\Gamma \rhd \lambda x{:}\sigma.M : \sigma \to \tau]\!]^{\mathcal{A}}_{\eta_{\mathcal{A}}}\, a,\; App_{\mathcal{B}}\, [\![\Gamma \rhd \lambda x{:}\sigma.N : \sigma \to \tau]\!]^{\mathcal{B}}_{\eta_{\mathcal{B}}}\, b)$$

Proposition 5.2. *Any admissible logical relation on lambda applicative structures is a pre-logical relation.*

Proof. Admissibility plus the reverse implication in the definition of logical relations gives the property in the definition of pre-logical relations. □

Corollary 5.3. *Any logical relation on combinatory algebras is a pre-logical relation.*

Proof. Logical relations on combinatory algebras are always admissible. □

To understand why the composition of logical relations $\mathcal{R}$ over $\mathcal{A}$ and $\mathcal{B}$ and $\mathcal{S}$ over $\mathcal{B}$ and $\mathcal{C}$ might not be a logical relation, it is instructive to look at examples. When composition fails, the problem is often that the interpretation of some function type in $\mathcal{B}$ has "too few values". But even if we take logical relations over full type hierarchies, where all possible values of function types are present, composition can fail because the required "missing link" in $\mathcal{B}$ is not a function:

Example 5.4. Let Σ contain just two type constants, b and b'. Consider three full type hierarchies $\mathcal{A}, \mathcal{B}, \mathcal{C}$ which interpret b and b' as follows: $[\![b]\!]^{\mathcal{A}} = \{*\} = [\![b']\!]^{\mathcal{A}}$; $[\![b]\!]^{\mathcal{B}} = \{*\}$ and $[\![b']\!]^{\mathcal{B}} = \{\circ, \bullet\}$; $[\![b]\!]^{\mathcal{C}} = \{\circ, \bullet\} = [\![b']\!]^{\mathcal{C}}$. Let $\mathcal{R}$ be the logical relation over $\mathcal{A}$ and $\mathcal{B}$ induced by $R^b = \{\langle *, *\rangle\}$ and $R^{b'} = \{\langle *, \circ\rangle, \langle *, \bullet\rangle\}$ and let $\mathcal{S}$ be the logical relation over $\mathcal{B}$ and $\mathcal{C}$ induced by $S^b = \{\langle *, \circ\rangle, \langle *, \bullet\rangle\}$ and $S^{b'} = \{\langle \circ, \circ\rangle, \langle \bullet, \bullet\rangle\}$. $\mathcal{S} \circ \mathcal{R}$ is not a logical relation because it does not relate the identity function in $[\![b]\!]^{\mathcal{A}} \to [\![b']\!]^{\mathcal{A}}$ to the identity function in $[\![b]\!]^{\mathcal{C}} \to [\![b']\!]^{\mathcal{C}}$. The problem is that the only two functions in $[\![b]\!]^{\mathcal{B}} \to [\![b']\!]^{\mathcal{B}}$ are $\{* \mapsto \circ\}$ and $\{* \mapsto \bullet\}$, and $\mathcal{S}$ does not relate these to the identity in $\mathcal{C}$. □

Proposition 5.5. *The composition $\mathcal{S} \circ \mathcal{R}$ of pre-logical relations $\mathcal{R}$ over $\mathcal{A}, \mathcal{B}$ and $\mathcal{S}$ over $\mathcal{B}, \mathcal{C}$ is a pre-logical relation over $\mathcal{A}, \mathcal{C}$.*

Proof. A proof from the definition is not at all straightforward, but Lemma 4.1 says that pre-logicality is equivalent to a property of relations that is obviously closed under composition. □

Composition of relations is definable in terms of product, intersection and projection:

$$\mathcal{S} \circ \mathcal{R} = \pi_{1,3}(\mathcal{A} \times \mathcal{S} \ \cap \ \mathcal{R} \times \mathcal{C})$$

Closure of pre-logical relations under these operations is a more basic property than closure under composition, and is not specific to binary relations. We have:

Proposition 5.6. *Pre-logical relations are closed under intersection, product, projection, restriction to a substructure, permutation and $\forall$. (Here, if $\mathcal{R} \subseteq \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ then $\forall\mathcal{R} \subseteq \mathcal{A}_2 \times \cdots \times \mathcal{A}_n$ is defined by $(\forall\mathcal{R})^\sigma = \{\langle a_2, \ldots, a_n\rangle \mid \forall a_1 \in [\![\sigma]\!]^{\mathcal{A}_1}. \langle a_1, a_2, \ldots, a_n\rangle \in \mathcal{R}^\sigma\}$.) Logical relations are closed under product, permutation and $\forall$ but not under intersection, projection or restriction to a substructure.* □

Other classes of relations satisfy different closure properties. For instance, I-relations are obviously closed under product, permutation and $\forall$, and it is easy to see that they are also closed under intersection. They are not closed under composition by Example 5.4, and since composition is definable in terms of product, intersection and projection it follows that they are not closed under projection. Failure of closure under restriction to a substructure follows from decidability and undecidability results in [Loa9?].

A consequence of closure under intersection is that given a property P of relations that is preserved under intersection, there is always a *least* pre-logical relation satisfying P. We then have (see Example 3.4 above):

Proposition 5.7. *The least pre-logical predicate over a given lambda Σ-applicative structure contains exactly those elements that are the values of closed Σ-terms.*

In a signature with no term constants, a logical relation may be constructed by defining a relation R on base types and using the definition to *"lift" R* inductively to higher types. The situation is different for pre-logical relations: there are in general many pre-logical liftings of a given R, one being of course its lifting to a logical relation (provided this gives an admissible relation). But since the property of lifting a given R is preserved under intersection, the least

pre-logical lifting of R is also a well-defined relation. Similarly for the least pre-logical *extension* of a given family of relations, for any signature. Notice that lifting/extending a given family of relations to a logical relation is problematic for signatures containing higher-order constants. Further ways of defining pre-logical relations are indicated by the examples at the end of Sect. 3.

It is easy to see that pre-logical relations are not closed under union. And even in a signature with no term constants, the class of pre-logical relations that lift a given relation R on base types cannot be endowed with a lattice structure in general. But the only logical relation in this class is one of its maximal elements under inclusion.

6 Pre-logical Relations via Composition of Logical Relations

Our weakening of the definition of logical relations may appear to be *ad hoc*, but for extensional structures it turns out to be the minimal weakening that is closed under composition. There are variants of this result for several different classes of models. We give the version for Henkin models.

Proposition 6.1. *Let $\mathcal{A}$ and $\mathcal{B}$ be Henkin models and let $\mathcal{R}$ be a pre-logical relation over $\mathcal{A}$ and $\mathcal{B}$. Then $\mathcal{R}$ factors into a composition of three logical relations over Henkin models.*

Proof sketch. We construct a Henkin model $\mathcal{A}[X]$ from $\mathcal{A}$ by taking $[\![\sigma]\!]^{\mathcal{A}[X]}$ to be the set of closed terms of type σ over a signature obtained by adding a constant for each element of $\mathcal{A}$ and a set of indeterminates of each type to Σ, quotiented by the congruence induced by the equality in $\mathcal{A}$. Take normal forms under evaluation in $\mathcal{A}$ up to $\beta\eta$-congruence as representatives of these congruence classes. We construct $\mathcal{B}[X]$ from $\mathcal{B}$ in the same way. We then define a logical relation $\mathcal{R}[X]$ on $\mathcal{A}[X]$ and $\mathcal{B}[X]$ by relating normal forms iff they are the same modulo $\mathcal{R}$. To see that $\mathcal{R}[X]$ is a logical relation, suppose that we have $f \in [\![\sigma \to \tau]\!]^{\mathcal{A}[X]}$ and $g \in [\![\sigma \to \tau]\!]^{\mathcal{B}[X]}$ such that whenever $R[X]^\sigma(a, b)$, we have $R[X]^\tau(App_{\mathcal{A}[X]}\ f\ a, App_{\mathcal{B}[X]}\ g\ b)$. Let $a = b = x$, an indeterminate. We know that $f\ x$ and $g\ x$ are the same normal form modulo $\mathcal{R}$, and hence $R[X]^{\sigma\to\tau}(f, g)$. It is easy to see that lifting the embedding on base types between $\mathcal{A}$ and $\mathcal{A}[X]$ to a logical relation is the embedding, and likewise for $\mathcal{B}$. Finally, we have to show that if $a \in [\![\sigma]\!]^{\mathcal{A}}$ and $b \in [\![\sigma]\!]^{\mathcal{B}}$ then $R^\sigma(a, b)$ iff $R[X]^\sigma([a], [b])$. This demonstrates that $\mathcal{R}$ is the composition of the embedding of $\mathcal{A}$ in $\mathcal{A}[X]$, $\mathcal{R}[X]$, and the inverse of the embedding of $\mathcal{B}$ in $\mathcal{B}[X]$. □

Corollary 6.2. *The class of pre-logical relations on Henkin models is the closure under composition of the class of logical relations on such structures.* □

This gives the following lambda-definability result:

Corollary 6.3. *Let $\mathcal{A}$ be a Henkin model and $a \in [\![\sigma]\!]^{\mathcal{A}}$. Then $\langle a, a\rangle$ belongs to all relations over $\mathcal{A}\times\mathcal{A}$ obtained by composing logical relations iff $a = [\![M : \sigma]\!]^{\mathcal{A}}$ for some closed Σ-term $M : \sigma$.*

Proof. By Corollary 6.2 and a binary version of Prop. 5.7. □

Corollary 6.2 does not hold if we restrict ourselves to considering just finite full type hierarchies: given an element a of a finite structure, it turns out to be co-r.e. if the pair $\langle a, a\rangle$ belongs to all binary relations which are obtainable by closing logical relations under intersection and projection (and hence by closure under composition), while by Prop. 5.7 and [Loa9?] it is not co-r.e. if $\langle a, a\rangle$ belongs to all binary pre-logical relations. In the case of arbitrary full type hierarchies, the question is open: the proof of Prop. 6.1 fails if we take a full type hierarchy in place of $\mathcal{A}[X]$, and we conjecture that Corollary 6.2 does not hold.

For non-extensional structures the notion of pre-logical relations is not the minimal weakening that gives closure under composition. The following variant is the minimal weakening for this case.

Definition 6.4. *An algebraic relation is* extensional *if whenever $R^{\sigma\to\tau}(f, g)$, f is extensionally equal to f' and g is extensionally equal to g', we have $R^{\sigma\to\tau}(f', g')$.*

All pre-logical relations over extensional structures are automatically extensional, and all logical relations over applicative structures (even non-extensional ones) are automatically extensional as well.

Proposition 6.5. *Let $\mathcal{A}$ and $\mathcal{B}$ be combinatory algebras and let $\mathcal{R}$ be an extensional pre-logical relation over $\mathcal{A}$ and $\mathcal{B}$. Then $\mathcal{R}$ factors into a composition of three logical relations.*

Proof. *As for Prop. 6.1 except that we need to take the extensional collapse over $\mathcal{A}$ and $\mathcal{B}$ respectively in the construction of $\mathcal{A}[X]$ and $\mathcal{B}[X]$. The fact that $\mathcal{R}$ is extensional is needed to show that the embeddings are logical relations.* □

Corollary 6.6. *The class of extensional pre-logical relations on combinatory algebras is the closure under composition of the class of logical relations on such structures.* □

These results may suggest that our definition of pre-logical relations on non-extensional structures should be strengthened by requiring the relation to be extensional, but this would make the reverse implication of the Basic Lemma fail. So although the notion of extensional pre-logical relations is the minimal weakening that gives closure under composition, these are stronger than necessary to give the Basic Lemma.

At the end of Sect. 5, we argued that intersection, projection etc. are more basic than composition of relations. Here are some counterparts to the above results in those terms. First, we can regard a pre-logical predicate $\mathcal{A}$ over a lambda applicative structure $\mathcal{B}$ as a *substructure* of $\mathcal{B}$, written $\mathcal{A} \subseteq \mathcal{B}$. Then, by analogy, we can call a logical predicate $\mathcal{A}$ over $\mathcal{B}$ a *logical substructure* of $\mathcal{B}$, written $\mathcal{A} \preceq \mathcal{B}$.

Proposition 6.7. *Suppose that $\mathcal{A} \subseteq \mathcal{B}$. Then:*

1. *If $\mathcal{A}$ is non-empty then there exist $\mathcal{C}, \mathcal{A}'$ such that $\mathcal{A}' \preceq \mathcal{B} \times \mathcal{C}$ and $\mathcal{A} = \pi_1(\mathcal{A}')$.*
2. *There exist $\mathcal{A}', \mathcal{B}'$ such that $\mathcal{A} \preceq \mathcal{A}' \preceq \mathcal{B}' \succeq \mathcal{B}$ and $\mathcal{A} = \mathcal{A}' \cap \mathcal{B}$.*

Proof. *Let $\mathcal{A} \subseteq \mathcal{B}$.*

1. *Take* $\mathcal{C} = \mathcal{A}[X]$*, where* $\mathcal{A}[X]$ *is as in the proof of Prop. 6.1, and* $\mathcal{A}' = \mathcal{A} \times \mathcal{A}[X]$*. Since* $\mathcal{A}$ *is non-empty, we have* $\mathcal{A} \times \mathcal{A}[X] \preceq \mathcal{B} \times \mathcal{A}[X]$ *and trivially* $\mathcal{A} = \pi_1(\mathcal{A} \times \mathcal{A}[X])$*.*
2. *Take* $\mathcal{A}' = \mathcal{A}[X]$ *and* $\mathcal{B}' = \mathcal{B}[X]$*.* □

If $\mathcal{A} \subseteq \mathcal{B}$ then there exists $\mathcal{B}'$ such that $\mathcal{A} \preceq_I \mathcal{B}' \succeq \mathcal{B}$, where $\mathcal{A} \preceq_I \mathcal{B}$ means that $\mathcal{A}$ is an I-predicate over $\mathcal{B}$.

It would be interesting to continue the above investigations to understand more fully the expressive power of classes of relations generated by logical relations under various operations, particularly on full type hierarchies.

7 Representation Independence and Data Refinement

Logical relations have been applied to explain the fact that the behaviour of programs does not depend on the way that data types are represented, but only on what can be observed about them using the operations that are provided. "Behaviour of programs" is captured by the notion of observational equivalence.

Definition 7.1. *Let* $\mathcal{A}$ *and* $\mathcal{B}$ *be lambda* Σ*-applicative structures and let OBS, the* observable types*, be a subset of Types*(B)*. Then* $\mathcal{A}$ *is* observationally finer *than* $\mathcal{B}$ *with respect to OBS, written* $\mathcal{A} \leq_{OBS} \mathcal{B}$*, if for any two closed terms* $M, N : \sigma$ *for* $\sigma \in OBS$ *such that* $[\![M : \sigma]\!]^{\mathcal{A}} = [\![N : \sigma]\!]^{\mathcal{A}}$ *we have* $[\![M : \sigma]\!]^{\mathcal{B}} = [\![N : \sigma]\!]^{\mathcal{B}}$*.*

$\mathcal{A}$ *and* $\mathcal{B}$ *are* observationally equivalent *with respect to OBS, written* $\mathcal{A} \equiv_{OBS} \mathcal{B}$*, if* $\mathcal{A} \leq_{OBS} \mathcal{B}$ *and* $\mathcal{B} \leq_{OBS} \mathcal{A}$*.*

It is usual to take *OBS* to be the "built-in" types for which equality is decidable, for instance *bool* and/or *nat*. Then $\mathcal{A}$ and $\mathcal{B}$ are observationally equivalent iff it is not possible to distinguish between them by performing computational experiments. Mitchell gives the following representation independence result:

Theorem 7.2 ([Mit96]). *Let* Σ *be a signature that includes a type constant nat, and let* $\mathcal{A}$ *and* $\mathcal{B}$ *be Henkin models, with* $[\![nat]\!]^{\mathcal{A}} = [\![nat]\!]^{\mathcal{B}} = \mathbb{N}$*. If there is a logical relation* $\mathcal{R}$ *over* $\mathcal{A}$ *and* $\mathcal{B}$ *with* R^{nat} *the identity relation on natural numbers, then* $\mathcal{A} \equiv_{\{nat\}} \mathcal{B}$*. Conversely, if* $\mathcal{A} \equiv_{\{nat\}} \mathcal{B}$*,* Σ *provides a closed term for each element of* $\mathbb{N}$*, and* Σ *only contains first-order functions, then there is a logical relation* $\mathcal{R}$ *over* $\mathcal{A}$ *and* $\mathcal{B}$ *with* R^{nat} *the identity relation.* □

The following example (Exercise 8.5.6 in [Mit96]) shows that the requirement that Σ contains only first-order functions is necessary.

Example 7.3. Let Σ have type constant *nat* and term constants $0, 1, 2, \ldots : nat$ and $f : (nat \to nat) \to nat$. Let $\mathcal{A}$ be the full type hierarchy over $[\![nat]\!]^{\mathcal{A}} = \mathbb{N}$ with $0, 1, 2, \ldots$ interpreted as usual and $[\![f]\!]^{\mathcal{A}}(g) = 0$ for all $g : \mathbb{N} \to \mathbb{N}$. Let $\mathcal{B}$ be like $\mathcal{A}$ but with $[\![f]\!]^{\mathcal{B}}(g) = 0$ if g is computable and $[\![f]\!]^{\mathcal{B}}(g) = 1$ otherwise. Since the difference between $\mathcal{A}$ and $\mathcal{B}$ cannot be detected by evaluating terms, $\mathcal{A} \equiv_{\{nat\}} \mathcal{B}$. But there is no logical relation over $\mathcal{A}$ and $\mathcal{B}$ which is the identity relation on *nat*: if $\mathcal{R}$ is logical then $R^{nat \to nat}(g, g)$ for any $g : \mathbb{N} \to \mathbb{N}$, and then $R^{nat}(App_{\mathcal{A}}\, [\![f]\!]^{\mathcal{A}}\, g, App_{\mathcal{B}}\, [\![f]\!]^{\mathcal{B}}\, g)$, which gives a contradiction if g is non-computable. □

We will strengthen this result by showing that pre-logical relations characterize observational equivalence for *all* signatures. We also generalize to arbitrary sets of observable types but this is much less significant. This characterization is obtained as a corollary of the following theorem which is a strengthening of Lemma 8.2.17 in [Mit96], again made possible by using pre-logical relations in place of logical relations.

Theorem 7.4. *Let $\mathcal{A}$ and $\mathcal{B}$ be lambda Σ-applicative structures and let $OBS \subseteq Types(B)$. Then $\mathcal{A} \leq_{OBS} \mathcal{B}$ iff there exists a pre-logical relation over $\mathcal{A}$ and $\mathcal{B}$ which is a partial function on OBS.*

Proof. $\Leftarrow$*: Suppose that $\mathcal{R}$ is a pre-logical relation over $\mathcal{A}$ and $\mathcal{B}$ which is a partial function on OBS and let $[\![M : \sigma]\!]^{\mathcal{A}} = [\![N : \sigma]\!]^{\mathcal{A}}$ for $\sigma \in OBS$. Apply the Basic Lemma to both sides and use the fact that $\mathcal{R}^{\sigma}$ is a partial function to get $[\![M : \sigma]\!]^{\mathcal{B}} = [\![N : \sigma]\!]^{\mathcal{B}}$.*

$\Rightarrow$*: Take the relation defined in Example 3.7.* □

(Mitchell's Lemma 8.2.17 is the "if" direction for Henkin models where $OBS = Types(B)$ but $\mathcal{R}$ is required to be logical rather than just pre-logical.)

Corollary 7.5. *Let $\mathcal{A}$ and $\mathcal{B}$ be lambda Σ-applicative structures and let $OBS \subseteq Types(B)$. Then $\mathcal{A} \equiv_{OBS} \mathcal{B}$ iff there exists a pre-logical relation over $\mathcal{A}$ and $\mathcal{B}$ which is a partial function on OBS in both directions.* □

Corollary 7.6. *Let Σ be a signature that includes a type constant nat and let $\mathcal{A}$ and $\mathcal{B}$ be lambda Σ-applicative structures with $[\![nat]\!]^{\mathcal{A}} = [\![nat]\!]^{\mathcal{B}} = \mathbb{N}$ such that Σ provides a closed term for each element of $\mathbb{N}$. There is a pre-logical relation $\mathcal{R}$ over $\mathcal{A}$ and $\mathcal{B}$ with R^{nat} the identity relation on natural numbers iff $\mathcal{A} \equiv_{\{nat\}} \mathcal{B}$.* □

Example 7.7. Revisiting Example 7.3, the pre-logical relation constructed in Example 3.7 has the required property, and it does not relate non-computable functions since they are not lambda definable. □

In accounts of data refinement in terms of logical relations such as Sect. 2 of [Ten94], the fact that logical relations do not compose conflicts with the experience that data refinements *do* compose in real life. Example 5.4 can be embellished to give refinements between data structures like lists and sets for which the logical relations underlying the refinement steps do not compose to give a logical relation, yet the data refinements involved do compose at an intuitive level. This failure to justify the soundness of *stepwise* refinement is a serious flaw. If pre-logical relations are used in place of logical relations, then the fact that the composition of pre-logical relations is again a pre-logical relation (Prop. 5.5) explains why stepwise refinement is sound. This opens the way to further development of the foundations of data refinement along the lines of [ST88], but we leave this to a separate future paper, see Sect. 10.

8 Other Applications

There are many other applications of logical relations. Take for instance the proof of strong normalization of $\lambda^{\rightarrow}$ in [Mit96]: one defines an admissible logical

predicate on a lambda applicative structure of terms by lifting the predicate on base types consisting of the strongly normalizing terms to higher types, proves that the predicate implies strong normalization, and then applies the general version of the Basic Lemma to give the result. The pattern for proofs of confluence, completeness of leftmost reduction, etc., is the same, sometimes with logical relations in place of logical predicates. There are also constructions that do not involve the Basic Lemma because the relations defined are not logical relations, but that include proofs following the same lines as the proof of the Basic Lemma. Examples include Gandy's proof that the hereditarily strict monotonic functionals model λ_I terms [Gan80], Plotkin's proof that lambda terms satisfy any I-relation [Plo80], and Jung and Tiuryn's proof that lambda terms satisfy any Kripke logical relation with varying arity at each arity (Theorem 3 of [JT93]).

All of these can be cast into a common mould by using pre-logical relations rather than logical relations. If a relation or predicate on a lambda applicative structure is logical and admissible, then it is pre-logical, and then the Basic Lemma for pre-logical relations gives the result. Plotkin's, Jung and Tiuryn's, and Gandy's relations can be shown to be pre-logical (in Gandy's case with respect to λ_I), see Examples 3.8 and 3.9 respectively, and then the application of the Basic Lemma for pre-logical relations gives the result in these cases as well. In each case, however, the interesting part of the proof is not the application of the Basic Lemma (or the argument that replaces its application in the case of Gandy, Plotkin, and Jung and Tiuryn) but rather the construction of the relation and the proof of its properties. The point of the analysis is not to say that this view makes the job easier but rather to bring forward the common pattern in all of these proofs, which is suggestive of a possible methodology for such proofs.

Definition 8.1. *A family of binary relations* $\{R^\sigma \subseteq [\![\sigma]\!]^{\mathcal{A}} \times [\![\sigma]\!]^{\mathcal{A}}\}_{\sigma \in Types(B)}$ *over a* Σ*-applicative structure* $\mathcal{A}$ *is a* partial equivalence relation *(abbreviated* PER*) if it is symmetric and transitive for each type.*

Proposition 8.2. *Let* $\mathcal{R}$ *be a PER on a* Σ*-applicative structure* $\mathcal{A}$ *which is algebraic. Define the* quotient *of* $\mathcal{A}$ *by* $\mathcal{R}$*, written* $\mathcal{A}/\mathcal{R}$*, as follows:*

- $[\![\sigma]\!]^{\mathcal{A}/\mathcal{R}} = [\![\sigma]\!]^{\mathcal{A}}/R^\sigma$*, i.e. the set of* $\mathcal{R}$*-equivalence classes of objects* $a \in [\![\sigma]\!]^{\mathcal{A}}$ *such that* $R^\sigma(a, a)$*.*
- $App^{\sigma,\tau}_{\mathcal{A}/\mathcal{R}}\,[f]_{\mathcal{A}/\mathcal{R}}\,[a]_{\mathcal{A}/\mathcal{R}} = [App^{\sigma,\tau}_{\mathcal{A}}\, f a]_{\mathcal{A}/\mathcal{R}}$
- $[\![c]\!]^{\mathcal{A}/\mathcal{R}} = [c]_{\mathcal{A}/\mathcal{R}}$ *for each constant* $c : \sigma$ *in* Σ*.*

Then:

1. *Let* $\mathcal{A}$ *be a lambda applicative structure. Then* $\mathcal{A}/\mathcal{R}$ *is a lambda applicative structure iff* $\mathcal{R}$ *is pre-logical.*
2. *Let* $\mathcal{A}$ *be a combinatory algebra. Then* $\mathcal{A}/\mathcal{R}$ *is a combinatory algebra iff* $\mathcal{R}$ *is pre-logical.*
3. $\mathcal{A}/\mathcal{R}$ *is an extensional applicative structure iff its restriction to the substructure of* $\mathcal{A}$ *consisting of the elements in* $Dom(\mathcal{R})$ *is a logical relation.* □

The last part of the above proposition says that one application of logical relations, that is their use in obtaining extensional structures by quotienting non-extensional structures i.e. the so-called *extensional collapse*, requires a relation that is logical (on a substructure) rather than merely pre-logical.

The above proposition allows us to prove completeness for different classes of structures using the traditional technique of quotienting an applicative structure of terms by a suitable relation defined by provability in a calculus. For non-extensional structures, this is not possible using logical relations because the relation defined by provability is pre-logical or algebraic rather than logical.

At this point one could develop a theory analogous to that of homomorphisms, quotients and substructures in universal algebra, but we refrain from doing this here. One would expect analogues of the usual theorems relating these three notions.

9 Beyond $\lambda^{\rightarrow}$ and Applicative Structures

Up to now we have been working in $\lambda^{\rightarrow}$, the simplest version of typed lambda calculus. We will now briefly indicate how other type constructors could be treated so as to obtain corresponding results for extended languages.

As a template, we shall discuss the case of product types. The syntax of types is extended by adding the type form $\sigma \times \tau$ and the syntax of terms is extended by adding pairing $\langle M, N\rangle$ and projections $\pi_1\ M$ and $\pi_2\ M$. If we regard these as additional term constants in the signature, e.g. $\langle\cdot,\cdot\rangle : \sigma \rightarrow \tau \rightarrow \sigma \times \tau$ for all σ, τ, rather than as new term forms, then the definition of pre-logical relations remains the same: the condition on constants says that e.g. $R^{\sigma\rightarrow\tau\rightarrow\sigma\times\tau}([\![\langle\cdot,\cdot\rangle]\!]^{\mathcal{A}}, [\![\langle\cdot,\cdot\rangle]\!]^{\mathcal{B}})$ and this is all that is required. For models that do not satisfy surjective pairing, this is weaker than the corresponding condition on logical relations, namely

- $R^{\sigma\times\tau}(a, b)$ iff $R^{\tau}(\pi_1\ a, \pi_1\ b)$ and $R^{\tau}(\pi_2\ a, \pi_2\ b)$.

The treatment of sum types $\sigma + \tau$ is analogous.

A type constructor that has received less attention in the literature is (finite) powerset, $\mathcal{P}(\sigma)$. For lack of space we do not propose a specific language of terms to which one could apply the paradigm suggested above, but we claim that the notion of pre-logical relations over full type hierarchies would be extended to powersets by the addition of the following condition:

- $R^{\mathcal{P}(\sigma)}(\alpha, \beta)$ iff $\forall a \in \alpha, \exists b \in \beta . R^{\sigma}(a, b)$ and $\forall b \in \beta, \exists a \in \alpha . R^{\sigma}(a, b)$.

Note that this is the same pattern used in defining bisimulations. The extension for other kinds of models remains a topic for future work.

Various other kinds of types can be considered, including inductive and co-inductive data types, universally and existentially quantified types, and various flavours of dependent types. We have not yet considered these in any detail, but we are confident that for any of them, one could take any existing treatment of logical relations and modify it by weakening the condition on functions as above without sacrificing the Basic Lemma. We expect that this would even yield improved results as it has above, but this is just speculation.

A different dimension of generalization is to consider models having additional structure — e.g. Kripke applicative structures [MM91], pre-sheaf models or cartesian closed categories — for which logical relations have been studied. We have not yet examined the details of this generalization but it appears that a corresponding weakening of the definition would lead to analogues of the results above, cf. [PPS98].

10 Conclusions and Directions for Future Work

Our feeling is that by introducing the notion of pre-logical relation we have, metaphorically and a little immodestly, removed a "blind spot" in the existing intuition of the use and scope of logical relations and related techniques. This is not to say that some specialists in the field have not previously contemplated generalizations similar to ours, but they have not carried the investigation far enough. We believe that in this paper we have exposed very clearly the fact that in many situations the use of logical relations is unnecessarily restrictive. Using pre-logical relations instead, we get improved statements of some results (e.g. Theorem 7.4 and its corollaries), we encompass constructions that had previously escaped the logical paradigm (e.g. Example 3.9), and we isolate the necessary and sufficient hypotheses for many arguments to go through (e.g. Lemma 4.1).

Throughout the paper we have indicated possible directions of future investigation, e.g. with respect to richer type theories. It is plausible that sharper characterizations of representation independence will appear in many different type contexts.

But probably the area where the most benefits will be achieved will be that of the foundations of data refinement. Here we think that a more comprehensive explanation of data refinement would be obtained by combining an account in terms of pre-logical relations with the first-order algebraic treatment in [ST88] which we would expect to extend smoothly to higher-order. Among other improvements, this would result in a non-symmetric refinement relation, giving a better fit with the real-life phenomenon being modelled.

There is a vast literature on logical relations in connection with areas like parametricity, abstract interpretation, etc. A treatment of these topics in terms of pre-logical relations is likely to be as fruitful and illuminating as it has proved to be for the classical example of simply-typed lambda calculus presented here.

Acknowledgements: Thanks to Samson Abramsky, Martin Hofmann, Jo Hannay, Yoshiki Kinoshita, John Mitchell, Peter O'Hearn, Gordon Plotkin, John Power and Ian Stark for helpful comments.

References

[Abr90] S. Abramsky. Abstract interpretation, logical relations, and Kan extensions. *Journal of Logic and Computation* 1:5–40 (1990).

[Gan80] R. Gandy. Proofs of strong normalization. In: *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, 457–477. Academic Press (1980).

[JT93] A. Jung and J. Tiuryn. A new characterization of lambda definability. *Proc. TLCA'93*. Springer LNCS 664, 245–257 (1993).

[KOPTT97] Y. Kinoshita, P. O'Hearn, J. Power, M. Takeyama and R. Tennent. An axiomatic approach to binary logical relations with applications to data refinement. *Proc. TACS'97*, Springer LNCS 1281, 191–212 (1997).

[Loa9?] R. Loader. The undecidability of λ-definability. Church Memorial volume, to appear (199?).

[Mit90] J. Mitchell. Type Systems for Programming Languages. Chapter 8 of *Handbook of Theoretical Computer Science, Vol B.* Elsevier (1990).

[Mit96] J. Mitchell. *Foundations for Programming Languages.* MIT Press (1996).

[MM91] J. Mitchell and E. Moggi. Kripke-style models for typed lambda calculus. *Annals of Pure And Applied Logic* 51:99–124 (1991).

[Plo80] G. Plotkin. Lambda-definability in the full type hierarchy. In: *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, 363–373. Academic Press (1980).

[PPS98] G. Plotkin, J. Power and D. Sannella. A compositional generalisation of logical relations. Draft report, `http://www.dcs.ed.ac.uk/home/dts/pub/laxlogrel.ps` (1998).

[ST88] D. Sannella and A. Tarlecki. Toward formal development of programs from algebraic specifications: implementations revisited. *Acta Informatica* 25:233–281 (1988).

[Sch87] O. Schoett. Data Abstraction and the Correctness of Modular Programming. Ph.D. thesis CST-42-87, Univ. of Edinburgh (1987).

[Sch90] O. Schoett. Behavioural correctness of data representations. *Science of Computer Programming* 14:43–57 (1990).

[Sta85] R. Statman. Logical relations and the typed lambda calculus. *Information and Control* 65:85–97 (1985).

[Ten94] R. Tennent. Correctness of data representations in Algol-like languages. In: *A Classical Mind: Essays in Honour of C.A.R. Hoare*. Prentice Hall (1994).