

A Translation of OCaml GADTs into Coq

Anonymous Author(s)

Abstract

Proof assistants based on dependent types are powerful tools for building certified software. In order to verify programs written in a different language, however, a representation of those programs in the proof assistant is required. When that language is sufficiently similar to that of the proof assistant, one solution is to use a *shallow embedding* to directly encode source programs as programs in the proof assistant. One challenge with this approach is ensuring that any semantic gaps between the two languages are accounted for. In this paper, we present *GSet*, a mixed embedding that bridges the gap between OCaml GADTs and inductive datatypes in Coq. This embedding retains the rich typing information of GADTs while also allowing pattern matching with impossible branches to be translated without additional axioms. We formalize this with GADTML, a minimal calculus that captures GADTs in OCaml, and gCIC, an impredicative variant of the Calculus of Inductive Constructions. Furthermore, we present the translation algorithm between GADTML and gCIC, together with a proof of the soundness of this translation. We have integrated this technique into **coq-of-ocaml**, a tool for automatically translating OCaml programs into Coq. Finally, we demonstrate the feasibility of our approach by using our enhanced version of **coq-of-ocaml** to translate a portion of the Tezos code base into Coq.

1 Introduction

Interactive proof assistants based on dependent type theory are powerful tools for program verification. These tools have been used to certify large and complex systems, including compilers [22], operating systems [20, 21], file systems [9], and implementations of cryptographic protocols [3]. While impressive, each of these efforts effectively constructed a new implementation of the system from scratch, as opposed to verifying an existing implementation. This points to an important hurdle to the adoption of proof assistants— in order to use an interactive theorem prover to certify programs written in different languages, users must first encode those programs in the language of the proof assistant.

A key challenge in this scenario is bridging the gap between the language of the source program and that of the proof assistant. In the case that the two are quite different, the standard solution is to employ a *deep embedding*, i.e. representing the abstract syntax trees of source programs as a data type in the proof assistant [4]. While flexible, this strategy demands considerable machinery, including a formalization of the semantics of the language inside the proof assistant, typically accompanied by an additional reasoning mechanism and proof automation, e.g. a program logic [7].

Thus, the formalization of the language semantics become part of the trusted code base (TCB) of any program verified using this approach.

When the languages are semantically similar, e.g. Haskell and Coq, an alternative strategy is to *shallowly embed* source programs in the target language. Recent efforts have shown how to automate this translation [12, 30], reducing user burden. A shallow embedding gives users access to all the built-in verification tooling of the proof assistant, and naturally inherits any further improvements made to the proof assistant. The semantics of translated programs are that of the proof assistant, and do not require extending the TCB. Instead, users rely on the translation itself to preserve the semantics of the source program. Since a key appeal of using an interactive proof assistant to verify programs are their minimal trusted code base, it is vital to ensure that the translation safely bridges any semantic gaps between the source and target languages, e.g. when translating from a partial language to a total one.

Even when the languages are quite similar, though, subtle discrepancies can exist that make a direct translation impossible. As an example, some OCaml functions over generalized algebraic datatypes (GADTs) [18] do not have a direct analogue in Gallina, the functional programming language of Coq [33], despite the fact that Coq’s inductive datatypes can be thought of as a generalization of GADTs. To see why, consider the following OCaml program:

```
type _ udu =
  | Unit : unit udu
  | Double_unit : (unit * unit) udu

let unit_twelve (x : unit udu) =
  match x with
  | Unit -> 12
```

The `udu` datatype is indexed by a type that varies according to the constructor used to build a value, in this case `unit` and `unit*unit`, respectively. The utility of this extra type information can be seen in the subsequent definition of the `unit_twelve` function. Observe that `Double_unit` can never be used to build a value of type `unit udu`, and thus corresponds to an *impossible branch*, i.e. a case that is never encountered at run-time. As a convenience, OCaml allows users to elide patterns for impossible branches. While this particular example is quite simple, GADTs are commonly used to encode rich type information: e.g. embedding type information into the type of syntax trees so that only well-formed expressions can be built.

A naive transliteration of this program into Gallina immediately encounters a problem:

```

111 Inductive udu : Set → Set :=
112   | Unit : udu unit
113   | Double_unit : udu (unit * unit).
114
115 Definition unit_twelve (x : udu unit) : nat :=
116   match x with
117   | Unit ⇒ 12
118   end.

```

Coq rejects the definition of `unit_twelve` as missing a case for `Double_unit`, as Gallina requires that `match` statements provide an exhaustive set of patterns. Adding a default pattern does not improve matters,

```

123 Definition unit_twelve (x : udu unit) : nat :=
124   match x with
125   | Unit ⇒ 12
126   | _ ⇒ _
127   end.

```

as Coq now complains that it cannot infer an instantiation of the body for the default case. While we could certainly provide a dummy value for this simple example, constructing a value of a given type in Coq is impossible for many polymorphic functions, e.g. the `get_head : 'a list -> 'a` function. An alternative solution is to equip Coq with the necessary typing information to *prove* that this branch is nonsensical, i.e. to derive a proof of `False` for this case. From here, we can appeal to the principle of explosion¹ to derive a dummy value. One way to do so is to use the convoy pattern [10] to augment the pattern match so that information about the type indices of the discriminée is propagated to each of the branches:

```

142 Program Definition unit_twelve' (x : udu unit) : nat :=
143   match x in udu T return T = unit → nat with
144   | Unit ⇒ fun h ⇒ 12
145   | Double_unit ⇒ fun h ⇒ _
146   end eq_refl.

```

This definition employs the `Program` command [29] so that we can use Coq's interactive proof mode to derive the proof of `False`. Promisingly, the resulting goal includes the assumption $H : \text{unit} * \text{unit} = \text{unit}$, which encodes the desired information about the type index of x . Unfortunately, we are no better off than before, as it is impossible to derive `False` from this assumption without additional axioms! The most straightforward way to prove that two types are not equal is via a cardinality argument, i.e. showing that the two types have a different number of elements. This is clearly not the case here, as `unit * unit` and `unit` are both singleton sets containing `(tt, tt)` and `tt` respectively. Moreover, these two types are equivalent [32]: if `unit * unit <> unit` were derivable in Coq, it would imply that the univalence axiom is inconsistent with the underlying type theory, an unwelcome outcome for fans of Homotopy Type Theory [34]. In other

¹The principle of explosion states that from falsehood, anything follows.

words, what was an impossible branch in OCaml could be possible in Coq if we use this straightforward embedding!

Alternatively, we might consider tweaking `unit_twelve` to use dependent pattern matching to allow the type of `match` to vary according to the index of x :

```

171 Definition unit_twelve' (x : udu unit) : nat :=
172   match x in udu T return (match T with
173     | unit ⇒ nat
174     | _ ⇒ unit
175   end) with
176   | Unit ⇒ 12
177   | Double_unit ⇒ tt
178   end.

```

The idea here is to have impossible branches return values of a type that is easily inhabited, e.g. `unit`. Unfortunately, Coq also rejects this definition, as case analysis on types is not allowed.

Yet another solution is to implement the missing branches using an axiom of the form: `unreachable_branch: forall {A}, A`. This is the approach previously adopted by **coq-of-ocaml**, a translator from OCaml programs to Coq [12]. While this approach permits `unit_twelve` to be translated, this comes at the cost of admitting an obviously unsound axiom to the trusted code base, relying on the translation to ensure that it is used safely.

In this paper, we propose an alternative approach that does not rely on the use of unsafe axioms. Our solution implements a mixed embedding [11] of GADTs in Gallina using a distinguished universe for GADT indices, which we call `GSet`. At its core, `GSet` is a universe whose members are both injective and disjoint. This could be accomplished by adding a new sort to the Calculus of Inductive Constructions (CIC), similar to the `SProp` sort that has recently been added to Coq [19], but we adopt a simpler approach of making `GSet` a datatype in `Set` instead, being careful with the translation of GADTs to use `GSets` in a way that ensures their indices are both injective and disjoint.

In summary, this paper makes the following contributions:

- We present a translation from GADTML, a formalization of OCaml with GADTs, to `GCIC`, a variant of `CIC`. Our approach translates impossible branches without any use of axioms.
- We prove that our translation is type-preserving when applied to programs that do not use user-defined type families as indices.
- We have integrated our approach into **coq-of-ocaml**². We evaluate our approach by translating a portion of the Tezos code base, removing a number of axioms required by the previous implementation.

We begin by illustrating our approach with a motivating example of `GSet` in action. Sections 3 and 4 then present a

²The implementation is part of the current release of **coq-of-ocaml** and is available at <https://github.com/formal-land/coq-of-ocaml>.

formalization of our translation and its metatheory, using a minimal functional language with GADTs (GADTML) as the source language, and a variant of the Calculus of Inductive Constructors (CIC) as the target language. Section 5 discusses our implementation of this translation as part of **coq-of-ocaml**³, and discuss its application to a real-world OCaml codebase. We then conclude with a discussion of related work and future directions.

2 An Overview of GSet

In order to properly translate OCaml clients of GADTs to Coq, we adopt a mixed embedding for the type indices of GADTs which provide similar assurances about impossible branches. The key insight is that while user-defined datatypes are not guaranteed to be disjoint in Coq, the constructors of an inductive datatype are. Thus, by adopting a deep embedding for the type indices of GADTs, we can force them to be distinct:

```
Inductive GSet : Set :=
| G_arrow : GSet → GSet → GSet
| G_tuple : GSet → GSet → GSet
| G_tconstr : nat → Set → GSet.
```

This type identifies three main kinds of OCaml types: an GADT index is either a function type, a tuple, or a labeled base type. The key intuition is that every element of this type is provably unique, modulo disjoint labels. We chose these three types for readability of the generated code and simplicity of the translation. Using GSets for GADT indices, we can finally correctly translate the impossible branch of `unit_twelve`:

```
Definition G_unit := G_tconstr 0 unit.
```

```
Inductive udu : GSet → Set :=
| Unit : udu G_unit
| Double_unit : udu (G_tuple G_unit G_unit).
```

```
Definition unit_twelve (x : udu G_unit) : int :=
match x in udu s0
return s0 = G_unit → int with
| Unit ⇒ fun eq0 ⇒ ltac:(subst; exact 12)
| _ ⇒ fun (neq : G_tuple G_unit G_unit = G_unit) ⇒
ltac:(discriminate)
end eq_refl.
```

The case for `Double_unit` now assumes

```
G_tuple G_unit G_unit = G_unit
```

which contradicts the semantics of inductive datatypes in Coq. Thus, we are able to automatically discharge this branch via `discriminate` using Coq's support for tactics in terms. By carefully propagating equalities on the indices of GADTs indexed by GSet, we are able to similarly disregard a large class of impossible branches when using **coq-of-ocaml** to

³The supplementary material includes an in-depth walkthrough of **coq-of-ocaml**.

translate OCaml programs. A key intuition underlying our approach is that these equalities can be used to reify the unification algorithm used by OCaml when typing `match` expressions.

This is not the whole story, however, as clients of GADTs also make use of the extra typing information to enhance their own typing guarantees. The canonical example of this is having an interpreter vary its return type based on the type index of an expression encoded as a GADT:

```
type _ term =
| T_Lift : 'a -> 'a term
| T_Int : int -> int term
| T_Bool : bool -> bool term
| T_Add : int term * int term -> int term
| T_Pair : 'a term * 'b term -> ('a * 'b) term

let rec eval : type a. a term -> a = function
| T_Lift x -> x
| T_Int n -> n
| T_Bool b -> b
| T_Add (x, y) -> (eval x) + (eval y)
| T_Pair (t1, t2) -> (eval t1, eval t2)
```

Here, each `term` expression is augmented with its type: the integer literal `T_Int 1` has the type `int term`, for example, while the boolean `T_Bool true` has type `bool term`. In addition to prohibiting nonsensical terms such as `T_Add (T_Bool true) (T_Int 1)`, these indices allow clients of `term` to vary their signature accordingly. Thus, in addition to ensuring that `eval` is only applied to semantically meaningful expressions, it also guarantees that it returns a tuple when applied to an expression of type `(int, bool) term`, for example. In order to provide appropriate types when translating such programs, we need a denotation of a index as a type in Coq.

In order to do so, we utilize the `decodeG` function, which uses the type parameter of a `G_tconstr` to interpret an index in GSet:

```
Fixpoint decodeG (s : GSet) : Set :=
match s with
| G_tconstr s t ⇒ t
| G_arrow t1 t2 ⇒ decodeG t1 → decodeG t2
| G_tuple t1 t2 ⇒ (decodeG t1) * (decodeG t2)
end.
```

Equipped with this function, we can now produce Coq versions of both `term` and `eval` with the expected types.

```
Inductive term : GSet → Set :=
| T_Lift : forall {a : GSet}, decodeG a → term a
| T_Int : int → term G_nat
| T_Bool : bool → term G_bool
| T_Add : term G_int → term G_int → term G_int
| T_Pair : forall {a b : GSet}, term a → term b
→ term (G_tuple a b).

Fixpoint eval {a : GSet} (function_parameter : term a)
: decodeG a :=
match function_parameter with
```

331 $s ::= \forall a.s \mid t$ Types
 332 $t, u ::= a \mid t \rightarrow t \mid t * t \mid T \bar{t}$ Monotype
 333 $e ::= x \mid \lambda x : t.e \mid e e$ Expression
 334 $\quad \mid \Lambda a.e \mid e[t] \mid (e, e)$
 335 $\quad \mid \text{match } e \text{ with } \overline{K \bar{x} \rightarrow e'}$
 336 $dcl ::= \text{type } T \bar{a} := \overline{K : \forall \bar{a}b. \bar{t} \rightarrow T \bar{a}}$ ADT Declaration
 337 $\quad \mid \text{gadt } G \bar{a} := \overline{K : \forall \bar{b}. \bar{t} \rightarrow G \bar{v}}$ GADT Declaration
 338 $p ::= \overline{dcl}; e$ Program

Figure 1. GADTML Syntax

344 $\mid T_Lift \ v \Rightarrow v$
 345 $\mid T_Int \ n \Rightarrow n$
 346 $\mid T_Bool \ b \Rightarrow b$
 347 $\mid T_Add \ x \ y \Rightarrow Z.add \ (eval \ x) \ (eval \ y)$
 348 $\mid T_Pair \ t1 \ t2 \Rightarrow ((eval \ t1), (eval \ t2))$
 349 end.

350 Note how the pattern for T_Lift uses `decodeG`, so that $T_Lift \ ()$
 351 is translated as $T_Lift \ (a := G_unit) \ ()$. Relying on `GSet` and
 352 `decodeG`, we are able to retain the ability to elide impossible
 353 branches when embedding OCaml GADTs into Coq without
 354 sacrificing the rich typing information of GADT clients, all
 355 while producing Coq programs that are syntactically similar
 356 to their OCaml counterparts.

3 GADTML and gCIC

359 In this section, we present GADTML, a minimal functional
 360 language with GADTs, and gCIC, our variant of CIC. The
 361 next section uses these calculi to formalize our translation. In
 362 a later section, we show how to bridge the gap between the
 363 formalism presented in this section and the implementation.

365 **GADTML.** GADTML is the source language of our compiler,
 366 and its syntax is defined in Figure 1. GADTML extends
 367 System F with tuples, user defined ADTs and GADTs, and
 368 pattern matching. We write GADTML terms in blue to easily
 369 contrast with CIC terms, which are colored in red. We use
 370 the notation $e[t]$ for type applications, uppercase lambdas
 371 are used for type abstractions, e.g. $\Lambda a.e$, and (e_1, e_2) represents
 372 a tuple. Overlines are used to represent sequences, e.g.
 373 \overline{K} . A GADTML program consists of a sequence of datatype
 374 declarations followed by an expression. There are two kinds
 375 of datatypes: ADTs, and GADTs; ADTs only builds homogeneous
 376 datatypes, whereas GADTs allows for a finer-grained
 377 polymorphism. Although every ADT can also be written as
 378 a GADT, we keep them separate to illustrate the challenges
 379 of translating GADTs to Coq. We use G to represent GADTs
 380 and T to represent regular ADTs.

381 The kinding and typing rules for GADTML are presented
 382 in Figure 2. These are largely identical to their counterparts
 383 in System F, with the addition of an extra context Σ . This
 384 context is used to keep track of type constructors for each

$$\frac{\boxed{\Sigma; \Gamma \vdash t : *}}{\text{type } T \bar{a} := \overline{K : \forall \bar{a}b. \bar{t} \rightarrow T \bar{a}} \in \Sigma} \frac{\Sigma; \Gamma \vdash \bar{u} : *}{\Sigma; \Gamma \vdash T \bar{u} : *} \quad (\text{KADT})$$

$$\frac{\boxed{\Sigma; \Gamma \vdash e : t}}{\text{gadt } G \bar{a} := \overline{K : \forall \bar{b}. \bar{t} \rightarrow G \bar{v}} \in \Sigma} \frac{\Sigma; \Gamma \vdash \bar{u} : *}{\Sigma; \Gamma \vdash G \bar{u} : *} \quad (\text{KGADT})$$

$$\frac{\Sigma; \Gamma \vdash e : T \bar{u} \quad \Sigma; \Gamma \vdash t : *}{\text{type } T \bar{a} := \overline{K : \forall \bar{a}b. \bar{t} \rightarrow T \bar{a}} \in \Sigma} \left\{ \frac{\Sigma; \Gamma, a, b, x_i : t_i \vdash e'_i : t}{\Sigma; \Gamma \vdash \text{match } e \text{ with } \overline{K_i \bar{x}_i \rightarrow e'_i} : t} \right\}_{K_i} \quad (\text{TYMATCH})$$

$$\frac{\Sigma; \Gamma \vdash e : G \bar{u} \quad \Sigma; \Gamma \vdash t : *}{\text{gadt } G \bar{a} := \overline{K : \forall \bar{b}. \bar{t} \rightarrow G \bar{v}} \in \Sigma} \left\{ \frac{\Sigma; \sigma_i(\Gamma, \bar{b}, x_i : t_i) \vdash e'_i : \sigma_i(t)}{\sigma_i \equiv \text{unifies}(\bar{u}, \bar{v}_i) \not\equiv \perp} \right\}_{K_i} \quad (\text{TYGMATCH})$$

Figure 2. Selected Kinding and Typing Rules for GADTML

418 declared datatype. The type context Γ is a telescope contain-
 419 ing type variables $a \in \Gamma$ and mappings of variables to their
 420 corresponding types $(x : t) \in \Gamma$. For simplicity, we assume
 421 that every type variable introduced in the type context has
 422 a fresh name.

423 The kind system includes the KADT and KGADT rules for
 424 type constructors, while the type system adds the typing
 425 rules TYMATCH and TYGMATCH for `match` expressions. The
 426 other rules are as expected and therefore elided, the complete
 427 definition of the kinding and typing rules can be found in
 428 the supplementary material.

429 The kinding rule KADT for type constructors $T \bar{u}$ states
 430 that T must be declared in the context Σ and that each u_i must
 431 also be well-kinded. It is implicit in this rule that the length
 432 of \bar{u} must agree with the number of declared parameters \bar{a} .
 433 The kinding rule for GADTs KGADT behaves similarly, the
 434 only difference is that the return type of the constructors
 435 can differ, i.e. for an ADT the constructors must always build
 436 a $T \bar{a}$, whereas GADTs can build $G \bar{v}$, for any well typed list
 437 of terms \bar{v} .

438 The typing rule TYMATCH for case analysis on ADTs re-
 439 quires that there must be a well-typed branch for each one

441 $T, e ::= x \mid \lambda x : A. e \mid e e \mid T \bar{v}$ Expressions
 442 $\mid \forall (a : A), t \mid \text{Set}$
 443 $\mid \text{let } (x : t) = \bar{e} \text{ in } e$
 444 $\mid \text{match } e \text{ in } T \bar{a} \text{ return } t \text{ with}$
 445 $\quad \mid K \bar{x} \Rightarrow e' \text{ end}$
 446 $\text{decl} ::= \frac{\text{Inductive } T \Xi : \Delta \rightarrow \text{Set} := \text{Inductive Types}}{\mid K : \Delta \rightarrow T \bar{v}}$
 447 $\text{prog} ::= \text{decl}; e$ Program

Figure 3. gCIC Syntax

452 of the declared constructors K_i in Σ of the expression being analysed. The corresponding rule for GADTs is more interesting: it only requires patterns for those constructors whose signatures are compatible with the type of the expression being analyzed. More precisely, this assumption uses the standard unification [28] algorithm to try to unify the signature of each constructor with the required type: if unification fails, the branch is impossible and can be safely elided; otherwise the resulting unifier σ is used to type the body of the pattern e_i . This rule also relies on the auxiliary function for doing type substitution in contexts, the definition is as expected and provided in the supplementary material. Notice that unification of GADTs is undecidable [17], thus we present a simple algorithm in the supplementary material.

465 In summary, the typing rule for pattern matching on GADTs states that a *match* expression has type t if:

- 466 • The type of the match t is well kinded;
- 467 • The discriminée e has type $T \bar{u}$, which must be well kinded;
- 468 • T must be declared in Σ with constructors \bar{K} , each of which constructs a $T \bar{v}$;
- 469 • Each K_i that can be unified with $T \bar{u}$ via a unifier σ_i must appear as a pattern. The body of the corresponding pattern e_i must have type $\sigma_i(t)$ in the context substituted with σ_i .

478 **gCIC.** The target language of our translation is gCIC, a variant of CIC equipped with impredicative Set⁴ and let bindings. We focus our attention to gCIC's treatment of inductive datatypes; interested readers can see Paulin-Mohring [26] for a more detailed treatment of CIC.

483 Figure 3 presents the syntax of gCIC, which consists of a single construct for types and terms, and another construct for type family declarations. There is no syntactic distinction between types and expressions, as is standard in dependently typed languages. We use uppercase letters A and T to emphasize that an expression is conceptually a type, and lowercase letters e to emphasize that an object is a term. gCIC expressions include variables a, b, x, y , lambda abstractions, applications, universal quantification, the type of all

493 ⁴We use impredicative Set to simplify the translation by avoiding the vexing details of universes. For more details, see the supplementary material.

496 $\Sigma; \Gamma \vdash e : t$
 497
 498 $\frac{\text{Inductive } T \Xi : \Delta \rightarrow \text{Set} := \mid K : \Delta \rightarrow T \bar{v} \in \Sigma}{\Sigma; \Gamma \vdash K_i : \Delta_i \rightarrow T \bar{v}_i}$ (CTyKONS)
 499
 500
 501 $\frac{\text{Inductive } T \Xi : \Delta \rightarrow \text{Set} := \mid K : \Delta \rightarrow T \bar{v} \in \Sigma}{\Sigma; \Gamma \vdash \bar{u} : \Xi \quad \Sigma; \Gamma \vdash \bar{v} : \Delta}$
 502
 503
 504 $\Sigma; \Gamma \vdash T \bar{u} \bar{v} : \text{Set}$ (CTyTyFAM)
 505
 506 $\Sigma; \Gamma \vdash e : T \bar{u}$
 507 $\Sigma; \Gamma, \bar{a} : \Delta \vdash t : s$
 508 $\text{Inductive } T \Xi : \Delta \rightarrow \text{Set} := \mid K : \Delta \rightarrow T \bar{v} \in \Sigma$
 509 $\{ \Sigma; \Gamma, \bar{x}_i : \Delta_i \vdash e'_i : t[\bar{u}_i/\bar{a}] \}_{K_i}$
 510 $\Sigma; \Gamma \vdash \text{match } e \text{ in } T \bar{a} \text{ return } t \text{ with } \mid K \bar{x} \Rightarrow e' \text{ end} : t[\bar{u}/\bar{a}]$ (CTyMATCH)
 511
 512
 513
 514
 515
 516
 517

Figure 4. Selected Typing Rules for gCIC

518 types *Set* (including itself), type families $T \bar{u}$, let bindings, and case analysis.

519 Adopting standard practice, we use arrows for non-dependent function types. We use a, b to emphasize when a variable is treated as a type variable, and t, s, τ to emphasize when an expression is conceptually a type. gCIC includes explicit syntax for instantiating inductive datatypes in order to simplify the presentation of our translation. We elide non-dependent motive of *match* expressions.

526 Inductive type families consist of a named datatype T and its constructors \bar{K} . Each type declaration uses two telescopes [13], Ξ for the non-varying indices—i.e., the parameters—of a type, and Δ for the indices that do vary, i.e., the arity. By convention, we use the letters u and v for the parameters of a type $T \bar{u}$; v , in particular, is used for the indices in the return type of a constructor: $K_i : \Delta_i \rightarrow T \bar{v}$.

534 Figure 4 presents the typing rules for gCIC, which are largely standard [26]. The typing rule for *match* expressions (CTyMATCH) requires a pattern for each constructor, in contrast to the TyMATCH rule of GADTML, which allows impossible branches to be elided. In addition, this rule uses the supplied motive *in* $T \bar{a}$ *return* t to ensure that the body of each pattern e_i has the expected type of $t[\bar{u}_i/\bar{a}]$. Motives are required because unification is undecidable in the presence of inductive types [24].

4 Translating GADTML into gCIC

544 As discussed in Section 2, a sound translation from GADTML to gCIC needs to deal with the semantic mismatches between how each language deals with pattern matching. In contrast to gCIC, GADTML's typing rule for *match* expressions permit both motives and impossible branches to be elided. This

551 enables, for example, the following GADTML program to be
552 well-typed⁵:

```
553 gadt term a =
554   | T_Lift : forall a. a -> term a
555   | T_Int  : int  -> term int
556   | T_Bool : bool -> term bool
557   | T_Pair : forall l r.
558     term l * term r -> term (l * r)
559
560 λ (e : term int) =>
561   match e with
562   | T_Lift x -> x
563   | T_Int n -> n
```

564 An embedding of this program in gCIC must supply both
565 an appropriate motive for the `match` expression and provide
566 bodies for the missing impossible branches. Our solution to
567 both issues is to modify the definition of `term` to use the type
568 `GSet` for its indices, instead of `Set`. This datatype allows us
569 to provide a dependent motive that equips each branch with
570 exactly the typing information provided by unification in
571 the `TyGMATCH` rule. In the case of reachable branches, our
572 translation uses this information to “cast” the body to the
573 expected dependent type. For impossible branches, this infor-
574 mation allows us to derive a proof of `False`; from this proof,
575 we apply the principle of explosion to provide a “default”
576 body for these patterns.

577 To accomplish this, we have implemented a translation
578 from GADTML to gCIC consisting of three distinct phases:

- 579 1. **Transpilation**: First, we generate a potentially ill-
580 typed gCIC program from a GADTML program, gather-
581 ing information about which types need to be mi-
582 grated to `GSet` along the way.
- 583 2. **Embedding**: Using the information from the previous
584 phase, we update the intermediate gCIC program to
585 use `GSet` indices based on the information gathered
586 by the previous phase.
- 587 3. **Repair**: Finally, we build the proof terms needed to
588 ensure reachable branches are well-typed and to rule
589 out any impossible branches.

590 Before diving into the details of each phase, we begin by il-
591 lustrating the output of each phase on the GADTML program
592 from above.

593 **Transpilation**. The first step of our translation produces
594 the following **ill-typed** gCIC term:

```
596 Inductive term : GSet → Set :=
597   | T_Lift : ∀ (a : Set), a → term a
598   | T_Int  : int  → term int
599   | T_Bool : bool → term bool
600   | T_Pair : ∀ (l : Set),
601     forall (r : Set), term l * term r → term (l * r)
602
603 λ (e : term int).
```

604 ⁵For simplicity, this example assumes definitions of `int` and `bool`.
605

```
606 match e in term c return c = int → int with
607 | T_Lift a x → λ (a = int). x
608 | T_Int n → λ (int = int). n
609 | T_Bool b → λ (bool = int). False
610 | T_Pair l r p → λ (l * r = int). False
611 end eq_refl
```

612 While quite similar to our input program, we can already
613 observe several key differences. The definition of `term`, for
614 example, is indexed on `GSet`, the main `match` statement now
615 includes a motive with an equality about the type index of
616 the discriminatee, and it furthermore includes branches for
617 each constructor of `term`. We note that the latter two changes
618 rely on some auxiliary definitions. These correspond to items
619 included in the standard library of Coq, namely `eq`, `False`,
620 `prod`, `nat`, and `bool`. These definitions are as expected — e.g.,
621 `eq` is the type of equality proofs and has a single constructor
622 `eq_refl`, while `False` is an uninhabited datatype — so we elide
623 them from our example. For now, the translation uses `False`
624 as a signal that later phases need to build the required proof
625 term. We also elide the definition of `GSet` and its decoding
626 function `decodeG`, both of which are equivalent to those
627 presented in Section 2. Our translation depends on other
628 functions commonly available in Coq, e.g., the recursion
629 principles for `eq` (`eq_rec`) and `False` (`False_rec`); our example
630 program elides the (completely standard) definitions of these
631 functions.

632 The translation also generates the following set of `GSet`
633 constraints; these track which type variables should live in
634 `GSet` by marking them with Δ :

$$635 \xi = \{(a : \Delta), (l : \Delta), (r : \Delta)\} \quad 636$$

637 This information is used by the next phase to help embed
638 each of these type variables into `GSet` in a well-typed way.

639 **Embedding**. From this intermediate program, the next
640 phase produces the following (also ill-typed) term:

```
641 Inductive term : GSet → Set :=
642   | T_Lift : ∀ (a : GSet), decodeG a → term a
643   | T_Int  : int  → term (G_tconstr 0 int)
644   | T_Bool : bool → term (G_tconstr 1 bool)
645   | T_Pair : ∀ (l : GSet), ∀ (r : GSet),
646     term l * term r → term (G_tuple l r)
647
648 λ (e : term int).
649   match e in term c return c = G_tconstr 0 int → int with
650   | T_Lift a x → λ (h : a = G_tconstr 0 int). x
651   | T_Int n → λ (h : G_tconstr 0 int = G_tconstr 0 int). n
652   | T_Bool b → λ (h : G_tconstr 1 bool = G_tconstr 0 int). False
653   | T_Pair l r p → λ (h : G_tuple l r = G_tconstr 0 int). False
654   end eq_refl
```

655 Note that all the type variables tagged with Δ in ξ now have
656 the type `GSet`. Any occurrence of these variables outside of
657 an index of `term` has been wrapped with a call to the (elided)
658 `decodeG` function, e.g. as in the first parameter of the `T_Lift`
659 constructor. Finally, each constructor now produces a `term`
660

with the right `GSet` index: `T_Int` now produces a value of type T (`G_tconstr 0 int`). The integer argument of `G_tconstr` uniquely identifies its corresponding type, using the position in the declaration context. In the aforementioned example, `0` marks the position of `int` in the context Σ while `bool` is tagged with `1`. After this phase, all the datatypes declarations are well typed, but it still remains to ensure that `match` expressions are well typed.

Repair. The last phase results in the following well-typed program⁶, by either casting the body of a reachable pattern to the appropriate term, or by supplying a proof of `False` to provide to `False_ind` when the branch is impossible.

```

669 λ (e : term int).
670   match e in term c return c = G_tconstr 0 int → int with
671   | T_lift a x → λ (h : a = G_tconstr 0 int).
672     eq_rec A (G_tconstr 0 int) (λ y ⇒ decodeG y → int)
673     (λ (z : decodeG (G_tconstr 0 int)) ⇒ z) a (eq_sym h) x
674   | T_Int n → λ (h : G_tconstr 0 int = G_tconstr 0 int). n
675   | T_Boolean b → λ (h : G_tconstr 1 bool = G_tconstr 0 int).
676     let (h1 : 1 = 0); (h2 : bool = int) := K_inj h
677     in False_ind (conflict h1)
678   | T_Pair l r p → λ (h : G_tuple 1 r = G_tconstr 0 int).
679     False_ind (conflict h)
680   end eq_refl

```

The body of the pattern for `T_lift` now utilizes the equality provided by the translation of `match` to “cast” its result to the expected type (via an application of the standard recursion principle for equality `eq_rec`). Similarly, both `T_Boolean` and `T_Pair` are impossible branches, so this phase uses the supplied equality to synthesize the required proof of `False`. This proof relies on two key properties of the constructors of inductive datatypes. First, that they are *injective*, which we abbreviate as $K_{inj} : h : K \bar{e}_1 = K \bar{e}_2 \rightarrow \bar{e}_1 = \bar{e}_2$. Second, that they are *disjoint*, which we abbreviate as $conflict : K_i \bar{e}_1 = K_j \bar{e}_2$ (where $K_i \neq K_j$). Our implementation of this translation uses the tactics `inversion` and `discriminate` to construct the proofs of both K_{inj} and `conflict` on demand. Having seen the results of the three phases of our translation on a simple example, we now proceed to a detailed presentation of each phase.

4.1 Transpilation Phase

In order to translate a program, we must also translate its type. More precisely, to translate a type t from a source language to a type t in a target language, we want an algorithm $[1, 5] \Sigma; \Gamma \vdash t : * \rightsquigarrow t$, meaning that under the declaration context Σ , and under the variable context Γ , the type t is well-kinded in the source language and is transpiled to t in the target language.

When translating programs with GADTs, however, we also need to identify which type variables should be translated into `Set` and which ones should be translated into `GSet`.

⁶The type declarations are already well-typed after the previous phase, so we elide them here.

$$\begin{array}{c}
\boxed{\Sigma; \Gamma \vdash t : * \rightsquigarrow_g t \mid \xi} \\
\hline
\frac{\Sigma; \Gamma \vdash a : *}{\Sigma; \Gamma \vdash a : * \rightsquigarrow_* a \mid \{a : *\}} \text{ (TyTRANSVAR)} \\
\hline
\frac{\Sigma; \Gamma \vdash a : *}{\Sigma; \Gamma \vdash a : * \rightsquigarrow_{\Delta} a \mid \{a : \Delta\}} \text{ (TyTRANSGSETVAR)} \\
\hline
\frac{\Sigma; \Gamma, a \vdash t : * \rightsquigarrow_* t \mid \xi}{\Sigma; \Gamma \vdash \forall a. t : * \rightsquigarrow_* \forall (a : \text{Set}), t \mid \xi} \text{ (TyTRANSALL)} \\
\hline
\frac{\text{gadt } G \bar{a} := \mid K : \forall \bar{b}. \bar{t} \rightarrow G \bar{v} \in \Sigma}{\Sigma; \Gamma \vdash u_i : * \rightsquigarrow_{\Delta} u_i \mid \xi_i, \text{ for each } u_i \in \bar{u}} \\
\quad \xi = \sqcup \xi_i \\
\hline
\Sigma; \Gamma \vdash G \bar{u} : * \rightsquigarrow_g G \bar{u} \mid \xi \text{ (TyTRANSGADT)} \\
\hline
\frac{\text{type } T \bar{a} := \mid K : \forall \bar{b}. \bar{t} \rightarrow T \bar{a} \in \Sigma}{\Sigma; \Gamma \vdash u_i : * \rightsquigarrow_* u_i \mid \xi_i, \text{ for each } u_i \in \bar{u}} \\
\quad \xi = \sqcup \xi_i \\
\hline
\Sigma; \Gamma \vdash T \bar{u} : * \rightsquigarrow_g T \bar{u} \mid \xi \text{ (TyTRANSADT)}
\end{array}$$

Figure 5. Selected Type Transpilation Rules

In order to achieve this, we track if we are currently translating an index of a GADT type constructor or not.

Formally, our translation has the form $\Sigma; \Gamma \vdash t : * \rightsquigarrow_g t \mid \xi$. The subscript g tracks if we are currently under a GADT type constructor or not, and the `GSet` constraint ξ tracks which variables should inhabit `GSet` and which ones should inhabit `Set`. We use the notation $\{a : *\}$, when a should inhabit `Set`, and $\{a : \Delta\}$ when a should inhabit `GSet`. Analogously, when translating a GADT index, we mark $g = \Delta$, otherwise $g = *$. For example, if $\text{gadt } G \bar{a} := \mid K : \forall \bar{b}. \bar{t} \rightarrow G \bar{v} \in \Sigma$, then $\Sigma; a \vdash G a : * \rightsquigarrow_* G a \mid \{a : \Delta\}$, since a is used as an index of the GADT, G ; and the algorithm tracks this via the constraint $\xi = \{a : \Delta\}$.

We define a join operation $\xi_1 \sqcup \xi_2$ such that $\langle \xi, \sqcup \rangle$ forms a join-semilattice; such that $\{a : *\} \sqcup \{a : \Delta\} = \{a : \Delta\}$, and therefore $\{a : *\} \leq \{a : \Delta\}$. For different variables it behaves as regular set union $\{a : *\} \sqcup \{b : \Delta\} = \{(a : *), (b : \Delta)\}$. This ensures that all type variables used as `GSet` will be appropriately marked as such.

Figure 5 lists a subset of the rules defining out the translation of types; the complete set of rules can be found in the supplementary material. The heart of the translation are the rules `TyTRANSGADT` and `TyTRANSADT`. The latter states that in order to translate a type $G \bar{u}$, where G is declared as a GADT, we first translate each index u_i at `GSet`, i.e. we set $g = \Delta$. The translation of each u_i yields a translated u_i and `GSet` constraint ξ_i . The final result of translating $G \bar{u}$ is $G \bar{u}$ and the join of all the sets of set of `GSet` constraints $\xi = \sqcup \xi_i$.

$$\begin{array}{c}
\boxed{\Sigma; \Gamma \vdash e : t \rightsquigarrow e \mid \xi} \\
\hline
\frac{\Sigma; \Gamma \vdash x : t \quad \Sigma; \Gamma \vdash t \rightsquigarrow_* t \mid \xi}{\Sigma; \Gamma \vdash x : t \rightsquigarrow x \mid \xi} \text{ (TRANSVAR)} \\
\hline
\frac{\Sigma; \Gamma, x : t_1 \vdash e : t \rightsquigarrow e \mid \xi_e \quad \Sigma; \Gamma \vdash t_1 : * \rightsquigarrow_* t_1 \mid \xi_t}{\Sigma; \Gamma \vdash \lambda(x : t_1).e : t_1 \rightarrow t \rightsquigarrow \lambda(x : t_1).e \mid \xi} \text{ (TRANSLAM)} \\
\hline
\frac{\Sigma; \Gamma \vdash e : \forall a.t \rightsquigarrow e \mid \xi_1 \quad \Sigma; \Gamma \vdash s : * \rightsquigarrow_{\xi_1(a)} s \mid \xi_2}{\Sigma; \Gamma \vdash e[s] : t[s/a] \rightsquigarrow e s \mid \xi} \text{ (TRANSAPP)} \\
\hline
\frac{\Sigma; \Gamma, a \vdash e : t \rightsquigarrow e \mid \xi}{\Sigma; \Gamma \vdash \Lambda a.e : \forall a.t \rightsquigarrow \lambda(a : \text{Set}).e \mid \xi} \text{ (TRANSKLAM)} \\
\hline
\frac{\Sigma; \Gamma \vdash e_1 : t_2 \rightarrow t_1 \rightsquigarrow e_1 \mid \xi_1 \quad \Sigma; \Gamma \vdash e_2 : t_2 \rightsquigarrow e_2 \mid \xi_2}{\Sigma; \Gamma \vdash e_1 e_2 : t_1 \rightsquigarrow e_1 e_2 \mid \xi} \text{ (TRANSAPP)} \\
\hline
\frac{\text{type } T \bar{a} := | K : \forall \bar{a} \bar{b}. \bar{t} \rightarrow T \bar{a} \in \Sigma \quad \Sigma; \Gamma \vdash e : T \bar{u} \rightsquigarrow e \mid \xi_e \quad \{ \Sigma; \Gamma, \bar{a}, \bar{b}, x_i : t_i \vdash e'_i : t \rightsquigarrow e'_i \mid \xi_i \}_{K_i} \quad \xi = (\sqcup \xi_i) \sqcup \xi_e}{\Sigma; \Gamma \vdash \text{match } e \text{ with } | K \bar{x} \rightarrow e \text{ end} : t \rightsquigarrow \text{match } e \text{ with } | K \bar{x} \Rightarrow e' \text{ end} \mid \xi} \text{ (TRANSMATCH)} \\
\hline
\frac{\text{gadt } G \bar{a} := | K : \forall \bar{b}. \bar{t} \rightarrow G \bar{v} \in \Sigma \quad \Sigma; \Gamma \vdash e : G \bar{u} \rightsquigarrow e \mid \xi_e \quad \Sigma; \Gamma \vdash G \bar{u} \rightsquigarrow_* G \bar{u} \mid \xi_u \quad \Sigma; \Gamma, \bar{a}, \bar{b} \vdash \bar{v} : * \rightsquigarrow_{\Delta} \bar{v} \mid \xi_v \quad \xi = (\sqcup \xi_i) \sqcup \xi_e \sqcup \xi_u \sqcup \xi_v \quad \left\{ \begin{array}{l} \Sigma; \sigma_i(\Gamma, \bar{a}, \bar{b}, x_i : t_i) \vdash e'_i : \sigma_i(t) \rightsquigarrow e'_i \mid \xi_i \\ \text{if } \sigma_i \equiv \text{unifies}(\bar{u}, \bar{v}_i) \neq \perp \end{array} \right\}_{K_i} \quad \text{if } \text{unifies}(\bar{u}, \bar{v}_i) \equiv \perp}{\Sigma; \Gamma \vdash \text{match } e \text{ with } | K \bar{x} \rightarrow e' \text{ end} : t \rightsquigarrow | K \bar{x} \Rightarrow \lambda(\bar{h} : \bar{v} = \bar{u}).e' \text{ end eq_refl} \mid \xi} \text{ (TRANSGMATCH)}
\end{array}$$

Figure 6. Expression Transpilation

The rule `TYTRANSADT` is similar, but instead we translate the indices at `Set`, i.e. $g = *$.

The remaining rules are largely as expected. If a variable is being translated at `Set` then the `TYTRANSVAR` rule records that $\{a : *\}$. Otherwise the rule `TYTRANSGETSETVAR` applies, and the constraint $\{a : \Delta\}$ is recorded. Finally, type abstractions are always translated into `Set`, as can be seen in the `TYTRANSALL` rule. A later phase will update this to reflect the information gathered by the `GSet` constraint. To see why we wait to update this term until a later phase consider the following type: $\forall a. T a \rightarrow G a$, the proper translation should be $\forall (a : \text{GSet}), T (\text{decode } a) \rightarrow G a$, however, we would have to first translate the right-hand-side of the arrow (i.e. $G a$) to know that $\{a : \Delta\}$ and be able to translate $T a$ into $T (\text{decode } a)$.

4.1.1 Expression Transpilation. Figure 6 defines the type directed translation of `GADTML` expressions into `GCIC` expressions. Most of the rules follow the typing rules, with some additional tracking of set of `GSet` constraints. The most interesting rules are `TRANSAPP`, `TRANSMATCH`, and `TRANSGMATCH`,

The `TRANSAPP` rule translates type applications $e[s]$, where e has type $\forall a.t$. It first translates e and the constraint

associated a will be used to determine if s should be translated at `GSet` or at `Set`.

The `TRANSMATCH` rule translates `match` expressions of the form `match e with | K $\bar{x} \rightarrow e_i$ end`, where the type of the discriminatee e is $T \bar{u}$ and T is an ADT. It first translates the discriminatee e into e , then it proceeds to translate every branch expression e_i into their respective e_i , with the proper constructor variables in the context $\bar{a}, \bar{b}, \bar{x}_i$. It also returns the join of all of the generated sets of `GSet` constraints ξ .

The `TRANSGMATCH` rule translates matching expressions of the form `match e with | K $\bar{x} \rightarrow e'$` when the discriminatee e has type $G \bar{u}$ and G is declared as a GADT. It translates the discriminatee into e , its type $G \bar{u}$ into $G \bar{u}$, and the return type t into t .

In order to capture the information provided by unification in the `TYGMATCH` rule, `TRANSGMATCH` exposes the equalities between the indices of the discriminatee and the indices of the constructors $\bar{v} = \bar{u}$, in each branch of the match. This is accomplished by using the motive `in G \bar{c} return ($\bar{c} = \bar{u}$) $\rightarrow t$` , and having each branch take the equalities $\lambda(\bar{h} : \bar{v} = \bar{u}).e_i$ as arguments. These proofs are provided at the end of the match with the appropriate number of `eq_refl`s, ensuring that the type of the translated match is t , as expected.

Impossible branches are elided in GADTML, but they must be provided in gCIC. To achieve this, the translation checks if each branch is possible by unifies(\bar{u}, \bar{v}_i). If this fails then we set the body of the branch as **False**, signaling to a later phase that a proof of **False** is required. If it succeeds, then we translate a unified version of the body e_i . As always, this rule returns the join of the constraints generated by each subexpression.

4.2 Embedding Phase

The embedding phase uses the embedding function $g[-]_{\xi}^{\Gamma}$ that is defined in Figure 7. This phase takes as input a gCIC term and returns another gCIC term, with type variables translated into **GSet** when necessary. As before, g tracks if the embedding is being done inside a GADT type constructor, Γ tracks the variable types necessary for the next phase, and ξ stores the set of **GSet** constraints produced by the first phase of the translation.

Similar to in the type translation, $g = \Delta$ denotes that the embedding is being performed inside a GADT type constructor, and $g = *$ otherwise. In other words, g flips into Δ when embedding the indices of a GADT type constructor, i.e. $*[G \bar{u}]_{\xi}^{\Gamma} = G^{\Delta}[\bar{u}]_{\xi}^{\Gamma}$ and flips back to $*$ when embedding indices of ADT type constructors, i.e. $\Delta[T \bar{u}]_{\xi}^{\Gamma} = T^*[\bar{u}]_{\xi}^{\Gamma}$.

Formally, embedding is a partial function defined over the structure of gCIC terms. It is only applied after the transpilation phase, and hence is only defined on the range of transpilation, which is a subset the gCIC language. As one example, $T(\lambda(x : t).e)$ can never be generated by the transpilation, and therefore embedding is not defined on this term.

To embed an ADT type constructor at **Set**, i.e. $*[T \bar{u}]_{\xi}^{\Gamma}$, we simply embed its indices: $T^*[\bar{u}]_{\xi}^{\Gamma}$. On the other hand, to embed this type at **GSet**, we use **G_tconstr**, and record its position in the declaration signature $\#\Sigma(T)$. Embedding a GADT is similar, with the only difference that the indices will be embedded at **GSet**, i.e. $G^{\Delta}[\bar{u}]_{\xi}^{\Gamma}$. Assigning a unique key to each type constructor is paramount for ensuring injectivity and disjointness of type constructors, which is crucial to the next phase.

To embed a **match** expression, we embed both the discriminée and return of the motive. To finish translating the branches of the match, the next phase will use information from the typing context Γ to repair the body of each match to have the correct type.

The other rules are largely as expected. Arrows and tuples are also embedded into **GSet** when necessary. The indices of equations are always translated with $g = \Delta$ because they are only generated by the transpiler to compare GADT indices. Universally quantified variables are now migrated to **GSet** when they are marked in the set of **GSet** constraints. The context Γ is also extended when embedding lambda terms and

universal quantifiers, as this information will be necessary by the repair phase.

4.3 Repair Phase

The last translation step repairs the body of **match** expressions so that they are well-typed. It does so via the repair function \vdash_s defined in Figure 8. This function takes as input a term e , a target type t and a context Γ , and outputs a term e^{\dagger} that has type t under the context Γ (Lemma 1). It recurses over the tail of the typing context Γ , and terminates when it either reaches a non-equation in Γ , or when a contradiction found.

The first two rules of \vdash_s perform a type cast when they find an equation on a variable x . The rules behave similarly, the only difference being that if x is found on the left of the equation the symmetry property of equations is first applied via the function **eq_sym**. To perform this cast, the algorithm first gathers all variables $\bar{z} : \bar{u}$ in Γ in which x appears in u_i . It then casts the body by recursively applying \vdash_s to e with all occurrences of x substituted by τ , including occurrences in the target type t and in the context Γ . The cast is built using the function **eq_rec**, which has the type:

$$\Sigma; \vdash \text{eq_rec} : \forall(A : \text{Set})(x : A)(P : A \rightarrow \text{Set}), \\ Px \rightarrow \forall(y : A), x = y \rightarrow Py$$

Building this substitution recursively is possible because the previous type marks the exact positions at which the rewrite will happen. This can be seen in the third argument supplied to **eq_rec**: $\lambda(y : A). (\bar{u} \rightarrow t)[x/y]$, which indicates that the expression being cast has type $\bar{u}[x/y] \rightarrow t[x/y]$. This allows the recursive call to access to each $z_0 : \bar{u}[x/\tau]$, after y is instantiated to τ . Substitution then replaces each z in Γ with its respective z_0 , i.e. $\Gamma[\bar{z}_0/\bar{z}]$. Since \bar{z} captures all variables that mentions x , and they have been substituted by z_0 , which doesn't mention x , we can safely remove it from the context, via $\Gamma[\bar{z}_0/\bar{z}] - \{x\}$.

When an equation over constructors is encountered, the function first check if they are the same constructor. If this is the case, the repair algorithm uses K_{inj} , the injectivity rule for constructors, and continues the type substitution recursively. If the constructors are not the same, it has reached a contradiction, and the term can be replaced by the aforementioned **conflict** term. The repair function also introduces function variables $\lambda(x : t)$ into the context, and furthermore ignores trivial equations.

Concretely, this algorithm behaves as an inverse function of the substitution of type variables. We can see this in Lemma 1, which states that for any term e that has type $t[\tau/x]$ under a context with the same substitution $\Gamma[\tau/x]$, applying the repair algorithm using the equality $h : \tau = x$ to the e yields a well-typed term without the substitution.

Lemma 1 (Repair step is the inverse of substitution). *If $\Sigma; \Gamma[\tau/x] \vdash e : t[\tau/x]$ and $\Gamma, h : x = \tau \vdash_s e : t = e^{\dagger}$ then*

991	$*[Set]_{\xi}^{\Gamma} = Set$	$*[T \bar{u}]_{\xi}^{\Gamma} = T * [\bar{u}]_{\xi}^{\Gamma}$	1046
992	$\Delta[Set]_{\xi}^{\Gamma} = GSet$	$\Delta[T \bar{u}]_{\xi}^{\Gamma} = G_tconstr (\#\Sigma(T)) (T * [\bar{u}]_{\xi}^{\Gamma})$	1047
993			1048
994	$*[a]_{\xi}^{\Gamma} = \begin{cases} \text{decode } a & \text{if } (a : \Delta) \in \xi \\ a & \text{otherwise} \end{cases}$	$*[G \bar{u}]_{\xi}^{\Gamma} = G \Delta[\bar{u}]_{\xi}^{\Gamma}$	1049
995		$\Delta[G \bar{u}]_{\xi}^{\Gamma} = G_tconstr (\#\Sigma(G)) (G \Delta[\bar{u}]_{\xi}^{\Gamma})$	1050
996			1051
997	$\Delta[a]_{\xi}^{\Gamma} = a, \quad \text{if } (a : \Delta) \in \xi$	$*[\bar{u} = \bar{v}]_{\xi}^{\Gamma} = \Delta[\bar{u}]_{\xi}^{\Gamma} = \Delta[\bar{v}]_{\xi}^{\Gamma}$	1052
998			1053
999	$\Delta[t_1 \rightarrow t_2]_{\xi}^{\Gamma} = G_arrow \Delta[t_1]_{\xi}^{\Gamma} \Delta[t_2]_{\xi}^{\Gamma}$	$*[\forall(a : Set), t]_{\xi}^{\Gamma} = \forall(a : GSet), * [t]_{\xi}^{\Gamma, (a : GSet)}$, if $(a : \Delta) \in \xi$	1054
1000		$*[\forall(a : Set), t]_{\xi}^{\Gamma} = \forall(a : Set), * [t]_{\xi}^{\Gamma, (a : Set)}$, otherwise	1055
1001	$*[t_1 \rightarrow t_2]_{\xi}^{\Gamma} = * [t_1]_{\xi}^{\Gamma} \rightarrow * [t_2]_{\xi}^{\Gamma}$		1056
1002	$*[t_1 * t_2]_{\xi}^{\Gamma} = * [t_1]_{\xi}^{\Gamma} * * [t_2]_{\xi}^{\Gamma}$	$*[\lambda(x : t_1), t]_{\xi}^{\Gamma} = \lambda(x : * [t_1]_{\xi}^{\Gamma}), * [t]_{\xi}^{\Gamma, (x : * [t_1]_{\xi}^{\Gamma})}$	1057
1003			1058
1004	$\Delta[t_1 * t_2]_{\xi}^{\Gamma} = G_tuple \Delta[t_1]_{\xi}^{\Gamma} \Delta[t_2]_{\xi}^{\Gamma}$	$^g[e_1 e_2]_{\xi}^{\Gamma} = ^g[e_1]_{\xi}^{\Gamma} ^g[e_2]_{\xi}^{\Gamma}$	1059
1005			1060
1006			1061
1007			1062
1008			1063
1009			1064
1010			1065
1011			1066
1012			1067
1013			1068
1014			1069
1015			1070
1016			1071
1017			1072
1018			1073
1019			1074
1020			1075
1021			1076
1022			1077
1023			1078
1024			1079
1025			1080
1026			1081
1027			1082
1028			1083
1029			1084
1030			1085
1031			1086
1032			1087
1033			1088
1034			1089
1035			1090
1036			1091
1037			1092
1038			1093
1039			1094
1040			1095
1041			1096
1042			1097
1043			1098
1044			1099
1045			1100

Figure 7. Embedding Function

1014	$\Gamma, h : x = \tau \vdash_s e : t \triangleq$		
1015	take all $(\bar{z} : \bar{u}) \in \Gamma, \text{ s.t } x \in u,$		
1016	$\text{eq_rec } A \tau (\lambda (y : A). (\bar{u} \rightarrow t)[x/y])$		
1017	$(\lambda (z_0 : u[\tau/x]). \Gamma[z_0/z] - \{x\} \vdash_s e[z_0/z] : t[\tau/x])$		
1018	$x (\text{eq_sym } h) \bar{z}$		
1019			
1020	$\Gamma, h : x = \tau \vdash_s e : t \triangleq$		
1021	take all $(\bar{z} : \bar{u}) \in \Gamma, \text{ s.t } x \in u,$		
1022	$\text{eq_rec } A \tau (\lambda (y : A). (\bar{u} \rightarrow t)[x/y])$		
1023	$(\lambda (z_0 : u[\tau/x]). \Gamma[z_0/z] - \{x\} \vdash_s e[z_0/z] : t[\tau/x])$		
1024	$x h \bar{z}$		
1025			
1026	$\Gamma, h : K \bar{x} = K \bar{y} \vdash_s e : t \triangleq$	$\text{let } (\overline{h : x = y}) := K_{inj} h \text{ in}$	
1027		$\Gamma, (h : x = y) \vdash_s e : t$	
1028			
1029	$\Gamma, h : K_1 \bar{x} = K_2 \bar{y} \vdash_s e : t \triangleq$	$\text{if } K_1 \neq K_2,$	
1030		$\text{False_ind (conflict } h)$	
1031			
1032	$\Gamma \vdash_s \lambda(x : t'). e : t' \rightarrow t \triangleq$	$\Gamma, (x : t') \vdash_s e : t$	
1033			
1034	$\Gamma, h : \tau = \tau \vdash_s e : t \triangleq$	$\Gamma \vdash_s e : t$	
1035			
1036	$\Gamma \vdash_s e : t \triangleq$	e , if the head of Γ is not	
1037		an equation	
1038			
1039			

Figure 8. Repair Function

$\Sigma; \Gamma, h : x = \tau \vdash e^{\dagger} : t$, assuming that Γ contains no equations.

Proof. Direct from the definition of \vdash_s and the type of `eq_rec`. See the supplementary material for the full proof. \square

4.4 Soundness of the Translation

In order to show that our translation is sound, we prove both that type translation preserves kinding and that expression translation preserves typing.

Theorem 1 establishes that if a type t is a well-kinded type in GADTML, then its translation t is also well-kinded under a translated context, after the embedding and repair phases. For this it is also necessary to define the translation of contexts: $\vdash \Sigma \rightsquigarrow \Sigma \mid \xi_{\Sigma}$ and $\Sigma \vdash \Gamma \rightsquigarrow \Gamma$, these definitions are straightforward and elided from our presentation due to space constraints, but they can be found in the supplementary material.

Also note the contexts can also be embedded using the generated set of `GSet` constraints, i.e. $[\Sigma]_{\xi_{\Sigma}}; [\Gamma]_{\xi}$. The definition of these embeddings are a straightforward application of the embedding algorithm to contexts. The translation of declaration contexts $\vdash \Sigma \rightsquigarrow \Sigma \mid \xi_{\Sigma}$ generates its own set of `GSet` constraints since the variable information on each datatype constructor is local.

Theorem 1 (Type Translation Preserves Kinding). *If $\Sigma; \Gamma \vdash t : * \rightsquigarrow_g t \mid \xi$ and $\vdash \Sigma \rightsquigarrow \Sigma \mid \xi_{\Sigma}$ and $\Sigma \vdash \Gamma \rightsquigarrow \Gamma$ then $[\Sigma]_{\xi_{\Sigma}}; [\Gamma]_{\xi} \vdash ^g [t]_{\xi}^{\Gamma} : ^g [Set]_{\xi}^{\Gamma}$*

Proof. By induction on the derivation of the type translation $\Sigma; \Gamma \vdash t : * \rightsquigarrow_g t \mid \xi$. \square

Our second theorem establishes that the translation of a well-typed GADTML term produces well-typed `gCIC` term.

We note here that the repair function isn't strong enough to handle nested user-defined type constructors of the form $G(T \bar{u})$, as the current formulation of `GSet` can only embed one level of type constructors with a unique key, as seen in `G_tconstr`.

Theorem 2 (Expression Translation Preserves Typing). *If $\Sigma; \Gamma \vdash e : t \rightsquigarrow e \mid \xi$ and $\Sigma; \Gamma \vdash t : * \rightsquigarrow_* t \mid \xi_t$ and $\vdash \Sigma \rightsquigarrow \Sigma \mid \xi_\Sigma$ and $\Sigma \vdash \Gamma \rightsquigarrow \Gamma$ then $[\Sigma]_{\xi_\Sigma}; [\Gamma]_{\xi} \vdash * [e]_{\xi}^{\Gamma} : * [t]_{\xi}^{\Gamma}$, assuming that e doesn't have pattern matchings over datatypes that uses user-defined types as indices*

Proof. By induction over the derivation of the transpilation of expressions $\Sigma; \Gamma \vdash e : t \rightsquigarrow e \mid \xi$. The more interesting case is that of the `TRANSMATCH` rule; the proof of this case is included in the supplementary material. \square

5 Implementation and Evaluation

We have implemented our translation of GADTs in `coq-of-ocaml`, a source-to-source compiler from OCaml to Coq. Our implementation closely follows the algorithm presented above, although there are two discrepancies worth mentioning. First, our translation supports mixing datatype and function declarations, in contrast to the algorithm, which requires type declarations to appear at the beginning of a program. In order to ensure that embedded types are unique, our implementation uses strings instead of numbers for identifiers, and uses the name of a type for this argument. Second, as described in Section 1, the repair phase of the algorithm inserts uses of the `discriminate` and `subst` tactics for the bodies of branches.

Notably, `coq-of-ocaml` handles a considerably larger subset of OCaml than `GADTML`, including many features that use type parameters, e.g. parametrized records, parametrized type synonyms and “grabbing” of existential variables. All of these represent another use of type variables that our implementation also carefully tracks and migrates to `GSet` when necessary. In addition, our implementation handles native types (e.g. `int`, `bool`, and `list`) and translates them as their equivalent counterpart in Coq's standard library. As the treatment of these base types is orthogonal to the translation of GADTs, we have opted to elide them from our formalization.

We have developed a set of micro-benchmarks showcasing each of the features needed to support `GSet`-indexed GADTs in `coq-of-ocaml`. They can be found in the folder `tests` of the supplementary material⁷, they are:

- `GSet_term.ml`: impossible branches and casts;
- `GSet_record.ml`: embedded records with parameters that are used as GADT indices;
- `GSet_existential.ml`: existential variables used as GADT indices;

⁷The micro-benchmarks are also included in the current version of `coq-of-ocaml`, but organized differently.

Function Name	OCaml LOC	Coq LOC
<code>reveal_case</code>	10	25
<code>transaction_case</code>	36	65
<code>origination_case</code>	30	47
<code>delegation_case</code>	11	31
<code>register_global_constant_case</code>	12	40
Total	99	208

Table 1. Size of translated `Operation_Repr` functions

- `GSet_record.ml`: regular records and irrefutable patterns;
- `GSet_ex_grab.ml`: grabbing of existential variables that are marked as `GSet`;

To evaluate the effectiveness of our approach, we have also used our extended version of `coq-of-ocaml` to translate a portion the Michelson interpreter, which is part of Tezos' code base. Michelson is a smart contract language that uses GADTs to ensure that operations are always applied to arguments of the expected type. Since Michelson can be used to manage real money, it is paramount that its interpreter is bug-free and reliable.

In order to evaluate our implementation, we picked a representative GADT from the Michelson interpreter, namely `manager_operation`. This datatype is responsible for managing some operations performed by the nodes and smart contracts of the Tezos protocol, and its definition can be found in `operation_repr.ml`

Before the implementation of the presented translation, `coq-of-ocaml` would translate impossible branches via a use of the axiom `gadt_unreachable_branch`. Using our translation, the updated version of `coq-of-ocaml` eliminated all uses of the `gadt_unreachable_branch` axiom in the five functions shown in Table 1. While the updated translation increased the size of the translated functions, e.g. by inserting type equalities in `match` statements, the small increase in code size has a clear benefit in terms of reducing the trusted code base by eliminating the use of axioms.

6 Related Work

The source language of our translation, `GADTML`, is similar to $\lambda_{2,G\mu}$, first presented in Xi et al. [35]. That calculus is an extension of System F with all the features of `GADTML` and more (e.g. fixpoints and let bindings). Although $\lambda_{2,G\mu}$ does not include GADTs directly, the authors show they can be derived from a surface language. While `GADTML` explicitly uses unification to type check pattern matching, $\lambda_{2,G\mu}$ instead solves constraints maintained in the type variable context. The authors also never discuss impossible branches. We argue that the alternative design choices of `GADTML` enable a cleaner presentation of our translation.

In a similar vein, Sulzmann et al. [31] presents System Fc, an extension of System F with type equality coercions. The authors show that their calculus can encode a plethora of

interesting language features, including GADTs. That work also presents a multi-step constraint-translation of a source language with GADTs into System Fc, and also shows that their translation is type-preserving. In addition to the different target languages, the key difference between our two approaches is that Sulzmann et al. [31] does not need to construct proof terms witnessing type casts and the infeasibility of impossible branches.

There have been a number of efforts using interactive proof assistants to verify programs written in mainstream functional programming languages. The most closely related example is **hs-to-coq** [30], a source-to-source translator from Haskell to Coq. Like **coq-of-ocaml**, **hs-to-coq** produces a shallow embedding of source programs in Coq. The tool is able to capture many advanced language features of Haskell, including typeclasses, records and guarded pattern matching, and it has been used to translate and verify several textbook examples as well as significant portions of Haskell's containers library [6]. **hs-to-coq** provides a best-effort approach to supporting GADTs in translated programs. For some datatypes, users can provide a specification file marking which arguments should be translated as indices, simplifying type argument inference. They then translate the indices directly, so they are forced to use an axiom to handle impossible branches, similar to the use of `unreachable_gadt_branch` in **coq-of-ocaml**. This axiom is also used to support incomplete patterns in Haskell functions. In principle, our translation algorithm is general enough to be incorporated into **hs-to-coq**, eliminating the need for this axiom when used for impossible branches.

CFML translates OCaml programs to Coq via characteristic formulae [8]. The key idea in this approach is to capture program behaviors via invariants expressed as higher-order formulae, which can then be expressed directly in the logic of Coq. Since this approach does not generate functions in Coq, it is capable of faithfully capturing the behaviors of non-terminating programs. The cost of this flexibility is that the translation loses much of the structure of the original program. Following the structure of the source program is an important design decision behind **coq-of-ocaml**.

OCaml-to-PVS Equivalence Validation (**OPEV**) [2] is a tool to validate translations between OCaml and PVS [25] programs by automatically generating a large number of test cases and automatically discharging them using PVS. This approach could help to address a different gap in the current implementation of **coq-of-ocaml**, as it currently relies on users to validate that the translated code matches the intended semantics of the OCaml source programs.

Cameleer [15, 27] takes as input OCaml programs annotated with specifications in the GOSPEL language and outputs verification conditions. These conditions are then discharged by Why3 [16], a deductive verification toolchain

that interfaces with several SMT solvers. In contrast to **coq-of-ocaml**, **Cameleer** relies on automated theorem provers to certify the correctness of OCaml programs.

Coq provides a mechanism to extract code [23] to OCaml, Haskell, Scheme and JSON. This is, in some sense, the inverse of the problem we address here, as we go from from a language with less expressive types to one with richer types. Thus, the extraction mechanism is tasked with safely *erasing* information, including any proof terms, as opposed to faithfully *preserving* type information. Spector-Zabusky et al. [30] proposes that extracting translated code and then testing its equivalence with the original program could greatly increase confidence in their results. Automatically validating the equivalence of roundtrip translations of translating code is an interesting direction for future work.

7 Future Work and Conclusion

One potential direction for future work is to provide a proof that the runtime behavior of the translated GADTML term is equivalent to the behavior of the generated gCIC term. In addition, the present definition of GSet is not expressive enough to translate some mutually recursive GADTs, due to positivity constraints in Coq. The current formulation of GSet is also currently not expressive enough to solve equations generated by GADT indexed by other user-defined type constructors, since these type constructors are injective in OCaml but not necessarily in CIC. As mentioned in the introduction, this could be more directly addressed by making GSet a new universe, similar to SProp, with is equipped with axioms encoding that all inhabitants of this new universe respects injectivity and disjointness of type constructors. A key technical challenge is developing restrictions that make this new universe sound, as injectivity of type constructors is known to be unsound in the presence of inductive types [14].

In this paper, we have presented *GSet*, a mixed embedding that bridges the gap between OCaml GADTs and inductive datatypes in Coq. This embedding retains the rich typing information of GADTs while also allowing case statements with impossible branches to be translated without additional axioms. We presented GADTML, a calculus that captures the essence of GADTs in OCaml and described a sound translation from GADTML to gCIC, a variant of CIC. We have implemented this technique in **coq-of-ocaml**, a tool for automatically translating OCaml programs into Coq. We have used this enhanced version of **coq-of-ocaml** to translate a portion of the OCaml interpreter for Michelson, the smart contract language of Tezos, into Coq, removing five axioms that were generated by previous versions of the tool.

References

- [1] Amal Ahmed and Matthias Blume. 2008. Typed Closure Conversion Preserves Observational Equivalence. *SIGPLAN Not.* 43, 9 (sep 2008), 157–168. <https://doi.org/10.1145/1411203.1411227>

- 1321 [2] Xiaoxin An, Amer Tahat, and Binoy Ravindran. 2020. A Validation
1322 Methodology for OCaml-to-PVS Translation. In *NASA Formal Meth-*
1323 *ods*, Ritchie Lee, Susmit Jha, Anastasia Mavridou, and Dimitra Gian-
1324 *nakopoulou* (Eds.). Vol. 12229. Springer International Publishing, -,
1325 207–221. https://doi.org/10.1007/978-3-030-55754-6_12 Series Title:
1326 Lecture Notes in Computer Science.
- 1327 [3] Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet,
1328 Markulf Kohlweiss, Jianyang Pan, Jonathan Protzenko, Aseem Rastogi,
1329 Nikhil Swamy, Santiago Zanella-Béguelin, and Jean Karim Zinzindou-
1330 houé. 2016. Implementing and Proving the TLS 1.3 Record Layer.
1331 Cryptology ePrint Archive, Report 2016/1178. <https://ia.cr/2016/1178>.
- 1332 [4] Richard J. Boulton, Andrew Gordon, Michael J. C. Gordon, John Har-
1333 rison, John Herbert, and John Van Tassel. 1992. Experience with
1334 Embedding Hardware Description Languages in HOL. In *Proceedings*
1335 *of the IFIP TC10/WG 10.2 International Conference on Theorem Provers*
1336 *in Circuit Design: Theory, Practice and Experience*. North-Holland Pub-
1337 lishing Co., NLD, 129–156.
- 1338 [5] William J. Bowman, Youyou Cong, Nick Rioux, and Amal Ahmed.
1339 2017. Type-Preserving CPS Translation of Σ and Π Types is Not Not
1340 Possible. *Proc. ACM Program. Lang.* 2, POPL, Article 22 (dec 2017),
1341 33 pages. <https://doi.org/10.1145/3158110>
- 1342 [6] Joachim Breitner, Antal Spector-Zabusky, Yao Li, Christine Rizkallah,
1343 John Wiegley, and Stephanie Weirich. 2018. Ready, Set, Verify! Ap-
1344 plying Hs-to-Coq to Real-World Haskell Code (Experience Report).
1345 *Proc. ACM Program. Lang.* 2, ICFP, Article 89 (July 2018), 16 pages.
1346 <https://doi.org/10.1145/3236784>
- 1347 [7] Qinxiang Cao, Lennart Beringer, Samuel Gruetter, J. Dodds, and A.
1348 Appel. 2018. VST-Floyd: A Separation Logic Tool to Verify Correctness
1349 of C Programs. *Journal of Automated Reasoning* 61 (2018), 367–422.
- 1350 [8] Arthur Charguéraud. 2010. Program Verification through Character-
1351 istic Formulae. In *Proceedings of the 15th ACM SIGPLAN International*
1352 *Conference on Functional Programming* (Baltimore, Maryland, USA)
1353 (ICFP '10). Association for Computing Machinery, New York, NY, USA,
1354 321–332. <https://doi.org/10.1145/1863543.1863590>
- 1355 [9] Haogang Chen, Daniel Ziegler, Tej Chajed, Adam Chlipala, M. Frans
1356 Kaashoek, and Nikolai Zeldovich. 2015. *Using Crash Hoare Logic for*
1357 *Certifying the FSCQ File System*. Association for Computing Machinery,
1358 New York, NY, USA, 18–37. <https://doi.org/10.1145/2815400.2815402>
- 1359 [10] Adam Chlipala. 2013. *Certified Programming with Dependent Types: A*
1360 *Pragmatic Introduction to the Coq Proof Assistant*. The MIT Press, -.
- 1361 [11] Adam Chlipala. 2021. Skipping the Binder Bureaucracy with Mixed
1362 Embeddings in a Semantics Course (Functional Pearl). *Proc. ACM*
1363 *Program. Lang.* 5, ICFP, Article 94 (Aug. 2021), 28 pages. <https://doi.org/10.1145/3473599>
- 1364 [12] Guillaume Claret. 2021. Coq of Ocaml. <https://github.com/clarus/coq-of-ocaml>. Accessed: 2021-09-09.
- 1365 [13] Jesper Cockx, Dominique Devriese, and Frank Piessens. 2016. Unifiers
1366 as Equivalences: Proof-Relevant Unification of Dependently Typed
1367 Data. In *Proceedings of the 21st ACM SIGPLAN International Conference*
1368 *on Functional Programming* (Nara, Japan) (ICFP 2016). Association
1369 for Computing Machinery, New York, NY, USA, 270–283. <https://doi.org/10.1145/2951913.2951917>
- 1370 [14] Leonardo de Moura. 2022. Coq of Ocaml. <https://github.com/leanprover/lean/issues/654>. Accessed: 2022-09-21.
- 1371 [15] Jean-Christophe Filliâtre, Léon Gondelman, Cláudio Lourenço, An-
1372 drei Paskevich, Mário Pereira, Simão Melo de Sousa, and Aymeric
1373 Walch. 2020. A Toolchain to Produce Verified OCaml Libraries. (Jan.
1374 2020). <https://hal.archives-ouvertes.fr/hal-01783851> working paper
1375 or preprint.
- 1376 [16] Jean-Christophe Filliâtre and Andrei Paskevich. 2013. Why3 – Where
1377 Programs Meet Provers. In *Programming Languages and Systems*,
1378 Matthias Felleisen and Philippa Gardner (Eds.). Springer Berlin Hei-
1379 delberg, Berlin, Heidelberg, 125–128.
- 1380 [17] Jacques Garrigue and Jacques Le Normand. 2017. GADTs and Ex-
1381 haustiveness: Looking for the Impossible. *Electronic Proceedings in*
1382 *Theoretical Computer Science* 241 (Feb. 2017), 23–35. <https://doi.org/10.4204/EPTCS.241.2>
- 1383 [18] Jacques Garrigue and Jacques Le Normand. 2011. Adding GADTs to
1384 OCaml the direct approach. (2011), 29.
- 1385 [19] Gaëtan Gilbert, Jesper Cockx, Matthieu Sozeau, and Nicolas Tabareau.
1386 2019. Definitional Proof-Irrelevance without K. *Proc. ACM Program.*
1387 *Lang.* 3, POPL, Article 3 (jan 2019), 28 pages. <https://doi.org/10.1145/3290316>
- 1388 [20] Ronghui Gu, Jérémie Koenig, Tahina Ramananandro, Zhong Shao,
1389 Xiongnan (Newman) Wu, Shu-Chun Weng, Haozhong Zhang, and
1390 Yu Guo. 2015. Deep Specifications and Certified Abstraction Layers.
1391 *SIGPLAN Not.* 50, 1 (Jan. 2015), 595–608. <https://doi.org/10.1145/2775051.2676975>
- 1392 [21] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick,
1393 David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt,
1394 Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and
1395 Simon Winwood. 2009. SeL4: Formal Verification of an OS Ker-
1396 nel. In *Proceedings of the ACM SIGOPS 22nd Symposium on Oper-*
1397 *ating Systems Principles* (Big Sky, Montana, USA) (SOSP '09). As-
1398 sociation for Computing Machinery, New York, NY, USA, 207–220.
1399 <https://doi.org/10.1145/1629575.1629596>
- 1400 [22] Xavier Leroy. 2009. Formal Verification of a Realistic Compiler. *Com-*
1401 *munic. ACM* 52, 7 (July 2009), 107–115. <https://doi.org/10.1145/1538788.1538814>
- 1402 [23] Pierre Letouzey. 2008. Extraction in coq: An overview. In *Conference*
1403 *on Computability in Europe*. Springer, 359–369.
- 1404 [24] Conor McBride. 2000. Dependently typed functional programs and
1405 their proofs. (2000).
- 1406 [25] Sam Owre, John M Rushby, and Natarajan Shankar. 1992. PVS: A pro-
1407 totype verification system. In *International Conference on Automated*
1408 *Deduction*. Springer, 748–752.
- 1409 [26] Christine Paulin-Mohring. 2015. Introduction to the calculus of induc-
1410 tive constructions.
- 1411 [27] Mário Pereira and António Ravara. 2021. Cameleer: A Deductive
1412 Verification Tool for OCaml. In *Computer Aided Verification*, Alexandra
1413 Silva and K. Rustan M. Leino (Eds.). Springer International Publishing,
1414 Cham, 677–689.
- 1415 [28] Benjamin C. Pierce. 2002. *Types and Programming Languages*. MIT
1416 Press.
- 1417 [29] Matthieu Sozeau. 2007. Subset Coercions in Coq. In *Types for Proofs and*
1418 *Programs*, Thorsten Altenkirch and Conor McBride (Eds.). Springer
1419 Berlin Heidelberg, Berlin, Heidelberg, 237–252.
- 1420 [30] Antal Spector-Zabusky, Joachim Breitner, Christine Rizkallah, and
1421 Stephanie Weirich. 2018. Total Haskell is Reasonable Coq. In *Pro-*
1422 *ceedings of the 7th ACM SIGPLAN International Conference on Cer-*
1423 *tified Programs and Proofs* (Los Angeles, CA, USA) (CPP 2018). As-
1424 sociation for Computing Machinery, New York, NY, USA, 14–27.
1425 <https://doi.org/10.1145/3167092>
- 1426 [31] Martin Sulzmann, Manuel M. T. Chakravarty, Simon Peyton Jones,
1427 and Kevin Donnelly. 2007. System F with Type Equality Coercions.
1428 In *Proceedings of the 2007 ACM SIGPLAN International Workshop on*
1429 *Types in Languages Design and Implementation* (Nice, Nice, France)
1430 (TLDI '07). Association for Computing Machinery, New York, NY, USA,
1431 53–66. <https://doi.org/10.1145/1190315.1190324>
- 1432 [32] Nicolas Tabareau, Éric Tanter, and Matthieu Sozeau. 2018. Equiv-
1433 alences for Free: Univalent Parametricity for Effective Transport.
1434 *Proc. ACM Program. Lang.* 2, ICFP, Article 92 (July 2018), 29 pages.
1435 <https://doi.org/10.1145/3236787>
- 1436 [33] The Coq Development Team. 2021. *The Coq Proof Assistant*. <https://doi.org/10.5281/zenodo.4501022>
- 1437 [34] The Univalent Foundations Program. 2013. *Homotopy Type Theory:*
1438 *Univalent Foundations of Mathematics*. <https://homotopytypetheory.org/>

1431	org/book, Institute for Advanced Study.	Orleans, Louisiana, USA) (<i>POPL '03</i>). Association for Computing Machinery, New York, NY, USA, 224–235. https://doi.org/10.1145/604131.604150	1486
1432	[35] Hongwei Xi, Chiyen Chen, and Gang Chen. 2003. Guarded Recursive Datatype Constructors. In <i>Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages</i> (New		1487
1433			1488
1434			1489
1435			1490
1436			1491
1437			1492
1438			1493
1439			1494
1440			1495
1441			1496
1442			1497
1443			1498
1444			1499
1445			1500
1446			1501
1447			1502
1448			1503
1449			1504
1450			1505
1451			1506
1452			1507
1453			1508
1454			1509
1455			1510
1456			1511
1457			1512
1458			1513
1459			1514
1460			1515
1461			1516
1462			1517
1463			1518
1464			1519
1465			1520
1466			1521
1467			1522
1468			1523
1469			1524
1470			1525
1471			1526
1472			1527
1473			1528
1474			1529
1475			1530
1476			1531
1477			1532
1478			1533
1479			1534
1480			1535
1481			1536
1482			1537
1483			1538
1484			1539
1485			1540