Gradually Structured Data*

STEFAN MALEWSKI, University of Chile, Chile

MICHAEL GREENBERG, Stevens Institute of Technology, USA

ÉRIC TANTER, University of Chile & IMFD, Chile

Dynamically-typed languages offer easy interaction with ad hoc data such as JSON and S-expressions; statically-typed languages offer powerful tools for working with structured data, notably *algebraic datatypes*, which are a core feature of typed languages both functional and otherwise. Gradual typing aims to reconcile dynamic and static typing smoothly. The gradual typing literature has extensively focused on the computational aspect of types, such as type safety, effects, noninterference, or parametricity, but the application of graduality to data structuring mechanisms has been much less explored. While row polymorphism and set-theoretic types have been studied in the context of gradual typing, algebraic datatypes in particular have not, which is surprising considering their wide use in practice. We develop, formalize, and prototype a novel approach to gradually structured data with algebraic datatypes. Gradually structured data bridges the gap between traditional algebraic datatypes and flexible data management mechanisms such as tagged data in dynamic languages, or polymorphic variants in OCaml. We illustrate the key ideas of gradual algebraic datatypes through the evolution of a small server application from dynamic to progressively more static checking, formalize a core functional language with gradually structured data, and establish its metatheory, including the gradual guarantees.

CCS Concepts: • **Theory of computation** → **Type structures; Program semantics.**

Additional Key Words and Phrases: gradual typing, algebraic datatypes, semi-structured data

ACM Reference Format:

Stefan Malewski, Michael Greenberg, and Éric Tanter. 2021. Gradually Structured Data. *Proc. ACM Program. Lang.* 5, OOPSLA, Article 126 (October 2021), 29 pages. <https://doi.org/10.1145/3485503>

1 INTRODUCTION

Most matters in programming can either be entirely resolved at runtime, or count on the help of static type checking to enforce basic guarantees. Data structuring mechanisms are no exception. Dynamically-typed programming languages like JavaScript, Scheme, and Python use ad-hoc, semi-structured datatypes to support a prototype-based approach to development. The archetypal example is S-expressions, thanks to which one can simply represent data as a list with a ‘tag’ symbol at its head indicating the kind of data, followed by arbitrary elements. As prototypes grow, though, it can be challenging to safely evolve semi-structured datatypes. It is very easy to miss an uncommon path and forget to add support for new parts of the datatype or update support for changed ones. Statically-typed languages excel at this kind of maintenance, pointing programmers

*S. Malewski is funded by ANID / Scholarship Program / Beca de doctorado nacional/2021 - 21210982. É. Tanter is partially funded by the ANID FONDECYT Regular Project 1190058 and the Millennium Science Initiative Program: code ICN17_002. M. Greenberg did some of this work while at Pomona College and with support from Harvard University.

Authors’ addresses: Stefan Malewski, PLEIAD Lab, Computer Science Department (DCC), University of Chile, Chile; Michael Greenberg, Stevens Institute of Technology, USA; Éric Tanter, PLEIAD Lab, Computer Science Department (DCC), and Instituto Milenio Fundamento de los Datos (IMFD), University of Chile & IMFD, Chile.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/10-ART126

<https://doi.org/10.1145/3485503>

directly to the changed cases. Indeed, static datatype definitions are more rigid but bring stronger guarantees, such as ensuring that all possible alternatives of a given datatype have been considered. In any case, statically-typed languages must also deal with semi-structured data, which can be found all over the web [Buneman 1997], in formats such as JSON, S-XML, or indeed XML itself. Dealing with semi-structured data then requires falling back on alternatives such as string manipulations or generic structures with reflection or runtime type coercions in order to access specific attributes. Notably, the OCaml programming language supports *polymorphic variants* [Garrigue 1998] as a more expressive and statically safe mechanism. Polymorphic variants are however completely detached syntactically from the standard way of structuring data, namely *algebraic datatypes*. For instance, a case expression cannot mix both algebraic datatypes and polymorphic variants. In addition, the type system machinery is also widely different. This strict separation of worlds complicates evolving code from one mechanism to the other.

This tension in data structuring mechanisms is reminiscent of the general tension between static and dynamic type checking, which has attracted a large body of work, in particular in gradual typing [Siek and Taha 2006]. Gradual typing allows for the smooth integration of static and dynamic typing, supporting both extremes as well as the continuum between them. Gradual typing has been explored for a variety of language features, such as *subtyping* [Garcia et al. 2016; Siek and Taha 2007; Takikawa et al. 2012], *references* [Herman et al. 2010; Siek et al. 2015b; Toro and Tanter 2020], *effects* [Bañados Schwerter et al. 2014, 2016], *ownership* [Sergey and Clarke 2012], *typestate* [Garcia et al. 2014; Wolff et al. 2011], *session types* [Igarashi et al. 2017b], *refinements* [Lehmann and Tanter 2017], *type inference* [Garcia and Cimini 2015; Vazou et al. 2018] *parametricity* [Ahmed et al. 2011, 2017; Igarashi et al. 2017a; New et al. 2020; Toro et al. 2019], etc.

However, the interaction of gradual typing with data structuring mechanisms has been rather scarce. Notable exceptions are *row polymorphism* [Garcia et al. 2016; Sekiyama and Igarashi 2020], *set-theoretic types* [Castagna and Lanvin 2017; Castagna et al. 2019] and *union types* [Siek and Tobin-Hochstadt 2016; Toro and Tanter 2017]. For instance, Sekiyama and Igarashi [2020] devise a language with row polymorphism and gradual records. In row types, every variant is its own type, similar to Typed Racket [Tobin-Hochstadt and Felleisen 2008]. Row types are very flexible, allowing for unknown variants, but they do not directly support the more conventional definitions of nominal datatypes with a static set of variants.

This work explores the application of gradual typing to nominal algebraic datatypes as found in many popular languages such as OCaml, Haskell, Scala, Elm, Rust, etc. We consider both closed and open datatypes as found in e.g. Scala. We show that gradualizing such a language yields a novel, expressive approach to gradually structured data: a simple type system allows for imprecise static information while keeping a single namespace for all constructors, whether or not they are statically declared. Gradual languages trade off between checking invariants statically and at runtime. When adding algebraic datatypes to a language, the guarantees we are after are that all pattern matches are complete, all constructors are statically known, and every constructor belongs to at most one datatype. As we “go gradual”, these properties are no longer statically guaranteed. Care must be taken to ensure that runtime checks allow moving programs seamlessly across the static/dynamic spectrum; that is, the resulting language must not only satisfy type soundness, but the gradual guarantees [Siek et al. 2015a] as well.

Contributions. This article presents the following contributions:

- We illustrate the use of gradually structured data in a novel language, GSD (for Gradually Structured Data), through the evolution of a web API of a simple arithmetic interpreter (Section 2).

- In addition to the standard *unknown type*, we identify the need for two novel gradual types related to datatypes: the *unknown datatype* and the *unknown open datatype*.
- We present a core statically-typed language with extensible, nominal algebraic datatypes called λ_D (Section 3) and develop its corresponding gradual language $\lambda_{D?}$ (Section 4).
- We develop the runtime semantics of $\lambda_{D?}$ using an evidence-based intermediate language, following the simplified approach to AGT proposed by Toro et al. [2019].
- We prove that $\lambda_{D?}$ satisfies the expected criteria of [Siek et al. 2015a] for gradually-typed languages (Section 4.5).
- We derive $\lambda_{D?}$ from λ_D following the Abstracting Gradual Typing (AGT) methodology [Garcia et al. 2016], thereby providing another case in favor of this systematic approach to gradual language design. Using AGT ensures that the semantics of $\lambda_{D?}$ is obtained systematically from that of λ_D , and that the gradual guarantees are easy to uphold.
- We provide an implementation of GSD, a practical language built on top of $\lambda_{D?}$, together with a number of illustrative examples, available at <https://pleiad.cl/gsd>.

In addition to briefly describing the implementation, Section 5 compares GSD and related languages. Section 6 discusses other related work, and Section 7 concludes.

2 GRADUALLY STRUCTURED DATA IN ACTION

We now illustrate the use of gradually structured data as supported in our prototype language GSD (for Gradually Structured Data). We walk through a four-step development scenario for the web API of an arithmetic interpreter in GSD. The first version is fully dynamic, the fourth is almost completely statically typed. Along the way we highlight the most important features related to gradually structured data and the GSD language.

The Basic Arithmetic Server (BAS) is a simple web API for doing arithmetic calculations. Like most contemporary web APIs, it communicates by sending and receiving JSON messages. Also like most contemporary web APIs, we implement it with a prototype-based approach: we build a core of functionality in the server, but the server evolves as its clients do. There is no fixed “protocol” up front, only conventions. As the API becomes stable, the protocol becomes more fixed. Practically speaking, this means the development starts out mostly untyped, and that developers only add static types as parts of the development stabilize.

Version 1. The first version of BAS written in GSD supports addition and subtraction (Figure 1). The server itself is the `serve` function (line 1), which takes a JSON request and returns a result (or error) as a JSON string, using `handleRequest` (line 2) to actually perform the request.

To handle a request (lines 3-5), first the API key is checked with `withValidKey` and then pattern matching determines what the request actually is. In this first version, BAS only supports addition (Plus) and subtraction (Minus). All other operations fail with a string error message (Fail). The definition of `withValidKey` (line 6) conveniently uses a direct field access (instead of pattern matching) to extract the key attribute of the given request; we omit the definition of `isValidKey`, which could be as complicated as a database call or as simple as a checksum.

The `fromJSON` function parses the JSON input to *constructed data*. Constructed data is similar to an S-expression: it consists of a constructor name as a tag, followed by zero or more arguments. In GSD, `fromJSON` expects its input to be an object with just one field. The field’s name becomes the constructor (i.e. tag) and the field’s value the constructor arguments. For example:

```
> fromJSON '{"Plus": {"key":10, "x":1, "y":2}}'
Plus { key = 10, x = 1, y = 2 }
> fromJSON '{"Sqrt":{"key":10, "x":{"Frac":{"numerator":11, "denominator":12}}}}'
Sqrt { key = 10, x = Frac { numerator = 11, denominator = 12 } }
```

```

1  serve jsonReq = toJSON (handleRequest (fromJSON jsonReq))
2  handleRequest request =
3    withValidKey request (\r ⇒ match r with
4      Plus key x y ⇒ Success {r = (x + y)}
5      Minus key x y ⇒ Success {r = (x - y)}
6      _ ⇒ Fail {r = "Error: unknown command"})
7  withValidKey r action = if isValidKey r.key -- key validation details omitted
8    then action r else Fail {r = "Error: invalid key"}

```

Fig. 1. BAS Version 1: dynamically-typed, supporting addition, subtraction, and API keys

To seamlessly support this, constructors have labeled parameters; arguments with non-primitive types use the field name as the outermost tag (e.g., `Frac`). Keeping fields and tags ensures no information is lost parsing JSON strings into data. Dually, `toJSON` serializes a data value to a JSON string. Both of these functions are built-ins of GSD, because constructor name generation is not first class.

The dynamically-typed BAS Version 1 handles API requests appropriately (assuming 10 is a valid key):

```

> serve '{"Plus": {"key":10, "x":1, "y":2}}'
'{"Success": 3}'
> serve '{"Times": {"key":10, "x":1, "y":2}}'
'{"Fail": "Error: unknown command"}'

```

BAS Version 1 uses unstructured data: there are no declared datatypes. If the programmer were to type `Foil` instead of `Fail` in the error case, there would be no static error, and `toJSON` would send a confusing message to a client.

Adding Types. GSD is gradually typed. Unannotated binders are considered to have the unknown type $?$. As usual, $?$ is the least precise type, and is consistent with any other type. Recall that in gradual typing, type precision ($\sqsubseteq$) characterizes the amount of static information conveyed by a gradual type. For instance, $\text{Int} \rightarrow \text{Int} \sqsubseteq \text{Int} \rightarrow ? \sqsubseteq ? \rightarrow ? \sqsubseteq ?$. Two types are *consistent* with each other there is a way to fill in their $?$ parts to reach an equal type, e.g. $\text{Int} \sim ?$ but $\text{Int} \not\sim \text{Int} \rightarrow ?$. Concretely, this means that a variable of type $?$ can be used in any position regardless of its expected type, and any value can flow to such a variable as well.

But adding datatypes to a language that only supports $?$ would only allow for the definition of constructors with unknown parameter types, it would not support *unclassified data*. Furthermore, having only the unknown type at hand would make it hard to precisely characterize when pattern matching can proceed: at runtime, the discriminatee must be some constructed data, no matter from which datatype. In fact, this data may be from a statically-defined algebraic datatype, or unclassified.

To precisely characterize the minimal shape of values that can be pattern-matched, GSD introduces a new gradual type, the unknown datatype $?_{\square}$. This type can be understood as the “ground type” of data values (which can be eliminated by pattern matching), just like $? \rightarrow ?$ is the ground type of functions (which can be eliminated by function application). Therefore, pattern matching is only well-typed if the discriminatee has a type consistent with $?_{\square}$: pattern matching on an expression of type $? \rightarrow ?$ or Int is a static type error. Likewise, it is a static type error to try to use a $?_{\square}$ -typed expression in another elimination form, such as a function application, or primitive operation like addition.

Additionally, GSD must assign a type to unclassified data. Using $?_{\square}$ would be too imprecise to be satisfactory. Indeed, it would allow unclassified data to be optimistically considered as part of

```

data Response = Success {x : ?} | Fail {msg : String}
open data Error
open data Request
serve : String → String
serve jsonReq = toJSON (handleRequest (fromJSON jsonReq))
handleRequest : Request → Response
handleRequest request =
  withValidKey request (\r : Request ⇒ match r with
    Plus key x y ⇒ Success (x + y)
    Minus key x y ⇒ Success (x - y)
    -           ⇒ Fail (msg CommandError))
withValidKey : Request → (Request → Response) → Response
withValidKey r action = if isValidKey r.key
  then action r else Fail (msg InvalidKeyError)
msg : Error → String
msg err = match err with CommandError ⇒ "unknown command"
                          InvalidKeyError ⇒ "invalid key"
                          -           ⇒ "unknown error"

```

Fig. 2. BAS Version 2: explicit response structure, but undetermined request and error structure

closed datatypes, which one expects to remain closed. When statically-typed code matches on an expression of a datatype with, say, two variants, the pattern match is expected to be exhaustive and not let unclassified data through. Therefore, GSD introduces another gradual type, the unknown open datatype $?_{\square}$, which is more precise than the unknown datatype $?_{\blacksquare}$, and is only consistent with *open* datatypes. Unclassified data has type $?_{\square}$ since it is considered as possibly inhabiting any *open* datatype. In summary, in GSD we have $?_{\square} \sqsubseteq ?_{\blacksquare} \sqsubseteq ?$. Also, datatypes are syntactically tagged with their openness: $D_{\square}$ refers to an open datatype, while $D_{\blacksquare}$ refers to a closed one. Given a closed datatype $D_{\blacksquare}$, we have $D_{\blacksquare} \sqsubseteq ?_{\blacksquare}$ but $D_{\blacksquare} \not\sim ?_{\square}$.

In Figure 1, matching on the r variable on line 2 is well-typed: r has type $?$, which is consistent with $?_{\blacksquare}$. At runtime, when unclassified data (of type $?_{\square}$) flows into r , the matching reduces successfully because $?_{\square}$ is more precise than $?_{\blacksquare}$. Importantly, if `fromJSON` returned a number instead of a data value, the match would fail at runtime. This is because pattern matches require the type of the discriminatee to be at least as precise as $?_{\blacksquare}$, and `Int` is not. The type of `fromJSON`, $\text{String} \rightarrow ?$, allows it to return both integers and data values.

Version 2. In BAS Version 2, the developers start to firm up some of their definitions as the application evolves (Figure 2), using GSD’s support for datatypes. First, the `Response` datatype is explicitly declared as a *closed* datatype with two alternatives, one for `Success` and one for `Failure`. A *closed* datatype is like a standard OCaml or Haskell algebraic datatype, in that all constructors are specified in place. Contrastingly, the developer remains uncertain of which requests and errors might occur, so the `Error` and `Request` datatypes are declared *open*. An *open* datatype admits the declaration of new variants after the datatype declaration.¹ With these datatype declarations in hand, the programmer can now give explicit types to most functions, such as `handleRequest`, of type `Request → Response`. BAS Version 2 has narrowed down many of its representations, even

¹Object-oriented class extensibility typically defaults to open, with keywords like `final` indicating closed types. Scala supports algebraic datatypes with *case classes*. They are open by default (i.e. new variants can be added in other files/modules), and Scala provides the keyword `sealed` for closed case classes, requiring all variants to be locally defined, as in a typical functional programming language like Haskell or OCaml.

```

data Response = Success { x : ? } | Fail { msg : String }
open data Error = CommandError | InvalidKeyError      -- Declaring possible
open data Request = Plus { key : Int, x : ?, y : ? } -- constructors
                  | Minus { key : Int, x : ?, y : ? } --
serve : String → String
serve jsonReq = toJSON (handleRequest (fromJSON jsonReq))
handleRequest : Request → Response
handleRequest request =
  withValidKey request (\r : Request ⇒ match r with
    Plus key x y ⇒ Success (x + y)
    Minus key x y ⇒ Success (x - y)
    Not key x     ⇒ Success (not x)
    _             ⇒ Fail (msg CommandError))

withValidKey : Request → (Request → Response) → Response
withValidKey r action = if isValidKey r.key
  then action r else Fail (msg InvalidKeyError)
msg : Error → String
msg err = match err with CommandError ⇒ "unknown command"
                  InvalidKeyError ⇒ "invalid key"
                  _                ⇒ "unknown error"

```

Fig. 3. BAS Version 3: partially-specified datatypes

as it leaves some open—like `Error` and `Request`. Errors are no longer mere strings, but structured data—the new `msg` function pattern matches on them to produce string-based error messages. In GSD, open datatypes can be optimistically inhabited by *any* constructed data, so it is valid to match a value of type `Error` with patterns of some unanticipated constructor names. The programmer accounts for this with a catch-all `_` case. Programmers familiar with polymorphic variants in OCaml will recognize this intermediate structuring of data. The difference at this point is that in GSD there is no polymorphic structural type inference as in OCaml, just simple gradual types. Many errors can still happen at runtime in GSD! For example, `handleRequest (Plus {key=1, x=False, y=7})` fails at runtime when `handleRequest` tries to evaluate `False + 7` and `handleRequest (Plus {key=1, x=False, y=7, z=0})` produces a `CommandError`, because there is no pattern matching a `Plus` constructor with four arguments.

Version 3. BAS Version 3 makes the `Error` and `Request` datatypes more static (Figure 3). For instance, `Plus` and `Minus` have been promoted to official constructors of the `Request` datatype. These explicit declarations fix the internal structure of these declared constructors: now `Plus {key=1, x=False, y=7, z=0}` is a *static* type error, because the constructor is used with the wrong number of arguments. Also, because `Plus` is now a declared constructor of the `Request` datatype, `Plus` is no longer a possible constructor of another open datatype: passing `Plus k x y` to `errorMsg` would yield a *static* type error.

Note that the types of the `x` and `y` fields of both `Plus` and `Minus` have type `?` at this stage. So not every such error is static: `handleRequest (Plus {key=1, x=False, y=7})` fails at runtime when trying to add `False` and `7`.

Observe that `Request` and `Error` are still declared **open**, so they can still be optimistically inhabited with unclassified data, and pattern matches on values of these types can handle arbitrary constructors. Here, for instance, the programmer has added experimental support for a `Not` operation.


```

data Response = Success {x : Data} | Fail {msg : String}
data Data = N {x : Int} | B {x : Bool}
data Error = InvalidKeyError          -- | CommandError is no longer possible
data Request = Plus {key : Int, x : Int, y : Int}
              | Minus {key : Int, x : Int, y : Int}
              | Not {key : Int, x : Bool}

serve : String → String
serve jsonReq = toJSON (handleRequest (fromJSON jsonReq))
handleRequest : Request → Response
handleRequest request =
  withValidKey request (\r : Request ⇒ match r with
    Plus key x y ⇒ Success (N (x + y))
    Minus key x y ⇒ Success (N (x - y))
    Not key x     ⇒ Success (B (not x))
    -- Fail case is ruled out

withValidKey : Request → (Request → Response) → Response
withValidKey r action = if isValidKey r.key
  then action r else Fail (msg InvalidKeyError)
msg : Error → String
msg err = match err with InvalidKeyError ⇒ "invalid key"

```

Fig. 4. BAS Version 4: fully specified datatypes

Version 4. Finally, BAS Version 4 is almost fully statically typed (Figure 4). Every datatype is **closed**, and all unknown types ? have been replaced with static types. For instance, because the type of x is statically declared to be `Int`, the expression `Plus {key=1, x=False, y=7}` is now ill-typed. Similarly, `handleRequest (Times {key=1, x=3, y=7})` now yields a static type error, because `Times` is not a valid constructor of `Request`. By closing all the datatypes, no catch-all cases are needed anymore, which eliminates some runtime errors in pattern matches: the `Error` datatype shrinks accordingly.

But every statically-typed language must eventually confront the outside world: `fromJSON` still introduces a term of type `?`. The “parse, don’t validate” approach popular in the statically-typed functional programming community suggests writing a wrapper around `fromJSON` to ensure that we only process appropriate `Requests`; we omit that development here to not belabor the point.

Summary. The table below summarizes the possible errors in each stage of the evolution of the BAS running example implemented in GSD:

	Unknown command	Invalid key	Unknown error	Invalid JSON	Unknown field	Wrong argument	Ill-formed Plus
Version 1	⚬	⚬	⚬	⚬	⚬	⚬	✓
Version 2	⚬	⚬	⚬	⚬	⚬	⊖	✓
Version 3	⚬	⚬	⚬	⚬	⚬	⊖	⊖
Version 4	⊖	⚬	⊖	⚬	⊖	⊖	⊖

The first three columns correspond to the errors defined in the datatype `Error`. The other are errors thrown by the interpreter: *Invalid JSON* is thrown if `fromJSON` is called with an invalid string, *Unknown field* when r in `withValidKey` does not have the field `key`, *Wrong argument* represents errors caused by calling a function with arguments of the wrong type, and *Ill-formed Plus* refers to the possibility of constructing `Plus` with multiple arities and labels. In the table, ✓ means “runs without errors”, ⚬ means “fails with a runtime error”, and ⊖ means “fails with a static error”.

As illustrated here, GSD is a simple language design for gradually structured data, which combines well-known open/closed datatype declarations with standard gradual types, further enriched with two new gradual types ($?_{\blacksquare}$ as the ground type of constructed data, and $?_{\square}$ as the unknown open datatype). As a result, GSD accommodates a wide range of evolution and type strengthening scenarios, all within the same frame of reference, which seems appealing in practice.

The following two sections develop the theory underlying GSD: Section 3 presents the static language we consider as a starting point, and Section 4 exposes the gradual language derived from it using the Abstracting Gradual Typing methodology (AGT) [Garcia et al. 2016].

3 ALGEBRAIC DATATYPES, STATICALLY

We begin by describing a statically-typed language λ_D : a call-by-value simply-typed lambda calculus with algebraic datatypes. In the literature, models of algebraic datatypes come primarily in two flavors: a simplified algebraic model, with pairs ((e_1, e_2) , `fst`, `snd`) and disjoint sums (`inl`, `inr`, `match`) and unit (the 0-ary or nullary product, `()`), and a general, flexibly structural model in terms of row types [Wand 1987]. We do not adopt either of these common models, opting instead for an explicit model of named datatypes with associated constructors and a general notion of pattern matching, in the spirit of GHC’s core calculus [Sulzmann et al. 2007]. We are interested in building a model that can simulate some of the dynamics and pragmatics of contemporary languages.

To this end, two key features of λ_D are worth highlighting upfront. First, λ_D supports *open* datatypes, which can be extended with new constructors. Recall that some statically-typed languages support similar features: Scala has both open and closed variants; OCaml has both extensible and polymorphic variants, though polymorphic variants are syntactically separated from standard closed datatypes. Second, λ_D ’s static semantics is parameterized over *matching strategies*, to account for variation in real languages: Haskell does not even *warn* on incomplete matches by default; OCaml lets incomplete matches off with a warning, likewise for Scala with *sealed case classes*; Elm rejects incomplete matches.

3.1 Syntax

The syntax of λ_D extends the simply-typed lambda calculus with algebraic datatypes in the style of statically-typed functional languages like Haskell and OCaml (Figure 5).² Types T are conventional, including base types B (such as `Int` or `String`), function types $T_1 \rightarrow T_2$, and named algebraic datatypes $D_{\blacksquare}$. Recall that each datatype is tagged with its openness: $D_{\square}$ denotes an open datatype D , while $D_{\blacksquare}$ denotes a closed one. Datatype contexts Δ map datatype names to *constructor sets* C , which may be empty (e.g., the uninhabited “void” type). Recall that closed datatypes have a fixed set of variants given at definition time, while open ones can be extended with additional variants. Constructor contexts Ξ map *constructor names* c to ordered, labeled products of types, i.e., records. Record labels l should be used at most once for each constructor, but can be reused across them.

The expressions of the language include the usual variables, constants k , lambda abstraction, applications, and static type ascriptions (which play a key role in the gradualization with AGT). We add three syntactic forms to support datatypes: constructor applications, field accesses, and pattern matching.

²Some notational conventions are defined for conciseness. We write $\bar{l} = \bar{e}$ to mean $l_1 = e_1, \dots, l_n = e_n$ and $\bar{p} \mapsto \bar{e}$ to mean $p_1 \mapsto e_1; \dots; p_n \mapsto e_n$. When necessary we add a superscript next to the overline to make explicit the number of elements being repeated, eg. $\bar{l}^i = \bar{e}^i$ means $l_1 = e_1, \dots, l_i = e_i$. Likewise, we write $(\bar{x} : \bar{T})$, or more explicitly $(\bar{x}^n : \bar{T}^n)$, to mean $(x_1 : T_1) \dots (x_n : T_n)$. We write $\{\bar{p}\}$, or $\{\bar{p}^n\}$, to mean the set $\{p_1, \dots, p_n\}$. In the typing judgments we also differentiate between definitional ($\doteq$) and propositional equality ($=$).

Datatypes			Types		
Datatype names	D	$\in$ DTName	Base types	$B ::= \text{Int} \mid \dots$	
Openness	$\blacksquare$	$::= \square \mid \blacksquare$	Types	$T ::= B \mid T \rightarrow T \mid D_{\blacksquare}$	
Constructor names	c	$\in$ CtorName	Terms		
Constructor sets	C	$\in$ Ctors $\subseteq$ CtorName			
Labels	l	$\in$ Label	Expressions	$e ::= x \mid k \mid e e \mid \lambda x:T. e$ $\mid e :: T \mid c \{ \bar{l} = \bar{e} \} \mid e.l$ $\mid \text{match } e \text{ with } \{ \bar{p} \mapsto \bar{e} \}$	
Contexts			Identifiers	$x \in \text{Var}$	
Datatype ctxs	Δ	$::= \cdot \mid \Delta, D_{\blacksquare} : C$	Constants	$k ::= 0 \mid 1 \mid \dots \mid + \mid \dots$	
Constructor ctxs	Ξ	$::= \cdot \mid \Xi, c : (\bar{l} : \bar{T})$	Patterns	$p ::= c \bar{x}$	
Type ctxs	Γ	$::= \cdot \mid \Gamma, x : T$	Values	$v ::= k \mid \lambda x:T. e \mid c \{ \bar{l} = \bar{v} \}$	
Programs			Errors	$\text{err} ::= \text{error}_M \mid \text{error}_A$	
Declaration	decl	$::= \text{data } D_{\blacksquare} \text{ with } \Xi$ $\mid \text{extend } D_{\square} \text{ with } \Xi$			
Programs	P	$::= \overline{\text{decl}}_i e$			

Fig. 5. λ_D syntax

The *constructor application* $c \{l_1 = e_1, \dots, l_n = e_n\}$ applies the constructor c with field l_i set to (the value of) e_i . The *field access* $e.l$ extracts field l from the data value resulting from evaluating e . Finally, a *pattern match* $\text{match } e \text{ with } \{p_1 \mapsto e_1; \dots; p_n \mapsto e_n\}$ compares the value of the discriminatee e against each pattern p_i in turn; when a pattern matches, the parts of the discriminatee are bound by p_i , and e_i is evaluated with these bindings. As a modeling compromise, we do not allow nested patterns in the formalism. However, top-level patterns are sufficiently expressive to encode nested patterns. Although having positional constructor arguments may be simpler, we adopt labeled arguments because working with semi-structured data is a key potential application of GSD. Labeled arguments support encoding JSON objects directly, without increasing the gap between GSD and the formal calculi. We include field access explicitly because it is (a) a useful feature in real languages, but (b) challenging to encode faithfully using only pattern matching in the presence of open datatypes. Finally, the outermost syntactic structures are programs. A program P is a sequence of declarations followed by an expression. There are two kinds of declarations *decl*: datatype declarations and datatype extensions. A datatype declaration $\text{data } D_{\blacksquare} \text{ with } c_1 : (\bar{l}_1 : \bar{T}_1), \dots, c_n : (\bar{l}_n : \bar{T}_n)$ extends the datatype context with the datatype $D_{\blacksquare}$ associated to the constructors $c_1, \dots, c_n$. It also extends the constructor context with each constructor c_i and its signature $(\bar{l}_i : \bar{T}_i)$. The *datatype extension* $\text{extend } D_{\square} \text{ with } c_1 : (\bar{l}_1 : \bar{T}_1), \dots, c_n : (\bar{l}_n : \bar{T}_n)$ extends the set of constructors associated with the open datatype $D_{\square}$ in the datatype context with the new constructors $c_1, \dots, c_n$. Like with datatype declarations, the constructor context is extended with every constructor and the labeled product of types associated with it.

3.2 Static Semantics

The static language λ_D enjoys a mostly conventional static semantics, with ordinary rules for type well-formedness, type equality, and term typing (Figure 7). Anticipating the use of AGT to derive the gradual language, we define the static semantics a little more abstractly than usual: side conditions in rules are explicit predicates and partial functions (Figure 8), so that we can use Galois connections to derive the gradual versions (Section 4.2). The judgments concerning types themselves are conventional, requiring that datatypes be well formed in the datatype context Δ ,

Context well-formedness

$$\begin{array}{c}
\frac{}{\vdash \cdot} \quad \Delta\text{-EMPTY} \qquad \frac{\boxed{\vdash \Delta} \quad \boxed{\Delta \vdash \Xi} \quad \boxed{\Delta; \Xi \vdash \Gamma}}{\vdash \Delta, D_{\blacksquare} : C} \quad \Delta\text{-EXT} \\
\\
\frac{\vdash \Delta}{\Delta \vdash \cdot} \quad \Xi\text{-EMPTY} \qquad \frac{\Delta \vdash \Xi \quad \forall i, j. 1 \leq i, j \leq n, i \neq j \implies l_i \neq l_j \quad \Delta \vdash \bar{T} \quad \bar{T} = \text{fty}_{\Delta, \Xi}(\bar{l}, \text{cty}_{\Delta}(c))}{\Delta \vdash \Xi, c : (\bar{l} : \bar{T})} \quad \Xi\text{-EXT} \\
\\
\frac{\Delta \vdash \Xi}{\Delta; \Xi \vdash \cdot} \quad \Gamma\text{-EMPTY} \qquad \frac{\Delta; \Xi \vdash \Gamma \quad \Delta \vdash T}{\Delta; \Xi \vdash \Gamma, x : T} \quad \Gamma\text{-EXT}
\end{array}$$

Program well-formedness $\boxed{P \text{ ok}}$

$$\overline{\text{decl}}, e \text{ ok} \Leftrightarrow \Delta \vdash \Xi \wedge \Delta; \Xi; \cdot \vdash e : T \quad \text{where } \Delta \doteq \text{data-ctx}(\overline{\text{decl}}) \text{ and } \Xi \doteq \text{ctor-ctx}(\overline{\text{decl}})$$

Context construction

$$\begin{aligned}
\text{data-ctx}(\overline{\text{decl}}) &= \text{fold}(\text{data-add-decl}, \cdot, \overline{\text{decl}}) \\
\text{ctor-ctx}(\overline{\text{decl}}) &= \text{fold}((\lambda(d, \Xi). \Xi, \text{decl-ctors}(d)), \cdot, \overline{\text{decl}}) \\
\\
\text{data-add-decl}(\text{data } D_{\blacksquare} \text{ with } \Xi, \Delta) &= \Delta, D_{\blacksquare} : \text{dom}(\Xi) \quad \text{if } \{D_{\square}, D_{\blacksquare}\} \cap \text{dom}(\Delta) = \emptyset \\
\text{data-add-decl}(\text{extend } D_{\square} \text{ with } \Xi, \Delta) &= \Delta, D_{\square} : \text{dom}(\Xi \cup \Xi') \quad \text{if } \Delta(D_{\square}) = \Xi \text{ and } \Xi \cap \Xi' = \emptyset \\
\text{data-add-decl}(_, _) &= \perp \\
\\
\text{decl-ctors}(\text{data } D_{\blacksquare} \text{ with } \Xi) &= \Xi \\
\text{decl-ctors}(\text{extend } D_{\square} \text{ with } \Xi) &= \Xi
\end{aligned}$$

Fig. 6. λ_D : well-formedness of programs and contexts, and context construction

which in turn requires that Δ be well formed. Context well-formedness is formally complicated—we build up three contexts well-formedness judgments on top of each other. Despite the formal complexity, the context rules are conceptually simple (Figure 6). First, the datatype context Δ is well formed when each datatype is assigned a disjoint set of constructors (Δ -EXT). Next, given a datatype context Δ , the constructor context Ξ is well formed (Ξ -EXT) when each constructor (a) belongs to some datatype D (cty_{Δ}), (b) has arguments with disjoint labels, (c) each label is mapped to a well formed, closed type, and (d) every other constructor in the same datatype that uses the same label uses the same type at that label. Given a datatype context Δ and constructor context Ξ , the type context Γ is well formed if the types inside of it are well formed (Γ -EXT). Finally, the judgment $P \text{ ok}$ asserts the well-formedness of the program P (Figure 6). A program is well formed when the contexts constructed from its declarations are well formed, and its top-level term is well typed. Contexts are constructed from the sequence of declarations in a program (Figure 6). A datatype context Δ is constructed by the partial function data-ctx , which maps datatypes to the constructors that appear in their datatype declarations and datatype extensions. Datatype context construction fails if the resulting context is ill-formed: either there are multiple declarations for the same datatype, or a datatype is extended before it is declared. Using the function ctor-ctx , a constructor context Ξ is built by collecting every constructor signature in the declarations.

Typing rules

		$\Delta; \Xi; \Gamma \vdash e : T$
$\frac{\Delta; \Xi \vdash \Gamma \quad T \doteq \Gamma(x)}{\Delta; \Xi; \Gamma \vdash x : T}$	T-VAR	
		$\frac{}{\Delta; \Xi; \Gamma \vdash k : \text{ty}(k)}$ T-CONST
$\frac{\Delta; \Xi; \Gamma, x : T \vdash e : T'}{\Delta; \Xi; \Gamma \vdash \lambda x:T. e : T \rightarrow T'}$	T-LAM	$\frac{\Delta; \Xi; \Gamma \vdash e_1 : T_1 \quad \Delta; \Xi; \Gamma \vdash e_2 : T_2 \quad T_2 = \text{dom}(T_1)}{\Delta; \Xi; \Gamma \vdash e_1 e_2 : \text{cod}(T_1)}$ T-APP
$\frac{\Delta; \Xi; \Gamma \vdash e : T \quad T = T'}{\Delta; \Xi; \Gamma \vdash e :: T' : T'}$	T-ASCRIBE	$\frac{T \doteq \text{cty}_{\Delta}(c) \quad \text{isdata}(T) \quad \Delta; \Xi; \Gamma \vdash \bar{e} : \bar{T} \quad \bar{T} = \text{lty}_{\Xi}(\bar{l}, c) \quad \text{satisfylab}_{n\Xi}(c, l_1 \times \dots \times l_n)}{\Delta; \Xi; \Gamma \vdash c \{ \bar{l}^n = \bar{e}^n \} : T}$ T-CTOR
$\frac{\Delta; \Xi; \Gamma \vdash e : T \quad \text{isdata}(T)}{\Delta; \Xi; \Gamma \vdash e.l : \text{fty}_{\Delta, \Xi}(l, T)}$	T-ACCESS	$\frac{\Delta; \Xi; \Gamma \vdash e : T \quad \text{isdata}(T) \quad \text{valid}_{\Delta, \Xi}(\{\bar{p}\}, T) \quad \forall i \in [1, n] \ (x_{i1} : T_{i1}) \times \dots \times (x_{im_i} : T_{im_i}) \doteq \text{parg}_{m_i \Xi}(p_i)}{\Delta; \Xi; \Gamma, x_{i1} : T_{i1}, \dots, x_{im_i} : T_{im_i} \vdash e_i : T_i}$ T-MATCH
		$\Delta; \Xi; \Gamma \vdash \text{match } e \text{ with } \{ \bar{p}^n \mapsto \bar{e}^n \} : \text{equate}_n(T_1, \dots, T_n)$

Reductions

		$e \rightarrow e \text{ or } \mathbf{err}$
(R-Beta)	$(\lambda x:T. e) v \rightarrow e [v/x]$	
(R-Delta)	$k v \rightarrow \delta(k, v)$	
(R-AscErase)	$v :: T \rightarrow v$	
(R-Match)	$\text{match } c \{ \bar{l} = \bar{v} \} \text{ with } \{ \bar{p} \mapsto \bar{e} \} \rightarrow \begin{cases} e_k [\bar{v}/\bar{x}_k] & k \text{ smallest } \mathbb{N} \text{ s.t. } c = c_k \text{ where } c_k \bar{x}_k \doteq p_k \\ \mathbf{error}_M & \text{there is no } k \text{ s.t. } c = c_k \end{cases}$	
(R-Access)	$(c \{ \bar{l} = \bar{v} \}).l \rightarrow \begin{cases} v_k & l_k = l \\ \mathbf{error}_A & \text{otherwise} \end{cases}$	

Frames

		$e \mapsto e \text{ or } \mathbf{err}$
$E ::= \square :: T \mid \square e \mid v \square \mid \square.l \mid c \{ \bar{l} = \bar{v}, l = \square, \bar{l} = \bar{e} \} \mid \text{match } \square \text{ with } \{ \bar{p} \mapsto \bar{e} \}$		
$\frac{e \rightarrow e'}{e \mapsto e'}$	R $\rightarrow$	$\frac{e \mapsto e'}{E[e] \mapsto E[e']}$ RE
$\frac{e \rightarrow \mathbf{err}}{e \mapsto \mathbf{err}}$	RERR	$\frac{e \mapsto \mathbf{err}}{E[e] \mapsto \mathbf{err}}$ RERR

Fig. 7. λ_D : typing and dynamic semantics

The term typing rules for λ_D are standard (Figure 7). The rules are written with explicit type predicates and functions (see T-App for example) in order to ease gradualization, following the recommendation of [Garcia et al. 2016]. In addition to the usual lambda calculus rules, there are rules specific to datatypes—T-CTOR, T-ACCESS and T-MATCH—as well as an ascription rule T-ASCRIBE. Constructor application is well typed (T-CTOR) when (a) its type is a datatype, (b) the labels l_i match what we know about the constructor, and (c) the subterms at each label are of the correct type for that label. A field access $e.l$ is well typed (T-ACCESS) when the type of e is a datatype with at least one constructor for which l is valid. Pattern matches are well typed (T-MATCH) when (a) the type of the term being matched is a datatype, (b) the branches have the same type in a context extended with the type of the pattern bindings, and (c) the patterns are valid with respect to the type of the

$$\begin{array}{ll}
\text{dom} : \text{Type} \rightarrow \text{Type} & \text{cty}_{\Delta} : \text{Ctors} \rightarrow \text{Type} \\
\text{dom}(T_1 \rightarrow T_2) = T_1 & \text{cty}_{\Delta}(c) = \bigcap \{D_{\blacksquare} \in \text{dom}(\Delta) \mid c \in \Delta(D_{\blacksquare})\} \\
\text{dom}(_) = \perp & \\
\\
\text{cod} : \text{Type} \rightarrow \text{Type} & \text{lty}_{\Xi} : \text{Label} \times \text{Ctors} \rightarrow \text{Type} \\
\text{cod}(T_1 \rightarrow T_2) = T_2 & \text{lty}_{\Xi}(l, c) = T_i \quad (l_1 : T_1) \times \dots \times (l_n : T_n) = \Xi(c) \wedge l = l_i \\
\text{cod}(_) = \perp & \text{lty}_{\Xi}(l, c) = \perp \quad \text{otherwise} \\
\\
\text{equate}_n : \text{Type}^n \rightarrow \text{Type} & \text{fty}_{\Delta, \Xi} : \text{Label} \times \text{Type} \rightarrow \text{Type} \\
\text{equate}_n(T, \dots, T) = T & \text{fty}_{\Delta, \Xi}(l, D_{\blacksquare}) = \bigcap \{\text{lty}_{\Xi}(l, c) \mid c \in \Delta(D_{\blacksquare})\} \\
\text{equate}_n(_, \dots, _) = \perp & \text{fty}_{\Delta, \Xi}(l, T) = \perp \\
\\
\sqcap : \text{Type} \times \text{Type} \rightarrow \text{Type} & \text{ctors} : \Delta \rightarrow \mathcal{P}(\text{Ctors}) \\
T \sqcap T = T & \text{ctors}(\Delta) = \bigcup \{\Delta(D_{\blacksquare}) \mid D_{\blacksquare} \in \text{dom}(\Delta)\} \\
_ \sqcap _ = \perp & \text{isdata}(T) \Leftrightarrow T \in \{D_{\blacksquare} \mid D_{\blacksquare} \in \text{DTName}\} \\
\\
\text{pctor} : \text{Pattern} \rightarrow \text{Ctors} & \\
\text{pctor}(c \bar{x}) = c & \\
\\
\text{parg}_{n\Xi} : \text{Pattern} \rightarrow (\text{Var} : \text{Type})^n & \\
\text{parg}_{n\Xi}(c \bar{x}^n) = \begin{cases} (x_1 : T_1) \times \dots \times (x_n : T_n) & (l_1 : T_1) \times \dots \times (l_m : T_m) \doteq \Xi(c) \wedge n = m \\ \perp & c \notin \text{dom}(\Xi) \vee n \neq m \end{cases} & \\
\\
\text{satisfylab}_{n\Xi}(c, l_1 \times \dots \times l_n) \Leftrightarrow \text{permutation}(l_1 \dots l_n, l'_1 \dots l'_m) \quad (l'_1 : T_1) \times \dots \times (l'_m : T_m) \doteq \Xi(c) & \\
\\
\text{valid}_{\Delta, \Xi}(P, D_{\blacksquare}) \Leftrightarrow \bigcup_{p \in P} \text{pctor}(p) \star \Delta(D_{\blacksquare}) \quad \text{where } \star \text{ is: } \subseteq \text{ if Sound, } = \text{ if Exact, or } \supseteq \text{ if Complete.} &
\end{array}$$

Fig. 8. λ_D : type functions and predicates for datatype-related rules

term being matched. λ_D 's static semantics accommodates different interpretations of when pattern matches are valid—that is, when a list of patterns is considered sufficiently exhaustive (see below).

The typing rules use a number of partial type functions (Figure 8): dom and cod compute the domain and codomain, respectively, of a given type (used in T-APP); equate_n computes the *meet* of a collection of types (used in T-MATCH); cty_{Δ} computes the type of a constructor (used in Ξ -EXT and T-CTOR); $\text{parg}_{n\Xi}$ computes the bindings, with their expected types, of a pattern (used in T-MATCH); $\text{fty}_{\Delta, \Xi}$ computes the type a label has for a specific type (used in Ξ -EXT and T-ACCESS); lty_{Ξ} computes the type of a label for a specific constructor (used in T-CTOR and $\text{fty}_{\Delta, \Xi}$); We also use two total functions: pctor computes the constructor name of a pattern (used in T-MATCH via $\text{valid}_{\Delta, \Xi}$; see below); ctors computes the constructor set of a datatype (used in Δ -EXT, $\text{fty}_{\Delta, \Xi}$, cty_{Δ} and $\text{valid}_{\Delta, \Xi}$); We use some type predicates beyond the type equality in T-ASCRIBE: $\text{satisfylab}_{n\Xi}$ checks that a constructor has a specific set of labels (used in T-CTOR); isdata checks if a type is a datatype (used in T-CTOR, T-ACCESS and T-MATCH); and, finally, $\text{valid}_{\Delta, \Xi}$ determines whether or not a set of patterns is sufficient (used in T-MATCH).

Matching Strategies. The language λ_D is parametric with respect to the meaning of valid pattern matches. We introduce three possible *matching strategies* that characterize valid matches ($\text{valid}_{\Delta, \Xi}$, Figure 8): sound, complete, and exact. *Sound* matching corresponds to Haskell's policy: every pattern in the match expression must correspond to at least one constructor in the datatype,

but partial matches are allowed. *Complete* matching does not allow for partial matches, but case branches can have patterns that do not match any constructor in the type. We are not aware of any statically-typed language that lets one write extra cases using impossible constructor names. Even so, allowing dead branches is harmless in the static system, and as we will see, the choice proves advantageous as we gradualize the type system. *Exact* matching is the combination of sound and complete matching: every single constructor is accounted for, with none missing and none extra. The exact regime is used for instance in Coq and Elm.

For example, using the concrete syntax of our implementation, the pattern matches below are sound exact and complete with respect to A respectively:

	<i>-- Sound</i>	<i>-- Exact</i>	<i>-- Complete</i>
<code>data A = Foo Bar</code>	<code>match x with</code>	<code>match x with</code>	<code>match x with</code>
<code>data B = Baz</code>	<code>Foo ⇒ ...</code>	<code>Foo ⇒ ...</code>	<code>Foo ⇒ ...</code>
		<code>Bar ⇒ ...</code>	<code>Bar ⇒ ...</code>
			<code>Baz ⇒ ...</code>

As we will see, AGT allows us to derive meaningful gradual variants of λ_D for each strategy.

3.3 Dynamic Semantics

The reduction rules of λ_D are mostly standard (Figure 7). We use evaluation frames with call-by-value semantics with errors; subterms are reduced from left to right and evaluation may fail. Values are the usual suspects: constants, functions, and constructors applied to values (Figure 5).

There are five reduction rules: R-Beta reduces a function application by substitution, R-Delta reduces an operator application by means of the δ function (which we assume to be adequately typed), R-AscErase drops ascriptions on values, R-Match performs pattern matching on datatype values, and R-Access performs partial field access on a constructor.

Note that field access is partial—i.e., it can fail at runtime, producing **error_A**. Field access $e.l$ fails when e reduces to a constructor without l . To illustrate this, suppose we define a datatype of integer lists ($\text{List} : \{\text{Cons}, \text{Nil}\} \in \Delta$ and $\{\text{Cons} : (hd : \text{Int}) \times (tl : \text{List}), \text{Nil} : \langle \rangle\} \subseteq \Xi$). Then $\text{Nil}.hd$ would be well typed but fail at runtime, because Nil does not have hd as a field.

Depending on which matching strategy is used to typecheck the program ($\text{valid}_{\Delta, \Xi}$, Section 3.2), pattern matching may fail at runtime. Specifically, complete and exact pattern matching are guaranteed not to fail at runtime, but sound pattern matching might, producing **error_M**.

3.4 Metatheory

The exact type safety property that λ_D enjoys depends on the choice of matching strategy. As we have just seen, in the exact and complete interpretations of $\text{valid}_{\Delta, \Xi}$, programs can only fail with errors caused by a field access error (**error_A**). But, if $\text{valid}_{\Delta, \Xi}$ is only meant to be sound, programs can also fail by a pattern match error (**error_M**). To encode this dependency, we define the function errors , which maps a matching strategy $m \in \{\text{Sound}, \text{Exact}, \text{Complete}\}$ to the set of errors that might be produced: $\text{errors}(\text{Sound}) = \{\text{error}_M, \text{error}_A\}$ and $\text{errors}(m) = \{\text{error}_A\}$ otherwise.

THEOREM 3.1 (TYPE SAFETY OF λ_D). *When using matching strategy m , if $\Delta; \Xi; \cdot \vdash e : T$ then either $e \Downarrow v$ with $\Delta; \Xi; \cdot \vdash v : T$, or $e \Downarrow \text{err}$ where $\text{err} \in \text{errors}(m)$.*

4 GRADUALLY STRUCTURED DATA

We now gradualize λ_D following the Abstracting Gradual Typing methodology (AGT) [Garcia et al. 2016], yielding the gradual language $\lambda_{D?}$. This section focuses on the statics of $\lambda_{D?}$. After a brief overview of AGT, we introduce gradual types and the induced notion of (im)precision, and then lift

the statics semantics of λ_D to account for imprecise types. We then develop the runtime semantics of λ_D ? via an intermediate language, and conclude with the metatheory of λ_D ?

4.1 Overview of AGT

The first step of the Abstracting Gradual Typing methodology (AGT) [Garcia et al. 2016] is to define the syntax of gradual types and give them meaning through a concretization function to the set of static types they denote. For instance, the unknown type represents any type, and a precise type (constructor) represents the equivalent static type (constructor). For example, $\text{Int} \rightarrow ?$ denotes the set of all function types from Int to any static type.

The key point of AGT is that once the meaning of gradual types is agreed upon, there is no space for ad hoc design in the static semantics of the language. First, the notion of precision between gradual types coincides with set inclusion of the denoted static types. Then, the abstract interpretation framework provides us with the *definitions* of type predicates and functions over gradual types, for which we can then find equivalent inductive or algorithmic *characterizations*.

In particular, a predicate on static types induces a counterpart on gradual types through *existential* lifting: two gradual types are related if there exist static types in their respective concretization that are related. For instance, consistency is the existential lifting of type equality. Lifting type functions requires an abstraction function, establishing a Galois connection: a lifted function is the abstraction of the result of applying the static function to all the denoted static types.

For obtaining the dynamic semantics of a gradual language, AGT augments a consistent judgment (such as consistency or consistent subtyping) with the *evidence* of *why* such a judgment holds. Then, reduction mimics proof reduction of the type preservation argument of the static language, combining evidences through steps of *consistent transitivity*, which either yield a more precise evidence, or fail if the evidences to combine are incompatible.³ A failure of consistent transitivity corresponds to a cast error in a traditional cast calculus.

In the rest of this section, we apply this methodology to λ_D in order to systematically derive λ_D ?, providing more details about the AGT approach as we go along.

4.2 Semantics of Gradual Types

The first step in building λ_D ? is to define a denotation of gradual types as static types; this denotation is called a *concretization*. Equipped with concretization, one can then lift static type predicates (such as equality) to operate on gradual types (yielding consistency). Then, using the dual operation of *abstraction*, static type functions (like dom , equate_n , and $\text{valid}_{\Delta, \Xi}$) can be lifted, yielding the gradual type system.

The concretization function γ gives meaning to a gradual type by producing the non-empty set of static types that it denotes (Figure 9). Gradual types are noted by the metavariable G :

$$G ::= B \mid G \rightarrow G \mid D_{\blacksquare} \mid ? \mid ?_{\blacksquare} \mid ?_{\square}$$

As usual, the unknown type denotes any static type. We deal with gradually structured data, and so new forms of gradual types are required. First, we introduce the unknown datatype $?_{\blacksquare}$, which stands for any datatype $D_{\blacksquare}$. Values of type $?_{\blacksquare}$ can be eliminated through pattern matching and field access. We say that $?_{\blacksquare}$ is the *ground type* for datatypes, both open and closed, just like $? \rightarrow ?$ is the ground type of functions, which can be eliminated through application. Additionally, to support gradually structured data, it is helpful to account for “free-floating” constructors, which do not (yet) belong to any statically defined datatype: the unknown open datatype $?_{\square}$, represents any *open* datatype.

³We refer to the evidence of a consistent judgment as a countable entity. Therefore, we use the plural *evidences*, following the accepted use in academic English [Oxford 2021], instead of writing *pieces of evidence* or *evidence objects*.

Galois Connection for Gradual Types

$$\begin{array}{ll}
\gamma : \text{GType} \rightarrow \mathcal{P}^*(\text{Type}) & \alpha : \mathcal{P}^*(\text{Type}) \rightarrow \text{GType} \\
\gamma(B) = \{B\} & \alpha(\{B\}) = B \\
\gamma(G_1 \rightarrow G_2) = \{T_1 \rightarrow T_2 \mid T_1 \in \gamma(G_1), T_2 \in \gamma(G_2)\} & \alpha(\{\overline{T_{i1} \rightarrow T_{i2}}\}) = \alpha(\{\overline{T_{i1}}\}) \rightarrow \alpha(\{\overline{T_{i2}}\}) \\
\gamma(D_{\blacksquare}) = \{D_{\blacksquare}\} & \alpha(\{D_{\blacksquare}\}) = D_{\blacksquare} \\
\gamma(?_{\square}) = \{D_{\square} \mid D \in \text{DTName}\} & \alpha(\{\overline{D_{\square i}}\}) = ?_{\square} \\
\gamma(?_{\blacksquare}) = \{D_{\blacksquare} \mid D \in \text{DTName}\} & \alpha(\{\overline{D_{\blacksquare i}}\}) = ?_{\blacksquare} \\
\gamma(?) = \text{Type} & \alpha(\{\overline{T_i}\}) = ?
\end{array}$$

Type Precision

$$\begin{array}{llll}
\frac{}{B \sqsubseteq B} \text{ P-BASE} & \frac{G_{11} \sqsubseteq G_{21} \quad G_{21} \sqsubseteq G_{22}}{G_{11} \rightarrow G_{12} \sqsubseteq G_{21} \rightarrow G_{22}} \text{ P-ARROW} & \frac{\blacksquare_1 = \blacksquare_2}{D_{\blacksquare 1} \sqsubseteq D_{\blacksquare 2}} \text{ P-DATA} & \boxed{G \sqsubseteq G} \\
\frac{}{G \sqsubseteq ?} \text{ P-?} & \frac{}{?_{\blacksquare} \sqsubseteq ?_{\blacksquare}} \text{ P-?}_{\blacksquare} & \frac{}{?_{\square} \sqsubseteq ?_{\square}} \text{ P-?}_{\square} & \\
\frac{}{D_{\blacksquare} \sqsubseteq ?_{\blacksquare}} \text{ P-?}_{\blacksquare}\text{R} & \frac{}{?_{\square} \sqsubseteq ?_{\blacksquare}} \text{ P-?}_{\square}\text{L} & \frac{}{D_{\square} \sqsubseteq ?_{\square}} \text{ P-?}_{\square}\text{R} & \\
\end{array}$$

Type Consistency

$$\begin{array}{llll}
\frac{}{B \sim B} \text{ C-BASE} & \frac{G_1 \sim G'_1 \quad G_2 \sim G'_2}{G_1 \rightarrow G_2 \sim G'_1 \rightarrow G'_2} \text{ C-ARROW} & \frac{D_{\blacksquare 1} = D_{\blacksquare 2}}{D_{\blacksquare 1} \sim D_{\blacksquare 2}} \text{ C-DATA} & \boxed{G \sim G} \\
\frac{}{?_{\square} \sim ?_{\square}} \text{ C-?}_{\square} & \frac{}{?_{\square} \sim D_{\square}} \text{ C-?}_{\square}\text{L} & \frac{}{D_{\square} \sim ?_{\square}} \text{ C-?}_{\square}\text{R} & \\
\frac{}{?_{\blacksquare} \sim ?_{\blacksquare}} \text{ C-?}_{\blacksquare} & \frac{}{?_{\blacksquare} \sim ?_{\blacksquare}} \text{ C-?}_{\blacksquare}\text{L} & \frac{}{?_{\blacksquare} \sim ?_{\blacksquare}} \text{ C-?}_{\blacksquare}\text{R} & \\
\frac{}{D_{\blacksquare} \sim ?_{\blacksquare}} \text{ C-DATAL} & & \frac{}{?_{\blacksquare} \sim D_{\blacksquare}} \text{ C-DATAR} & \\
\end{array}$$

Fig. 9. Concretization and abstraction functions, type precision and type consistency

Gradual types are drawn from static types and type constructors together with the three unknown types. As expected, a static type like `Int` or `Int  $\rightarrow$  Bool $_{\blacksquare}$` is a fully precise gradual type, which only denotes itself.

Because several type-level functions and predicates are indexed by contexts, we need to define concretization for both datatype and constructor contexts. The concretization functions (in supplementary material) take some context with gradual types and compute the set of static contexts that the gradual context represents. The concretization of a gradual datatype context produces a set of static contexts in which unclassified data (i.e., constructors not appearing statically in any datatype) are added to the constructor set of each open datatype. The concretization of a gradual constructor context is the pointwise concretization of the types in each constructor's definition. And to accommodate for unclassified data, the produced static constructor contexts are extended with arbitrary constructor definitions. Since different contexts can have different definitions for the same unclassified data, the type of the arguments of unclassified data is `?`. Hereafter, we mark static contexts with a subscript *S*, such as Δ_S , to avoid ambiguities.

Having defined the concretization of gradual types, one can derive the notion of *precision* between gradual types: a gradual type G_1 is less precise than another gradual type G_2 if its concretization is a superset of that of G_2 .

Definition 4.1 (Type Precision). $G \sqsubseteq G'$ if and only if $\gamma(G) \subseteq \gamma(G')$.

An equivalent inductive definition is given in Figure 9.

Note that for any open datatype $D_\square$, we have $D_\square \sqsubseteq ?_\square \sqsubseteq ?_\blacksquare \sqsubseteq ?$. On the other hand, $?_\square$ and $? \rightarrow ?$ are unrelated by precision.

Armed with concretization, the AGT framework gives us a direct way to lift type predicates from the static language: a predicate holds on gradual types if there *exist* types in the respective concretizations that satisfy the static predicate. For instance, the type equality judgment is lifted into type consistency. Two gradual types are said to be consistent if they have at least one static type in common in their denotations.

Definition 4.2 (Type consistency). $G \sim G' \Leftrightarrow \exists T \in \gamma(G), \exists T' \in \gamma(G'), T = T'$.

Equivalently, one can characterize consistency as $G \sim G' \Leftrightarrow \gamma(G) \cap \gamma(G') \neq \emptyset$. An equivalent inductive definition is given in Figure 9.

In addition to type equality, the statics of λ_D use two predicates specific to datatypes: `isdata` holds if a given type is a datatype, and `valid $\Delta, \Xi$` holds if a match expression is valid for a given type (recall that we consider three matching strategies, Figure 8). The consistent lifting of `isdata` is simply consistency with respect to $?_\blacksquare$: `isdata $\Delta$ (G)  $\Leftrightarrow G \sim ?_\blacksquare$` (Figure 10).

A match is valid if there exists some type in the concretization of the type of the discriminatee for which it is valid. Formally, `valid Δ, Ξ (P, G) $\Leftrightarrow \exists \Delta_S \in \gamma(\Delta), \exists \Xi_S \in \gamma(\Xi), T \in \gamma(G), \text{valid}_{\Delta, \Xi}(P, T)$. In practice this means that when matching on a term whose type is gradual, patterns must be valid for any datatype, in the case of $?_\square$ or $?_\blacksquare$, or any open datatype, for $?_\square$. Moreover, if one wants to make a match valid with respect to some open datatype (or $?_\square$), patterns must take into account unclassified data also: Exact and Complete matches require a default case. Furthermore, patterns of unclassified data do not interfere with validity in any strategy, when matching a term of an open datatype. On the other hand, if the discriminatee has a static type, every matching strategy behaves in the same way as in the static language (Exact and Complete matches cannot fail). Importantly, in contrast to the static language, no strategy can avoid matching errors. For example, given two closed datatypes A and B , each with a single constructor named A and B respectively, the expression match (A :: ?) with {B $\mapsto$ 0} both typechecks and fails at runtime for any matching strategy, because it will be valid Δ, Ξ for B . Complete matches are especially relevant in the gradual language. When the type of the discriminatee is $?_\blacksquare$, match expressions can match on constructors of multiple datatypes, enabling more dynamic programs. This is possible because unlike Sound and Exact matches, Complete matches are not restricted to patterns from a single datatype.`

Finally, to be able to lift type functions to operate on gradual types, we need to define abstraction as the counterpart of concretization. Given a set of static types, abstraction yields the most precise gradual type that denotes (at least) this set; its definition is straightforward (Figure 9). We can then establish that abstraction is both sound and optimal, yielding a Galois connection, an important property for AGT-derived languages. Equipped with abstraction, type functions such as `dom` and `equate $n$` are lifted into their consistent counterparts by first concretizing their inputs and abstracting the collection of their possible outputs. Figure 10 shows the equivalent definitions by cases. For example, consider the consistent lifting of the functions `cty $\Delta$` . The first case of `cty $\Delta$` behaves like the static function `cty $\Delta$` : constructors declared with a datatype have it as their type. The second case assigns the unknown open datatype $?_\square$ to unclassified data as long as there is an open datatype in context. If only closed datatypes have been declared, unclassified data cannot be typed and the

$$\begin{array}{ll}
\widetilde{\text{dom}} : \text{GType} \rightarrow \text{GType} & \widetilde{\text{lt}}_{\Xi} : \text{Label} \times \text{Ctors} \rightarrow \text{GType} \\
\widetilde{\text{dom}}(G_1 \rightarrow G_2) = G_1 & \widetilde{\text{lt}}_{\Xi}(l, c) = \begin{cases} ? & c \notin \text{dom}(\Xi) \\ G_i & (\bar{l} : \bar{G}) \doteq \Xi(c) \wedge l = l_i \\ \perp & (\bar{l} : \bar{G}) \doteq \Xi(c) \wedge \nexists i. l = l_i \end{cases} \\
\widetilde{\text{dom}}(?) = ? & \\
\widetilde{\text{dom}}(_) = \perp & \\
\\
\widetilde{\text{cod}} : \text{GType} \rightarrow \text{GType} & \widetilde{\text{fty}}_{\Delta, \Xi} : \text{Label} \times \text{GType} \rightarrow \text{GType} \\
\widetilde{\text{cod}}(G_1 \rightarrow G_2) = G_2 & \widetilde{\text{fty}}_{\Delta, \Xi}(l, D_{\blacksquare}) = \bigsqcap_{c \in \Delta(D_{\blacksquare})} \widetilde{\text{lt}}_{\Xi}(l, c) \\
\widetilde{\text{cod}}(?) = ? & \widetilde{\text{fty}}_{\Delta, \Xi}(l, D_{\square}) = \begin{cases} \bigsqcap_{c \in \Delta(D_{\square})} \widetilde{\text{lt}}_{\Xi}(l, c) & \text{hasLabel}_{\Delta, \Xi}(D_{\square}, l) \\ ? & \neg \text{hasLabel}_{\Delta, \Xi}(D_{\square}, l) \end{cases} \\
\widetilde{\text{cod}}(_) = \perp & \widetilde{\text{fty}}_{\Delta, \Xi}(l, ?_{\square}) = \bigsqcup_{D_{\square} \in \text{dom}(\Delta)} \widetilde{\text{fty}}_{\Delta, \Xi}(l, D_{\square}) \\
\\
\widetilde{\text{cty}}_{\Delta} : \text{Ctors} \rightarrow \text{GType} & \widetilde{\text{fty}}_{\Delta, \Xi}(l, ?_{\blacksquare}) = \bigsqcup_{D_{\blacksquare} \in \text{dom}(\Delta)} \widetilde{\text{fty}}_{\Delta, \Xi}(l, D_{\blacksquare}) \\
\widetilde{\text{cty}}_{\Delta}(c) = \begin{cases} D_{\blacksquare} & \exists D_{\blacksquare}. c \in \Delta(D_{\blacksquare}) \\ D_{\square} & \exists! D_{\square} \in \text{dom}(\Delta) \\ ?_{\square} & \exists D_{\square} \in \text{dom}(\Delta) \\ \perp & \end{cases} & \widetilde{\text{fty}}_{\Delta, \Xi}(l, ?_{\blacksquare}) = \bigsqcup_{D_{\blacksquare} \in \text{dom}(\Delta)} \widetilde{\text{fty}}_{\Delta, \Xi}(l, D_{\blacksquare}) \\
& \widetilde{\text{fty}}_{\Delta, \Xi}(l, ?_{\square}) = \widetilde{\text{fty}}_{\Delta, \Xi}(l, ?_{\blacksquare}) \\
& \widetilde{\text{fty}}_{\Delta, \Xi}(l, G) = \perp \\
\\
& \widetilde{\text{equate}}_n : \text{GType}^n \rightarrow \text{GType} \\
& \widetilde{\text{equate}}_n(G_1, \dots, G_n) = \bigsqcap \{G_1, \dots, G_n\} \\
\\
\widetilde{\text{parg}}_{g, \Xi} : \text{Pattern} \rightarrow (\text{Var} : \text{GType})^n & \\
\widetilde{\text{parg}}_{g, \Xi}(c \vec{x}^n) = \begin{cases} (x_1 : ?) \times \dots \times (x_n : ?) & c \notin \text{dom}(\Xi) \\ (x_1 : G_1) \times \dots \times (x_n : G_n) & (l_1 : G_1) \times \dots \times (l_m : G_m) \doteq \Xi(c) \wedge n = m \\ \perp & (l_1 : G_1) \times \dots \times (l_m : G_m) \doteq \Xi(c) \wedge n \neq m \end{cases} \\
\\
\text{satisfylabels}_{n, \Xi}(c, l_1 \times \dots \times l_n) \Leftrightarrow c \notin \Xi \vee \text{permutation}(l_1 \dots l_n, l'_1 \dots l'_m) \\
\text{where } (l'_1 : G_1) \times \dots \times (l'_m : G_m) \doteq \Xi(c) \\
\\
\text{valid}_{\Delta, \Xi}(P, G) \Leftrightarrow \exists \Delta_S \in \gamma(\Delta), \exists \Xi_S \in \gamma(\Xi), \exists T \in \gamma(G), \text{valid}_{\Delta, \Xi}(P, T) \\
\\
\text{hasLabel}_{\Delta, \Xi}(D_{\blacksquare}, l) \Leftrightarrow \exists c \in \Delta(D_{\blacksquare}), \exists (l', G') \in \Xi(c), l = l' \quad \widetilde{\text{isdata}}(G) \Leftrightarrow G \sim ?_{\blacksquare}
\end{array}$$

Fig. 10. $\lambda_{D?}$: consistent type functions and predicates

function fails. Likewise, $\widetilde{\text{lt}}_{\Xi}$ computes the gradual type of label for a given constructor. Every field of unclassified data is typed with the unknown type $?$. The fields of regular constructors have the type that is declared in the constructor context. And as expected, fields not present in a constructor cannot be typed. Finally, $\widetilde{\text{equate}}_n$ computes the meet of its inputs. The precision meet of two gradual types is defined as $G_1 \sqcap G_2 = \alpha(\gamma(G_1) \cap \gamma(G_2))$ [Garcia et al. 2016]. Likewise, the precision join—used only in the algorithmic characterization of $\widetilde{\text{fty}}_{\Delta, \Xi}$ —is defined as $G_1 \sqcup G_2 = \alpha(\gamma(G_1) \cup \gamma(G_2))$. We omit the rather tedious and uninformative algorithmic definitions of join and meet.

4.3 Static Semantics of $\lambda_{D?}$

Armed with the definitions of type predicates and functions systematically derived from the meaning of gradual types, the static semantics of the gradual language $\lambda_{D?}$ follow directly. The typing rules (Figure 11) mirror the rules of the static language, except that all predicates and functions are replaced by their consistent counterparts. The only interesting rules are G-Ascribe, G-Ctor, G-Access and G-Match. An ascription $e :: G$ is well-typed (G-Ascribe) when the type of e is consistent with G . Constructor application is well-typed (G-Ctor) when (a) its type is *consistent*

Typing rules

$$\begin{array}{c}
\boxed{\Delta; \Xi; \Gamma \vdash e : G} \\
\frac{\Delta; \Xi \vdash \Gamma \quad G \doteq \Gamma(x)}{\Delta; \Xi; \Gamma \vdash x : G} \text{G-VAR} \qquad \frac{}{\Delta; \Xi; \Gamma \vdash k : \text{ty}(k)} \text{G-CONST} \\
\frac{\Delta; \Xi; \Gamma, x : G \vdash e : G'}{\Delta; \Xi; \Gamma \vdash \lambda x : G. e : G \rightarrow G'} \text{G-LAM} \qquad \frac{\Delta; \Xi; \Gamma \vdash e_1 : G \quad \Delta; \Xi; \Gamma \vdash e_2 : G' \quad G' \sim \widetilde{\text{dom}}(G)}{\Delta; \Xi; \Gamma \vdash e_1 e_2 : \widetilde{\text{cod}}(G)} \text{G-APP} \\
\frac{\Delta; \Xi; \Gamma \vdash e : G \quad G \sim G'}{\Delta; \Xi; \Gamma \vdash e :: G' : G'} \text{G-ASCRIBE} \qquad \frac{G \doteq \widetilde{\text{cty}}_{\Delta}(c) \quad G \sim ?_{\blacksquare} \quad \Delta; \Xi; \Gamma \vdash \bar{e} : \bar{G} \quad \bar{G} \sim \widetilde{\text{lty}}_{\Xi}(\bar{l}, c) \quad \text{satisfylabels}_{n\Xi}(c, l_1 \times \dots \times l_n)}{\Delta; \Xi; \Gamma \vdash c \{ \bar{l}^n = \bar{e}^n \} : G} \text{G-CTOR} \\
\frac{\Delta; \Xi; \Gamma \vdash e : G \quad G \sim ?_{\blacksquare}}{\Delta; \Xi; \Gamma \vdash e.l : \widetilde{\text{fty}}_{\Delta, \Xi}(l, G)} \text{G-ACCESS} \qquad \frac{\Delta; \Xi; \Gamma \vdash e : G \quad G \sim ?_{\blacksquare} \quad \widetilde{\text{valid}}_{\Delta, \Xi}(\{ \bar{p} \}, G) \quad \forall i \in [1, n] \ (x_{i1} : G_{i1}) \times \dots \times (x_{im_i} : G_{im_i}) \doteq \text{parg}_{m_i \Xi}(p_i)}{\Delta; \Xi; \Gamma, x_{i1} : G_{i1}, \dots, x_{im_i} : G_{im_i} \vdash e_i : G_i} \text{G-MATCH} \\
\frac{}{\vdash \cdot} \Delta\text{-EMPTY} \qquad \frac{\boxed{\vdash \Delta} \quad \boxed{\Delta \vdash \Xi} \quad \boxed{\Delta; \Xi \vdash \Gamma} \quad \vdash \Delta \quad \{D_{\blacksquare}, D_{\square}\} \cap \text{dom}(\Delta) = \emptyset \quad \text{ctors}(\Delta) \cap C = \emptyset}{\vdash \Delta, D_{\blacksquare} : C} \Delta\text{-GEXT} \\
\frac{\vdash \Delta}{\Delta \vdash \cdot} \Xi\text{-EMPTY} \qquad \frac{\Delta \vdash \Xi \quad \forall i, j. 1 \leq i, j \leq n, i \neq j \implies l_i \neq l_j \quad \Delta \vdash \bar{G} \quad \bar{G} \sim \widetilde{\text{fty}}_{\Delta, \Xi}(\bar{l}, \widetilde{\text{cty}}_{\Delta}(c))}{\Delta \vdash \Xi, c : (\bar{l} : \bar{G})} \Xi\text{-GEXT} \\
\frac{\Delta \vdash \Xi}{\Delta; \Xi \vdash \cdot} \Gamma\text{-EMPTY} \qquad \frac{\Delta; \Xi \vdash \Gamma \quad \Delta; \Xi \vdash G}{\Delta; \Xi \vdash \Gamma, x : G} \Gamma\text{-GEXT}
\end{array}$$

Context well-formedness

$$\begin{array}{c}
\boxed{\vdash \Delta} \qquad \boxed{\Delta \vdash \Xi} \qquad \boxed{\Delta; \Xi \vdash \Gamma} \\
\frac{}{\vdash \cdot} \Delta\text{-EMPTY} \qquad \frac{\vdash \Delta \quad \{D_{\blacksquare}, D_{\square}\} \cap \text{dom}(\Delta) = \emptyset \quad \text{ctors}(\Delta) \cap C = \emptyset}{\vdash \Delta, D_{\blacksquare} : C} \Delta\text{-GEXT} \\
\frac{\vdash \Delta}{\Delta \vdash \cdot} \Xi\text{-EMPTY} \qquad \frac{\Delta \vdash \Xi \quad \forall i, j. 1 \leq i, j \leq n, i \neq j \implies l_i \neq l_j \quad \Delta \vdash \bar{G} \quad \bar{G} \sim \widetilde{\text{fty}}_{\Delta, \Xi}(\bar{l}, \widetilde{\text{cty}}_{\Delta}(c))}{\Delta \vdash \Xi, c : (\bar{l} : \bar{G})} \Xi\text{-GEXT} \\
\frac{\Delta \vdash \Xi}{\Delta; \Xi \vdash \cdot} \Gamma\text{-EMPTY} \qquad \frac{\Delta; \Xi \vdash \Gamma \quad \Delta; \Xi \vdash G}{\Delta; \Xi \vdash \Gamma, x : G} \Gamma\text{-GEXT}
\end{array}$$

Program well-formedness

$$\boxed{\text{P ok}} \\
\overline{\text{decl}}, e \text{ ok} \Leftrightarrow \Delta \vdash \Xi \wedge \Delta; \Xi; \cdot \vdash e : G \quad \text{where } \Delta \doteq \text{data-ctx}(\overline{\text{decl}}) \text{ and } \Xi \doteq \text{ctor-ctx}(\overline{\text{decl}})$$

Fig. 11. λ_D : static semantics

with a datatype, (b) the labels l_i match what we know about the constructor, and (c) the subterms at each label have a type that is consistent with the correct type for that label. A field access $e.l$ is well-typed (G-Access) when the type of e is consistent with a datatype with at least one constructor for which l is valid.

Pattern matches are well-typed (G-Match) when (a) the type of the term being matched is consistent with a datatype, (b) the types of the branches in a context extended with the type of the pattern bindings are consistent with each other, and (c) the patterns are valid with respect to the type of the term being matched.

Notions of well-formedness for contexts and programs follow directly from the static definitions, using lifted functions and predicates to deal with gradual types (Figure 11). Context construction is unchanged from λ_D .

4.4 Dynamic Semantics of $\lambda_{D?}$

In order to define the dynamic semantics of $\lambda_{D?}$, we follow the approach of [Toro et al. \[2019\]](#) and introduce an auxiliary language $\lambda_{D?}^\epsilon$, which is just a variant of $\lambda_{D?}$ in which every ascription carries the *evidence* that supports its validity. Evaluation of $\lambda_{D?}$ consists of an elaboration to $\lambda_{D?}^\epsilon$ (akin to a cast insertion translation [[Siek and Taha 2006](#)], but following the AGT framework), and then the actual reduction of the elaborated $\lambda_{D?}^\epsilon$ term. We first explain the syntax, static and dynamic semantics of $\lambda_{D?}^\epsilon$, and then come back to the evaluation of $\lambda_{D?}$ and complete its metatheory.

Syntax and Static Semantics of $\lambda_{D?}^\epsilon$. The syntax of $\lambda_{D?}^\epsilon$ differs from that of $\lambda_{D?}$ in the introduction of *evidence* in ascriptions and values (Figure 12). Intuitively, evidences justify consistent judgments locally, and during reduction, evidences are combined through a partial operation called *consistent transitivity*. If the combination is successful, reduction proceeds; otherwise a runtime type error is reported, denoting the clash between two mutually-incompatible local justifications.

The metavariable ϵ ranges over evidences: in a language where the consistency notion is symmetric, evidence can be represented simply by one gradual type.⁴ For instance, the evidence that $? \rightarrow \text{Int}$ is consistent with $\text{Bool}_\blacksquare \rightarrow ?$ is the gradual type $\text{Bool}_\blacksquare \rightarrow \text{Int}$, i.e. the *meet* (relative to precision) of both types. Evidence-augmented consistency judgments are written $\epsilon \Vdash G \sim G'$.

An ascription $\epsilon e :: G$ carries the evidence that the actual type of e is consistent with G . A value v in $\lambda_{D?}^\epsilon$ is a *raw value* u ascribed to a gradual type, together with the supporting evidence. For convenience, we write ϵ_G to denote the obvious reflexive evidence that G is consistent with itself (i.e. $\epsilon_G = G$).

Finally, $\lambda_{D?}^\epsilon$ introduces a new kind of runtime errors, **error_T**, which correspond to runtime type errors, witnessed when reduction requires combining incompatible evidences (more below).

The typing rules of $\lambda_{D?}^\epsilon$ are almost identical to those of $\lambda_{D?}$ (Figure 12). In the ascription rule, consistency is replaced by evidence-augmented type consistency; observe how the evidence supporting the consistency judgment is held in the ascription term itself. This is the key “runtime tracking” mechanism of AGT. Additionally, the use of consistency in other rules is replaced by type precision when the consistent judgment has one fixed gradual type. For example, in $\epsilon G\text{-Access}$, $G \sim ?_\blacksquare$ is replaced by $G \sqsubseteq ?_\blacksquare$.

Reduction Semantics. The reduction semantics of $\lambda_{D?}^\epsilon$ is described in Figure 12. The ascription rule $\epsilon R\text{-AscErase}$, standard in evidence-based reduction semantics, is key to understand the mechanisms of runtime type checking in this technical setting. The rule describes how an ascription surrounding a value reduces to a single value if the two evidences can be combined through the *consistent transitivity* operator. Consistent transitivity is the key runtime operator in evidence-based semantics. It is the runtime type tracking mechanism, playing the dual role of type tags and casts in other presentations of gradual languages; likewise, a failure of consistent transitivity corresponds to a cast error [[Garcia et al. 2016](#)].

In our context where evidences are just gradual types, the general definition of consistent transitivity [[Garcia et al. 2016](#)] boils down to $\epsilon_{G_1} \circ \epsilon_{G_2} = \alpha(\gamma(G_1) \cap \gamma(G_2))$. Note that this definition coincides with the precision meet between gradual types introduced earlier: $\epsilon_{G_1} \circ \epsilon_{G_2} = G_1 \sqcap G_2$.

In rule $\epsilon R\text{-AscErase}$, ϵ_1 justifies that G_u , the type of the raw value, is consistent with G_1 , while ϵ_2 justifies $G_1 \sim G_2$. Composition of these evidences via consistent transitivity, if defined, justifies that $G_u \sim G_2$. If consistent transitivity is undefined, the reduction steps to **error_T**. For example, since $\text{Int} \sqcap \text{Bool}_\blacksquare$ is not defined: $\epsilon_{\text{Bool}_\blacksquare} (\epsilon_{\text{Int}} (\epsilon_{\text{Int}} 1 :: \text{Int}) :: ?) :: \text{Bool}_\blacksquare \longrightarrow \epsilon_{\text{Bool}_\blacksquare} (\epsilon_{\text{Int}} 1 :: ?) :: \text{Bool}_\blacksquare \longrightarrow \text{error}_T$. Crucially, consistent transitivity ensures that unclassified data does not “infiltrate” a closed

⁴In a language with (consistent) subtyping, evidences are typically represented by a pair of gradual types [[Garcia et al. 2016](#)], likewise in languages where the consistency relation is oriented [[Toro et al. 2019](#)].

<i>Evidence</i>	$\varepsilon \in \text{GType}$	<i>Values</i>	$v ::= \varepsilon u :: G$
<i>Expressions</i>	$e ::= v \mid x \mid e e \mid \varepsilon e :: G \mid c \{\bar{l} = \bar{v}\} \mid \text{match } e \text{ with } \{\bar{p} \mapsto \bar{e}\}$	<i>Raw Values</i>	$u ::= k \mid \lambda x:G. e \mid c \{\bar{l} = \bar{v}\}$
		<i>Errors</i>	$\text{err} ::= \text{error}_M \mid \text{error}_A \mid \text{error}_T$
Typing rules		$\Delta; \Xi; \Gamma \vdash e : G$	
$\frac{\Delta; \Xi \vdash \Gamma \quad G \doteq \Gamma(x)}{\Delta; \Xi; \Gamma \vdash x : G} \varepsilon\text{G-VAR}$		$\frac{}{\Delta; \Xi; \Gamma \vdash k : \text{ty}(k)} \varepsilon\text{G-CONST}$	
$\frac{\Delta; \Xi; \Gamma, x : G \vdash e : G'}{\Delta; \Xi; \Gamma \vdash \lambda x:G. e : G \rightarrow G'} \varepsilon\text{G-LAM}$		$\frac{\Delta; \Xi; \Gamma \vdash e_1 : G \quad \Delta; \Xi; \Gamma \vdash e_2 : \widetilde{\text{dom}}(G)}{\Delta; \Xi; \Gamma \vdash e_1 e_2 : \widetilde{\text{cod}}(G)} \varepsilon\text{G-APP}$	
$\frac{\Delta; \Xi; \Gamma \vdash e : G \quad \varepsilon \Vdash G \sim G'}{\Delta; \Xi; \Gamma \vdash \varepsilon e :: G' : G'} \varepsilon\text{G-ASCRIIBE}$		$\frac{\Delta; \Xi; \Gamma \vdash \bar{e} : \widetilde{\text{ltv}}_{\Xi}(\bar{l}, c) \quad \text{satisfylabels}_{n, \Xi}(c, l_1 \times \dots \times l_n)}{\Delta; \Xi; \Gamma \vdash c \{\bar{l}^n = \bar{e}^n\} : \widetilde{\text{ctv}}_{\Delta}(c)} \varepsilon\text{G-CTOR}$	
$\frac{\Delta; \Xi; \Gamma \vdash e : G \quad G \sqsubseteq ?_{\blacksquare}}{\Delta; \Xi; \Gamma \vdash e.l : \widetilde{\text{fty}}_{\Delta, \Xi}(l, G)} \varepsilon\text{G-ACCESS}$		$\frac{\Delta; \Xi; \Gamma \vdash e : G' \quad G' \sqsubseteq ?_{\blacksquare} \quad \widetilde{\text{valid}}_{\Delta, \Xi}(\{\bar{p}\}, G') \quad \forall i \in [1, n] (x_{i1} : G_{i1}) \times \dots \times (x_{imi} : G_{imi}) \doteq \text{parg}_{m_i, \Xi}(p_i)}{\Delta; \Xi; \Gamma, x_{i1} : G_{i1}, \dots, x_{imi} : G_{imi} \vdash e_i : G} \varepsilon\text{G-MATCH}$	
$\frac{}{\Delta; \Xi; \Gamma \vdash \text{match } e \text{ with } \{\bar{p} \mapsto \bar{e}\} : G} \varepsilon\text{G-MATCH}$			
Reductions		$e \longrightarrow e \text{ or } \text{err}$	
$(\varepsilon\text{R-Beta}) \quad (\varepsilon_1 (\lambda x:G_{11}. e) :: G_1 \rightarrow G_2) (\varepsilon_2 u :: G_1) \longrightarrow$		$\begin{cases} \text{cod}(\varepsilon_1)(e[(\varepsilon_2 \circ \text{dom}(\varepsilon_1)) u :: G_{11}/x]) :: G_2 \\ \text{error}_T & \text{if } \varepsilon_2 \circ \text{dom}(\varepsilon_1) \text{ is not defined} \end{cases}$	
$(\varepsilon\text{R-Delta}) \quad (\varepsilon_1 k :: G_1 \rightarrow G_2) (\varepsilon_2 u :: G_1) \longrightarrow$		$\text{cod}(\varepsilon_1) \delta(k, u) :: G_2$	
$(\varepsilon\text{R-AscErase}) \quad \varepsilon_2 (\varepsilon_1 u :: G_1) :: G_2 \longrightarrow$		$\begin{cases} (\varepsilon_1 \circ \varepsilon_2) u :: G_2 \\ \text{error}_T & \text{if } \varepsilon_1 \circ \varepsilon_2 \text{ is not defined} \end{cases}$	
$(\varepsilon\text{R-Match}) \quad \text{match } \varepsilon c \{\bar{l} = \bar{v}\} :: G \text{ with } \{\bar{p} \mapsto \bar{e}\} \longrightarrow$		$\begin{cases} e_k[\bar{v}/\bar{x}_k] & k \text{ smallest } \mathbb{N} \text{ s.t. } c = c_k \\ & \text{where } c_k \bar{x}_k \doteq p_k \\ \text{error}_M & \text{there is no } k \text{ s.t. } c = c_k \end{cases}$	
$(\varepsilon\text{R-Access}) \quad (\varepsilon c \{\bar{l} = \bar{v}\} :: G).l \longrightarrow$		$\begin{cases} (\varepsilon_k \circ \varepsilon_{G'}) u_k :: G' & l_k = l \\ & \text{where } \varepsilon_k u_k :: G_k \doteq v_k \\ & \text{and } G' \doteq \widetilde{\text{fty}}_{\Delta, \Xi}(l_k, G) \\ \text{error}_T & \text{if } \varepsilon_k \circ \varepsilon_{G'} \text{ is not defined} \\ \text{error}_A & \text{otherwise} \end{cases}$	
Frames		$e \mapsto e \text{ or } \text{err}$	
$E ::= \varepsilon \square :: G \mid \square e \mid v \square \mid \square.l \mid c \{\bar{l} = \bar{v}, l = \square, \bar{l} = \bar{e}\} \mid \text{match } \square \text{ with } \{\bar{p} \mapsto \bar{e}\}$			
$\frac{e \longrightarrow e'}{e \mapsto e'} \varepsilon\text{R} \longrightarrow$		$\frac{e \mapsto e'}{E[e] \mapsto E[e']} \varepsilon\text{RE}$	
$\frac{e \longrightarrow \text{err}}{e \mapsto \text{err}} \varepsilon\text{RERR}$		$\frac{e \mapsto \text{err}}{E[e] \mapsto \text{err}} \varepsilon\text{RERR}$	

Fig. 12. $\lambda_{D?}^{\varepsilon}$: Syntax, typing and dynamic semantics

datatype at runtime, because $\varepsilon_{D_{\blacksquare}} \circ \varepsilon_{?_{\square}} = D_{\blacksquare} \sqcap ?_{\square}$ is only defined if $D_{\blacksquare}$ is an *open* datatype (in which case it is equal to $D_{\square}$).

Rule ε R-BETA reduces a function application as usual if consistent transitivity between the evidence of the argument and the domain of the function's evidence is defined; otherwise it steps to **error_T**. Rule ε R-Delta takes the raw value from the ascription, performs the primitive operation and returns the resulting raw value wrapped in an ascription, using the codomain of the operation's evidence.

The novel rules for dealing with datatypes are derived similarly. Rule ε R-Match deals with pattern matching: it performs an ordinary pattern match and binding on the underlying raw value, selecting the first matching clause. Thanks to the translation function, the evidences of the discriminatee's arguments will be at least as precise as the type expected by the pattern. A consistent transitivity check is not necessary. Rule ε R-Access reduces successfully only if consistent transitivity between the evidence of the value and that of the expected type of the access is defined; otherwise it steps to **error_T**.

Evaluation of $\lambda_{D?}$. As mentioned earlier, the evaluation of a $\lambda_{D?}$ term first elaborates the term to $\lambda_{D?}^{\varepsilon}$, written $\Delta; \Xi; \cdot \vdash e \rightsquigarrow e_{\varepsilon} : G$, and then reduces this internal term. For a $\lambda_{D?}$ term e , we write $\Delta; \Xi \vdash e \Downarrow v$ (resp. $e \Downarrow \mathbf{err}$) if $\Delta; \Xi; \cdot \vdash e \rightsquigarrow e_{\varepsilon} : G$ and $e_{\varepsilon} \mapsto^* v$ (resp. $e_{\varepsilon} \mapsto^* \mathbf{err}$). We often write $e \Downarrow v$ instead of $\Delta; \Xi \vdash e \Downarrow v$ for brevity. The elaboration of $\lambda_{D?}$ terms to $\lambda_{D?}^{\varepsilon}$ terms, provided in supplementary material, is straightforward [Toro and Tanter 2020], and very similar to a standard cast insertion translation [Siek and Taha 2006]. It inserts ascriptions on every raw value and on every term whose typing derivation relies on a consistent judgment, synthesizing the corresponding evidence that supports the ascription.

To illustrate, evaluating term $(\text{Foo } \{x = 2\}).x + 1$, where Foo is unclassified data, goes as follows:

$$\begin{aligned}
 & (\text{Foo } \{x = 2\}).x + 1 \\
 \rightsquigarrow & (\varepsilon_{\text{Int}} (\varepsilon_{?_{\square}} \text{Foo } \{x = \varepsilon_{\text{Int}} (\varepsilon_{\text{Int}} 2 :: \text{Int}) :: ?\} :: ?_{\square}).x :: \text{Int}) + (\varepsilon_{\text{Int}} 1 :: \text{Int}) \\
 \longrightarrow & (\varepsilon_{\text{Int}} (\varepsilon_{?_{\square}} \text{Foo } \{x = \varepsilon_{\text{Int}} 2 :: ?\} :: ?_{\square}).x :: \text{Int}) + (\varepsilon_{\text{Int}} 1 :: \text{Int}) & \varepsilon_{\text{Int}} \circ \varepsilon_{\text{Int}} = \varepsilon_{\text{Int}} \\
 \longrightarrow & (\varepsilon_{\text{Int}} (\varepsilon_{\text{Int}} 2 :: ?) :: \text{Int}) + (\varepsilon_{\text{Int}} 1 :: \text{Int}) & \varepsilon_{\text{Int}} \circ \varepsilon_{?} = \varepsilon_{\text{Int}} \\
 \longrightarrow & (\varepsilon_{\text{Int}} 2 :: \text{Int}) + (\varepsilon_{\text{Int}} 1 :: \text{Int}) & \varepsilon_{\text{Int}} \circ \varepsilon_{\text{Int}} = \varepsilon_{\text{Int}} \\
 \longrightarrow & \varepsilon_{\text{Int}} 3 :: \text{Int}
 \end{aligned}$$

As another example, the transformation and evaluation of $(\text{Foo } \{x = 2\}).x :: ?_{\blacksquare}$ goes as follows:

$$\begin{aligned}
 & (\text{Foo } \{x = 2\}).x :: ?_{\blacksquare} \\
 \rightsquigarrow & \varepsilon_{?_{\blacksquare}} (\varepsilon_{?_{\square}} \text{Foo } \{x = \varepsilon_{\text{Int}} (\varepsilon_{\text{Int}} 2 :: \text{Int}) :: ?\} :: ?_{\square}).x :: ?_{\blacksquare} \\
 \longrightarrow & \varepsilon_{?_{\blacksquare}} (\varepsilon_{?_{\square}} \text{Foo } \{x = \varepsilon_{\text{Int}} 2 :: ?\} :: ?_{\square}).x :: ?_{\blacksquare} & \varepsilon_{\text{Int}} \circ \varepsilon_{\text{Int}} = \varepsilon_{\text{Int}} \\
 \longrightarrow & \varepsilon_{?_{\blacksquare}} (\varepsilon_{\text{Int}} 2 :: ?) :: ?_{\blacksquare} & \varepsilon_{\text{Int}} \circ \varepsilon_{?} = \varepsilon_{\text{Int}} \\
 \longrightarrow & \mathbf{error}_{\text{T}} & \varepsilon_{\text{Int}} \circ \varepsilon_{?_{\blacksquare}} \text{ is not defined}
 \end{aligned}$$

4.5 Metatheory of $\lambda_{D?}$

We now turn to the expected properties of $\lambda_{D?}$ [Siek et al. 2015a]. First, $\lambda_{D?}$ is type safe. Unlike type safety of λ_D , in general in $\lambda_{D?}$ programs can fail with any error **err**, irrespective of the selected matching strategy.

THEOREM 4.3 (TYPE SAFETY). *If $\Delta; \Xi; \cdot \vdash e : G$, then either $e \Downarrow v$ with $\Delta; \Xi; \cdot \vdash v : G$, or $e \Downarrow \mathbf{err}$.*

The $\lambda_{D?}$ type system is equivalent to the λ_D type system on *static programs*. Specifically, we write P static for a well-formed program P that uses neither imprecise types nor unclassified data. Let $\vdash_S$ denote the typing judgment of λ_D .

THEOREM 4.4 (STATIC EQUIVALENCE FOR STATIC PROGRAMS). *Let $P \doteq \overline{\text{decl}}, e_S$ s.t. P static, $\Delta_S \doteq \text{data-ctx}(\overline{\text{decl}})$ and $\Xi_S \doteq \text{ctor-ctx}(\overline{\text{decl}})$. Then $\Delta_S; \Xi_S; \cdot \vdash_S e_S : T \Leftrightarrow \Delta_S; \Xi_S; \cdot \vdash e_S : T$.*

Additionally, for static programs, the dynamic semantics of $\lambda_{D?}$ is equivalent to that of λ_D . Let $\Downarrow_S$ denote the evaluation judgment of λ_D .

THEOREM 4.5 (DYNAMIC EQUIVALENCE FOR STATIC PROGRAMS). *Let $P \doteq \overline{\text{decl}}, e_S$ s.t. P static, $\Delta_S \doteq \text{data-ctx}(\overline{\text{decl}})$ and $\Xi_S \doteq \text{ctor-ctx}(\overline{\text{decl}})$. Then $\Delta_S; \Xi_S \vdash e_S \Downarrow_S u \Leftrightarrow \Delta_S; \Xi_S \vdash e_S \Downarrow \textcolor{violet}{\epsilon}T u :: T$.*

The equivalence is proven by weak bisimulation relating λ_D terms with $\lambda_{D?}^\epsilon$ terms that potentially include ascriptions in sub-expressions.

Finally, $\lambda_{D?}$ satisfies the gradual guarantees [Siek et al. 2015a]. First, a well-typed program remains well-typed when made less precise. (Type precision is naturally extended to terms and contexts.)

THEOREM 4.6 (STATIC GRADUAL GUARANTEE). *If $\Delta; \Xi; \cdot \vdash e : G$, $\Delta \sqsubseteq \Delta'$, $\Xi \sqsubseteq \Xi'$, and $e \sqsubseteq e'$, then $\Delta'; \Xi'; \cdot \vdash e' : G'$ where $G \sqsubseteq G'$.*

Second, a program that runs without errors still does when made less precise.

THEOREM 4.7 (DYNAMIC GRADUAL GUARANTEE). *Suppose $\Delta; \Xi; \cdot \vdash e : G$ and $\Delta'; \Xi'; \cdot \vdash e' : G'$ with $\Delta \sqsubseteq \Delta'$, $\Xi \sqsubseteq \Xi'$, and $e \sqsubseteq e'$. If $e \Downarrow v$ then $e' \Downarrow v'$ where $v \sqsubseteq v'$.*

5 IMPLEMENTATION AND COMPARISON TO EXISTING LANGUAGES

We support the GSD language (Section 2) with core calculi, but core calculi alone do not make a language! In addition to the `toJSON` and `fromJSON` builtins, our GSD examples make use of a variety of extensions to pattern matching. Such extensions are critical enablers of a variety of dynamic idioms, as seen in `flatten` [Fagan 1991; Greenberg 2019]. This section first discusses the proof-of-concept implementation of GSD, and then compares GSD to existing approaches.

5.1 Implementation

We provide a reference interpreter for GSD, written in Haskell. An online demo of the interpreter can be found at <https://pleiad.cl/gsd>. The goal of the interpreter is not to be a practical and efficient implementation of GSD, but rather a reference implementation and exploration tool for further research. In particular, the interpreter outputs an interactive execution trace, which can display evidences and ascriptions from the intermediate language.

The interpreter implementation is very close to the formal calculi described in this article. GSD is composed of two languages: `GSDCore` and `GSDEv`, corresponding to $\lambda_{D?}$ and $\lambda_{D?}^\epsilon$, respectively. Interpretation consists of four main phases: parsing and desugaring source code into `GSDCore`, type checking `GSDCore`, elaborating `GSDCore` to `GSDEv`, and reducing `GSDEv`. The implementations of the typechecker and elaborator directly follow the corresponding rules for $\lambda_{D?}$. In order to collect step-by-step reduction traces, the final reduction phase is implemented as a CEK machine [Felleisen and Friedman 1986], closely following the reduction rules for $\lambda_{D?}^\epsilon$. The Galois connection used to derive the various elements of the gradual language is not used directly in any part of the interpreter; instead, consistent predicates and functions are implemented using their equivalent inductive characterizations.

The performance of gradually-typed languages is a research area in itself [Bauman et al. 2017; Campora et al. 2018; Greenman et al. 2019; Herman et al. 2010; Kuhlenschmidt et al. 2019; Siek and Wadler 2010; Siek et al. 2015b; Takikawa et al. 2016], and we have so far not explored how to best implement gradually-structured data. First of all, the dynamic semantics derived with AGT as presented here are not space efficient: evidence transitivity checks accumulate in tail calls and

	Typed	Declared	Unclassified	Default	Values	Arity	Reflect	Multiple	Uniform	Exhaustive	? _□
Scheme	D	struct	✓	✓	✓	✓	✓	N/A	×	×	✓
Typed Racket	G	ctor, sets	✓	✓	✓	✓	✓	✓	×	declared	✓
CDuce	G	sets	✓	✓	✓	✓	✓	✓	✓	✓	✓
OCaml	S	types	✓	✓	✓	×	×	variants	×	declared	×
Scala	S	✓	sets	✓	✓	×	✓	✓	N/A	sealed	×
Rust	S	✓	×	✓	✓	✓	×	×	N/A	✓	×
Elm	S	✓	×	✓	✓	×	×	×	N/A	✓	×
GSD	G	✓	✓	✓	NYI	NYI	NYI	? _□	✓	declared	✓

Fig. 13. Language feature comparison. Languages are loosely ordered from dynamic to static, with GSD last. N/A means “not applicable”; NYI means “not yet implemented”. Typed: *Dynamic*, *Static*, or *Gradual* types. Declared: statically declared datatypes; Scheme declares structs but not symbols in S-expressions; OCaml’s polymorphic variants are ad-hoc; others declare only constructors (ctor) or sets of types (sets). Unclassified: ad hoc constructors, like S-expressions or polymorphic variants. Default: catch-all cases. Values: match on non-constructors. Arity: match on some prefix of positional arguments or selection of record fields. Reflect: match on type of value. Multiple: constructors can belong to more than one declared type. Uniform: no distinction between declared and ad hoc constructors. Exhaustive: static checking of exhaustiveness. ?_□: type characterizing data.

are only resolved after a call returns. Recent work shows how to achieve space efficiency with AGT [Bañados Schwerter et al. 2021; Toro and Tanter 2020], and the technique directly applies to the semantics presented here. The novel implementation challenge introduced by GSD is the representation of datatypes. Static datatypes can often be given very efficient representations, but open datatypes and dynamically-generated constructors will require more information at runtime.

Additionally, while the decision of having a fixed global context simplifies the formalization, it is a strong assumption for a real programming language. For instance, implementing separate compilation would be challenging.

5.2 Comparison to Existing Languages

We now review the broad range of approaches to pattern matching that already exist (Figure 13). A brief disclaimer: we bias the listed features to those relevant to GSD; a full of analysis of data structuring mechanisms would be a serious endeavor in itself.

Scheme. Scheme is the mainstream programming language most closely resembling the lambda calculus. Scheme lacks algebraic datatype definitions. The language comes with lists and a notion of symbol sufficient to encode tagged data: S-expressions are lists where the first element is a symbol (representing the name of a constructor) and possibly more values (representing arguments to that constructor). Extensions add features like record types [Kelsey 1999] and pattern matching [Godek 2020]. Even with these extensions, the standard approach for semi-structured data in Scheme is to use S-expressions.

Typed Racket. Typed Racket is a dynamic-first gradual type system for Racket, a Scheme-like language [Tobin-Hochstadt et al. 2014]. Typed Racket’s static checking uses *occurrence typing* to statically guarantee type safety and exhaustive checking [Tobin-Hochstadt and Felleisen 2008]. The type system is based on union types, which accommodate unclassified data as in Scheme, i.e., tagged data as S-expressions. Typed Racket encodes open datatypes using union types and the type `Any`, comparable to `?`. Constructors and datatypes are declared separately. Every constructor is

in some sense a singleton type, and there is no restriction on how many types a constructor can belong to.

CDuce. CDuce is a pure functional programming language with set-theoretic types [Benzaken et al. 2003], which excels at working with unstructured data. CDuce makes no distinction between defined and undefined constructors. In fact, there is no way to statically define constructors at all! Constructors are encoded as tagged data like in Scheme. As in Typed Racket, each constructor has its own singleton type, and it can belong to multiple named datatypes. CDuce allows the same constructor name to be used with different argument types, and programmers can match against types. CDuce’s set-theoretic types go well beyond Typed Racket: intersection and negation allow for *very* precise types. For an example, see Greenberg [2019].

OCaml. In addition to standard algebraic datatypes, OCaml also supports two relaxed notions of datatype: polymorphic variants [Garrigue 1998], a disciplined static approach to unclassified data; and extensible variants [Leroy et al. 2021]. Extensible variants, unsurprisingly, are variant types that can be extended with new constructors and use the conventional (static) constructor syntax. The `exn` type (used for exceptions) is a built-in example of an extensible variant, for which programmers can introduce new constructors.

There are two main distinctions between GSD and OCaml’s polymorphic variants: (a) how polymorphic variants inhabit types and (b) how polymorphic variants can be matched. In GSD, the closest corresponding idea is $?_{\square}$, which represents open datatypes. There is no way in GSD to name a type of some particular bounded set of unclassified constructors, but OCaml can approximate $?_{\square}$ by using lower bounds ($\sim$ in Figure 13): $[>\sim \text{Foo}]$ represents all types that have at least the constructor $\sim \text{Foo}$. OCaml reasons well about such bounded sets: functions can be annotated with the exact set of polymorphic variants they expect, or with upper or lower bounds; these bounds can even be inferred. In contrast to unclassified data in GSD, polymorphic variants do not mix with regular variants in a pattern match. Such a distinction can be an advantage: typos of declared constructors are caught as syntax errors by the compiler. GSD makes no such distinction, handling any constructed data in the same match expression, irrespective of whether its classified or not—allowing for a smoother transition from a dynamic program to a more static one at the expense of some static checking.

We do not include Haskell in the table, but its affordances are similar to OCaml. Haskell lacks polymorphic and extensible variants, but features a `Dynamic` type that provides for a limited but highly structured notion of reflection.

Scala. Scala is a statically typed programming language that combines functional and object-oriented programming [Odersky et al. 2006]. Scala’s *case classes* are akin to algebraic datatypes, and can be matched on. Classes are open by default and can be extended with multiple case classes. Pattern matches on open classes cannot be checked for exhaustiveness, since case classes can be extended from outside their declaring file. But case classes marked as *sealed* can only be extended from within their source file—in this case, the compiler checks for exhaustiveness. GSD makes similar tradeoffs between allowing openness and checking exhaustiveness. While Scala does not support ad-hoc constructors, it offers great extensibility for datatypes: the programmer can extend one datatype with another; set-theoretic types like union (of arbitrary types) and intersection (of traits) cover some ad hoc use cases.

Rust. Rust is a safe systems programming language [Klabnik and Nichols 2018]. Its type system uses substructural types to enforce an ownership discipline—these powerful features have no real bearing in our setting. Rust’s `enums` resemble algebraic datatypes, and are significantly more expressive than `enums` in C or C++. Like C and C++, though, Rust offers ways for programmers to

control the representation of enums, e.g., assigning them integer values. Rust pattern matching is expressive, with convenient forms for matching only parts of a datatype as well as keywords for managing memory dereference as part of pattern matches. Because of its low-level target, Rust doesn't offer much in the way of runtime reflection.

Elm. Elm is a pure functional language that compiles to JavaScript [Czaplicki and Chong 2013]. Elm takes static checking seriously: incomplete pattern matches are an error. Of the static languages considered, Elm offers the least flexibility: Elm offers no reflection of any kind; release builds forbid debug logging and exceptions; the built-in definitions and libraries force users to handle errors with option types. We do not include them in the table, but Coq, Agda, and other total languages are similar to Elm.

6 RELATED WORK

We now expand on two lines of related work: work on datatypes, and work on gradual types. Algebraic datatypes are a fixture of typed functional programming—and they are increasingly found in other languages, too [Klabnik and Nichols 2018; Odgersky et al. 2006]. Several ways of extending datatypes with new variants have been explored beyond those discussed above, including in the gradual typing literature.

Garrigue [1998] uses polymorphic variants to encode extensible datatypes: programmers can match on variants that are not defined in any datatype, bounding the accepted variants of a function at the type level. Zenger and Odgersky [2001] define a way of modeling extensible algebraic datatypes with defaults in an object-oriented language. Datatypes are modeled as classes and every extension, with potentially multiple variants, as a subclass of the extended datatype. It is said to be “with defaults”, by the way it implements pattern matching. When doing pattern matching on a subclass, if no pattern holds then the default case does a supercall. Datatypes are extensible: programmers introduce new variants using inheritance and interfaces. These extensions are said to be linear, since each extension brings only the variants of its superclass. Linearity ensures that pattern matches cannot fail at runtime. But because extensions are all statically declared, they cannot account for dynamically discovered unclassified data.

Gradual typing per se is relatively recent, but attempts to unify dynamically and statically typed programming go back at least to the 1990s. We have already discussed several systems in Section 5; we discuss a few more here.

Garcia et al. [2016] develop *gradual rows*, which are rows with possible extra unknown fields. Extending our system with gradual rows would have complicated the formalism, but would have allowed for not only extensible datatypes, but constructors with gradual arity. Sekiyama and Igarashi [2020] describe a gradual language with row types and row polymorphism. They support gradual variants: like GSD and unlike polymorphic variants, there is no special syntax for static constructors. They can bound input or output variants of a function using row polymorphism [Wand 1991]. Row types are very expressive, but they come with more metatheoretical baggage. They use scoped labels [Leijen 2017], which lead to an operational semantics that does not directly correspond to an efficient implementation.

Siek and Tobin-Hochstadt [2016], Castagna and Lanvin [2017] and [Toro and Tanter 2017] describe gradual languages with union types, under different forms. Union types are a set-theoretic alternative to algebraic datatypes. As seen in the discussion of Typed Racket and CDuce (Section 5), union types are often more flexible than algebraic approaches: a single constructor can appear in multiple types, and type unions can mix primitive values like numbers or booleans with constructors. We don't support true unions or pattern matching on anything but constructors in GSD, but we can simulate some aspects of union types using the unknown datatype `?_`, where `match` expressions

have cases for the relevant variants. GSD has no way to statically differentiate between $\text{Bool} \cup \text{Foo}$ and $\text{Bool} \cup \text{Bar}$: both are simply $?_{\square}$. However, $\lambda_{D?}$ facilitates working with unclassified data that is explicitly not yet defined: even with negation, it is challenging to concisely express the idea of “constructor that does not appear statically in closed datatypes”.

Jafery and Dunfield [2017] develop gradual datasort refinements with sum types to eliminate pattern matching errors. The range of graduality they support is much more static than what we study here: the dynamic language is ML-like, and the static language has datasort refinements, i.e., type-level reasoning about subsets of datatype constructors. Our notion of $\text{valid}_{\Delta, \Xi}$ is related; it would be interesting to combine our work and theirs for a “full spectrum” system that ranges from no guarantees (Scheme-like) to some guarantees (ML-like) to strong guarantees (datasort refinements).

7 CONCLUSION

Gradual typing aims to reconcile dynamic and static approaches, but has not yet confronted a critical, defining feature of statically-typed programming languages: algebraic datatypes. After defining a simple calculus with support for nominal algebraic datatypes, we use AGT [Garcia et al. 2016] to derive $\lambda_{D?}$, a core calculus for GSD, a language for gradually structured data. Our design hinges on carefully separating open and closed datatypes, and introducing two new unknown types, $?_{\square}$ as the ground type of data, and $?_{\square}$, the unknown open datatype. Gradually structured data lets programmers handle data at different levels of static precision: from ad hoc and semi-structured or “tagged” data all the way up to fully statically defined algebraic datatypes. GSD achieves all of this while using a very simple type system. In addition to the metatheory of $\lambda_{D?}$, we implemented a reference interpreter for GSD and illustrated how GSD handles the evolution of a basic web server.

There is much left to be done. Our implementation of GSD is an interpreter, not a compiler, and we have not yet attempted to improve space or time performance. We build our account of name generation into primitives like `fromJSON`, and we can imagine much more robust ways of constraining the shape of unclassified data to, e.g., some particular number or pattern of arguments. Finally, real languages support interesting extensions of datatypes, like GADTs and typeclasses. Accounting for these features is important but challenging. Overall, we show that even with simple types, gradually structured data is an expressive and flexible approach.

REFERENCES

- Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. 2011. Blame for All. In *Proceedings of the 38th annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2011)*. ACM Press, Austin, Texas, USA, 201–214. <https://doi.org/10.1145/1926385.1926409>
- Amal Ahmed, Dustin Jamner, Jeremy G. Siek, and Philip Wadler. 2017. Theorems for Free for Free: Parametricity, with and Without Types. *Proceedings of the ACM on Programming Languages* 1, ICFP (Sept. 2017), 39:1–39:28. <https://doi.org/10.1145/3110283>
- Felipe Bañados Schwerter, Alison M Clark, Khurram A Jafery, and Ronald Garcia. 2021. Abstracting gradual typing moving forward: precise and space-efficient. *Proceedings of the ACM on Programming Languages* 5, POPL (Jan. 2021), 1–28. <https://doi.org/10.1145/3434342>
- Felipe Bañados Schwerter, Ronald Garcia, and Éric Tanter. 2014. A Theory of Gradual Effect Systems. In *Proceedings of the 19th ACM SIGPLAN Conference on Functional Programming (ICFP 2014)*. ACM Press, Gothenburg, Sweden, 283–295. <https://doi.org/10.1145/2692915.2628149>
- Felipe Bañados Schwerter, Ronald Garcia, and Éric Tanter. 2016. Gradual Type-and-Effect Systems. *Journal of Functional Programming* 26 (Sept. 2016), 19:1–19:69. <https://doi.org/10.1017/S0956796816000162>
- Spenser Bauman, Carl Friedrich Bolz-Tereick, Jeremy Siek, and Sam Tobin-Hochstadt. 2017. Sound gradual typing: Only mostly dead. *Proceedings of the ACM on Programming Languages* 1, OOPSLA (Oct. 2017), 1–24. <https://doi.org/10.1145/3133878>
- Véronique Benzaken, Giuseppe Castagna, and Alain Frisch. 2003. CDuce: An XML-Centric General Purpose Language. In *Proceedings of the 8th ACM SIGPLAN Conference on Functional Programming (ICFP 2003)*. ACM Press, Uppsala, Sweden.

<https://doi.org/10.1145/944705.944711>

- Peter Buneman. 1997. Semistructured Data. In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (Principles of Database Systems)*. ACM Press, New York, NY, USA, 117–121. <https://doi.org/10.1145/263661.263675>
- John Peter Campora, Sheng Chen, and Eric Walkingshaw. 2018. Casts and costs: Harmonizing safety and performance in gradual typing. *Proceedings of the ACM on Programming Languages* 2, ICFP (Sept. 2018), 1–30. <https://doi.org/10.1145/3236793>
- Giuseppe Castagna and Victor Lanvin. 2017. Gradual Typing with Union and Intersection Types. *Proceedings of the ACM on Programming Languages* 1, ICFP (Sept. 2017), 41:1–41:28. <https://doi.org/10.1145/3110285>
- Giuseppe Castagna, Victor Lanvin, Tommaso Petrucciani, and Jeremy G. Siek. 2019. Gradual typing: a new perspective. *Proceedings of the ACM on Programming Languages* 3, POPL (Jan. 2019), 16:1–16:32. <https://doi.org/10.1145/3290329>
- Evan Czaplicki and Stephen Chong. 2013. Asynchronous Functional Reactive Programming for GUIs. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2013)*. ACM Press, Seattle, Washington, USA, 411–422. <https://doi.org/10.1145/2499370.2462161>
- Mike Fagan. 1991. *Soft typing: An approach to type checking for dynamically typed languages*. Ph.D. Dissertation. Rice University. <https://scholarship.rice.edu/handle/1911/16439>
- Matthias Felleisen and Daniel P. Friedman. 1986. Control Operators, the SECD-Machine, and the λ -calculus. In *Proceedings of IFIP Working Conference on Formal Descriptions of Programming Concepts*, Martin Wirsing (Ed.). Ebberup, Denmark, 193–219.
- Ronald Garcia and Matteo Cimini. 2015. Principal Type Schemes for Gradual Programs. In *Proceedings of the 42nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2015)*. ACM Press, Mumbai, India, 303–315. <https://doi.org/10.1145/2775051.2676992>
- Ronald Garcia, Alison M. Clark, and Éric Tanter. 2016. Abstracting Gradual Typing. In *Proceedings of the 43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2016)*, Rastislav Bodik and Rupak Majumdar (Eds.). ACM Press, St Petersburg, FL, USA, 429–442. <https://doi.org/10.1145/2837614.2837670> See erratum: <https://www.cs.ubc.ca/~rxg/agt-erratum.pdf>.
- Ronald Garcia, Éric Tanter, Roger Wolff, and Jonathan Aldrich. 2014. Foundations of Typestate-Oriented Programming. *ACM Transactions on Programming Languages and Systems* 36, 4, Article 12 (Oct. 2014), 12:1–12:44 pages. <https://doi.org/10.1145/2629609>
- Jacques Garrigue. 1998. Programming with polymorphic variants. In *ML Workshop*, Vol. 13. Baltimore, 7.
- Panicz Maciej Godek. 2020. SRFI 200: Pattern Matching. <https://srfi.schemers.org/srfi-200/srfi-200.html>
- Michael Greenberg. 2019. The Dynamic Practice and Static Theory of Gradual Typing. In *3rd Summit on Advances in Programming Languages (SNAPL 2019) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 136)*, Benjamin S. Lerner, Rastislav Bodik, and Shriram Krishnamurthi (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Providence, RI, USA, 6:1–6:20. <https://doi.org/10.4230/LIPIcs.SNAPL.2019.6>
- Ben Greenman, Asumu Takikawa, Max S New, Daniel Feltey, Robert Bruce Findler, Jan Vitek, and Matthias Felleisen. 2019. How to evaluate the performance of gradual type systems. *jfp* 29 (2019). <https://doi.org/10.1017/S0956796818000217>
- David Herman, Aaron Tomb, and Cormac Flanagan. 2010. Space-efficient gradual typing. *Higher-Order and Sympolic Computation* 23, 2 (June 2010), 167–189. <https://doi.org/10.1007/s10990-011-9066-z>
- Atsushi Igarashi, Peter Thiemann, Vasco T. Vasconcelos, and Philip Wadler. 2017b. Gradual Session Types. *Proceedings of the ACM on Programming Languages* 1, ICFP (Sept. 2017), 38:1–38:28. <https://doi.org/10.1145/3110282>
- Yuu Igarashi, Taro Sekiyama, and Atsushi Igarashi. 2017a. On Polymorphic Gradual Typing. *Proceedings of the ACM on Programming Languages* 1, ICFP (Sept. 2017), 40:1–40:29. <https://doi.org/10.1145/3110284>
- Khurram A. Jafery and Joshua Dunfield. 2017. Sums of Uncertainty: Refinements Go Gradual. In *Proceedings of the 44th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2017)*. ACM Press, Paris, France, 804–817. <https://doi.org/10.1145/3009837.3009865>
- Richard Kelsey. 1999. SRFI 9: Defining Record Types. <https://srfi.schemers.org/srfi-9/srfi-9.html>
- Steve Klabnik and Carol Nichols. 2018. *The Rust Programming Language*. No Starch Press.
- Andre Kuhlenschmidt, Deyaaeldeen Almahallawi, and Jeremy G Siek. 2019. Toward efficient gradual typing for structural types via coercions. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2019)*. ACM Press, Phoenix, AZ, USA, 517–532. <https://doi.org/10.1145/3314221.3314627>
- Nico Lehmann and Éric Tanter. 2017. Gradual Refinement Types. In *Proceedings of the 44th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2017)*. ACM Press, Paris, France, 775–788. <https://doi.org/10.1145/3009837.3009856>
- Daan Leijen. 2017. Type Directed Compilation of Row-Typed Algebraic Effects. In *Proceedings of the 44th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2017)*. ACM Press, Paris, France, 486–499. <https://doi.org/10.1145/3093333.3009872>

- Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. 2021. The OCaml system, release 4.12, Documentation and user's manual. (27 Jan. 2021). <https://ocaml.org/releases/4.12/ocaml-4.12-refman.pdf>
- Max S. New, Dustin Jamner, and Amal Ahmed. 2020. Graduality and Parametricity: Together Again for the First Time. *Proceedings of the ACM on Programming Languages* 4, POPL (Jan. 2020), 46:1–46:32. <https://doi.org/10.1145/3371114>
- Martin Odersky, Philippe Altherr, Vincent Cremet, Burak Emir, Sebastian Maneth, Stéphane Micheloud, Nikolay Mihaylov, Michel Schinz, Erik Stenman, and Matthias Zenger. 2006. *An overview of the Scala programming language*. Technical Report LAMP-REPORT-2006-001. EPFL. <https://www.scala-lang.org/docu/files/ScalaOverview.pdf>
- Oxford. 2021. *Oxford Advanced Learner's Dictionary* (10th ed.). Oxford University Press. Evidence.
- Taro Sekiyama and Atsushi Igarashi. 2020. Gradual Typing for Extensibility by Rows. (Jan. 2020). Workshop on Gradual Typing (WGT).
- Ilya Sergey and Dave Clarke. 2012. Gradual Ownership Types. In *Proceedings of the 21st European Symposium on Programming Languages and Systems (ESOP 2012) (Lecture Notes in Computer Science, Vol. 7211)*, Helmut Seidl (Ed.). Springer-Verlag, Tallinn, Estonia, 579–599. https://doi.org/10.1007/978-3-642-28869-2_29
- Jeremy Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In *Proceedings of the Scheme and Functional Programming Workshop*. 81–92.
- Jeremy Siek and Walid Taha. 2007. Gradual Typing for Objects. In *Proceedings of the 21st European Conference on Object-oriented Programming (ECOOP 2007) (Lecture Notes in Computer Science, 4609)*, Erik Ernst (Ed.). Springer-Verlag, Berlin, Germany, 2–27. https://doi.org/10.1007/978-3-540-73589-2_2
- Jeremy Siek and Sam Tobin-Hochstadt. 2016. The Recursive Union of Some Gradual Types. In *Wadler Festschrift (Lecture Notes in Computer Science, Vol. 9600)*. Springer-Verlag, 388–410. https://doi.org/10.1007/978-3-319-30936-1_21
- Jeremy Siek and Philip Wadler. 2010. Threesomes, with and without blame. In *Proceedings of the 37th annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2010)*. ACM Press, Madrid, Spain, 365–376. <https://doi.org/10.1145/1707801.1706342>
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015a. Refined Criteria for Gradual Typing. In *1st Summit on Advances in Programming Languages (SNAPL 2015) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 32)*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Asilomar, California, USA, 274–293. <https://doi.org/10.4230/LIPIcs.SNAPL.2015.274>
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, Sam Tobin-Hochstadt, and Ronald Garcia. 2015b. Monotonic References for Efficient Gradual Typing. In *Proceedings of the 24th European Symposium on Programming Languages and Systems (ESOP 2015) (Lecture Notes in Computer Science, Vol. 9032)*, Jan Vitek (Ed.). Springer-Verlag, London, UK, 432–456. https://doi.org/10.1007/978-3-662-46669-8_18
- Martin Sulzmann, Manuel MT Chakravarty, Simon Peyton Jones, and Kevin Donnelly. 2007. System F with type equality coercions. In *Proceedings of the 2007 ACM SIGPLAN international workshop on Types in languages design and implementation*. acm, Nice, France, 53–66. <https://doi.org/10.1145/1190315.1190324>
- Asumu Takikawa, Daniel Feltey, Ben Greenman, Max S. New, Jan Vitek, and Matthias Felleisen. 2016. Is sound gradual typing dead?. In *Proceedings of the 43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2016)*, Rastislav Bodík and Rupak Majumdar (Eds.). ACM Press, St Petersburg, FL, USA, 456–468. <https://doi.org/10.1145/2837614.2837630>
- Asumu Takikawa, T. Stephen Strickland, Christos Dimoulas, Sam Tobin-Hochstadt, and Matthias Felleisen. 2012. Gradual Typing for First-Class Classes. In *Proceedings of the 27th ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA 2012)*. ACM Press, Tucson, AZ, USA, 793–810. <https://doi.org/10.1145/2398857.2384674>
- Sam Tobin-Hochstadt and Matthias Felleisen. 2008. The Design and Implementation of Typed Scheme. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2008)*. ACM Press, San Francisco, CA, USA, 395–406. <https://doi.org/10.1145/1328438.1328486>
- Sam Tobin-Hochstadt, Vincent St-Amour, and Eric Dobson. 2014. The Typed Racket Reference. (2014).
- Matias Toro, Elizabeth Labrada, and Éric Tanter. 2019. Gradual Parametricity, Revisited. *Proceedings of the ACM on Programming Languages* 3, POPL (Jan. 2019), 17:1–17:30. <https://doi.org/10.1145/3290330>
- Matias Toro and Éric Tanter. 2017. A Gradual Interpretation of Union Types. In *Proceedings of the 24th Static Analysis Symposium (SAS 2017) (Lecture Notes in Computer Science, Vol. 10422)*. Springer-Verlag, New York City, NY, USA, 382–404. https://doi.org/10.1007/978-3-319-66706-5_19
- Matias Toro and Éric Tanter. 2020. Abstracting Gradual References. *Science of Computer Programming* 197 (Oct. 2020), 1–65. <https://doi.org/10.1016/j.scico.2020.102496>
- Niki Vazou, Éric Tanter, and David Van Horn. 2018. Gradual Liquid Type Inference. *Proceedings of the ACM on Programming Languages* 2, OOPSLA (Nov. 2018), 132:1–132:25. <https://doi.org/10.1145/3276502>
- Mitchell Wand. 1987. Complete Type Inference for Simple Objects. In *Proceedings of the Symposium on Logic in Computer Science (LICS 1987)*. IEEE Computer Society Press, Ithaca, New York, USA, 37–44.

- Mitchell Wand. 1991. Type inference for record concatenation and multiple inheritance. *Information and Computation* 93, 1 (1991), 1–15. [https://doi.org/10.1016/0890-5401\(91\)90050-C](https://doi.org/10.1016/0890-5401(91)90050-C) Selections from 1989 IEEE Symposium on Logic in Computer Science.
- Roger Wolff, Ronald Garcia, Éric Tanter, and Jonathan Aldrich. 2011. Gradual Typestate. In *Proceedings of the 25th European Conference on Object-oriented Programming (ECOOP 2011) (Lecture Notes in Computer Science, Vol. 6813)*, Mira Mezini (Ed.). Springer-Verlag, Lancaster, UK, 459–483. https://doi.org/10.1007/978-3-642-22655-7_22
- Matthias Zenger and Martin Odersky. 2001. Extensible Algebraic Datatypes with Defaults. In *Proceedings of the 6th ACM SIGPLAN Conference on Functional Programming (ICFP 2001)*. ACM Press, Florence, Italy, 241–252. <https://doi.org/10.1145/507635.507665>