

# Compositional Embeddings of Domain-Specific Languages

ANONYMOUS AUTHOR(S)

A common approach to defining domain-specific languages (DSLs) is via a direct embedding into a host language. There are several well-known techniques to do such embeddings, including *shallow* and *deep* embeddings. However, such embeddings come with various trade-offs in existing programming languages. Owing to such trade-offs, many embedded DSLs end up using a mix of approaches in practice, requiring a substantial amount of code, as well as some advanced coding techniques.

In this paper, we show that the recently proposed *Compositional Programming* paradigm and the CP language provide improved support for embedded DSLs. In CP, we obtain a new form of embedding, which we call a *compositional embedding*, that has most of the advantages of both shallow and deep embeddings. On the one hand, compositional embeddings enable various forms of linguistic reuse that are characteristic of shallow embeddings, including the ability to reuse host-language optimizations in the DSL and add new DSL constructs easily. On the other hand, similarly to deep embeddings, compositional embeddings support definitions by pattern matching or dynamic dispatching (including dependent interpretations, transformations, and optimizations) over the abstract syntax of the DSL and have the ability to add new interpretations. We illustrate an instance of compositional embeddings with a DSL for document authoring called  $\text{E}\chi\text{T}$ . The DSL is highly flexible and extensible, allowing users to create various non-trivial extensions easily. For instance,  $\text{E}\chi\text{T}$  supports various extensions that enable the production of wiki-like documents,  $\text{\LaTeX}$  documents, vector graphics or charts. The viability of compositional embeddings for  $\text{E}\chi\text{T}$  is evaluated with three applications.

## 1 INTRODUCTION

A common approach to defining domain-specific languages (DSLs) is via a direct embedding into a host language. This approach is used in several programming languages, such as Haskell, Scala, and Racket. In those languages, various DSLs – including pretty printers [Hughes 1995; Wadler 2003], parser combinators [Leijen and Meijer 2001], and property-based testing frameworks [Claessen and Hughes 2000] – are defined as embedded DSLs. There are a few techniques for such embeddings, including the well-known *shallow* and *deep* embeddings [Boulton et al. 1992].

Unfortunately, shallow and deep embeddings come with various trade-offs in existing programming languages. Such trade-offs have been widely discussed in the literature [Gibbons and Wu 2014; Rompf et al. 2012; Scherr and Chiba 2014]. On the one hand, the strengths of shallow embeddings are in providing *linguistic reuse* [Krishnamurthi 2001], exploiting meta-language optimizations, and allowing the addition of new DSL constructs easily. On the other hand, deep embeddings shine in enabling the definition of complex semantic interpretations and optimizations over the abstract syntax tree (AST) of the DSL, and they enable adding new semantic interpretations easily. Regarding such trade-offs, Svenningsson and Axelsson [2015] made the following striking comment:

*“The holy grail of embedded language implementation is to be able to combine the advantages of shallow and deep in a single implementation.”*

While progress has been made in embedded language implementation, the holy grail is still not fully achieved in existing programming languages. Owing to the trade-offs between shallow and deep embeddings, many realistic embedded DSLs end up using a mix of both approaches in practice or more advanced forms of embeddings. For instance, there have been several approaches [Jovanović et al. 2014; Rompf et al. 2012; Svenningsson and Axelsson 2015] promoting the use of shallow embeddings as the frontend of the DSL to enable linguistic reuse, while deep embeddings are used as the backend for added flexibility in defining semantic interpretations. While such approaches manage to alleviate some of the trade-offs, they require translations between the two embeddings, a substantial amount of code, and some advanced coding techniques. Alternatively, more advanced embedding techniques, such as *tagless-final embeddings* [Carette et al. 2009; Kiselyov 2010], *data types à la*

50 *carte* [Swierstra 2008], *polymorphic embeddings* [Hofer et al. 2008] and *object algebras* [Oliveira  
51 and Cook 2012], are able to eliminate some of the trade-offs too. In particular, those approaches  
52 eliminate the trade-offs with respect to extensibility, facilitating both the addition of new DSL  
53 constructs and semantic interpretations. However, being quite close to shallow embeddings, those  
54 approaches lack some important capabilities, such as the ability to define complex interpretations  
55 and the use of (nested) pattern matching to express semantic interpretations and transformations  
56 easily and modularly.

57 In this paper, we show that the recently proposed *Compositional Programming* paradigm [Zhang  
58 et al. 2021] and the CP language provide improved programming language support for embedded  
59 DSLs. Compositional Programming provides an alternative way to define data structures compared  
60 to algebraic data types in functional programming or class hierarchies in object-oriented program-  
61 ming. Compositional interfaces and traits play a role similar to algebraic data types and definitions  
62 by pattern matching in functional programs. However, unlike algebraic data types and definitions  
63 by pattern matching, compositional interfaces and traits are modularly extensible. *Intersection*  
64 *types* [Oliveira et al. 2016] enable the composition of interfaces, while *nested composition* [Bi et al.  
65 2018] enables the composition of traits with pattern matching definitions. Thus, the CP language  
66 does not suffer from the infamous *Expression Problem* [Wadler 1998]. In addition, CP comes with  
67 several mechanisms to express *modular dependencies*, allowing powerful forms of dependency  
68 injection and complex semantic interpretations.

69 With those programming language features, we obtain a new form of embedding called a *compo-*  
70 *sitional embedding*, with nearly all of the advantages of both shallow and deep embeddings. On the  
71 one hand, compositional embeddings enable various forms of linguistic reuse that are characteristic  
72 of shallow embeddings, including the ability to reuse host-language optimizations in the DSL and  
73 easily add new DSL constructs. On the other hand, similarly to deep embeddings, compositional  
74 embeddings support definitions by pattern matching or dynamic dispatching, including optimiza-  
75 tions and transformations over the abstract syntax of the DSL, as well as the ability to add new  
76 interpretations. In short, we believe that compositional embeddings come very close to the holy  
77 grail of embedded language implementation desired by Svenningsson and Axelsson [2015].

78 We reimplemented the CP language and made it available as an in-browser interpreter. Using this  
79 new implementation of CP, we illustrate an instance of compositional embeddings with a DSL for  
80 document authoring called *ExT* (**Ex**ten**sible** **T**ypesetting). The DSL is highly flexible and extensible,  
81 allowing users to create various non-trivial extensions. For instance, *ExT* supports extensions that  
82 enable the production of wiki-like documents,  $\LaTeX$  documents, SVG charts, or computational  
83 graphics like fractals. Our largest application is Minipedia: a mini version of Wikipedia. Minipedia  
84 includes wiki pages for several states and cities and some tables that list countries by population or  
85 area. The *ExT* DSL, as well as its various extensions and applications, is available online.

86 In summary, the contributions of this paper are:

- 87
- 88 • **Compositional embeddings.** We show that the previously proposed Compositional Pro-  
89 gramming [Zhang et al. 2021], together with the CP language, enables a new form of  
90 embedding, which we call a compositional embedding. This paper explicitly reveals that  
91 compositional embeddings have most of the advantages of shallow and deep embeddings.
- 92 • **Comparison of embedding techniques.** We present a detailed comparison between com-  
93 positional embeddings and various techniques used in embedded language implementations,  
94 including shallow, deep, hybrid, and polymorphic embeddings.
- 95 • **The *ExT* DSL.** As a realistic instance of compositional embeddings, we present a DSL for  
96 document authoring called *ExT*. *ExT* is extremely flexible and customizable by users, with  
97 many features implemented in a modular way.
- 98

- **Implementation and applications of ExT.** We have built several applications with ExT, three of which are discussed in more detail in this paper. The largest application is Minipedia, and the other two applications illustrate computational graphics like fractals and a document extension for charts. The implementation of ExT and its applications are available at:

<https://plground.org>

## 2 EMBEDDINGS OF DSLS

In this section, we give some background on existing approaches to embedded DSLs and evaluate their strengths and drawbacks. This will be useful for the comparison in Section 3. We focus on *shallow* and *deep* embeddings [Boulton et al. 1992], which are the two main alternative forms of embeddings. We also discuss some other forms of embeddings near the end of this section. To illustrate all these embeddings, we present programs in Haskell, which is well known for its good support for embedded DSLs and has a syntax close to the CP language.

### 2.1 A Simple Region DSL

Inspired by Hudak [1998] and Hofer et al. [2008], we consider a simple region DSL for plane geometry. To illustrate the challenges in developing such a DSL, we consider five separate steps in the development, which illustrate various desired features, such as linguistic reuse and the ease of adding features for software evolution.

- (1) **Initial system.** We start with a small region language with five constructs: `circle` for creating a circular region with a given radius, `outside` for a complement to a certain region, two set operators `union` and `intersect`, and finally `translate` for moving a region by a given vector. The initial interpretation is to simply calculate the size of a region.
- (2) **Linguistic reuse.** We create a region with heavy sharing to assess the difficulty of reusing host-language optimizations.
- (3) **Extensibility with new constructs.** We extend the region language with three more constructs: `univ` for the universal region that contains all points, `empty` for the empty region that contains no points, and `scale` for resizing a region by two scale factors in a vector.
- (4) **Extensibility with new operations.** We add a new interpretation that checks whether a given point resides in a region.
- (5) **Complex interpretations, transformations, and optimizations.** Last but not least, we discuss three kinds of more complex interpretations that illustrate transformations and dependent interpretations over the AST of the region language.

### 2.2 Initial System

Fig. 1 shows the definitions of the initial DSL, including the language interface and a simple interpretation for shallow and deep embeddings.

*Language interfaces.* In the shallow embedding, `Region` is directly mapped to its semantics, and all the five language constructs are implemented as denotation functions. Thus, the language interface (essentially the signature of the denotation functions) is entangled with the definition of a particular semantic domain. In the deep embedding, `Region` is defined as an algebraic data type<sup>1</sup>. Algebraic data types only capture the abstract syntax, thus enabling the language interface to be separated from any concrete semantics.

<sup>1</sup>We use the notation of *generalized* algebraic data types [Peyton Jones et al. 2006] to make type annotations clear.

```

148 data Vector = Vector { x :: Double, y :: Double } deriving Show
149
150 type Region = Int
151
152 circle :: Double → Region
153 circle _ = 1
154 outside :: Region → Region
155 outside a = a + 1
156 union :: Region → Region → Region
157 union a b = a + b + 1
158 intersect :: Region → Region → Region
159 intersect a b = a + b + 1
160 translate :: Vector → Region → Region
161 translate _ a = a + 1
162
163 (a) A shallow embedding

```

```

data Region where
164 Circle :: Double → Region
165 Outside :: Region → Region
166 Union :: Region → Region → Region
167 Intersect :: Region → Region → Region
168 Translate :: Vector → Region → Region
169
170 size :: Region → Int
171 size (Circle _) = 1
172 size (Outside a) = size a + 1
173 size (Union a b) = size a + size b + 1
174 size (Intersect a b) = size a + size b + 1
175 size (Translate _ a) = size a + 1
176
177 (b) A deep embedding

```

Fig. 1. The initial system for the region DSL.

*Interpretations and semantics.* The initial semantics is an abstract interpretation, which calculates the size of a region. In the shallow embedding, the semantics is defined together with the language interface since we must specify some concrete type (**Int** in this case) that instantiates **Region**. In contrast, to create a new semantic interpretation in the deep embedding, we define a function `size` using pattern matching over the **Region** data type.

We can see that the two embeddings work quite differently: the shallow embedding encodes semantics in region constructors, and the interpretation function from **Region** to **Int** is merely the identity function; the deep embedding does nothing concerning constructors but hands over the work to the interpretation function `size`.

### 2.3 Linguistic Reuse and Meta-Language Optimizations

An important concern for embedded DSLs is *linguistic reuse* [Krishnamurthi 2001]. One of the key selling points of embedded DSLs is that they can reuse much of the infrastructure of the host language and inherit various optimizations that are available in the host language. However, there are important differences between shallow and deep embeddings with respect to linguistic reuse. As widely observed in previous work [Jovanović et al. 2014; Rompf et al. 2012; Svenningsson and Axelsson 2015], shallow embeddings make linguistic reuse easier. To illustrate this, we can create a region that contains a series of repeated subregions with horizontally aligned circles:

```

185 circles = go 20 (2 ** 18)
186   where go :: Int → Double → Region
187         go 0 offset = circle 1
188         go n offset = let shared = go (n - 1) (offset / 2)
189                       in union (translate Vector { x = -offset, y = 0 } shared)
190                              (translate Vector { x =  offset, y = 0 } shared)
191
192
193
194
195
196

```

The region `circles` is defined via an auxiliary recursive function `go`. In the shallow embedding, `shared` is of type **Int**, and using `shared` twice in the `let`-body will avoid interpreting the region twice. The code above also works for the deep embedding, assuming some smart constructors such as `circle = Circle`. However, in the deep embedding, `shared` is an AST of type **Region**. In this case, the AST will be duplicated in the `let`-body, and later when we interpret the region, we need to

197 traverse the same AST twice, duplicating the size calculation. Since `go` is recursive, it will lead to a  
 198 lot of sharing for the shallow approach and a lot of repetition for the deep approach. As a result,  
 199 the evaluation of `circles` is instantaneous in the shallow approach, whereas in the deep approach,  
 200 it basically does not terminate in a reasonable amount of time. In short, in shallow embeddings, we  
 201 are able to preserve sharing by naturally exploiting the sharing optimizations present in the host  
 202 language, but sharing is lost in deep embeddings.

203 We should remark that this particular issue regarding sharing (which is just an instance of  
 204 linguistic reuse) is well known [Gill 2009; Kiselyov 2011; Oliveira and Löh 2013]. There have been  
 205 several proposed techniques that enable preserving sharing in deeply embedded DSLs. However,  
 206 all those techniques add extra complexity to the DSL encoding, and they may even make the DSL  
 207 harder to use. This is in contrast to the shallow approach, where no extra work is required by the  
 208 programmer and host-language features are naturally reused.

## 209 2.4 Adding More Language Constructs

211 As part of evolving the DSL with additional features, we may decide to add more constructs for  
 212 regions. To evaluate how easy it is to add more language constructs, we add three more language  
 213 constructs for the universal or empty region as well as a scaling operator. Shallow embeddings  
 214 make such extensions very easy. We only need to add three new denotation functions:

```
215 univ :: Region          empty :: Region          scale :: Vector → Region → Region
216 univ = 1                empty = 1                scale _ a = a + 1
```

217 Notwithstanding the modular extension in shallow embeddings, it is awkward to add language  
 218 constructs in deep embeddings. We have to modify the algebraic data type `Region` and all related  
 219 interpretation functions to add new cases. This is not modular because we must change the existing  
 220 code. Even if we have access to previous definitions, it is inevitable to recompile all the code that  
 221 depends on `Region`.

222 In essence, once we start adding new features, we are quickly faced with an instance of the  
 223 *Expression Problem* [Wadler 1998]. Shallow embeddings make adding constructs easy, whereas  
 224 adding new constructs in deep embeddings is more difficult and non-modular. As we shall see next,  
 225 when adding new semantic interpretations, we have a dual problem.

## 227 2.5 Adding a New Interpretation

228 Size is too abstract to describe a region, so let us add a new interpretation that checks whether  
 229 a region contains a given point. For example, a circular region of radius  $r$  contains  $(x, y)$  if and  
 230 only if  $x^2 + y^2 \leq r^2$ . Although adding constructs is awkward in deep embeddings, adding a new  
 231 interpretation function is trivial:

```
232 contains :: Region → Vector → Bool
233 Circle   r `contains` p = p.x ** 2 + p.y ** 2 ≤ r ** 2
234 Outside  a `contains` p = not (a `contains` p)
235 Union    a b `contains` p = a `contains` p || b `contains` p
236 Intersect a b `contains` p = a `contains` p && b `contains` p
237 Translate Vector {..} a `contains` p = a `contains` Vector { x = p.x - x, y = p.y - y }
```

238 The code is mostly straightforward<sup>2</sup>. The only minor remark is that we use a record wildcard  
 239 (`Vector {..}`) in the case of `Translate` to bring record fields ( $x$  and  $y$ ) into scope.

241 However, there is no easy modular way to support multiple interpretations in shallow embeddings.  
 242 We need to completely change the definition of `Region` and remap all the language constructs to a  
 243 new semantic domain. The issue of semantic extension in shallow embeddings is dual to that of

244 <sup>2</sup>Two GHC language extensions are required here: *OverloadedRecordDot* and *RecordWildCards*.

246 language extension in deep embeddings. To address the tension between the two dimensions, some  
 247 alternative embeddings are proposed, such as tagless-final embeddings [Carette et al. 2009; Kiselyov  
 248 2010] and polymorphic embeddings [Hofer et al. 2008], though they still have some significant  
 249 drawbacks, which will be revealed shortly.

## 251 2.6 Dependencies, Complex Interpretations, and Domain-Specific Optimizations

252 One area where deeply embedded DSLs shine is in enabling more complex kinds of interpretations.  
 253 These interpretations may, for example, enable transformations or rewritings on the AST, which  
 254 are helpful for writing domain-specific optimizations, among other things. For writing such com-  
 255 plex forms of interpretations, we often require multiple *dependent* functions defined by pattern  
 256 matching, and sometimes we may even need nested pattern matching. Such interpretations are very  
 257 challenging for shallow DSLs and are often the reason why DSL writers prefer deep embeddings.

258 *Dependent interpretations.* Let us start with a dependent interpretation that shows a text repre-  
 259 sentation of a region using a deep embedding:

```
260 text :: Region → String
261 text (Circle r) = "a circular region of radius " ++ show r
262 text (Outside a) = "outside a region of size " ++ show (size a)
263 text s@(Union _ _) = "the union of two regions of size " ++ show (size s) ++ " in total"
264 text s@(Intersect _ _) = "the intersection of two regions of size " ++ show (size s) ++ " in total"
265 text s@(Translate _ _) = "a translated region of size " ++ show (size s)
```

266 Note that the definition of text *depends* on the previously defined size function. Such a definition  
 267 is challenging in shallow embeddings. A workaround sometimes used in shallow embeddings is to  
 268 use tuples to define multiple interpretations simultaneously. For example, we can define size and  
 269 text together as:

```
270
271 type Region = (Int, String) -- (size, text)
272 circle r = (1, "a circular region of radius " ++ show r)
273 outside a = (fst a + 1, "outside a region of size " ++ show (fst a))
274 union a b = (size, "the union of two regions of size " ++ show size ++ " in total")
275   where size = fst a + fst b + 1
276 intersect a b = (size, "the intersection of two regions of size " ++ show size ++ " in total")
277   where size = fst a + fst b + 1
278 translate _ a = (size, "a translated region of size " ++ show size)
279   where size = fst a + 1
```

280 *Mutually dependent interpretations.* A similar, but more interesting, example is two interpretations  
 281 for checking universality and emptiness:

```
282 isUniv :: Region → Bool           isEmpty :: Region → Bool
283 isUniv Univ = True                isEmpty Empty = True
284 isUniv (Outside a) = isEmpty a     isEmpty (Outside a) = isUniv a
285 isUniv (Union a b) = isUniv a || isUniv b  isEmpty (Union a b) = isEmpty a && isEmpty b
286 isUniv (Intersect a b) = isUniv a && isUniv b  isEmpty (Intersect a b) = isEmpty a || isEmpty b
287 isUniv (Translate _ a) = isUniv a         isEmpty (Translate _ a) = isEmpty a
288 isUniv (Scale _ a) = isUniv a            isEmpty (Scale _ a) = isEmpty a
289 isUniv _ = False                    isEmpty _ = False
```

290 Unlike in the previous example, where only text depends on size, the two definitions are *mutually*  
 291 *recursive*, depending on each other. If we want to rewrite them using shallow embeddings, they  
 292 also have to be encoded as a pair to call each other via **fst** and **snd**:

```
293 type Region = (Bool, Bool) -- (isUniv, isEmpty)
294
```

```

295 univ      = (True,  False)      union    a b = (fst a || fst b, snd a && snd b)
296 empty    = (False, True)       intersect a b = (fst a && fst b, snd a || snd b)
297 circle  _ = (False, False)    translate _ a = (fst a,      snd a)
298 outside a = (snd a, fst a)     scale    _ a = (fst a,      snd a)

```

299 Generally speaking, using tuples to deal with dependencies is non-modular and awkward to use  
300 in that interpretations become tightly coupled. Adding more interpretations with dependencies,  
301 for example, would require a complete rewrite of the code: dependent interpretations have to be  
302 defined together with the interpretations that they depend on. More advanced encodings, such as  
303 tagless-final embeddings [Kiselyov 2010] or polymorphic embeddings [Hofer et al. 2008], provide  
304 ways to modularize *independent* interpretations into reusable units, but they cannot scale well to  
305 *dependent* interpretations and often have to resort to similar techniques using tuples like the code  
306 above. In Appendix A, we encode the two examples above in tagless-final embeddings and illustrate  
307 that they suffer from the same problems as shallow embeddings with respect to dependencies, both  
308 mutual and non-mutual.

309 *Nested pattern matching.* Another strength of deep embeddings is their ability to perform nested  
310 pattern matching, which is rather useful to inspect smaller constructs. Here we take an optimization  
311 that eliminates double negation for example. Nested pattern matching is used to check if an Outside  
312 is directly wrapped around an outer Outside. If so, both of them are eliminated; otherwise, eliminate  
313 is recursively called with inner constructs:  
314

```

315 eliminate :: Region → Region
316 eliminate (Outside (Outside a)) = eliminate a
317 eliminate (Outside      a)      = Outside (eliminate a)
318 eliminate (Union        a b)    = Union (eliminate a) (eliminate b)
319 eliminate (Intersect    a b)    = Intersect (eliminate a) (eliminate b)
320 eliminate (Translate    v a)    = Translate v (eliminate a)
321 eliminate (Scale        v a)    = Scale v (eliminate a)
322 eliminate                a      = a

```

323 The common use of nested pattern matching poses a second challenge to domain-specific optimiza-  
324 tions. Kiselyov [2010] discusses the “*seemingly impossible pattern-matching*” problem with nested  
325 patterns, which appear not directly encodable in tagless-final embeddings. Instead, he proposes  
326 that such algorithms can often be implemented in a tagless-final interpreter as context-dependent  
327 interpretations by essentially changing the algorithm. An extra parameter is added so the semantic  
328 type looks like  $\text{Ctx} \rightarrow \text{repr}$ . However, it is not obvious how to convert an algorithm with nested  
329 pattern matching into one based on contexts.

### 311 3 COMPOSITIONAL EMBEDDINGS

332 After looking back at the existing embedding techniques, this section shows that Compositional  
333 Programming enables a new form of DSL embedding called a *compositional embedding*. Composi-  
334 tional Programming is a statically typed modular programming paradigm implemented by the CP  
335 language. For the details of the semantics of CP, we refer the reader to work by Zhang et al. [2021].  
336 We have reimplemented CP as an in-browser interpreter, which can be directly run on a serverless  
337 web page. There are some small syntactic differences between our implementation and the original  
338 work, which we point out when necessary. Additionally, our implementation uses the call-by-need  
339 evaluation strategy, which optimizes the original call-by-name one.

340 Here, we only introduce the necessary language features of CP to illustrate compositional  
341 embeddings. With compositional embeddings, we can get many of the benefits of both shallow and  
342 deep embeddings without most of the drawbacks. We revisit the challenges illustrated in Section 2

344 using our approach in this section. A detailed comparison between compositional embeddings  
 345 and various existing forms of embeddings, including shallow, deep, hybrid, and polymorphic  
 346 embeddings, is presented at the end of this section.

347

### 348 3.1 Initial System, Revisited: Compositional Interfaces and Traits

349 As before, we start with an initial language interface inspired by Hudak [1998] and an abstract  
 350 interpretation that calculates the size of a region:

```
351 type Vector = { x : Double; y : Double };      type Size = { size : Int };
352 type HudakSig<Region> = {                    sz = trait implements HudakSig<Size> => {
353   Circle   : Double → Region;                (Circle   _).size = 1;
354   Outside  : Region → Region;                (Outside  a).size = a.size + 1;
355   Union    : Region → Region → Region;      (Union    a b).size = a.size + b.size + 1;
356   Intersect : Region → Region → Region;    (Intersect a b).size = a.size + b.size + 1;
357   Translate : Vector → Region → Region;     (Translate _ a).size = a.size + 1;
358 };                                            };
```

359 Since CP is also statically typed, a *compositional interface* (HudakSig) serves as the specification of  
 360 the DSL syntax, playing a similar role to algebraic data types in the deep embedding. The type  
 361 parameter Region is called the *sort* of the interface HudakSig. Sorts can be instantiated with different  
 362 concrete types that represent different semantic domains. Every constructor, which is capitalized,  
 363 must return a sort. Compositional interfaces are implemented by *traits* [Bi and Oliveira 2018;  
 364 Ducasse et al. 2006], which serve as the unit of code reuse in CP. Traits play a similar role to classes  
 365 in traditional object-oriented languages and are used to provide implementations for interpretations  
 366 of the region DSL. To implement the abstract interpretation of the DSL, we first define the semantic  
 367 domain type Size. Then we define a trait that implements HudakSig<Size>. In the body of the trait,  
 368 we use *method patterns*, such as (Circle \_).size = 1, to define the implementation. The method  
 369 patterns used in the trait body allow us to define interpretations that look very much like the  
 370 corresponding definition of size by pattern matching in Haskell in Fig. 1b.

### 371 3.2 Linguistic Reuse and Meta-Language Optimizations

372 The example of horizontally aligned circles in Section 2.3 can be translated to CP as:

```
373 sharing Region = trait [self : HudakSig<Region>] => {
374   circles = letrec go (n:Int) (offset:Double) : Region =
375     if n == 0 then Circle 1.0
376     else let shared = go (n-1) (offset/2.0)
377           in Union (Translate { x = -offset; y = 0.0 } shared)
378                 (Translate { x =  offset; y = 0.0 } shared)
379   in go 20 (pow 2.0 18);
380 };
381
```

382 Since the region DSL can have multiple interpretations, circles should be oblivious of any concrete  
 383 interpretation but only refer to the language interface. To achieve this, we use *self-type annotations*.  
 384 The trait sharing is parametrized by a type parameter Region and annotated with a self-type  
 385 HudakSig<Region>. Like Scala traits, CP traits can be annotated with self-types to express abstract  
 386 dependencies. This serves as an elegant way to inject region constructors. All of the constructors  
 387 formerly declared in the interface are available through self-references, for instance, self.Circle  
 388 (“self.” can be omitted for the sake of convenience). These constructors are *virtual* [Ernst et al.  
 389 2006; Madsen and Møller-Pedersen 1989]: they are not attached to a specific implementation.

390 With the self-type dependency on HudakSig<Region>, the static type checker of CP rejects a  
 391 direct instantiation like new sharing @Size (@ is the prefix for a type argument). Instead, we need  
 392



393 to compose sharing with a trait implementing `HudakSig<Size>`, using a *merge operator*<sup>3</sup> [Dunfield  
394 2014]:

```
395 test = new sharing @Size , sz;    -- test = new ((sharing @Size) , sz);
396 test.circles.size                -- Above is an equivalent definition with redundant parentheses.
```

397 The merge of the two traits (`sharing @Size` and `sz`) is still a trait. With self-type dependencies  
398 satisfied, the merged trait can be instantiated successfully, and we can call its method. According  
399 to call-by-need evaluation, the result of `shared` is stored for subsequent use once evaluated. Thanks  
400 to the linguistic reuse of the host-language `let` expressions and their optimizations, evaluating  
401 `circles.size` is able to exploit sharing and terminates almost immediately. This is similar to the  
402 Haskell approach in Section 2.3 using shallow embeddings.  
403

### 404 3.3 Adding New Language Constructs

405 In compositional embeddings, language constructs can be modularly added. We first create a new  
406 language interface inspired by Hofer et al. [2008] and then define a trait implementing it:  
407

```
408 type HoferSig<Region> = {                               sz' = trait implements HoferSig<Size> => {
409     Univ  : Region;                                     (Univ    ).size = 1;
410     Empty : Region;                                     (Empty  ).size = 1;
411     Scale : Vector → Region → Region;                 (Scale _ a).size = a.size + 1;
412 };                                                       };
```

413 To compose multiple interfaces, we use *intersection types*. To illustrate this, we create a repository  
414 of shapes that make use of language constructs from both interfaces:

```
415 type RegionSig<Region> = HudakSig<Region> & HoferSig<Region>;
416 repo Region = trait [Self : RegionSig<Region>] => {
417     ellipse = Scale { x = 4.0; y = 8.0 } (Circle 1.0);
418     universal = Union (Outside Empty) (Circle 1.0);
419 };
```

420 `RegionSig` is an intersection type combining the two interfaces into one. The definition of `repo` is  
421 similar to the definition of `sharing` in Section 3.2. It defines an ellipse, by using the `Scale` constructor  
422 from `HoferSig` and the `Circle` constructor from `HudakSig`. Similarly, it defines an essentially universal  
423 region by using constructors from both signatures.  
424

### 425 3.4 Adding New Interpretations

426 Besides language constructs, it is also trivial to add new semantic interpretations in compositional  
427 embeddings. Just like creating a new interpretation function in deep embeddings, we only need to  
428 define a new trait implementing a compositional interface:

```
429 type Eval = { contains : Vector → Bool };
430 eval = trait implements RegionSig<Eval> => {
431     (Circle    r).contains p = pow p.x 2 + pow p.y 2 ≤ pow r 2;
432     (Outside   a).contains p = !(a.contains p);
433     (Union     a b).contains p = a.contains p || b.contains p;
434     (Intersect a b).contains p = a.contains p && b.contains p;
435     (Translate {...} a).contains p = a.contains { x = p.x - x; y = p.y - y };
436     (Univ      ).contains _ = true;
437     (Empty     ).contains _ = false;
438     (Scale    {...} a).contains p = a.contains { x = p.x / x; y = p.y / y };
439 };
```

440 <sup>3</sup>We simplified the notation of the merge operator from double commas (`,,`) to a single comma (`,`).  
441




Fig. 2. A visualization of nested trait composition: inner components of traits are recursively composed.

Similarly to Haskell, CP supports record wildcards `{ . }`, which bring record fields into scope. We instantiate the sort with a function checking if a point resides in a region.

*Composing multiple interpretations.* Now that we have multiple interpretations, we use *nested trait composition*, provided by the merge operator, to compose all interpretations:

```

459 shapes = new repo @(Eval&Size) , eval , sz , sz'; -- nested trait composition
460 "The elliptical region of size " ++ toString shapes.ellipse.size ++ (if shapes.ellipse.contains {
461   x = 0.0; y = 0.0 } then " contains " else " does not contain ") ++ "the origin."
462 --> "The elliptical region of size 2 contains the origin."

```

We provide the final semantic domain (Eval&Size), as the type argument, for the trait `repo` and merge it with three other modularly defined traits. The trait composition is called *nested* [Bi et al. 2018] because not only different sets of constructors, but also different semantics nested in the same constructor are merged. In our case, both interpretations are available under the two sets of language constructs after nested composition, as visualized in Fig. 2. Here, the language constructs and their semantics form a two-level hierarchy of traits. Such composition of whole hierarchies can be traced back to *family polymorphism* [Ernst 2001]. With this powerful mechanism of nested composition, we can easily develop highly modular *product lines* [Apel et al. 2016] of DSLs, where DSL features can be composed *à la carte*.

### 3.5 Dependency Injection and Domain-Specific Optimizations

Now we illustrate CP's support for *modular dependent interpretations*. For space reasons, we skip the simpler example with text (which can be found in our online implementation) and focus on the more interesting example with `isUniv` and `isEmpty`. Similarly to the deep embedding in Haskell in Section 2.6, we define two compositional traits `chkUniv` and `chkEmpty`:

```

479 type IsUniv = { isUniv : Bool };           type IsEmpty = { isEmpty : Bool };
480
481 chkUniv = trait                            chkEmpty = trait
482   implements RegionSig<IsEmpty> => IsUniv => {   implements RegionSig<IsUniv> => IsEmpty => {
483     (Univ      ).isUniv = true;                   (Empty      ).isEmpty = true;
484     (Outside  a).isUniv = a.isEmpty;              (Outside  a).isEmpty = a.isUniv;
485     (Union    a b).isUniv = a.isUniv || b.isUniv; (Union    a b).isEmpty = a.isEmpty && b.isUniv;
486     (Intersect a b).isUniv = a.isUniv && b.isUniv; (Intersect a b).isEmpty = a.isUniv || b.isEmpty;
487     (Translate _ a).isUniv = a.isUniv;           (Translate _ a).isEmpty = a.isEmpty;
488     (Scale    _ a).isUniv = a.isUniv;           (Scale    _ a).isEmpty = a.isUniv;
489     (Scale    _ a).isUniv = false;               (Scale    _ a).isEmpty = false;
490   };

```

491 The first thing to note is that the sort of `RegionSig` is instantiated with two types separated by a fat  
 492 arrow ( $\Rightarrow$ ). This is how we *refine* interface types for dependency injection. For example, `chkUniv` does  
 493 not implement `RegionSig<IsEmpty>`, but we assume the latter is available. The static type checker of  
 494 CP guarantees that `chkUniv` is later merged with another trait that implements `RegionSig<IsEmpty>`.  
 495 That is why we can safely use `a.isEmpty` when implementing `(Outside a).isUniv`. The second trait  
 496 `chkEmpty` is just the dual of `chkUniv`. It is hard to define such interpretations modularly in traditional  
 497 approaches because they are dependent, but we make them compositional via dependency injection.  
 498 Lastly, `_.isUniv` and `_.isEmpty` are *default patterns*, which compensate for the remaining construc-  
 499 tors declared in the interface. For example, `_.isUniv = false` implies `Empty.isUniv = false` and  
 500 `Circle.isUniv = false`. Note that in CP, for modularity, patterns are *unordered*, unlike functional  
 501 languages like Haskell or ML. Therefore, default patterns are local to the traits where they are used.  
 502 These patterns can be understood as a concise way to fill the missing cases in the local trait.

503 *Delegated method patterns.* In some complicated transformations, nested pattern matching is  
 504 required to inspect smaller constructs. Nested pattern matching is not directly supported in CP yet.  
 505 However, there is a simple transformation that we can do whenever we would like to have nested  
 506 patterns: we can delegate the inner tasks to other methods whereby only top-level patterns are  
 507 needed. We call such a technique *delegated method patterns*.

508 Concerning the nested pattern matching in Section 2.6, we can implement it with two mutually  
 509 dependent methods. Since delegated method patterns are not reused, the two methods are defined  
 510 together for the sake of simplicity (but can always be separated into two traits):  
 511

```
512 type Eliminate Region = { eliminate : Region; delOutside : Region };
513 elim Region = trait [fself : RegionSig<Region>]
514   implements RegionSig<Region  $\Rightarrow$  Eliminate Region>  $\Rightarrow$  {
515     (Outside a).eliminate = a.delOutside;
516     (Union a b).eliminate = Union a.eliminate b.eliminate;
517     (Intersect a b).eliminate = Intersect a.eliminate b.eliminate;
518     (Translate v a).eliminate = Translate v a.eliminate;
519     (Scale v a).eliminate = Scale v a.eliminate;
520     [self].eliminate = self;
521     -- delegated method patterns:
522     (Outside a).delOutside = a.eliminate;
523     [self].delOutside = Outside self.eliminate;
524   };
```

524 The translation from nested pattern matching to delegated method patterns is straightforward. Our  
 525 code lifts nested patterns (`Outside (Outside a)`) to top-level delegated patterns. In this way, we do  
 526 not change the underlying algorithm at all. This is in contrast with approaches like `tagless-final`  
 527 embeddings that, as discussed in Section 2.6, seem unable to directly support nested patterns,  
 528 requiring different algorithms to achieve the same goal. Though our approach is not as convenient  
 529 as deep embeddings, it is noteworthy that we can just write the original algorithm *modularly*.

530 If we add new language constructs in the future, it is easy to augment the traits with more top-level  
 531 patterns, but extending nested patterns is difficult because there is no name to denote the *extension*  
 532 *point*. With delegated method patterns, the method name `delOutside` offers an extension point for  
 533 additional cases in the nested pattern matching. For instance, if we wish for an extension supporting  
 534 the empty region and a special case for eliminating a region outside the empty region (`eliminate`  
 535 `(Outside Empty)`), we can have a trait with a method pattern of the form `(Empty).delOutside = ...`  
 536 modularly defined in another trait.

537 Although we believe that the design of CP can be improved to better support similar logic  
 538 to nested patterns, there are important differences between traditional (closed) forms of pattern  
 539

Table 1. A comparison between different embedding approaches.

|  | SHALLOW    | DEEP | HYBRID         | POLY.          | COMP.          |
|--|------------|------|----------------|----------------|----------------|
| Transcoding free                       | ●          | ●    | ○              | ●              | ●              |
| Linguistic reuse                       | ●          | ○    | ●              | ●              | ●              |
| Language construct extensibility       | ●          | ○    | ◐ <sup>1</sup> | ●              | ●              |
| Interpretation extensibility           | ○          | ●    | ●              | ●              | ●              |
| Transformations and optimizations      | ○          | ●    | ●              | ◐ <sup>2</sup> | ◐ <sup>2</sup> |
| Linguistic reuse after transformations | <i>n/a</i> | ○    | ○              | ●              | ●              |
| Modular dependencies                   | ○          | ●    | ●              | ○              | ●              |
| Nested pattern matching                | ○          | ●    | ●              | ○              | ◐ <sup>3</sup> |

<sup>1</sup> The extensibility of language constructs is limited or precludes exhaustive pattern matching.

<sup>2</sup> Transformations require some ingenuity and are sometimes awkward to write.

<sup>3</sup> Nested pattern matching is implemented as delegated method patterns.

matching and open pattern matching [Zhang and Oliveira 2020]. For instance, patterns in CP are unordered for compositionality, but the order of patterns matters in conventional pattern matching. Thus the design of improved language support for nested patterns in CP, which could make the use of delegated method patterns more convenient, requires further research and is left for future work.

*Linguistic reuse after transformations.* Before finishing the discussion of modular transformations, let us revisit CP’s support for linguistic reuse: can we still have meta-language optimizations for free after complicated transformations? The answer is yes. To demonstrate this point, we can modify the definition of circles in Section 3.2 to have an inefficient Outside (Outside (Circle 1.0)). As usual, we compose all the necessary traits together to make sure no dependencies are missing:

```
test' = new sharing' @(Size & Eliminate Size) , sz , sz' , elim @Size;
test'.circles.eliminate.size
```

The evaluation terminates as quickly as before, meaning that meta-language optimizations are still performed even after a transformation.

We should remark that there are limits to the form of implicit sharing (and linguistic reuse) that shallow embeddings or compositional embeddings provide. For both shallow and compositional embeddings, implicit sharing will not work if the interpretation is a function, such as contains. In such cases, the sharing is still lost. Nevertheless, implicit sharing is an important feature that is often exploited in shallow DSLs. With compositional embeddings, we can take this idea further and make it work even after some transformations and optimizations have been applied. Moreover, it is also possible to adopt solutions with explicit sharing by modeling a Let construct in the DSL.

### 3.6 A Comparison between Different Embedding Approaches

In Table 1, we compare existing embedding approaches in the literature with compositional embeddings. Shallow and deep embeddings have been discussed in detail in Section 2, and the table summarizes the points that we have made. Hybrid approaches [Jovanović et al. 2014; Rompf et al. 2012; Svenningsson and Axelsson 2015] employ both embeddings together. However, transcoding from shallow to deep is generally needed, as both shallow and deep implementations must coexist side by side. There is a clear boundary between the two parts: linguistic reuse is available only in the shallow part, while complex transformations are available only in the deep part. Therefore, it is still hard to exploit host-language features after transformations. In work by Svenningsson and Axelsson, we cannot add more constructs to the deeply embedded AST but can only extend

Table 2. A briefer comparison with respect to the three dimensions of the Expression Problem.

|                | OOP | FP | POLY. | COMP. |
|----------------|-----|----|-------|-------|
| New constructs | ●   | ○  | ●     | ●     |
| New functions  | ○   | ●  | ●     | ●     |
| Dependencies   | ●   | ●  | ○     | ●     |

● modular (no code duplication needed)    ○ non-modular (code duplication needed)

shallowly embedded syntactic sugar. In Scala, if *open* case classes [Emir et al. 2007] are used for deep embeddings, ASTs are also extensible but do not ensure the exhaustiveness of pattern matching. Compositional embeddings offer a unified approach directly capable of all these functionalities.

Polymorphic embeddings [Hofer et al. 2008], tagless-final embeddings [Carette et al. 2009; Kiselyov 2010], and object algebras [Oliveira and Cook 2012] (denoted as POLY. in the table) also provide a unified encoding where most advantages of shallow and deep embeddings are available. However, they cannot deal with dependencies modularly. They encounter almost the same issue as in shallow embeddings, as described in Section 2.6. For example, we have to duplicate code of the size interpretation in the text interpretation, even if the size calculation is completely independent of the textual representation:

```
instance Region SizeAndText where
  circle r = ST { size = 1, text = "a circular region of radius " ++ show r }
  outside a = ST { size = a.size + 1, text = "outside a region of size " ++ show a.size }
  union a b = ST { size, text = "the union of two regions of size " ++ show size ++ " in total" }
  where size = a.size + b.size + 1
  .....
```

We refer curious readers to Appendix A for a detailed demonstration of how naive tagless-final embeddings cannot modularly handle dependencies in general.

Duplicating code every time we need a dependent interpretation is a serious problem since dependencies are extremely common. Nearly all non-trivial code will have dependencies. In functional programming, it is common to have functions defined by pattern matching that call other functions defined by pattern matching, which is where an important form of dependencies appears.

*The third dimension of the Expression Problem.* In essence, dependencies are a third dimension that is not discussed in the formulation of the Expression Problem by Wadler [1998]: whether dependencies require code duplication or not. If we take this into consideration, we can get a revised formulation of the Expression Problem, where instead of only two dimensions, we have a third dimension that evaluates the ability of an approach to deal with dependencies. With this extra dimension taken into account, what we get is Table 2. While tagless-final or other embeddings address the two original dimensions of extensibility modularly, they become non-modular with respect to dependencies, and programming with dependencies becomes significantly more awkward compared to conventional FP or OOP.

Note that we are not the first to observe the problem with dependencies. The problem above is known and discussed widely in the literature. In the context of polymorphic embeddings, Hofer et al. [2008] face the problem of dependencies and use the non-modular approach with tuples. Later Hofer and Ostermann [2010] adopt an extensible visitor pattern to improve modularity, but this results in much more boilerplate code. In his lecture notes on tagless-final interpreters, Kiselyov [2010] has a strong focus on discussing how to write non-compositional interpretations, which basically include operations with dependencies and nested pattern matching. He shows that non-compositional

638 programs can often be rewritten as programs that use contexts instead. The problem of modular  
639 dependencies has been widely discussed within the context of object algebras. Most recently, [Zhang  
640 and Oliveira \[2017, 2020\]](#) propose solutions to the problem using meta-programming techniques.  
641 [Zhang and Oliveira \[2019\]](#) further explore modular solutions in both Scala and Haskell, but their  
642 approaches still require a lot of boilerplate code. In Scala, it is due to a lack of language support for  
643 nested composition, whereas Haskell needs to encode dependencies via type classes to modularize  
644 dependent interpretations (see [Appendix B](#) for a similar solution). The drawbacks of the existing  
645 solutions above motivate the design of CP and compositional embeddings.

646 *Transformations.* It is widely believed that transformations and optimizations are much easier to  
647 write in deep embeddings [[Jovanović et al. 2014](#); [Scherr and Chiba 2014](#)]. We agree that it requires  
648 some ingenuity and is sometimes awkward to write complex transformations in compositional  
649 embeddings, but [Kiselyov \[2010\]](#) has demonstrated that several challenging transformations can be  
650 written with tagless-final embeddings. Compositional programming is a generalization of such forms  
651 of embeddings, and therefore all the transformations are possible with tagless-final embeddings are  
652 also possible in CP. Even for structurally non-local transformations, we can instantiate a sort with  
653 a function type (e.g. `RegionSig<Ctx→Region>`) and pass down the context when traversing the AST.  
654 Such contextual transformations can also be modular with the help of *polymorphic contexts* [[Zhang  
655 et al. 2021](#)]. We have implemented an example of common subexpression elimination inspired by  
656 [Kiselyov \[2011\]](#) (which can be found in our online implementation). Both our code and [Kiselyov’s](#)  
657 are not fully modular with respect to an auxiliary data structure (his implementation relies on  
658 a closed algebraic data type). In our case, this is because nested composition for recursive types  
659 is currently not supported, which relies on future theoretical progress of reconciling distributive  
660 subtyping and recursive types. Nevertheless, our approach is less entangled than tagless final  
661 embeddings since we can handle dependent transformations modularly.

## 663 4 E<sub>X</sub>T: A DSL FOR DOCUMENT AUTHORIZING

664 We have presented the advantages of compositional embeddings using a relatively small DSL. One  
665 may wonder if our approach scales to larger DSLs. As our answer to this question, we developed a  
666 larger DSL for document authoring called E<sub>X</sub>T.

667 E<sub>X</sub>T applies compositional embeddings to a document DSL, inspired by  $\LaTeX$  and Scribble [[Flatt  
668 et al. 2009](#)]. We have also done minor syntactic extensions to the CP language to support writing  
669 documents more naturally. Since documents usually involve large portions of text, it would be awk-  
670 ward to write such portions of text using string literals adopted by most programming languages.  
671 CP does not yet have facilities for syntactic extension that other languages (e.g. Racket, Haskell,  
672 or Coq) offer, so we have extended the parser directly. Thus, CP parses some document-specific  
673 syntax and desugars it to compositionally embedded fragments during parsing. Such a generative  
674 approach is similar to Racket macros [[Ballantyne et al. 2020](#)] and Template Haskell [[Sheard and  
675 Peyton Jones 2002](#)], which provide more flexible facilities for performing such desugaring. Never-  
676 theless, the surface syntax of E<sub>X</sub>T is essentially a set of lightweight syntactic sugar for its underlying  
677 compositional embeddings (similar, for instance, to Scribble in Racket).

### 679 4.1 Design Goals and Non-Goals

680 There are already a large number of document DSLs, so why shall we create yet another one? We  
681 explain why by identifying important design goals and non-goals of E<sub>X</sub>T.

682 *A document language for the web.* A first goal is to have a lightweight but powerful document  
683 language tailored mostly for the web. We view E<sub>X</sub>T as an alternative to Wikitext [[Wikipedia 2022b](#)],  
684 the markup language used by Wikipedia, which shares similar goals. Similar to Wikipedia pages,  
685  
686

687 users can directly edit ExT documents on a web page. But different from Parsoid [MediaWiki 2011],  
688 the official Wikitext parser that runs on the server side, our implementation can directly render  
689 documents on the browser side.

690 *General-purpose computation.* The majority of existing document DSLs (e.g. CommonMark [Mac-  
691 Farlane 2014], Textile [Allen 2002], and reStructuredText [Goodger 2002]) have a fixed set of  
692 language features and do not provide general-purpose computation. For instance, we may want to  
693 compute a table from some data, like a table listing the largest cities in descending order of popula-  
694 tion. For such kinds of tasks, it is useful to have mechanisms for general-purpose computation that  
695 allow us to *compute* documents rather than manually write them.

696 Although some other document DSLs provide language constructs that enable some computation,  
697 they still have significant limitations. For instance, Wikitext provides templates, which play a role  
698 similar to functions in conventional programming languages. However, templates cannot perform  
699 arbitrary computation as CP does. The documentation says, “A template can call itself but will stop  
700 after one iteration to prevent an infinite loop.” [Wikipedia 2022a] In other words, templates by  
701 themselves are incapable of recursion or loops, but in practice, loops are often needed in Wikipedia.  
702 For instance, `{{loop}}` has been used on approximately 99,000 pages [Wikipedia 2022c]. A Lua  
703 extension was introduced to bypass such restrictions, and the widely-used template `{{loop}}`  
704 invokes `string.rep` in Lua. Instead of handing it over to an external language, CP directly supports  
705 functions, and thus general-purpose computation is easily done in ExT, including recursion.

706 *Static type checking.* In contrast to many document languages, CP and ExT are statically type  
707 checked. Static type checking enables us to detect some errors that may arise when processing  
708 documents. In applications like wikis, it is very frequent that we need some form of structured data.  
709 States or cities, for example, may have infoboxes associated with them that list several pieces of data,  
710 such as areas, populations, official languages, etc. Using simple text to store such data can easily lead  
711 to simple errors. In addition, if we want to enable computed information as previously described,  
712 then it is useful to have a type system that allows us to write functions that take a well-defined  
713 form of structured data. Most document languages that we know of, including Wikitext, L<sup>A</sup>T<sub>E</sub>X, and  
714 Scribble, are not statically typed. We will discuss some examples where static typing is helpful in  
715 ExT in Section 4.4.

716 *Extensible meta-programming and customizability.* A benefit of modeling a DSL as an algebraic  
717 data type or a compositional interface is that we can easily write meta-functions over the abstract  
718 syntax of the DSL. For instance, in the context of documents, such meta-functions could include  
719 rendering to various backends (such as L<sup>A</sup>T<sub>E</sub>X or HTML) or computing a table of contents based  
720 on the document tree. In many external DSLs, such meta-functions are typically much harder to  
721 write, requiring us to directly modify the implementation. However, by embedding a DSL, the host  
722 language can act as a meta-language for the DSL and allow us to easily write meta-programs, for  
723 instance, as functions to process the AST by pattern matching.

724 In addition, the extensibility of compositional embeddings enables a form of extensible meta-  
725 programming: not only can we write meta-functions, but those meta-functions can be extended  
726 later to cater for new language constructs modularly defined by users. This enables DSL users to  
727 *modularly* extend the basic document language to add their own extensions, including infoboxes,  
728 vector graphics, etc. While Scribble does provide good facilities for syntactic extension and meta-  
729 programs, the document structure is fixed [Flatt and Barzilay 2009], so the DSL is not fully extensible.  
730 In short, ExT is designed to be extensible and highly customizable by users.

731 *Non-goals.* Currently, our implementation of CP is a fairly naive call-by-need interpreter, so the  
732 performance of ExT is not competitive with well-established tools like L<sup>A</sup>T<sub>E</sub>X. Furthermore, unlike  
733

|  |   |
|--|---|
| <pre> 736 \Section[ 737   Welcome to \Href("https://plground.org")[PLGround]! 738 ] 739 \Bold[CP] is a \Emph[compositional] programming 740 language. \\ There are \((1+1+1)\) key concepts in CP: 741 \Itemize[ 742   \Item[Compositional interfaces;] 743   \Item[Compositional traits;] 744   \Item[Method patterns;] 745   \Item[Nested trait composition.] 746 ] 747 748 (a) E<sub>X</sub>T code </pre> | <pre> <b>WELCOME TO PLGROUND!</b> CP is a <i>compositional</i> programming language. There are 4 key concepts in CP: <ul style="list-style-type: none"> <li>• Compositional interfaces;</li> <li>• Compositional traits;</li> <li>• Method patterns;</li> <li>• Nested trait composition.</li> </ul> </pre> |
| <pre> 749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765 766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782 783 784 </pre>   | <pre> (b) Rendered document </pre>  |

Fig. 3. A sample document illustrating E<sub>X</sub>T syntax.

Like L<sup>A</sup>T<sub>E</sub>X, which has extensive libraries developed over many years, there are no third-party libraries for E<sub>X</sub>T. Thus, E<sub>X</sub>T is *not* yet production-ready. Another non-goal is security. Since we never sanitize user-generated HTML, it is easy to perform JavaScript code injection on our website. It is off-topic to consider how to ensure web security here.

## 4.2 Syntax Overview

The syntax of E<sub>X</sub>T is designed for easy document authoring, as shown in Fig. 3. It is reminiscent of L<sup>A</sup>T<sub>E</sub>X, but still has significant differences. The similarity is that plain text can be directly written while commands start with a backslash. But we choose different brackets for command arguments:

- `\Cmd[...]` encloses document arguments. Square brackets indicate that a sequence of nested elements is recursively parsed. (Most commands in Fig. 3 are in such a form, for example, `\Section` and `\Bold`.)
- `\Cmd(...)` encloses positional arguments. Parentheses indicate that the whole construct is desugared to function application. (`\Href` in Fig. 3 is an example.)
- `\Cmd{...}` encloses labeled arguments. Braces indicate that these arguments are wrapped in a record. (As mentioned in Section 4.4, `\Line{ x1 = 0.0; y1 = 0.0; x2 = 4.6; y2 = 4.8 }` uses this form to distinguish different attributes.)

A command can be followed by an arbitrary number of different arguments in any order. This syntax is more flexible than the fixed form `@op[...] { ... }` in Scribble. Although both arguments are optional in Scribble, they can neither occur more than once nor be shuffled. In addition, there are two more useful E<sub>X</sub>T constructs. One is string interpolation with the form `\( ... )`, which has exactly the same syntax as the Swift programming language. We can see an example of string interpolation in Fig. 3: `\(1+1+1)`. It will be desugared to `Str (toString (1+1+1))`. This example will execute the code, which computes “4”, and its string form will be interpolated in the document. The other construct is double backslashes (`\\`), a shorthand for newline, which is borrowed from L<sup>A</sup>T<sub>E</sub>X. In E<sub>X</sub>T, it is equivalent to the `\Endl` command. There is also an example in Fig. 3, which inserts a newline character between the first and second sentences in the body.

*Desugaring.* As mentioned earlier, all E<sub>X</sub>T constructs are desugared to normal CP code. For example, the following two expressions are equivalent (the document DSL is enclosed within backticks):

```

`\\Entry("id"){ hidden = true; level = 1 }[This entry is \Emph[hidden]]` -- is desugared to:
Entry "id" { hidden = true; level = 1 } (Comp (Str "This entry is ") (Emph (Str "hidden")))

```



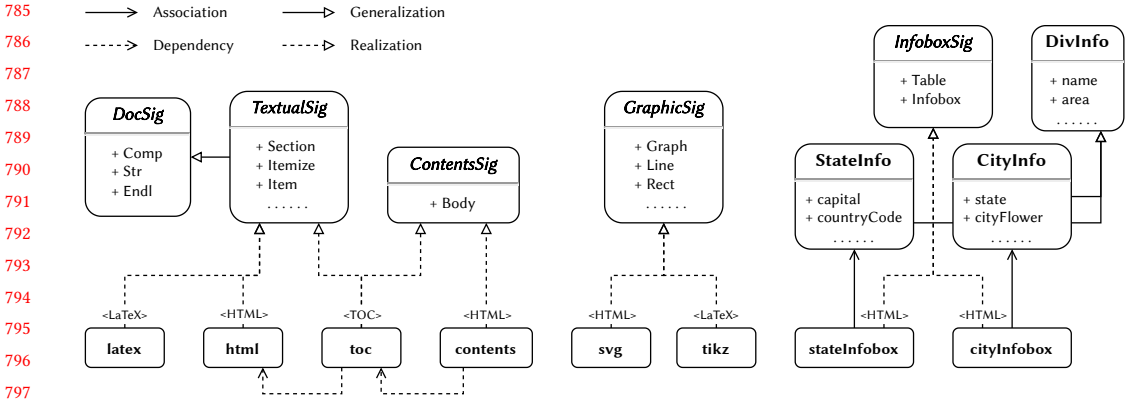


Fig. 4. A simplified diagram of ExT components.

Here, `Comp` and `Str` are two special commands that are assumed to be implemented by library authors. `Comp` is used for composing an AST of document elements, while `Str` is for plain text.

*Specification.* The grammar of ExT can be summarized in whole as follows (`<arg>*` means that `<arg>` may occur zero or more times):

```

786 <doc> ::= <str> | "\<esc>                <str> ::= text without "\" and "]"
787 <esc> ::= <cmd> <arg>* | "(" <exp> ")" | "\" <cmd> ::= identifier
788 <arg> ::= "[" <doc> "]" | "(" <exp> ")" | "{" <rcd> "}" <exp> ::= expression
789                                     <rcd> ::= record fields
    
```

### 4.3 Extensible Commands with Multi-Backend Semantics

ExT is so flexible that only a few commands are really compulsory. Additional commands and their semantics can be extended as needed. At the very beginning, we only need to define the aforementioned three special commands:

```

814 type DocSig<Element> = {
815   Comp : Element → Element → Element;   Str : String → Element;   Endl : Element;
816 };
    
```

However, for rich text in real-world documentation, these commands are not sufficient for various document elements. Adding new commands and their semantics is not hard at all in compositional embeddings, like what we have done in Section 3.3 and Section 3.4. Language interfaces can be separately declared and then intersected with each other. Besides extensible commands, their semantics are also retargetable. We can modularly create different traits for different backends, such as HTML and L<sup>A</sup>T<sub>E</sub>X:

```

824 type HTML = { html : String };
825 html = trait implements DocSig<HTML> => {
826   (Comp l r).html = l.html ++ r.html;
827   (Str s).html = s;
828   (Endl ).html = "<br>";
829 };
830
831 type LaTeX = { latex : String };
832 latex = trait implements DocSig<LaTeX> => {
833   (Comp l r).latex = l.latex ++ r.latex;
834   (Str s).latex = s;
835   (Endl ).latex = "\\\\";
836 };
    
```

*The components of ExT.* We have implemented several extensions of ExT, as shown in Fig. 4. In the upper part of the diagram, the capitalized names in a bold italic font (e.g. **DocSig**) are compositional interfaces, while those in a bold upright font (e.g. **DivInfo**) are normal types. The smaller boxes at

834 the bottom of the diagram are compositional traits. They implement different interfaces with the  
 835 sorts specified on the dashed arrows, (e.g. <HTML>). Specifically, TextualSig adds a set of common  
 836 commands for rich text, such as Section, Itemize, Item, etc. It extends DocSig and targets both  
 837 HTML and  $\LaTeX$ . GraphicSig is used to draw vector graphics (including lines, rectangles, circles,  
 838 etc.), which are supported in HTML and  $\LaTeX$  via SVG and TikZ respectively. InfoboxSig targets  
 839 only HTML and adds Wikipedia-like infoboxes (containing, for instance, information about areas  
 840 and populations of states or cities). All these compositional interfaces and traits can be combined  
 841 to form a heterogeneous composition. Notably, the implementation of ContentsSig, which is used  
 842 to generate a table of contents, introduces some non-trivial dependencies (see [Section 5.1](#)).

843

#### 844 4.4 Static Typing

845 A major difference of  $\text{E}\lambda\text{T}$  from other document DSLs is that it is statically type checked. With  
 846 static typing, potential type errors can be detected ahead of time, saving users' time for debugging.  
 847 For example, when drawing a line using SVG, we need to use the <line> element with four  
 848 numeric coordinates (x1, y1, x2, and y2). However, since additional information of an SVG element  
 849 is represented as XML attributes, all of the coordinates have to be quoted as strings instead of  
 850 numbers. Only after rendering it with a web browser and opening developer tools can one check if  
 851 attributes are valid. In  $\text{E}\lambda\text{T}$ , we model the line construct like this:

852

```
853 type GraphicSig<Graphic> = {
854   Line : { x1: Int; y1: Int; x2: Int; y2: Int } → Graphic;
855   -- and more constructors
856 };
```

856

857 Before rendering lines via SVG, the types of arguments are checked against the language interface,  
 858 and invalid values are rejected in advance. Note that the error messages are reported in terms of  
 859 the  $\text{E}\lambda\text{T}$  surface syntax, which is more friendly to users than some meta-programming techniques,  
 860 where errors are reported in terms of the generated code.

861

862 When modeling infoboxes and bibliography in [Section 5.1](#), we also make use of data types to  
 863 represent structured information. For example, we use DivInfo to model common information  
 864 required by all political divisions, as well as two specific types for states and cities respectively,  
 865 which extend DivInfo using intersection types:

865

```
866 type DivInfo = {
867   name      : String;
868   area      : Int;
869   population : Int;
870   languages : [String];
871 };
872
873 type StateInfo = DivInfo & {
874   countryCode : String;
875   religions   : [String];
876   -- and more fields
877 };
878
879 type CityInfo = DivInfo & {
880   cityFlower : String;
881   timeZone   : String;
882   -- and more fields
883 };
```

870

871 In this case, there is quite a bit of structure in the data. It statically describes what fields are allowed  
 872 and what types these fields have.

873

## 874 5 EVALUATION OF COMPOSITIONALLY EMBEDDED $\text{E}\lambda\text{T}$

875

876 In this section, we describe three applications of CP, compare them with alternatives in other docu-  
 877 ment languages, and evaluate the use of compositionally embedded  $\text{E}\lambda\text{T}$  for the three applications.

877

### 878 5.1 Minipedia

879

880 Minipedia is a mini document repository of states and cities, which reconstructs a small portion  
 881 of Wikipedia. Minipedia currently contains pages of the world's ten smallest countries and their  
 882 capitals, as well as a structured data file storing all facts about them used for infoboxes. There are

882

also pages consisting of sorted tables of microstates by area or population, which are computed using the information stored in the data file. Minipedia comprises over 4,000 lines of code in total, most of which is accounted for by document pages. Among the rest, the data file and E<sub>x</sub>T libraries account for around 300 and 200 lines of code, respectively.

*Drawbacks of Wikipedia.* Compared to our approach, Wikipedia and its markup language Wikitext have several drawbacks, as partly mentioned in [Section 4.1](#):

- All data in Wikitext is in string format. Wikitext does not differentiate data types like **Int**, **Bool**, etc. This makes type errors remain uncaught in Wikipedia infoboxes. For example, in the infobox of a country, the population field can be changed to a non-numeric meaningless value, and Wikipedia will not warn the editor at all.
- Statistical data is manually written on every Wikipedia page. This can easily result in data inconsistency for the same field across different pages, and even for different places on the same page. For example, the population size on a country page is often different from that in a statistical page listing all countries by population, owing to different data sources.
- Wikitext provides user-defined templates as a reusable unit of document fragments. They can be parametrized and play a similar role to functions. However, they do not natively support general-purpose computation like recursion. Along with the Lua extension and parser functions, Wikitext offers some computational power to Wikipedia documents, but it is not as easy to use as E<sub>x</sub>T is.

*Type safety and data consistency.* Minipedia has a structured data file containing the information of states and cities, each piece of which is modeled as a record. The record types have been shown in [Section 4.4](#). Every record for states or cities is type checked against their corresponding types and collected in two arrays in the data file:

```

808  tuvalu = {                                funafuti = {
809      name      = "Tuvalu";                name      = "Funafuti";
810      area      = 26;                       area      = 2;
811      population = 10645;                   population = 6320;
812      -- and more fields                    -- and more fields
813 } : StateInfo;                            } : CityInfo;
814 states = [tuvalu; {- and more states -}];  cities = [funafuti; {- and more cities -}];

```

Such a centralized data file forces different documents to read from the same data source and prevents data inconsistency. Not only infoboxes but also computed documents like sorted lists of countries can be created using the information from the data file. Moreover, if the area of Tuvalu, for example, is assigned a string value, there will be a type error before rendering.

*Writing Minipedia pages in E<sub>x</sub>T.* Writing Minipedia documents is enabled by the E<sub>x</sub>T libraries shown in [Fig. 4](#). Different features of E<sub>x</sub>T are defined in different libraries, and we can import whatever libraries we want when writing a document. For Minipedia, we only need HTML-related libraries. There are two modes for document authoring in Minipedia: “doc-only” and “program”. In the “doc-only” mode, with required libraries specified, documents are written directly in E<sub>x</sub>T without any wrapping code in CP, as shown in [Fig. 3](#). In the “program” mode, a document is created programmatically in a mixture of CP and E<sub>x</sub>T code. This mode makes general-purpose computation easier to write in a document.

For example, a sorted table of the smallest countries by area demonstrates general-purpose computation in the “program” mode. The program begins with an **open** directive, which imports definitions from the specified libraries. We sort the array of countries imported from Database using

932 a predefined function `sort` and then iterate it to generate  $\text{E}\lambda\text{T}$  code. As explained in [Section 3.2](#), we  
 933 use self-type annotations to inject dependencies on the library commands. The intersection type  
 934 `DocSig<T>&TableSig<T>` allows  $\text{E}\lambda\text{T}$  commands from both compositional interfaces to be used in the  
 935 trait `doc`:

```
936 open LibDoc; open LibTable; open Database;
937 sortedStates = sort states;
938 doc T = trait [self : DocSig<T>&TableSig<T>] => {
939   table = letrec rows (i:Int) : T = if i == 0 then `Trow[
940     \Theader[No.] \Theader[Country] \Theader[Area (km^2)]
941   ]` else let state = sortedStates!!(i-1) in `rows(i-1) \Trow[
942     \Tdata[\(i+1)] \Tdata[\(state.name)] \Tdata[\(state.area)]
943   ]` in `tbody[ \rows(#sortedStates) ]`;
944 };
945 document = new doc @HTML , html , table;
946 document.table.html
```

947 *Adding a table of contents (TOC).* We have introduced how to write Minipedia pages using  $\text{E}\lambda\text{T}$   
 948 libraries, but now let us go back and see how we implement a  $\text{E}\lambda\text{T}$  library. The library of TOC adds a  
 949 new command `\Body[...]`, which prepends a TOC to the document body. The difficulty here is that  
 950 HTML rendering depends on the TOC, whereas the TOC in turn depends on the HTML rendering  
 951 of section (and sub<sup>n</sup>section) titles. As we have demonstrated in [Section 3.5](#), we can modularly tackle  
 952 such complex dependencies by refining the interface types:

```
953 type ContentsSig<Element> = { Body : Element -> Element };
954 type TOC = { toc : String };
955
956 contents = trait implements ContentsSig<TOC => HTML> => {
957   [self]@(Body e).html = self.toc ++ e.html;
958 };
959 toc = trait implements TextualSig<HTML => TOC> & ContentsSig<TOC> => {
960   (Comp l r).toc = l.toc ++ r.toc;
961   (Section e).toc = listItem 0 e.html;
962   (SubSection e).toc = listItem 1 e.html;
963   (SubSubSection e).toc = listItem 2 e.html;
964   (Body e).toc = "<ul>" ++ e.toc ++ "</ul>";
965     ..toc = "";
966 };
```

967 With TOC specified in `ContentsSig<TOC => HTML>` and the self-reference specified in `[self]`, we  
 968 can call `self.toc` to generate a TOC when rendering HTML. Similarly, we can call `e.html` when  
 969 generating the TOC since the interface type is refined. A minor remark is that `listItem` is an  
 970 auxiliary function to generate an appropriate `<li>` wrapping for a given nesting level.

## 971 5.2 Fractals and Sharing

972 In the second application, we briefly introduce how to implement computational graphics like  
 973 fractals in  $\text{E}\lambda\text{T}$  with the help of linguistic reuse. Fractal drawing is all about repeating the same  
 974 pattern. This drives us to exploit meta-language sharing to avoid redundant computation. We  
 975 have implemented several well-known fractals, such as the Koch snowflake, the T-square, and the  
 976 Sierpiński carpet. For the sake of simplicity, we take the Sierpiński carpet as an example:  
 977

```
978 fractal T C = trait [self : DocSig<T> & GraphicSig<T><C> & ColorSig<C> & Draw T C]
979   implements Draw T C => {
```

980

```

981 draw {..} =
982   let center = Rect { x = x + width/3; y = y + height/3;
983                     width = width/3; height = height/3; color = White } in
984   if level == 0 then center else
985     let w = width/3 in let h = height/3 in let l = level-1 in
986       let shared = draw { x = x; y = y; width = w; height = h; level = l } in `\Group(id)[
987         \shared           \Translate{x=w;y=0}(shared)  \Translate{x=2*w;y=0}(shared)
988         \Translate{x=0;y=h}(shared)  \center           \Translate{x=2*w;y=h}(shared)
989         \Translate{x=0;y=2*h}(shared) \Translate{x=w;y=2*h}(shared) \Translate{x=2*w;y=2*h}(shared)
990       ]`
991 };

```

992 We make variables shared via the **let** expressions in CP. There are many uses in the code above but,  
 993 among them, shared accelerates fractal generation the most. The variable shared is later used eight  
 994 times to constitute the repeating patterns in the Sierpiński carpet. `\Translate` is a command declared  
 995 in `GraphicSig` for geometric translations. With automatic linguistic reuse in `ExT`, we only need to  
 996 generate a pattern once instead of repeating it eight times. Besides the efficiency in generating  
 997 fractals, our implementation of `\Translate` makes use of the `<use>` element in `SVG` to avoid the  
 998 exponential growth of output. Therefore, the duplication of `SVG` elements is also eliminated. These  
 999 optimizations are available for free thanks to the linguistic reuse in compositional embeddings.

### 1000 5.3 Customizing Charts

1001 In the last application, we illustrate how line charts and bar charts are modularly rendered using  
 1002 external data. Charts can be rendered into a document using the `ExT` language and its support  
 1003 for vector graphics. What is interesting about charts is that there are many alternative ways to  
 1004 present charts and customize them. The flexibility of the CP language, in terms of modularity and  
 1005 compositionality, is then very useful here. We will show how to adapt traditional object-oriented  
 1006 design patterns for compositional embeddings when modeling graphic components.

1007 *Charting stocks.* Taking stock prices as an example, we have a separate file containing data from  
 1008 big companies, as well as a few configuration items for chart rendering. The configuration includes  
 1009 the choice between lines and bars, whether to show borders or legends, what labels to display,  
 1010 and so on. Serving as an infrastructure, a base chart is created with the drawing functions for the  
 1011 caption and axes:

```

1012
1013
1014 type Base = { caption : HTML; xAxis : HTML; yAxis : HTML };
1015 baseChart (data : [Data]) = trait implements Base => {
1016   caption = `...`;  xAxis = `...`;  yAxis = `...`;
1017 };

```

1018 The simplified code above shows a rough sketch, where a base trait takes data as its parameter and  
 1019 implements the base interface. However, it does not implement the primary rendering function. As  
 1020 we have two options to visualize stock prices, we leave the choice to the STRATEGY pattern [Gamma  
 1021 et al. 1995].

1022 *STRATEGY pattern.* Since the configuration is unknown until run time, the rendering strategy  
 1023 cannot be hard-coded in the base trait. Instead, we define two strategy traits implementing the  
 1024 rendering, respectively for lines and bars. They constitute two variants of the base chart. It is not  
 1025 the programmer but the configurator that decides which variant of charts should be rendered. In  
 1026 the original STRATEGY pattern, a chart object delegates to a desired strategy object, to which its  
 1027 mutable field points. That mutable field can be changed from clients who use the chart. In CP,  
 1028  
 1029

1030 we just need to merge the base trait with the strategy we want, without the need for mutable  
1031 references:

```
1032 type Render = { render : HTML };
1033 lineStrategy = trait [self : Base] implements Render => {
1034   render = let lines = ... in `xAxis \yAxis \lines \caption`
1035 };
1036 barStrategy = trait [self : Base] implements Render => {
1037   render = let bars = ... in `xAxis \yAxis \bars \caption`
1038 };
1039 chart = baseChart data , if config.line then lineStrategy else barStrategy;
```

1040 After merging, chart is still a *reusable trait*. This is impossible in traditional object-oriented  
1041 languages, like Java, because they lack a dynamic way to compose classes at run time. In contrast,  
1042 such a dynamic trait composition is ubiquitous in CP. It is also easy to add a new strategy, like  
1043 switching to pie charts, and merge the base trait with the new one instead. If a client forgets to pick  
1044 any strategy, the type checker will notify them because the trait types before and after merging are  
1045 different. Moreover, another advantage of our approach over the original strategy pattern is that the  
1046 self-type annotations make methods in the base trait accessible to strategy traits. So caption, xAxis  
1047 and yAxis are directly shared with strategies, requiring no extra effort to pass them as arguments  
1048 when delegating. This remedies the “communication overhead between Strategy and Context”  
1049 caused by the original STRATEGY pattern [Gamma et al. 1995].  
1050

1051 *DECORATOR pattern.* Besides the STRATEGY pattern, the DECORATOR pattern [Gamma et al. 1995]  
1052 is also adapted for CP. Since decorations are also determined by external configuration at run time,  
1053 we need a modular way to add additional drawing processes to the chart trait dynamically. When  
1054 multiple decorators are enabled, their functionalities should be added. In the original DECORATOR  
1055 pattern, the decorator class should inherit the chart class and be instantiated with a reference to a  
1056 chart object. The decorator will store the chart object as a field and forward all methods to it. In  
1057 concrete decorators, the rendering function will be overridden to add decorations. The decorator  
1058 pattern sounds complicated in traditional object-oriented programming, but the same goal can be  
1059 achieved easily in CP using the *dynamic inheritance* provided by CP:  
1060

```
1061 type Chart = Base & Render;
1062 borderDecorator (chart : Trait<Chart>) = trait [self : Chart] implements Chart inherits chart => {
1063   override render = let sr = super.render in let border = ... in `sr \border`
1064 };
1065 legendDecorator (chart : Trait<Chart>) = trait [self : Chart] implements Chart inherits chart => {
1066   override caption = let legends = ... in `legends`
1067 };
1068 chart' = if config.border then borderDecorator chart else chart;
1069 chart'' = if config.legend then legendDecorator chart' else chart';
```

1070 A decorator takes a chart trait as the argument and creates another trait inheriting it dynamically. In  
1071 the implementation, we can override any methods as usual, and static type safety is still guaranteed.  
1072 It is worth noting that, in legendDecorator, only caption is overridden while render is just inherited.  
1073 In the original DECORATOR pattern, render will never be conscious of the overridden caption since  
1074 the decorator hands over control to chart after forwarding render. This is partly why Gamma  
1075 et al. [1995] warn readers that “a decorator and its component aren’t identical”. However, in CP,  
1076 chart.render can access the overridden caption through the late-bound self-reference. This solution  
1077 makes our code clean and easy to maintain.  
1078

1079 *True delegation via trait composition.* The chart application shows how other aspects of the  
 1080 modularity of CP are helpful to create highly customizable DSLs. We have shown how CP adapts  
 1081 and simplifies object-oriented design patterns. In general, we avoid complicated class hierarchies in  
 1082 the original versions and replace them with relatively simple trait composition. It showcases that we  
 1083 can get rid of verbose design patterns if a proper language feature fills in the gap. In both patterns,  
 1084 component adaptation is needed at run time, so delegation is at the heart of the challenges.

## 1085 6 DISCUSSION AND RELATED WORK

1086  
 1087 In this section, we discuss how the ideas of Compositional Programming and compositional  
 1088 embeddings can be more generally applied to other languages as well as related work.

### 1089 6.1 Encoding Compositional Embeddings in Other Languages

1090  
 1091 There are three key features that enable compositional embeddings in CP: *compositional interfaces*,  
 1092 *self-references*, and *nested composition*. All the three features can be encoded, with some effort and  
 1093 boilerplate code, in other programming languages, such as Haskell, Scala, or even Java.

1094 *Compositional interfaces* can be encoded indirectly using Haskell type classes, Scala traits, or  
 1095 Java interfaces. For example, the compositional interface on the left corresponds to the Scala trait  
 1096 for generalized object algebras [Oliveira et al. 2013] on the right:

```
1097 type HoferSig<Region> = {                               trait HoferSig[[In,Out]] {
1098   Univ  : Region;                                       def Univ  : Out
1099   Empty : Region;                                       def Empty : Out
1100   Scale : Vector → Region → Region;                   def Scale : (Vector,In) ⇒ Out
1101 };                                                       }
```

1102 Here we distinguish between types used in input positions and types used in output positions. This  
 1103 distinction is the key to the mechanism of dependency injection in CP [Zhang et al. 2021].

1104 *Self-references* are also essential to dependency injection in CP. Complex interpretations may  
 1105 depend on self dependencies other than child dependencies [Zhang et al. 2021], which has occurred  
 1106 in the example of the text interpretation:

```
1107 [self]@(Translate v a).text = "a translated region of size " ++ toString self.size;
```

1108  
 1109 The method pattern above refers to `self.size` instead of `a.size`, which is seldom directly supported  
 1110 in existing embedding techniques. This feature can be modeled via open recursion, or we may  
 1111 reuse the available self-references in some object-oriented languages like Scala.

1112 *Nested composition* can be expressed in existing languages by writing explicit composition  
 1113 operators. For example, a merge operator for composing `Eval` and `Count` will have such a type:

```
1114 merge : HoferSig<Eval> → HoferSig<Count> → HoferSig<Eval&Count>
```

1115  
 1116 In essence, a merge operator takes two interpretations and creates a single merged one. Meanwhile,  
 1117 this is where most boilerplate code would come from since we need to write a new implementation  
 1118 for every new signature of the merge operator. A workaround is to generate this kind of boilerplate  
 1119 code using meta-programming.

1120 A detailed description of encoding the key features in Scala can be found in work by Oliveira et al.  
 1121 [2013]. Inspired by Bi et al. [2019], we also demonstrate a Haskell encoding in Appendix B. With  
 1122 the help of quite a few GHC language extensions, dependent interpretations are defined modularly  
 1123 in a tagless-final style, and their composition is directed by type classes. However, self-references  
 1124 are still not supported in that Haskell encoding, so we cannot directly address self dependencies  
 1125 like `self.size`. Moreover, the boilerplate code of smart constructors is needed for every language  
 1126 construct in the region DSL, and explicit projections are inevitable when writing interpretations.

1128 Compositional embeddings are more compact and cleaner in CP thanks to its language-level support  
1129 for all the three key features.

1130

## 1131 6.2 Related Work

1132 *Modular embeddings.* There are quite a few existing approaches to modular embeddings, which  
1133 essentially exploit solutions to the *Expression Problem* [Wadler 1998] in existing languages. Such  
1134 approaches include, for example, *tagless-final embeddings* [Carette et al. 2009] and *data types à*  
1135 *la carte* [Swierstra 2008] in functional programming, *polymorphic embeddings* [Hofer et al. 2008]  
1136 and *object algebras* [Oliveira and Cook 2012] in object-oriented languages, just to name a few.  
1137 Section 3.6 already compares those approaches with compositional embeddings in detail. In essence,  
1138 the lack of sufficient programming language support for modularity makes it hard to model modular  
1139 dependencies in those approaches. Compositional embeddings offer an elegant solution to many  
1140 issues by exploiting the support for Compositional Programming in the CP language.

1141

1142 *Hybrid embeddings.* Svenningsson and Axelsson [2015] propose combining shallow and deep  
1143 embeddings by translating a shallowly embedded interface to a deeply embedded core language. In  
1144 this way, language interfaces can be shallowly extended, and multiple interpretations are possible  
1145 in the deep core. In addition, some features in the host language can be exploited in the shallow  
1146 part, which is called *deep linguistic reuse* in Scala-Virtualized [Rompf et al. 2012]. Jovanović et al.  
1147 [2014] further propose an automatic translation between shallow and deep embeddings using Scala  
1148 macros instead of a manual translation. Their Yin-Yang framework targets lightweight modular  
1149 staging [Rompf and Odersky 2010] as the deep embedding backend but completely conceals  
1150 internal encodings from the users. A similar but slightly different approach called *implicit staging* is  
1151 proposed by Scherr and Chiba [2014]. Implicit staging extracts an intermediate representation from  
1152 a shallowly embedded DSL and reintegrates it after some transformations. Their research prototype  
1153 works in Java through load-time reflection. Compared to such hybrid approaches, compositional  
1154 embeddings do not require the use of two different embeddings and the translations between them  
1155 while supporting most features from both shallow and deep embeddings.

1156

1157 *Generative programming.* As stated in Section 4, ExT performs a lightweight desugaring during  
1158 parsing to generate compositionally embedded fragments in CP. This is an *ad-hoc* generative  
1159 technique. Some generative approaches in other languages, such as Racket macros [Ballantyne et al.  
1160 2020] and Template Haskell [Sheard and Peyton Jones 2002], provide more flexible mechanisms  
1161 compared to ours. Nevertheless, in this work, our goal is not to develop generative programming  
1162 in CP. Instead, we try to keep the DSL as close to the underlying compositional embeddings as  
1163 possible so that there is an obvious one-to-one mapping between ExT and CP. Rather than syntactic  
1164 extensibility (i.e. the ability to extend the syntax of the source language), our focus is on semantic  
1165 support for DSLs. Our approach is different from some generative frameworks like EVF [Zhang  
1166 and Oliveira 2017] and Castor [Zhang and Oliveira 2020], which generate extensible visitors from  
1167 annotations via meta-programming. Although they also enable modular and extensible program-  
1168 ming language components, their code uses various non-standard annotations, whose semantics  
1169 may be unclear to programmers. What is worse, type checking is delayed in the aforementioned  
1170 generative frameworks so error messages are reported in terms of the generated code. This can be  
1171 quite confusing without detailed knowledge of how the generated code works. In contrast, in our  
1172 implementation, type checking is done in the source language, and error messages are reported in  
1173 terms of the original ExT code.

1174

1175 *Document DSLs.* There are quite a few existing DSLs for document authoring.  $\LaTeX$  is one of  
1176 the most commonly used document languages, and ExT's syntax is inspired by it. In contrast to

1176



1177  $\text{E}\bar{\text{x}}\text{T}$ ,  $\text{L}\bar{\text{T}}\bar{\text{E}}\bar{\text{X}}$  is an external DSL and does not have a static type system.  $\text{L}\bar{\text{T}}\bar{\text{E}}\bar{\text{X}}$  supports arbitrary  
1178 computation, but it is not easy to write meta-programs over the AST of  $\text{L}\bar{\text{T}}\bar{\text{E}}\bar{\text{X}}$ . Moreover,  $\text{L}\bar{\text{T}}\bar{\text{E}}\bar{\text{X}}$  is not  
1179 intended to render web documents, which is an important goal of  $\text{E}\bar{\text{x}}\text{T}$ . *Wikitext* [Wikipedia 2022b]  
1180 is widely used across wiki websites, including Wikipedia and Fandom. Wikitext has some design  
1181 limitations like the absence of type safety, data consistency, and general-purpose computation,  
1182 which have been elaborated in Section 5.1. *Skribe* [Gallesio and Serrano 2005] is a document DSL  
1183 built on top of the Scheme programming language and extends Scheme’s syntax from S-expressions  
1184 to *Sk-expressions* for easier text writing. Skribe makes use of Scheme macros and functions to  
1185 construct documents. This idea is inherited and popularized by *Scribble* [Flatt et al. 2009] in the  
1186 Racket community, and DrRacket offers very good tool support for Scribble. Regarding syntax,  
1187 a new *@-notation* is designed for Scribble, which is desugared to normal S-expressions.  $\text{E}\bar{\text{x}}\text{T}$ ’s  
1188 syntax is similar to Scribble’s but is more flexible. Based on the infrastructure provided by Racket,  
1189 Scribble supports arbitrary computation and linguistic reuse like  $\text{E}\bar{\text{x}}\text{T}$ . However, Scribble has a fixed  
1190 document structure [Flatt and Barzilay 2009], and its rendering function performs conventional  
1191 pattern matching on it. Thus, it is not easy to extend their core document constructs and functions  
1192 without modifying the original source code. Furthermore, there is no static typing in Scribble,  
1193 though it uses Racket’s *contract* system to perform run-time type checking.

1194

## 1195 7 CONCLUSION

1196 We have presented compositional embeddings and shown their advantages compared to existing  
1197 embedding techniques in terms of modularity and compositionality. CP’s support for dependency  
1198 injection, pattern matching, and nested composition makes *seemingly* non-compositional inter-  
1199 pretations feasible in compositional embeddings. Thus, dependent interpretations and complex  
1200 transformations can be modularly defined. These virtues of compositional embeddings lay down  
1201 the foundation for the  $\text{E}\bar{\text{x}}\text{T}$  DSL. The viability of compositionally embedded  $\text{E}\bar{\text{x}}\text{T}$  has been evalu-  
1202 ated with three applications: *Minipedia* shows its extensibility, type safety, data consistency, and  
1203 ability to modularly handle dependencies; *fractals* show that general-purpose computation and  
1204 linguistic reuse are both available; finally, *charts* show a simpler approach to delegation compared  
1205 to object-oriented design patterns. Last but not least, an in-browser implementation of CP and  $\text{E}\bar{\text{x}}\text{T}$   
1206 is publicly available, and all code in the present paper can be type checked and run online.

1207 Compared to other techniques, compositional embeddings require dedicated programming  
1208 language support. While we are working hard on improving CP and its support for Compositional  
1209 Programming, we believe that an equally important outcome of our work is to inform current and  
1210 future language designers about the benefits and applications of Compositional Programming and  
1211 compositional embeddings. We hope that the results of our work will lead to other programming  
1212 languages considering extensions that support Compositional Programming and compositional  
1213 embeddings as well.

1214

1215 *Future work.* Since CP and  $\text{E}\bar{\text{x}}\text{T}$  are research prototypes, they are not yet production-ready. Our  
1216 implementation of CP is an in-browser interpreter, which is not so fast as industrial-strength  
1217 language implementations. We expect to improve its dynamic semantics and write a compiler for  
1218 CP, targeting either JavaScript or WebAssembly.

1219 Currently,  $\text{E}\bar{\text{x}}\text{T}$  is evaluated with some relatively small-scale applications. To show that  $\text{E}\bar{\text{x}}\text{T}$  can  
1220 provide a realistic alternative to established wiki languages, we need to deploy  $\text{E}\bar{\text{x}}\text{T}$  on a much larger  
1221 scale. At the moment, Minipedia is a toy compared to ultra-large wiki websites, such as Wikipedia.  
1222 We hope to attract more people to try  $\text{E}\bar{\text{x}}\text{T}$  on our website and send feedback to us. We also need  
1223 pragmatic advice from wiki experts to make  $\text{E}\bar{\text{x}}\text{T}$  a realistic alternative to Wikitext or other wiki  
1224 markup languages.

1225

## REFERENCES

- 1226  
1227 Dean Allen. 2002. Textile Markup Language Documentation. <https://textile-lang.com>
- 1228 Sven Apel, Don Batory, Christian Kästner, and Gunter Saake. 2016. *Feature-Oriented Software Product Lines: Concepts and*  
1229 *Implementation*. Springer-Verlag.
- 1230 Michael Ballantyne, Alexis King, and Matthias Felleisen. 2020. Macros for Domain-Specific Languages. In *OOPSLA*.
- 1231 Xuan Bi and Bruno C. d. S. Oliveira. 2018. Typed First-Class Traits. In *ECOOP*.
- 1232 Xuan Bi, Bruno C. d. S. Oliveira, and Tom Schrijvers. 2018. The Essence of Nested Composition. In *ECOOP*.
- 1233 Xuan Bi, Ningning Xie, Bruno C. d. S. Oliveira, and Tom Schrijvers. 2019. Distributive Disjoint Polymorphism for Composi-  
1234 tional Programming. In *ESOP*.
- 1235 Richard Boulton, Andrew Gordon, Mike Gordon, John Harrison, John Herbert, and John Van Tassel. 1992. Experience with  
1236 Embedding Hardware Description Languages in HOL. In *TPCD*.
- 1237 Jacques Carette, Oleg Kiselyov, and Chung-chieh Shan. 2009. Finally Tagless, Partially Evaluated: Tagless Staged Interpreters  
1238 for Simpler Typed Languages. *J. Funct. Program.* 19, 5 (2009).
- 1239 Koen Claessen and John Hughes. 2000. QuickCheck: A Lightweight Tool for Random Testing of Haskell Programs. In *ICFP*.
- 1240 Stéphane Ducasse, Oscar Nierstrasz, Nathanael Schärli, Roel Wuyts, and Andrew P. Black. 2006. Traits: A Mechanism for  
1241 Fine-grained Reuse. *ACM Trans. Program. Lang. Syst.* 28, 2 (2006).
- 1242 Jana Dunfield. 2014. Elaborating Intersection and Union Types. *J. Funct. Program.* 24, 2-3 (2014).
- 1243 Burak Emir, Martin Odersky, and John Williams. 2007. Matching Objects with Patterns. In *ECOOP*.
- 1244 Erik Ernst. 2001. Family Polymorphism. In *ECOOP*.
- 1245 Erik Ernst, Klaus Ostermann, and William R. Cook. 2006. A Virtual Class Calculus. In *POPL*.
- 1246 Matthew Flatt and Eli Barzilay. 2009. Low-Level Scribble API: Structures And Processing. [https://docs.racket-lang.org/  
1247 scribble/core.html](https://docs.racket-lang.org/scribble/core.html)
- 1248 Matthew Flatt, Eli Barzilay, and Robert Bruce Findler. 2009. Scribble: Closing the Book on Ad Hoc Documentation Tools. In  
1249 *ICFP*.
- 1250 Erick Gallesio and Manuel Serrano. 2005. Scribe: a Functional Authoring Language. *J. Funct. Program.* 15, 5 (2005).
- 1251 Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. 1995. *Design Patterns: Elements of Reusable Object-Oriented*  
1252 *Software*. Addison-Wesley Professional.
- 1253 Jeremy Gibbons and Nicolas Wu. 2014. Folding Domain-Specific Languages: Deep and Shallow Embeddings (Functional  
1254 Pearl). In *ICFP*.
- 1255 Andy Gill. 2009. Type-Safe Observable Sharing in Haskell. In *Haskell@ICFP*.
- 1256 David Goodger. 2002. reStructuredText: Markup Syntax and Parser Component of Docutils. [https://docutils.sourceforge.io/  
1257 rst.html](https://docutils.sourceforge.io/rst.html)
- 1258 Christian Hofer and Klaus Ostermann. 2010. Modular Domain-Specific Language Components in Scala. In *GPCE*.
- 1259 Christian Hofer, Klaus Ostermann, Tillmann Rendel, and Adriaan Moors. 2008. Polymorphic Embedding of DSLs. In *GPCE*.
- 1260 Paul Hudak. 1998. Modular Domain Specific Languages and Tools. In *ICSR*.
- 1261 John Hughes. 1995. The Design of a Pretty-printing Library. In *AFP*.
- 1262 Vojin Jovanović, Amir Shaikhha, Sandro Stucki, Vladimir Nikolaev, Christoph Koch, and Martin Odersky. 2014. Yin-Yang:  
1263 Concealing the Deep Embedding of DSLs. In *GPCE*.
- 1264 Oleg Kiselyov. 2010. Typed Tagless Final Interpreters. In *SSGIP*.
- 1265 Oleg Kiselyov. 2011. Implementing Explicit and Finding Implicit Sharing in Embedded DSLs. In *DSL*.
- 1266 Shriram Krishnamurthi. 2001. *Linguistic Reuse*. Ph. D. Dissertation. Rice University.
- 1267 Daan Leijen and Erik Meijer. 2001. *Parsec: Direct Style Monadic Parser Combinators for the Real World*. Technical Report.
- 1268 John MacFarlane. 2014. CommonMark Spec. <https://spec.commonmark.org>
- 1269 Ole Lehrmann Madsen and Birger Møller-Pedersen. 1989. Virtual Classes: A Powerful Mechanism in Object-Oriented  
1270 Programming. In *OOPSLA*.
- 1271 MediaWiki. 2011. Parsoid. <https://www.mediawiki.org/wiki/Parsoid>
- 1272 Bruno C. d. S. Oliveira and William R. Cook. 2012. Extensibility for the Masses: Practical Extensibility with Object Algebras.  
1273 In *ECOOP*.
- 1274 Bruno C. d. S. Oliveira and Andres Löb. 2013. Abstract Syntax Graphs for Domain Specific Languages. In *PEPM@POPL*.
- Bruno C. d. S. Oliveira, Zhiyuan Shi, and João Alpuim. 2016. Disjoint Intersection Types. In *ICFP*.
- Bruno C. d. S. Oliveira, Tijs van der Storm, Alex Loh, and William R. Cook. 2013. Feature-Oriented Programming with  
Object Algebras. In *ECOOP*.
- Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Geoffrey Washburn. 2006. Simple Unification-Based Type  
Inference for GADTs. In *ICFP*.
- Tiark Rumpf, Nada Amin, Adriaan Moors, Philipp Haller, and Martin Odersky. 2012. Scala-Virtualized: Linguistic Reuse for  
Deep Embeddings. *High. Order Symb. Comput.* 25, 1 (2012).

- 1275 Tiark Rompf and Martin Odersky. 2010. Lightweight Modular Staging: A Pragmatic Approach to Runtime Code Generation  
1276 and Compiled DSLs. In *GPCE*.
- 1277 Maximilian Scherr and Shigeru Chiba. 2014. Implicit Staging of EDSL Expressions: A Bridge between Shallow and Deep  
1278 Embedding. In *ECOOP*.
- 1279 Tim Sheard and Simon Peyton Jones. 2002. Template Meta-programming for Haskell. In *Haskell@ICFP*.
- 1280 Josef Svenningsson and Emil Axelsson. 2015. Combining Deep and Shallow Embedding of Domain-Specific Languages.  
1281 *Comput. Lang. Syst. Struct.* 44 (2015).
- 1282 Wouter Swierstra. 2008. Data Types à la Carte (Functional Pearl). *J. Funct. Program.* 18, 4 (2008).
- 1283 Philip Wadler. 1998. The Expression Problem. Posted on the Java Genericity mailing list. [https://homepages.inf.ed.ac.uk/  
1284 wadler/papers/expression/expression.txt](https://homepages.inf.ed.ac.uk/wadler/papers/expression/expression.txt)
- 1285 Philip Wadler. 2003. A Prettier Printer. In *The Fun of Programming*.
- 1286 Wikipedia. 2022a. Help:Template. Retrieved April 2022 from <https://en.wikipedia.org/wiki/Help:Template>
- 1287 Wikipedia. 2022b. Help:Wikitext. Retrieved April 2022 from <https://en.wikipedia.org/wiki/Help:Wikitext>
- 1288 Wikipedia. 2022c. Template:Loop. Retrieved April 2022 from <https://en.wikipedia.org/wiki/Template:Loop>
- 1289 Weixin Zhang and Bruno C. d. S. Oliveira. 2017. EVF: An Extensible and Expressive Visitor Framework for Programming  
1290 Language Reuse. In *ECOOP*.
- 1291 Weixin Zhang and Bruno C. d. S. Oliveira. 2019. Shallow EDSLs and Object-Oriented Programming: Beyond Simple  
1292 Compositionality. *Art Sci. Eng. Program.* 3, 3 (2019).
- 1293 Weixin Zhang and Bruno C. d. S. Oliveira. 2020. Castor: Programming with Extensible Generative Visitors. *Sci. Comput.  
1294 Program.* 193 (2020).
- 1295 Weixin Zhang, Yaozhu Sun, and Bruno C. d. S. Oliveira. 2021. Compositional Programming. *ACM Trans. Program. Lang.  
1296 Syst.* 43, 3 (2021).
- 1297
- 1298
- 1299
- 1300
- 1301
- 1302
- 1303
- 1304
- 1305
- 1306
- 1307
- 1308
- 1309
- 1310
- 1311
- 1312
- 1313
- 1314
- 1315
- 1316
- 1317
- 1318
- 1319
- 1320
- 1321
- 1322
- 1323

## 1324 A NON-MODULAR DEPENDENCIES IN TAGLESS-FINAL EMBEDDINGS

```
1325 {-# LANGUAGE DuplicateRecordFields, NamedFieldPuns, OverloadedRecordDot, RecordWildCards #-}
```

### 1327 A.1 Modular Interpretations in Tagless-Final Embeddings

```
1329 data Vector = Vector { x :: Double, y :: Double } deriving Show
```

1330 A tagless-final embedding defines region constructors in a type class, instead of a closed algebraic  
1331 data type:

```
1333 class RegionHudak repr where
1334   circle    :: Double → repr
1335   outside   :: repr → repr
1336   union     :: repr → repr → repr
1337   intersect :: repr → repr → repr
1338   translate :: Vector → repr → repr
```

1339 Size can be modularly defined as an instance of the type class since it has no dependency:

```
1340 newtype Size = S { size :: Int }
1341
1342 instance RegionHudak Size where
1343   circle    _ = S 1
1344   outside   a = S $ a.size + 1
1345   union     a b = S $ a.size + b.size + 1
1346   intersect a b = S $ a.size + b.size + 1
1347   translate _ a = S $ a.size + 1
```

1348 But how about Text? As a first try, we might write:

```
1349 newtype Text = T { text :: String }
1350
1351 instance RegionHudak Text where
1352   circle r = T { text = "a circular region of radius " ++ show r }
1353   -- outside a = T { text = "outside a region of size " ++ show a.size }
1354   -- .....
```

1355 We will get a type error concerning `a.size` if we uncomment the line above, because `a` has type `Text`  
1356 and thus does not contain a field named `size`. So the problem is that, once we need operations that  
1357 have some dependencies on other operations, we get into trouble! But programs with dependencies  
1358 are common in practice. This is a serious limitation.

### 1360 A.2 A Non-Modular Workaround for Dependencies

1361 A simple workaround is to pack the two operations together and duplicate the code of size calcula-  
1362 tion. We have already described this approach for shallow embeddings, and the same workaround  
1363 works for tagless-final embeddings:

```
1364 data SizeAndText = ST { size :: Int, text :: String }
1365
1366 instance RegionHudak SizeAndText where
1367   circle r = ST { size = 1, text = "a circular region of radius " ++ show r }
1368   outside a = ST { size = a.size + 1
1369                 , text = "outside a region of size " ++ show a.size }
1370   union a b = ST { size, text = "the union of two regions of size " ++ show size ++ " in total" }
1371   where size = a.size + b.size + 1
```

1372

```

1373 intersect a b = ST { size, text = "the intersection of two regions of size " ++ show size ++ "
1374     in total" }
1375     where size = a.size + b.size + 1
1376 translate v a = ST { size, text = "a translated region of size " ++ show size }
1377     where size = a.size + 1

```

1378 However, this is not modular since we have to duplicate code for size calculation. If we have another  
 1379 operation that depends on size, we have to repeat the same code even again. This workaround is  
 1380 an *anti-pattern*, which violates the basic principle of software engineering.

### 1381 A.3 Modular Language Constructs

1382 Of course, new region constructors can be modularly added in a tagless-final embedding:

```

1384 class RegionHofer repr where
1385     univ  :: repr
1386     empty :: repr
1387     scale :: Vector → repr → repr

```

1388 We have to resort to the workaround again when we encounter mutual recursion:

```

1389 data UE = UE { isUniv :: Bool, isEmpty :: Bool }
1390
1391 instance RegionHudak UE where
1392     circle _ = UE { isUniv = False, isEmpty = False }
1393     outside a = UE { isUniv = a.isUniv, isEmpty = a.isUniv }
1394     union a b = UE { isUniv = a.isUniv || b.isUniv, isEmpty = a.isEmpty && b.isEmpty }
1395     intersect a b = UE { isUniv = a.isUniv && b.isUniv, isEmpty = a.isEmpty || b.isEmpty }
1396     translate _ a = UE { isUniv = a.isUniv, isEmpty = a.isEmpty }
1397
1398 instance RegionHofer UE where
1399     univ = UE { isUniv = True, isEmpty = False }
1400     empty = UE { isUniv = False, isEmpty = True }
1401     scale _ a = UE { isUniv = a.isUniv, isEmpty = a.isEmpty }

```

### 1402 A.4 Use Case

1403 We can use all constructors from RegionHudak and RegionHofer to create a region, as long as UE  
 1404 implements both type classes:

```

1406 region :: UE
1407 region = outside empty `union` circle 1
1408
1409 main :: IO ()
1410 main = do putStrLn $ "Univ: " ++ show (region.isUniv)
1411         putStrLn $ "Empty: " ++ show (region.isEmpty)

```

1412  
1413  
1414  
1415  
1416  
1417  
1418  
1419  
1420  
1421

## 1422 B MODULAR DEPENDENCIES IN TAGLESS-FINAL EMBEDDINGS

```
1423 {-# LANGUAGE ConstraintKinds, DataKinds, FlexibleInstances, GADTs, KindSignatures,
1424     MultiParamTypeClasses, RankNTypes, ScopedTypeVariables, TypeApplications, TypeOperators,
1425     UndecidableInstances #-}
```

### 1427 B.1 Generic Definitions for Records

1428 First of all, we need to define a type-indexed record as follows:

```
1430 data Record :: [*] → * where
1431   Nil  :: Record '[]
1432   Cons :: a → Record as → Record (a ': as)
```

1433 We also need a projection operation that finds an element by its type from the record:

```
1435 class a `In` as where
1436   project :: Record as → a
1437
1438 instance {-# OVERLAPPING #-} a `In` (a ': as) where
1439   project (Cons x _) = x
```

```
1440
1441 instance {-# OVERLAPPING #-} a `In` as ⇒ a `In` (b ': as) where
1442   project (Cons _ xs) = project xs
```

1443 Moreover, `All c s` is a data type whose term can be constructed only if all types in `s` implement the type class `c`:

```
1445 data All c :: [*] → * where
1446   AllNil  :: All c '[]
1447   AllCons :: c a ⇒ All c as → All c (a ': as)
```

### 1450 B.2 Region DSL Infrastructure

```
1451 data Vector = Vector { x :: Double, y :: Double }
```

1452 The same as before, we define region constructors in a type class:

```
1454 class Region0 r where
1455   univ_  :: r
1456   empty_ :: r
1457   circle_ :: Double → r
```

1458 Note that the constructors above are simpler because they do not have `r` in input positions. The encoding of other constructors is more ingenious because we need to consider how to inject dependencies. Here we employ `Record s` to encode the dependencies an interpretation relies on:

```
1462 class Region0 r ⇒ Region1 s r where
1463   outside_ :: Record s → r
1464   union_   :: Record s → Record s → r
1465   intersect_ :: Record s → Record s → r
1466   translate_ :: Vector → Record s → r
1467   scale_    :: Vector → Record s → r
```

1468 Furthermore, we create an auxiliary type class that constrains `s1` to satisfy the dependencies of all interpretations in `s2`:

1470

```

1471 class Region2 s1 s2 where
1472   modality :: All (Region1 s1) s2
1473
1474 instance Region2 s1 '[] where
1475   modality = AllNil
1476
1477 instance (Region1 s1 a, Region2 s1 as)  $\Rightarrow$  Region2 s1 (a ': as) where
1478   modality = AllCons modality
1479
1480 With the infrastructure above, we can define smart constructors for all kinds of regions. Each smart
1481 constructor returns a term of type Record s that composes all corresponding interpretations if s is
1482 self-contained (no interpretation in s has external dependencies):
1483
1484 univ :: forall s. Region2 s s  $\Rightarrow$  Record s
1485 univ = univ' (modality @s @s)
1486   where univ' :: All (Region1 s1) s2  $\rightarrow$  Record s2
1487         univ' AllNil      = Nil
1488         univ' (AllCons m) = Cons univ_ (univ' m)
1489
1490 empty :: forall s. Region2 s s  $\Rightarrow$  Record s
1491 empty = empty' (modality @s @s)
1492   where empty' :: All (Region1 s1) s2  $\rightarrow$  Record s2
1493         empty' AllNil      = Nil
1494         empty' (AllCons m) = Cons empty_ (empty' m)
1495
1496 circle :: forall s. Region2 s s  $\Rightarrow$  Double  $\rightarrow$  Record s
1497 circle = circle' (modality @s @s)
1498   where circle' :: All (Region1 s1) s2  $\rightarrow$  Double  $\rightarrow$  Record s2
1499         circle' AllNil      _ = Nil
1500         circle' (AllCons m) r = Cons (circle_ r) (circle' m r)
1501
1502 outside :: forall s. Region2 s s  $\Rightarrow$  Record s  $\rightarrow$  Record s
1503 outside = outside' (modality @s @s)
1504   where outside' :: All (Region1 s1) s2  $\rightarrow$  Record s1  $\rightarrow$  Record s2
1505         outside' AllNil      _ = Nil
1506         outside' (AllCons m) a = Cons (outside_ a) (outside' m a)
1507
1508 union :: forall s. Region2 s s  $\Rightarrow$  Record s  $\rightarrow$  Record s  $\rightarrow$  Record s
1509 union = union' (modality @s @s)
1510   where union' :: All (Region1 s1) s2  $\rightarrow$  Record s1  $\rightarrow$  Record s1  $\rightarrow$  Record s2
1511         union' AllNil      _ _ = Nil
1512         union' (AllCons m) a b = Cons (union_ a b) (union' m a b)
1513
1514 intersect :: forall s. Region2 s s  $\Rightarrow$  Record s  $\rightarrow$  Record s  $\rightarrow$  Record s
1515 intersect = intersect' (modality @s @s)
1516   where intersect' :: All (Region1 s1) s2  $\rightarrow$  Record s1  $\rightarrow$  Record s1  $\rightarrow$  Record s2
1517         intersect' AllNil      _ _ = Nil
1518         intersect' (AllCons m) a b = Cons (intersect_ a b) (intersect' m a b)
1519
1520 translate :: forall s. Region2 s s  $\Rightarrow$  Vector  $\rightarrow$  Record s  $\rightarrow$  Record s
1521 translate = translate' (modality @s @s)
1522   where translate' :: All (Region1 s1) s2  $\rightarrow$  Vector  $\rightarrow$  Record s1  $\rightarrow$  Record s2
1523         translate' AllNil      _ _ = Nil

```

```

1520     translate' (AllCons m) v a = Cons (translate_ v a) (translate' m v a)
1521
1522 scale :: forall s. Region2 s s => Vector -> Record s -> Record s
1523 scale = scale' (modality @s @s)
1524   where scale' :: All (Region1 s1) s2 -> Vector -> Record s1 -> Record s2
1525         scale' AllNil _ _ = Nil
1526         scale' (AllCons m) v a = Cons (scale_ v a) (scale' m v a)

```

1527 As shown above, there is a lot of boilerplate code for each language construct of the region DSL.

### 1528 B.3 Tagless-Final Embeddings

1529 Now we can write dependent interpretations in a tagless-final style. We take the more interesting  
 1530 example with mutual recursion for example:

```

1532 newtype IsUniv = U { isUniv :: Bool }
1533 newtype IsEmpty = E { isEmpty :: Bool }
1534
1535 instance Region0 IsUniv where
1536   circle_ _ = U False
1537   univ_     = U True
1538   empty_   = U False
1539
1540 instance (IsEmpty `In` s, IsUniv `In` s) => Region1 s IsUniv where
1541   outside_ a = U $ isEmpty (project a)
1542   union_   a b = U $ isUniv (project a) || isUniv (project b)
1543   intersect_ a b = U $ isUniv (project a) && isUniv (project b)
1544   translate_ _ a = U $ isUniv (project a)
1545   scale_     _ a = U $ isUniv (project a)
1546
1547 instance Region0 IsEmpty where
1548   circle_ _ = E False
1549   univ_     = E False
1550   empty_   = E True
1551
1552 instance (IsUniv `In` s, IsEmpty `In` s) => Region1 s IsEmpty where
1553   outside_ a = E $ isUniv (project a)
1554   union_   a b = E $ isEmpty (project a) && isEmpty (project b)
1555   intersect_ a b = E $ isEmpty (project a) || isEmpty (project b)
1556   translate_ _ a = E $ isEmpty (project a)
1557   scale_     _ a = E $ isEmpty (project a)

```

1556 We can easily declare dependencies using type constraints, and explicit projections help us find  
 1557 appropriate interpretations in the type-indexed record.

### 1558 B.4 Use Case

1559 Since '[IsUniv, IsEmpty]' is self-contained, we can use smart constructors to create a region:

```

1560 region :: Record '[IsUniv, IsEmpty]
1561 region = outside empty `union` circle 1
1562
1563
1564 main :: IO ()
1565 main = do let Cons u (Cons e Nil) = region
1566         putStrLn $ "Univ: " ++ show (isUniv u)
1567         putStrLn $ "Empty: " ++ show (isEmpty e)

```

1568