

Coverage Semantics for Dependent Pattern Matching

Joseph Eremondi¹[0000–0002–9631–4826] and Ohad Kammar²[0000–0002–2071–0929]

¹ Department of Computer Science, University of Regina, Regina, SK, Canada
jeremondi@uregina.ca

² Laboratory for the Foundations of Computer Science, University of Edinburgh,
Edinburgh, Scotland, United Kingdom
ohad.kammar@ed.ac.uk

Abstract. Dependent pattern matching is a key feature in dependently typed programming. However, there is a theory-practice disconnect: while many proof assistants implement pattern matching as primitive, theoretical presentations give semantics to pattern matching by elaborating to eliminators. Though theoretically convenient, eliminators can be awkward and verbose, particularly for complex combinations of patterns.

This work aims to bridge the theory-practice gap by presenting a direct categorical semantics for pattern matching, which does not elaborate to eliminators. This is achieved using sheaf theory to describe when sets of arrows (terms) can be amalgamated into a single arrow. We present a language with top-level dependent pattern matching, without specifying which sets of patterns are considered covering for a match. Then, we give a sufficient criterion for which pattern-sets admit a sound model: patterns should be in the canonical coverage for the category of contexts. Finally, we use sheaf-theoretic saturation conditions to devise some allowable sets of patterns. We are able to express and exceed the status quo, giving semantics for datatype constructors, nested patterns, absurd patterns, propositional equality, and dot patterns.

Keywords: semantics, dependent pattern matching, sheaf, coverage

1 Introduction

Pattern matching is a core feature in dependently typed programming. With pattern matching one can specify a function consuming an input by giving functions for every possible way that input might have been constructed. For dependent types, the defined function can be a universally quantified proof, giving a Curry-Howard analogue of proof-by-cases.

However, there is a disconnect between the theory and practice of dependent pattern matching. Many dependently typed languages take pattern matching as a built-in user-facing construct: Coq [4], Agda [22], and Idris [5] all contain a form of dependent pattern matching in their core calculi. However, most theoretical treatments of dependent types deal with *eliminators* [19]: primitive recursors

with result types dependent on a value of the eliminated type. Eliminators and pattern matching are equally expressive, with or without with Axiom K [12, 6].

While it is possible to express all pattern matches using eliminators, it is not always convenient. Some pattern matching features require lengthy translations when converting to eliminators, such as overlapping patterns, catch-all branches, or matching on multiple values at once. Moreover, languages differ in which pattern matches they allow, so every variant of pattern matching requires a new eliminator-translation to prove consistency. Even the implementations of dependently typed languages are restricted by eliminators, since most pattern matches are elaborated into case trees with a 1:1 correspondence between branches and constructors of an inductive type.

The contribution of this paper is to narrow the theory-practice divide with a highly general syntax (Section 2) and categorical semantics for dependent pattern matching (Section 3). The semantics is direct and generic: pattern matches are translated directly into semantic objects without desugaring to eliminators or case trees, and the semantics is parameterized over an abstract coverage specifying which sets of patterns one can match against. We investigate the vision set forth by Epigram [20, 21] to enable diverse pattern-matching abstractions that go beyond the list of constructors declared by a datatype. We focus on non-overlapping patterns, but give a potential road map to supporting overlap.

In constructing our generic semantics, we present a general sufficient criterion for when a coverage leads to well-defined pattern matches without compromising logical consistency (Section 4), mechanized in Lean [11]. We define this criterion, drawing on parallels between dependent pattern matching and the theory of sheaves on a site, discovering that it is sufficient for each allowed set of patterns to correspond to a cover in the canonical coverage for the semantic category. Moreover, we use elementary results from sheaf theory to describe a group of closure operations which preserve the canonicity of a coverage (Section 5). These give a simple and direct way to model common features like multi-value matches, nested patterns, and matching on propositional equality proofs. With the exception of recursion, we achieve feature parity with the original presentation of dependent pattern matching by Coquand [9]. We conclude with an illustrative example (Section 6) and related and future work (Section 7).

A major contribution of our work is expressing pattern matching in the language of categories and sheaves. Our approach is semantic: instead of elaborating pattern matching syntax, we take pattern matches as primitive and work directly in the semantic domain, avoiding the need to consider eliminators as a syntactic primitive. The connection has been implicit for many decades, but we make it formal. That said, we only assume basic knowledge of functors, pullbacks, and slice categories, and we present all the required sheaf theory.

2 COVERTT: The Source Language

We begin with a variant of Martin L f Type Theory, called COVERTT, parameterized over which sets of patterns can be matched against. Drawing from

sheaf-theory terminology, when a set of patterns is permissible on the left-hand side of a pattern match we call it a **cover**, and say it is **covering**. The set of all covers together is called a **coverage**. The main distinct feature of COVERTT is that it is parameterized over a coverage.

2.1 The Anatomy of a Datatype

As they are presented in most dependently typed languages, inductively defined datatypes conflate four different concepts. Consider the quintessential inductive family of length-indexed vectors:

$$\begin{aligned} \text{data } \text{Vec } (A : \mathcal{U}) : (m : \mathbb{N}) \rightarrow \mathcal{U} \text{ where} \\ \text{nil} : \text{Vec } A \ 0 \\ \text{cons} : (n : \mathbb{N}) \rightarrow A \rightarrow \text{Vec } A \ n \rightarrow \text{Vec } A \ (n + 1) \end{aligned}$$

This definition implicitly relies on the following concepts, the separation of which motivates the design COVERTT (as exhibited in Section 2.3).

- **Coproducts:** An inductive type behaves like the sum of its constructor types. For **Vec**, it behaves like $\mathbb{1} + (A \times \text{Vec } A \ n)$. In COVERTT, each inductive I is declared at the top level to have a finite collection of constructors $D_1^I \dots D_n^I$, each of which has a type ending in I .
- **Dependent fields:** each constructor has a dependent product type, so the arguments to a constructor are a curried dependent record, where the types of later fields can depend on the values of earlier ones. For **Vec**, in **cons** the return type and the type of the tail of the list depend on the earlier parameter n . In COVERTT, the type of a constructor is given as a telescope, where the types of later entries are allowed to depend on the values of previous entries.
- **Indexing:** each constructor implies a specific equation about the index values. For **Vec**, **nil** restricts that $m = 0$ and **cons** restricts that $m = n + 1$. In COVERTT, we use **Fording** [18], where each constructor has an identical return type, but may constrain type parameters using equality-proof fields. We treat equality as primitive and make it the only way to constrain indices of a constructor, simplifying the models of COVERTT.
- **Recursion:** inductive types can refer to themselves in fields, except to the left of a function arrow, a condition known as strict positivity. Typically, one can define structurally recursive functions over **Vec**. Here, **Vec** occurs as a field type for **cons**, which is allowed because it is not to the left of an arrow type. We omit recursion from COVERTT, as we believe it requires a separate toolkit of abstractions. In Section 7.2 we discuss its possible addition. In examples, we refer to some inductive types defined using self-reference, such as vectors or natural numbers.

Restricting to Top-Level Datatypes COVERTT only allows data types and pattern matches to be declared at the top level. The parameter and constructor types of datatypes must be typeable in the empty context, though they can refer to other data types. Likewise, the branches and motive of a pattern match

must be typeable in the empty context. These assumptions simplify the presentation of COVERTT while reflecting how datatypes are implemented in languages like Agda, where declarations in a non-empty context are desugared into top level declarations with extra parameters. Treating nested patterns and arbitrary coverages is an interesting problem even with top-level matches, and requires substantial technical developments even to handle ordinary matching.

2.2 Syntax and Typing

The syntax for COVERTT below is standard, except for the separation of concerns from Section 2.1. Figure 1 gives the syntax, along with typing rules for COVERTT.

Overline arrows denote sequences, while bold metavariables denote dependent sequences, e.g. substitutions. Variables are assigned types from the context. Typing for dependent functions and equality is standard, with rules for their

$$\begin{array}{l}
\text{TERM} \ni s, S, t, T ::= \quad x \mid \mathcal{U} \mid \Pi(x : S).T \mid \lambda x. t \mid t \ s \mid s =_T t \mid \mathbf{refl}_t \\
\quad \mid I \ \mathbf{t} \mid D^I \ \mathbf{t} \mid \mathbf{case} \ (\mathbf{t} : \Gamma) \ \mathbf{to} \ T \ \mathbf{of} \ \{\overline{\Delta_i. \mathbf{s}_i} \Rightarrow \overline{t}^i\} \\
\text{CTX} \ni \Gamma, \Delta, \Xi ::= \quad \cdot \mid \Gamma, (x : T) \\
\text{SUBST} \ni \mathbf{s}, \mathbf{t} ::= \quad \cdot \mid \mathbf{s} ; t
\end{array}$$

$\Gamma \vdash t : T \text{ (Well Typed Terms/Types)}$

$\{t_i\}_i \triangleright \Xi \text{ (Covering Patterns)}$

| | | | |
|--|---|---|--|
| $\frac{\text{TYPECONV} \quad \Gamma \vdash t : S \quad \Gamma \vdash S \equiv T : \mathcal{U}}{\Gamma \vdash t : T}$ | $\frac{\text{TYVAR} \quad \vdash \Gamma \text{ ctx} \quad (x : T) \in \Gamma}{\Gamma \vdash x : T}$ | $\frac{\text{TYPI} \quad \Gamma \vdash S : \mathcal{U} \quad \Gamma, (x : S) \vdash T : \mathcal{U}}{\Gamma \vdash \Pi(x : S).T : \mathcal{U}}$ | $\frac{\text{TYAPP} \quad \Gamma \vdash t : \Pi(x : S).T \quad \Gamma \vdash s : S}{\Gamma \vdash t \ s : [s/x]T}$ |
| $\frac{\text{TYLAM} \quad \Gamma, (x : S) \vdash t : T}{\Gamma \vdash \lambda x. t : \Pi(x : S).T}$ | $\frac{\text{TYREFL} \quad \Gamma \vdash t : T}{\Gamma \vdash \mathbf{refl}_t : t =_T t}$ | $\frac{\text{TYIND} \quad \text{Params}(I) := \Delta \quad \Gamma \vdash \mathbf{t} \Rightarrow \Delta}{\Gamma \vdash I \ \mathbf{t} : \mathcal{U}}$ | |
| $\frac{\text{TYEQ} \quad \Gamma \vdash T : \mathcal{U} \quad \Gamma \vdash s : T \quad \Gamma \vdash t : T}{\Gamma \vdash s =_T t : \mathcal{U}}$ | | $\frac{\text{TYCTOR} \quad \text{Fields}(D^I) := \Delta, \Xi \quad \Gamma \vdash \mathbf{s} \Rightarrow \Delta \quad \Gamma \vdash \mathbf{t} \Rightarrow [\mathbf{s}/\Delta]\Xi}{\Gamma \vdash D^I \ \mathbf{t} : I \ \mathbf{s}}$ | |
| $\frac{\text{TYCASE} \quad \vdash \Xi \text{ ctx} \quad \Gamma \vdash \mathbf{t}_{\text{scrut}} \Rightarrow \Xi \quad \Xi \vdash T_{\text{motive}} : \mathcal{U} \quad \overline{\vdash \Delta_i \text{ ctx}}^i \quad \overline{\Delta_i \vdash \mathbf{s}_i \Rightarrow \Xi}^i \quad \vdash \{\mathbf{s}_i\}_i \triangleright \Xi \quad \overline{\Delta_i \vdash t_i : [\mathbf{s}_i/\Xi]T_{\text{motive}}}^i}{\Gamma \vdash \mathbf{case} \ (\mathbf{t}_{\text{scrut}} : \Xi) \ \mathbf{to} \ T_{\text{motive}} \ \mathbf{of} \ \{\overline{\Delta_i. \mathbf{s}_i} \Rightarrow \overline{t_i}^i\} : [\mathbf{t}_{\text{scrut}}/\Xi]T_{\text{motive}}}$ | | | |

Fig. 1. COVERTT: Term Typing

| $\vdash \Gamma \text{ ctx}$ (<i>Well Formed Contexts</i>) | | $\Gamma \vdash \mathbf{t} \Rightarrow \Delta$ (<i>Well Formed Substitutions</i>) | |
|---|--|--|--|
| CTXNIL | CTXCONS | ENVNIL | ENVCONS |
| $\vdash \cdot \text{ ctx}$ | $\frac{\vdash \Gamma \text{ ctx} \quad \Gamma \vdash T : \mathcal{U}}{\vdash \Gamma, (x : T) \text{ ctx}}$ | $\frac{}{\Gamma \vdash \cdot \Rightarrow \cdot}$ | $\frac{\Gamma \vdash \mathbf{s} \Rightarrow \Delta \quad \Gamma \vdash t : [\mathbf{s}/\Delta]T}{\Gamma \vdash \mathbf{s}; t \Rightarrow \Delta, (x : T)}$ |

Fig. 2. Typing: Contexts and Substitutions

types, introduction, and elimination. The equality type is intensional, and there is no reflection rule by which propositional equalities can be made judgmental equalities. Nevertheless, COVERTT is consistent with models that identify all propositionally-equal terms. We have no J-axiom, instead using the coverage to specify how to match on refl_t . We have one universe which does not have a type, leaving a universe hierarchy for future work.

Contexts and Substitutions Figure 2 also specifies well-formedness rules for contexts and substitutions. Contexts are sequences of typed variables, where later types may refer to variables from earlier in the context. Substitutions are the inhabitants of contexts. The rules are like an iterated version of dependent pairs: the type of the later values in the substitutions may depend on the earlier values. We borrow the notation $\Gamma \vdash \mathbf{t} \Rightarrow \Delta$ from Hofmann [14], since such a substitution corresponds to a morphism $\Gamma \longrightarrow \Delta$ in the models we define. We use the notation $[\mathbf{s}/\Delta]t$ to denote the simultaneous substitution of the variables bound in Δ by the terms of \mathbf{s} in t . If $\Gamma \vdash \mathbf{s} \Rightarrow \Delta$ and $\Delta \vdash t : T$, then $\Gamma \vdash [\mathbf{s}/\Delta]t : [\mathbf{s}/\Delta]T$.

Datatype and Pattern Matching Syntax We assume a fixed collection of inductive type constructors, along with data constructors. For each datatype there is a fixed context of parameters $\text{Params}(I)$ and, for each constructor, fields $\text{Fields}(D^I)$, such that $\vdash \text{Params}(I), \text{Fields}(D^I)$, i.e., both are well-formed and the fields may depend on the parameters. The type $I \mathbf{t}$ denotes the type constructor I applied to parameters \mathbf{t} . $D^I \mathbf{t}$ is the data constructor D for the type I , given \mathbf{t} for its fields. We omit I in D^I when it is clear from context.

The pattern matching form is a nameless version of defining functions by multiple pattern matching clauses, as in Agda or Idris. The term

$$\text{case } (\mathbf{t} : \Xi) \text{ to } T \text{ of } \{\overrightarrow{\Delta_i \cdot \mathbf{s}_i \Rightarrow t_i}^i\}$$

denotes a match on **scrutinees** \mathbf{t} of context-type Ξ , producing a result of the **motive** type T , where T may refer to the variables bound in Ξ . The branches, indexed by i , each have a left-side pattern \mathbf{s}_i , which may contain pattern variables from the context Δ_i , and which produces a result t_i . The scrutinee is a substitution because pattern matching functions can take multiple arguments, and the types of later arguments can depend on the values of earlier ones.

Typing Pattern Matches The TYCASE typing rule (Fig. 1) is the most important rule. To type $\Gamma \vdash \text{case } (t_{\text{scrut}} : \Xi) \text{ to } T_{\text{motive}} \text{ of } \{\overline{\Delta_i. s_i} \Rightarrow t_i^i\}$, the scrutinee inhabits some Ξ . Because pattern matching is dependent, the motive T_{motive} is indexed over the scrutinee type. The pattern s_i for each branch inhabit the scrutinee type Ξ in the context of its pattern variables Δ_i . Each branch is typed against the context of its pattern variables, and inhabits the motive for the scrutinee value given by that branch’s pattern, e.g., s_i . Finally, the entire match inhabits the motive type, instantiated to the scrutinee. Critically, the motive T_{motive} and each branch t_i must be typeable in the closed context Δ_i , making no reference to Γ . This restriction matches practice: in Agda and Idris, pattern matching elaborates to top level declarations. We give an example in Section 2.3.

The Generic Coverage Relation In a match, the patterns $\overline{s_i}$ must be covering. At no point do we require the patterns of a match to correspond to constructors of an inductive type, or even that the matched upon type be inductive. Instead, we appeal to an arbitrary judgment $\{s_i\}_i \triangleright \Xi$, relating sets of substitutions to contexts. This relation is to be read as “the patterns $s_1 \dots s_n$ are a total decomposition of the context Ξ ”. This relation replaces the usual condition that there must be a case for each constructor of the datatype. It is the parameter by which we tune COVERTT, and can take many forms, from requiring a single scrutinee with exactly one branch per constructor, to allowing multiple scrutinees with arbitrary nested patterns and absurd branches omitted. Section 4.2 explores the conditions a coverage must satisfy to result in a well-behaved type theory. In general, we expect that a coverage will consist of a basis set of coverings, containing at least variables and constructors for each datatype, which can be closed under composition, concatenation, etc. For generality we do not require closure conditions, but we show in Section 5 they are always permissible.

Computational Rules Figure 3 gives definitional equality for COVERTT, omitting structural (reflexivity, symmetry, transitivity) and congruence rules. Rule EQAPP is the usual β -reduction. In EQMATCH, we reduce a match when the scrutinee is $[t_{\text{mat}}/\Delta_j]s_j$ for some $\Gamma \vdash t_{\text{mat}} \Rightarrow \Delta_j$, i.e., it is a pattern applied to values for the pattern variables of s_j . This substitution t_{mat} instantiates the pattern variables in t_j , which is the value of the entire match.

2.3 Example: Vectors

To see these constructs concretely, we show how length-indexed vectors from Section 2.1 would be represented in our system, along with a type safe head function. In COVERTT, vectors can be defined using labelled sums. Notice the equality-type fields which encode the index constraints via Fording.

$\text{Params}(\text{Vec}) := \cdot, (A : \mathcal{U}), (n : \mathbb{N})$

$\text{Fields}(\text{Nil}^{\text{Vec}}) := \cdot, (A : \mathcal{U}), (n : \mathbb{N}), (eq : n = 0)$

$\text{Fields}(\text{Cons}^{\text{Vec}}) := \cdot, (A : \mathcal{U}), (n : \mathbb{N}), (m : \mathbb{N}), (h : A), (t : \text{Vec } A \ m), (eq : n = m + 1)$

| | |
|--|---|
| $\Gamma \vdash t = s : T$ (<i>Term Definitional Equality</i>) | $\vdash \Gamma \equiv \Delta \text{ ctx}$ (<i>Equal Contexts</i>) |
| $\Gamma \vdash \mathbf{t} \Rightarrow \Delta$ (<i>Equal Substitutions</i>) | |

EQMATCH

$$\frac{\Gamma \vdash \mathbf{t}_{\text{mat}} \Rightarrow \Delta_j \quad \Xi \vdash T : \mathcal{U} \quad \frac{\vdash \Xi \text{ ctx} \quad \frac{\overline{\vdash \Delta_i \text{ ctx}}^i \quad j \in \{1 \dots i\}}{\Delta_i \vdash \mathbf{s}_i \Rightarrow \Xi^i} \quad \vdash \{\mathbf{s}_i\}_i \triangleright \Xi \quad \overline{\Delta_i \vdash t_i : [\mathbf{s}_i / \Xi] T}^i}{\Gamma \vdash \text{case } ([\mathbf{t}_{\text{mat}} / \Delta_i] \mathbf{s}_j : \Xi) \text{ to } T \text{ of } \{\overline{\Delta_i \cdot \mathbf{s}_i \Rightarrow t_i}^i\} \equiv [\mathbf{t}_{\text{mat}} / \Delta_j] t_j : [\mathbf{t}_{\text{mat}} / \Delta_j][\mathbf{s}_j / \Xi] T}$$

EQAPP

$$\frac{\Gamma, (x : S) \vdash t : \Pi(x : S).T \quad \Gamma \vdash s : S}{\Gamma \vdash (\lambda x. t) s \equiv [s/x]t : [s/x]T}$$

EQCTXNIL

$$\frac{}{\vdash \cdot \equiv \cdot \text{ ctx}}$$

SUBNIL

$$\frac{}{\Gamma \vdash \cdot \equiv \cdot : \cdot}$$

EQCTXCONS

$$\frac{\vdash \Gamma \equiv \Gamma' \text{ ctx} \quad \Gamma \vdash T \equiv T' : \mathcal{U}}{\vdash \Gamma, (x : T) \equiv \Gamma', (x : T') \text{ ctx}}$$

SUBCONS

$$\frac{\Gamma \vdash \mathbf{s} \equiv \mathbf{s}' : \Delta \quad \Gamma \vdash t \equiv t' : [\mathbf{s} / \Delta] T}{\Gamma \vdash \mathbf{s} \circledast t \equiv \mathbf{s}' \circledast t' : \Delta, (x : T)}$$

Fig. 3. Definitional Equality: Computational Rules

In Agda-style notation, a safe head function for vectors can be written:

```
head : (A : U) → (n : N) → Vec A (1 + n) → A
head A n (Cons h t) = h
```

In COVERTT this would be defined as:

```
head : Π(A : U)(n : N)(x : Vec A (n + 1)). A
head := λA. λn. λx. case (· ∘ n ∘ x : ·, (n : N), (x : Vec A (n + 1))) to A of {
  ·, (m : N), (h : A), (t : Vec A m). · ∘ m ∘ Cons A (m + 1) m h t reflm+1 ⇒ h}
```

That is, the function takes A , n , and x as type, number, and vector parameters. It then passes those parameters as the scrutinee of the pattern match, which is annotated with their types. The annotation is a telescope of types, so it introduces new names for the scrutinees. It happens that here we are passing names as the scrutinees, but this need not be the case, which is why we need new names for them. The result type is annotated as A . The match has a single branch, with a telescope of pattern variables m , h and t with their types. The pattern for the branch has $m + 1$ as the value for scrutinee n and Cons applied to its arguments for scrutinee x . Finally, to the right of \Rightarrow is the result for this case, which is h .

Implicit in this example is the need for the following to hold:

$$\begin{aligned} & \{ \cdot, (m : N)(h : A), (t : \text{Vec } A \ m). \cdot \circ m \circ \text{Cons } A \ (m + 1) \ m \ h \ t \ \text{refl}_{m+1} \} \\ & \triangleright \cdot, (n : N), (x : \text{Vec } A \ (n + 1)) \end{aligned}$$

That is, the pattern for the single branch needs to cover the scrutinee type. Deducing that these patterns are a valid coverage involves seeing that refl_{m+1} constrains the value of n and that Nil has an absurd type.

The goal of this paper is to define direct semantics that justify such deductions, and give a broad framework to define valid coverages.

3 Categorical Models of COVERTT

In this section, we translate the syntactic constructs of COVERTT into the language of **Categories with Families** (CwFs). CwFs correspond almost exactly to the syntactic structure of dependent type theory, but with syntactic substitution replaced by a semantic operation, and with implicit liftings between contexts made explicit.

3.1 Background: Categories with Families

We recapitulate the definition and notation of CwFs, a categorical model for dependent type theory that follows the syntax fairly closely. See Hofmann [14] for more details.

Recall that a **family** X is a pair (I_X, \underline{X}) consisting of a set I_X and an I -indexed sequence of sets $(\underline{X}_i)_{i \in I}$. A **map** of families $f : X \rightarrow Y$ is a pair (f_I, \underline{f}) consisting of a function $f_I : I_X \rightarrow I_Y$ and an I -indexed sequence of functions $(f_i : \underline{X}_i \rightarrow \underline{Y}_{f_i})_{i \in I}$. The category **Fam** has families as objects and maps of families as morphisms, with componentwise identity and composition structure.

A **basic CwF** \mathcal{C} is a pair (\mathcal{C}_o, F) consisting of a category \mathcal{C}_o and a functor $F : \mathcal{C}_o^{\text{op}} \rightarrow \mathbf{Fam}$. The functor F packs four pieces of structure which we'll unpack using the following notation:

- Objects $\Gamma \in \mathcal{C}_o$ are **contexts**, and morphisms $\theta : \Gamma \rightarrow \Delta$ are **substitutions**.
- For every context $\Gamma \in \mathcal{C}_o$, we denote the family $F\Gamma$ by $(\text{Ty}(\Gamma), \text{Tm}_\Gamma(_))$. We call the elements of its indexing set $\text{Ty}(\Gamma)$ the **types** in context Γ . For each type in $T \in \text{Ty}(\Gamma)$, we call the element of the component $\text{Tm}_\Gamma(T)$ the **terms** of type T in context Γ . We omit Γ when it is clear from context.
- For every substitution $\theta : \Gamma \rightarrow \Delta$, we have a map of families $F\theta : F\Gamma \leftarrow F\Delta$, and we use the same notation for both its components $F\theta = (_ \{\theta\}, _ \{\theta\})$ and call both **substitution functions**. The first component is the substitution function on **types** $_ \{\theta\} : \text{Ty}(\Gamma) \leftarrow \text{Ty}(\Delta)$. The second component is a sequence of substitution functions on the terms of each type $_ \{\theta\} : \text{Tm}_\Gamma(T \{\theta\}) \leftarrow \text{Tm}_\Delta(T)$.
- The functoriality of F amounts to the following four properties, which we call the **substitution lemma** for this CwF, where T ranges over $\text{Ty}(\Gamma)$, t ranges over $\text{Tm}_\Gamma(T)$:

$$\begin{aligned} T\{\text{id}_\Gamma\} &= T & t\{\text{id}_\Gamma\} &= t & T\{\theta \circ \sigma\} &= (T\{\theta\})\{\sigma\} & t\{\theta \circ \sigma\} &= (t\{\theta\})\{\sigma\} \\ & & & & & & & (\text{where } \Xi \xrightarrow{\theta} \Delta \xrightarrow{\sigma} \Gamma) \end{aligned}$$

A basic CwF includes only the bare bones of semantic models for a dependent type theory. In order to model dependent type theories of interests we need to equip them with additional structure.

We start with context extension. Let \mathcal{C} be a basic CwF, and assume \mathcal{C}_o has a terminal object \cdot . A **comprehension structure** $(\triangleright, \mathbf{p}, \mathbf{v})$ over \mathcal{C} consists of, for each context $\Gamma \in \mathcal{C}_o$ and type $T \in \text{Ty}(\Gamma)$:

- A context $\Gamma \triangleright T$, the context Γ **extended by** T .
- A substitution $\mathbf{p}_T : \Gamma \triangleright T \rightarrow \Gamma$, the **weakening** of Γ by T . We say that we **weaken** a type or a term by T when we apply the corresponding substitution function for the weakening \mathbf{p}_T .
- A term $\mathbf{v}_T \in \text{Tm}_{\Gamma \triangleright T}(T)$, the **variable** we extend the context Γ with.
- Moreover, for every substitution $\theta : \Delta \rightarrow \Gamma$ and term $t \in \text{Tm}_{\Delta}(T\{\theta\})$ there is a unique substitution $\langle \theta, t : T \rangle : \Delta \rightarrow \Gamma \triangleright T$ satisfying:

$$\mathbf{p}_T \circ \langle \theta, t \rangle = \theta \quad \mathbf{v}_T\{\langle \theta, t \rangle\} = t$$

We call this substitution the **extension** of the substitution θ by t . In the sequel we omit the type ascription and write $\langle \theta, t \rangle$ for $\langle \theta, t : T \rangle$ when T is clear from context. Likewise, we omit the subscript on \mathbf{p} and \mathbf{v} when clear.

A CwF is a basic CwF together with comprehension structure. For a CwF $\mathcal{C} = (\mathcal{C}_o, F)$, we refer to \mathcal{C}_o as \mathcal{C} when doing so will cause no ambiguity.

3.2 Sections, Slices, and Dependent Types

Let \mathcal{C} be a category and $\Gamma \in \mathcal{C}$ an object in it. Recall the slice category \mathcal{C}/Γ whose objects (T, d) consist of an object $T \in \mathcal{C}$ and a morphism $d : T \rightarrow \Gamma$. Morphisms $f : (T, d) \rightarrow (S, e)$ in the slice are morphisms $f : T \rightarrow S$ that lift d through e , i.e.: $e \circ f = d$. For example, a section of Γ is a morphism out of (Γ, id) in the slice \mathcal{C}/Γ , i.e., an object $T \in \mathcal{C}$, and a pair of morphisms $f : \Gamma \rightarrow T$ and $d : T \rightarrow \Gamma$ such that $d \circ f = \text{id}$.

We can use the slices of a category \mathcal{C} to model dependent type theory by requiring \mathcal{C} to be locally Cartesian closed (LCCC). The intuition behind this structure is that the object (T, d) in the slice \mathcal{C}/Γ represent types T in context Γ , and d represents the dependency of terms of this type on their context. We will not recapitulate the LCCC conditions explicitly. We will, however, spell the induced LCCC structure needed for a CwF, for two reasons. First, we indicate what we have formalized in LEAN. Second, we rely on this relationship between type families and slice objects in Section 4. It lets us use core results of sheaf theory to model dependencies in pattern matching.

Lemma 1. (\checkmark_{LEAN}) *Let \mathcal{C} be a CwF.*

- *The weakening $\mathbf{p}_T : (\Gamma \triangleright T) \rightarrow \Gamma$ encodes the type T as the slice object $((\Gamma \triangleright T), \mathbf{p}_T)$, in the sense that sections of Γ correspond to terms in context Γ :*

$$\text{Tm}_{\Gamma}(T) \cong \text{Hom}_{(\mathcal{C}/\Gamma)}\left((\Gamma, \text{id}), (\Gamma \triangleright T, \mathbf{p}_T)\right)$$

- The morphisms in \mathcal{C} encode indexed types—for all $\Gamma, \Delta \in \mathcal{C}$, $T \in \text{Ty}(\Gamma)$, and $\theta : \Delta \rightarrow \Gamma$:

$$\text{Tm}_\Delta(T\{\theta\}) \cong \text{Hom}_{(\mathcal{C}/\Gamma)}\left((\Delta, \theta), (\Gamma \triangleright T, \text{p}_T)\right)$$

For $t \in \text{Tm}_\Gamma(T)$ we write $\bar{t} : \Gamma \rightarrow \Gamma \triangleright T$ for $\langle id, t \rangle$, the corresponding p -section.

3.3 Semantic Type Formers and Closedness

CwFs give the core structure of type dependency, but give no indication of what types a model supports. Here we give semantic closedness conditions which postulate the existence of type and term constructors corresponding to common features of a dependent type theory.

Dependent Functions We say that a CwF \mathcal{C} supports dependent functions if it has:

- For each $S \in \text{Ty}(\Gamma)$ and $T \in \text{Ty}(\Gamma \triangleright S)$, a type $\Pi(S, T) \in \text{Ty}(\Gamma)$;
- For each $T \in \text{Ty}(\Gamma \triangleright S)$ and $t \in \text{Tm}(T)$, a term $\lambda(t) \in \text{Tm}(\Pi(S, T))$;
- For each $T \in \text{Ty}(\Gamma \triangleright S)$, $s \in \text{Tm}(S)$ and $t \in \text{Tm}(\Pi(S, T))$, a term $\text{App}(t, s) \in \text{Tm}(T\{\langle id, s \rangle\})$;

such that the usual structural substitution rules hold, as well as the β -reduction equality $\text{App}(\lambda(t), s) = t\{\langle id, s \rangle\}$. We could impose a version of an η rule, but this is not required for our results.

Equality A CwF \mathcal{C} supports propositional equality types [14] if, for every type $T \in \text{Ty}(\Gamma)$, there exists substitution-stable terms as follows:

- A type $\text{Id}(T) \in \text{Ty}(\Gamma \triangleright T \triangleright T\{\text{p}\})$;
- A morphism $\text{Refl}_T : \Gamma \triangleright T \rightarrow \Gamma \triangleright T \triangleright T\{\text{p}\} \triangleright \text{Id}(T)$, such that $\text{p} \circ \text{Refl} = \langle id, \text{v}_T \rangle$;
- For each $S \in \text{Ty}(\Gamma \triangleright T \triangleright T\{\text{p}\} \triangleright \text{Id}(T))$, a function (on sets) $J_{T,S}$ which is in $\text{Tm}(T\{\text{Refl}_T\}) \rightarrow \text{Tm}(S)$, where, for any $t \in \text{Tm}(S\{\text{Refl}_T\})$, $J(t)\{\text{Refl}_T\} = t$.

We have analogues of COVERTT terms using substitution. For $T \in \text{Ty}(\Gamma)$ and $s, t \in \text{Tm}(T)$:

- $\text{Id}(T, s, t) := \text{Id}(T)\{\langle \langle id, s \rangle, t \rangle\} \in \text{Tm}(\Gamma)$;
- $\text{Refl}_T(t) := \text{v}\{\text{Refl}_T \circ \bar{t}\} \in \text{Tm}(\text{Id}(T)\{\text{p} \circ \text{Refl}_T \circ \bar{t}\}) = \text{Tm}(\text{Id}(T, t, t))$.

We say that a CwF supports **extensional equality** when $\text{Tm}(\text{Id}(T, s, t)) \neq \emptyset$ iff $s = t$, such as in presheaf or set-theoretic models. Note that this condition does not require COVERTT to soundly model an equality reflection rule. Intensional equality in COVERTT can be modelled extensionally, so long as one does not augment COVERTT with axioms like univalence, which are inconsistent with equality reflection. We focus on extensional models, discussing alternatives in Section 7.2.

Labelled Variants Next we define what it means for a category to support **labelled variants**, so that we can interpret datatypes as coproducts of their constructor types. Labelled variants are simply coproducts with extra structure to mediate the type-context relationship. Assume:

- a fixed set of type constructors TYCON ;
- for each $I \in \text{TYCON}$, a set of data constructors DATACON_I ;
- for each $I \in \text{TYCON}$, an object $\text{Params}_I \in \mathcal{C}$ specifying types of arguments to the type constructor;
- for each $D^I \in \text{DATACON}_I$ a type $\text{Fields}_{D^I} \in \text{Ty}(\text{Params}_I)$ specifying the types of arguments to the data constructor.

The last condition will usually rely on having some sort of dependent pair to encode multiple fields. Then a CwF supports the labelled variants for I if there exists some $\text{TyCon}_I \in \text{Ty}(\text{Params}_I)$, such that for each $\theta : \Gamma \rightarrow \text{Params}_I$, there exists an isomorphism

$$\iota : \Gamma \triangleright \text{TyCon}_I\{\theta\} \cong \coprod_i (\Gamma \triangleright \text{Fields}_{D_i^I}\{\theta\}) \quad \text{s.t. } \forall i, \text{pFields}_{D_i^I}\{\theta\} = \text{pTyCon}_I\{\theta\} \circ \iota^{-1} \circ \text{inj}_i$$

That is, projecting out the context of from the fields' type is the same as converting into TyCon_I with ι and then projecting the context. Then, for each $t \in \text{Tm}(\text{Fields}_{D_i^I}\{\theta\})$, the arrow $\iota^{-1} \circ \text{inj}_i \circ \bar{t}$ is a section of $\text{pTyCon}_I\{\theta\}$, and hence denotes a term in $\text{Tm}(\text{TyCon}_I\{\theta\})$. We call this term $\text{DataCon}_{D_i^I}(t)$, since it denotes the i th data constructor applied to field values t .

3.4 Pattern Matching

Here we give a semantic presentation of COVERTT-style pattern matching. In the CwF framework, we can follow the definition of COVERTT. This serves as a statement of what we need to define pattern matching semantically. We show how to fulfill those requirements in Section 4.

Consider a semantic coverage relation \triangleright whose elements are sets containing arrows into Δ . We say \mathcal{C} supports matching over the semantic coverage \triangleright if, for every, $\Delta \in \mathcal{C}$, $T \in \text{Ty}(\Delta)$, index set \mathcal{I} , covering $\{\theta_i : \Delta_i \rightarrow \Delta\}_i \triangleright \Delta$ for $i \in \mathcal{I}$, branch results $t_i \in \text{Tm}(T\{\theta_i\})$ for $i \in \mathcal{I}$, and scrutinee $\theta : \Gamma \rightarrow \Delta$, there exists a term $\text{match}_{(\Delta_i, \theta_i, t_i)}^{\rightarrow i}(\theta) \in \text{Tm}(T\{\theta\})$ such that:

- for any $\sigma_i : \Gamma \rightarrow \Delta_i$, if $\theta = \theta_i \circ \sigma_i$, then $\text{match}_{(\Delta_i, \theta_i, t_i)}^{\rightarrow i}(\theta) = t_i\{\sigma_i\}$;
- The above choice commutes with substitution, i.e., for $\theta_1 : \Gamma_1 \rightarrow \Gamma_2$ and $\theta_2 : \Gamma_2 \rightarrow \Delta$, we have $\text{match}_{(\Delta_i, \theta_i, t_i)}^{\rightarrow i}(\theta_2 \circ \theta_1) = (\text{match}_{(\Delta_i, \theta_i, t_i)}^{\rightarrow i}(\theta_2))\{\theta_1\}$;

To see this concretely, consider the **head** function from Section 2.3 in the CwF structure of COVERTT. The scrutinee type Δ is $\cdot, (A : \mathcal{U}), (n : \mathbb{N}), (x : \text{Vec } A \ n)$. The cover is the singleton $\{\cdot \mathbin{\text{\textcircled{;}}} (m+1) \mathbin{\text{\textcircled{;}}} \text{Cons } m \ h \ t \ \text{refl}_{m+1}\}$, where syntactic extension $\mathbin{\text{\textcircled{;}}}$ corresponds to $\langle _, _ \rangle$. This pattern corresponds to θ_1 . The pattern

context $\Delta_1 := \cdot, (m : \mathbb{N}), (h : A), (t : \text{Vec } A \ m)$, so the pattern in the cover corresponds to an arrow $\cdot, (m : \mathbb{N}), (h : A), (t : \text{Vec } A \ m) \rightarrow \cdot, (A : \mathcal{U}), (n : \mathbb{N}), (x : \text{Vec } A \ n)$. There is one branch, whose result t_1 is $h : A$ (which matches the overall result because the result type is not dependent). Finally, the scrutinees are the variables $\cdot \ ; n \ ; x$, which have context-type Δ in the empty context.

The pattern matching condition says that, if \triangleright is supported by a model of COVERTT, and $\{\cdot \ ; (m+1) \ ; \text{Cons } m \ h \ t \ \text{refl}_{m+1}\} \triangleright \Delta$, then there exists a term $\text{match}_{\Delta_1, \theta_1, t_1}$ such that, for any m, h, t , applying the substitution yields t , i.e. $[\cdot \ ; (m+1) \ ; \text{Cons } m \ h \ t \ \text{refl}_{m+1} / \Delta] \text{match}_{\Delta_1, \theta_1, t_1} = t$. In the case of our term model, this is given by the pattern match:

$$\text{case } (x : \cdot \ ; n \ ; x : \cdot, (n' : \mathbb{N}), (x' : \text{Vec } A \ n')) \text{ to } A \text{ of } \{ \\ \cdot, (m : \mathbb{N})(h : A), (t : \text{Vec } A \ m). \cdot \ ; (m+1) \ ; \text{Cons } m \ h \ t \ \text{refl}_{m+1} \Rightarrow h\}$$

However, in other models, there may not be an obvious way to form $\text{match}_{\Delta_1, \theta_1, t_1}$. We provide a general way of forming such a term in Section 4.

3.5 Model compatibility and soundness

The correspondence between the type formers we have introduced and the constructs of COVERTT is fairly direct, but we must account for some technical details to make the connection formal.

Not every possible syntactic cover of COVERTT leads to a non-trivial model. For example, if $\{(x : 0). \text{inl } x\} \triangleright (\mathbb{0} + \mathbb{N})$, then we can write an inhabitant for the empty type:

$$\text{bad} : \mathbb{0} := \text{case } (\text{inr } 3 : (\mathbb{0} + \mathbb{N})) \text{ to } \mathbb{0} \text{ of } \{(x : 0). (\text{inl } x) \Rightarrow x\}$$

If we have a CwF structure on a category \mathcal{C} supporting functions, labelled variants, and pattern matching in the sense of Section 3.4 with a coverage relation \triangleright , then we can soundly model COVERTT in \mathcal{C} so long as the syntactic \triangleright is compatible with the semantic \triangleright .

In Fig. 4 we define partial translations $\llbracket \Gamma \rrbracket \in \mathcal{C}$, $\llbracket \Gamma \vdash T \rrbracket \in \text{Ty}(\Gamma)$, $\llbracket \Gamma \vdash t : T \rrbracket \in \text{Tm}_\Gamma(T)$. We omit the cases other than pattern matching, as they are standard. The translation for pattern matches checks if the patterns translate to a semantic cover, which is well founded because each pattern is syntactically smaller than the entire match.

We have the following soundness result, characterizing how the syntactic coverage must correspond to a semantic coverage supporting pattern matching.

Theorem 1 (soundness). *If every COVERTT syntactic cover $\{\Delta_i. \mathbf{s}_i\}_i \triangleright \Xi$ has $\llbracket \Delta_i \rrbracket$ and $\llbracket \mathbf{s}_i \rrbracket$ defined, and $\{\llbracket \mathbf{s}_i \rrbracket\}_i \triangleright \llbracket \Xi \rrbracket$, then $\llbracket _ \rrbracket$ is a total mapping on terms/types/environments/contexts that are well typed with respect to \triangleright . Moreover, the model is sound with respect to definitional equality.*

Proof. If all covers are compatible, then straightforward induction shows that the undefined case never arises on well-typed terms. The argument follows the standard CwF model of type theory, except for pattern matching, where the equations from Section 3.4 directly satisfy EQMATCH.

$$\begin{aligned}
\llbracket \cdot \rrbracket &= 1 & \llbracket \Gamma, (x : T) \rrbracket &= \llbracket \Gamma \rrbracket \triangleright \llbracket T \rrbracket \\
\llbracket \Gamma \vdash \cdot \Rightarrow \cdot \rrbracket &= !\llbracket \Gamma \rrbracket : \Gamma \rightarrow 1 & \llbracket \Gamma \vdash s ; t \Rightarrow \Delta, (x : T) \rrbracket &= \langle \llbracket s \rrbracket, \llbracket t \rrbracket \rangle : \llbracket \Gamma \rrbracket \rightarrow \llbracket \Delta \rrbracket \triangleright \llbracket T \rrbracket \\
\llbracket \Gamma \vdash \text{case } (t_{\text{scrut}} : \Xi) \text{ to } T_{\text{motive}} \text{ of } \overrightarrow{\{\Delta_i. s_i \Rightarrow t_i\}} \rrbracket & \\
&= \text{match}_{(\llbracket \Delta_i \rrbracket, \llbracket s_i \rrbracket, \llbracket t_i \rrbracket) \xrightarrow{i}}^i (\llbracket t_{\text{scrut}} \rrbracket) \in \text{Tm}(T\{\llbracket t_{\text{scrut}} \rrbracket\}) \\
&\text{when } \llbracket t_i \rrbracket \text{ are defined for all } i \text{ and } \{\llbracket s_i \rrbracket\}_i \triangleright \llbracket \Xi \rrbracket \\
\llbracket \Gamma \vdash \text{case } (t_{\text{scrut}} : \Xi) \text{ to } T_{\text{motive}} \text{ of } \overrightarrow{\{\Delta_i. s_i \Rightarrow t_i\}} \rrbracket &\text{ undefined otherwise}
\end{aligned}$$

Fig. 4. Model of Pattern Matching in COVERTT

4 Coverages and Sheaves to Model COVERTT

Theorem 1 lists conditions that ensure we can soundly interpret COVERTT. What categories and coverages can we find that fulfill our criteria?

In this section we connect some core concepts of sheaf theory to pattern matching. Specifically, we provide a sufficient condition for when a coverage on a category \mathcal{C} admits semantic pattern matching as in Section 3.4.

4.1 Coverages and Sheaves

Alongside our contribution we provide a brief introduction to sheaves and sites for completeness. Systematic overviews of sheaves are given, for example, by Johnstone [15, C2] or MacLane and Moerdijk [17].

Coverages and Sites A **sheaf-theoretic coverage** on a category \mathcal{C} is, for each $\Delta \in \mathcal{C}$, a set of subsets of $\text{Hom}_{\mathcal{C}}(_, \Delta)$, called **covers**, which fulfill the following closure condition [15, C2.1.1]: For each cover $\{f_i : \Gamma_i \rightarrow \Delta\}_{i \in 1 \dots n}$, and other morphism $g : \Xi \rightarrow \Delta$ there exists a Ξ -cover $\{h_j : \Xi_j \rightarrow \Xi\}_{j \in 1 \dots m}$ such that, for each j , there exists an f_i such $g \circ h_j$ lifts along f_i . That is, there exists a k_j making the diagram to the right commute.

$$\begin{array}{ccc}
\Xi_j & \xrightarrow{k_j} & \Gamma_i \\
\downarrow h_j & & \downarrow f_i \\
\Xi & \xrightarrow{g} & \Delta
\end{array}$$

This condition is weaker than requiring covers to be closed under pullback by any arrow. In particular, we do not require \mathcal{C} to have all pullbacks. If \mathcal{C} does have all pullbacks, one can saturate the coverage so that the property of being a cover is preserved under pullbacks.

A sheaf-theoretic coverage on \mathcal{C} gives a coverage J for each object $\Delta \in \mathcal{C}$. Arrows in a J -cover share a common codomain, so it is clear to which object a cover belongs. A **site** (\mathcal{C}, J) is a category \mathcal{C} equipped with a coverage J .

We refer to “sheaf theoretic coverages” specifically to distinguish them from coverages in the sense of Sections 2 and 3, i.e. the sets of patterns that we allow. Both denote the sets of morphisms/patterns that cover a given context, but we don’t require pattern coverages to fulfill the sheaf-theoretic closure conditions. In Section 7.2 we show that a constructive syntactic model cannot fulfill them.

Sheaves The next conceptual tools we need are those of a presheaf and a sheaf. A **presheaf** P is a functor $P : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$. Sheaf theorists think of presheaves as abstract collection of functions/terms, with $P\Delta \in \mathbf{Set}$ being the set of functions out of Δ /terms in context Δ . A **sheaf** is a collection that ‘thinks’ all pattern matches over every cover uniquely defines a term. Formulating sheaves precisely involves multiple nested quantifiers, and we break it down in stages.

For any presheaf P and cover $\theta_i : \Delta_i \rightarrow \Delta \in J$, a **matching family** is a collection $x_i \in P(\Delta_i)$ such that for every $\sigma : \Xi \rightarrow \Delta_i$ and $\sigma' : \Xi \rightarrow \Delta_j$, if $\theta_i \circ \sigma = \theta_j \circ \sigma'$, then $P(\sigma)(x_i) = P(\sigma')(x_j)$. I.e., a matching family assigns a P value for all arrows in a cover, while agreeing in overlapping cases.

An **amalgamation** of a matching family $\overrightarrow{x_i}$ over $\theta_i : \Delta_i \rightarrow \Delta \in J$ is an $x \in P\Delta$ such that $P(\theta_i)(x) = x_i$, i.e., it is a value in the covered object that is compatible with the matching family.

A **sheaf** on (\mathcal{C}, J) for a cover $\theta_i : \Delta_i \rightarrow \Delta \in J$ is a presheaf P such that every matching family has a unique amalgamation. A presheaf is a J -sheaf, or just a sheaf, when it is a sheaf for each cover in J . It is in this way that a J -sheaf is a presheaf that ‘thinks’ all covers admit pattern-matching.

Let $\mathbf{y} : \mathcal{C} \rightarrow (\mathcal{C}^{\text{op}} \rightarrow \mathbf{Set})$ denote the **Yoneda embedding** that maps each $\Gamma \in \mathcal{C}$ to the presheaf $\text{Hom}_{\mathcal{C}}(_, \Gamma)$. A coverage is **subcanonical** when for every $\Delta \in \mathcal{C}$, $\mathbf{y}\Delta$ is a sheaf. There is a largest such coverage $J_{\text{canonical}}$ —the **canonical coverage**. We say a cover is **canonical** when it is in the canonical coverage. Every representable is a sheaf for a canonical cover, though in Section 7.2 we disprove the converse: there are models which support pattern matching, so every representable is a sheaf for each allowed pattern, but where the allowed patterns do not fulfill the necessary conditions to be a sheaf-theoretic coverage.

4.2 Pattern Matching via Sheaves

The similarity between amalgamation and pattern matching is apparent, and was informally established by Coquand [9]: since morphisms in \mathcal{C} correspond to substitutions (sequences of terms) in COVERTT, the sheaf condition gives a way to merge arrows (branches) with the same codomain (return type). However, to model dependent pattern matching, we need to handle the dependency of the branch result type on the scrutinee’s value. Thankfully, slices give us the tools to model type dependency, and sheaf theory lets us convert subcanonical coverages on a category to coverages on a slice. The key properties, which we have mechanized in Lean [11], are as follows (see, e.g. Johnstone [15, C2.2.17]):

Theorem 2. ($\sqrt{\text{LEAN}}$) *If (\mathcal{C}, J) is a subcanonical site, then for $\Gamma \in \mathcal{C}$, the site $(\mathcal{C}/\Gamma, J_\Gamma)$ is subcanonical, where we define $\{f_i : (\Delta_i, \theta_i) \rightarrow (\Xi, \sigma)\}_i \in J_\Gamma$ if and only if $\{f_i : \Delta_i \rightarrow \Xi\}_i \in J$. In particular, if $\{f_i : \Delta_i \rightarrow \Gamma\}_i$ is canonical, then $\{f_i : (\Delta_i, f_i) \rightarrow (\Gamma, \text{id})\}_i$ is too.*

These properties are related to the **fundamental theorem of topos theory**, which says that a slice of a sheaf category is equivalent to a category of sheaves over the slice.

We now have what we need to state and prove the main result of this section: a criterion ensuring that a coverage can model pattern matching. The following theorem has been mechanized in the Lean 4 theorem prover [11]; work is underway to mechanize the model's soundness and the coverage building rules of Section 5.

Theorem 3. (\checkmark LEAN) *Consider a CwF \mathcal{C} and, for each $\Delta \in \mathcal{C}$, a relation $\{\theta_i\}_{i \triangleright \Delta}$ where the θ_i are disjoint monomorphisms into Δ . If all covers in \triangleright are canonical, then \mathcal{C} supports pattern matching (in the sense of Section 3.4).*

Proof. Let $(C, J_{\text{canonical}})$ be a canonical site with a CwF structure and a relation $\triangleright \subseteq J_{\text{canonical}}$. Consider a scrutinee type $\Delta : \mathcal{C}$, dependent result type $T \in \text{Ty}(\Delta)$, canonical cover $\{\theta_i : \Delta_i \rightarrow \Delta\}_{i \triangleright \Delta}$ of non-overlapping monos, branch results $t_i \in \text{Tm}(T\{\theta_i\})$, and scrutinee $\theta : \Gamma \rightarrow \Delta$. To construct $\text{match}_{(\Delta_i, \theta_i, t_i)}^i(\theta) \in \text{Tm}(T\{\theta\})$, we build $t'_{\text{match}} \in \text{Tm}_\Delta(T\{id\})$ with which we compose the scrutinee θ . We will show how the sheaf condition corresponds to the pattern match.

Matches as Arrows Recall from Lemma 1 that there is a **Set**-isomorphism $\text{Tm}_\Delta(T\{\theta\}) \cong \text{Hom}_{(\mathcal{C}/\Gamma)}((\Delta, \theta), (\Gamma \triangleright T, \mathbf{p}_T))$. To find a term in $\text{Tm}_\Delta(T\{id\})$, we use an arrow in $\text{Hom}_{(\mathcal{C}/\Delta)}((\Delta, id), (\Delta \triangleright T, \mathbf{p})) = \mathbf{y}(\Delta \triangleright T, \mathbf{p})(\Delta, id)$.

Pattern Sets as Slice Covers The patterns $\{\theta_i : \Delta_i \rightarrow \Delta\}_i$ correspond to a canonical \mathcal{C} -cover by our premise, so by Thm. 2 the cover $\{\theta'_i : (\Delta_i, \theta_i) \rightarrow (\Delta, id)\}_i$ is canonical in \mathcal{C}/Δ .

Branches as Matching Families The branch results of the pattern match form a matching family for $\mathbf{y}((\Gamma \triangleright T, \mathbf{p}))$. Our branches are $\overrightarrow{t_i \in \text{Tm}_{\Delta_i}(T\{\theta_i\})}^i$. By Lem. 1, this family yields a sequence $\overrightarrow{x_i \in \text{Hom}_{(\mathcal{C}/\Delta)}((\Delta_i, \theta_i), (\Delta \triangleright T, \mathbf{p}))}^i$, i.e., $\overrightarrow{x_i \in \mathbf{y}((\Delta \triangleright T, \mathbf{p}))((\Delta_i, \theta_i))}^i$, which is a matching family for the presheaf $\mathbf{y}(\Delta \triangleright T, \mathbf{p})$ and the cover $\{\theta_i : (\Delta_i, \theta_i) \rightarrow (\Delta, id)\}_i$.

Amalgamating Branches Because the cover is canonical, then $\mathbf{y}(\Delta \triangleright T, \mathbf{p})$ is a sheaf for it. The sheaf condition states that the above matching family has an amalgamation $x \in \mathbf{y}(\Delta \triangleright T, \mathbf{p})(\Delta, id)$, such that $\theta_i \circ x = x_i$. So Lem. 1 yields a term $t'_{\text{match}} \in \text{Tm}(T\{id\})$ such that $t'_{\text{match}}\{\theta_i\} = t_i$.

Equations and the Scrutinee Finally, given a scrutinee $\theta : \Gamma \rightarrow \Delta$, we choose $t'_{\text{match}}\{\theta\}$ as $\text{match}_{(\Delta_i, \theta_i, t_i)}^i(\theta) \in \text{Tm}(T\{\theta\})$. It is in $\text{Tm}_\Gamma(T\{\theta\})$, so it has the correct type. It satisfies the requisite equations. Indeed, since $t'_{\text{match}}\{\theta_i\} = t_i$, whenever $\theta = \theta_i \circ \sigma_i$ for some σ_i , we have $t'_{\text{match}}\{\theta\} = t'_{\text{match}}\{\theta_i \circ \sigma_i\} = (t'_{\text{match}}\{\theta_i\})\{\sigma_i\} = t_i\{\sigma_i\}$ just as Section 3.4 requires. For substitution, given $\theta = \theta_2 \circ \theta_1$, we have $t'_{\text{match}}\{\theta_2 \circ \theta_1\} = (t'_{\text{match}}\{\theta_2\})\{\theta_1\} = \text{match}_{(\Delta_i, \theta_i, t_i)}^i(\theta_2)\{\theta_1\}$.

The above construction lets us model pattern matching for any canonical cover, where the motive type corresponds to a representable sheaf. If J is subcanonical, then every representable is a sheaf, so we can define dependent pattern matching for any motive type. Moreover, the canonical coverage contains the covers from every subcanonical coverage, so it suffices that each allowed pattern set is a canonical cover. The disjointness and injectivity conditions ensure that the branches of a match follow the sheaf-theoretic definition of a matching family. One could instead require that branches agree on their overlap [7], which is suited to matching on real numbers [25].

We conclude this section by recalling a property characterizing subcanonical covers [15, C2.1.11], which we will utilize in Section 5.

Theorem 4. *A set of arrows $\{\theta_i : U_i \rightarrow U\}_i$ is canonical for \mathcal{C} if and only if, for every $\sigma : V \rightarrow U$ and every object $T \in \mathcal{C}$, the presheaf $\mathbf{y}(T)$ is a sheaf for the pulled back family $\{\sigma^*\theta_i : \sigma^*U_i \rightarrow V\}_i$.*

We can use this theorem to form basic subcanonical coverages. To build new coverages from old ones, we employ **saturation conditions**: operations on coverages which do not change which presheaves are sheaves for that coverage. This is of interest to us because the canonical coverage is invariant under every saturation: it is already the largest possible subcanonical coverage, so no covers can be added without changing its notion of sheaf. The next section gives several examples of useful saturation conditions.

5 Tools for Building Coverages

In this section, we take the abstract canonicity condition and derive concrete rules for forming canonical covers. This section justifies the idea that COVERTT can begin with a basic set of covers for each type and obtain a language of patterns by allowing nesting, variables, and multiple scrutinees. We provide base coverages, along with composition rules for building complex coverages from simpler ones. This recreates commonly supported features of dependent pattern matching: variables, constructors for labelled variants, pruning absurd branches, and matching on `refl` with inaccessible (`dot`) patterns. In Section 7.2 we discuss potential novel coverages that can be supported using coverage semantics.

5.1 Identity and Isomorphism

The most basic canonical covers are singletons consisting of an isomorphism. If two contexts are isomorphic, then moving from one to the other covers all cases.

Lemma 2. *A presheaf is a sheaf for a singleton cover containing an isomorphism $\iota : \Gamma \cong \Delta$, so every isomorphism is a canonical singleton cover.*³

³ In a category with pullbacks, sheaves may be defined for Grothendieck pretopologies, in which all isomorphisms definitionally yield singleton covers.

As a consequence, identity arrows are in singleton canonical covers:

Corollary 1. *The singleton cover $\{id : \Gamma \rightarrow \Gamma\}$ is canonical for any category, so a pattern consisting entirely of variables x_1, x_2, \dots, x_n can be safely included in any coverage for COVERTT.*

Supporting identity arrows is the bare minimum we need for pattern matching. They are the base case out of which other patterns are built, where no discrimination or computation happens at all. Likewise, a catch-all pattern, commonly written as an underscore ‘ $_$ ’, is just an unnamed variable that does not occur in the right-hand side of the branch.

More generally, the canonical coverage contains all isomorphisms. So we can devise a sound semantics for COVERTT where any isomorphism is a valid cover of a type. This allows for operations such as rearranging variables in a dependency-respecting way or re-bracketing nested sums and products.

Functions written by the programmer can even be used as patterns if they are isomorphisms, opening the door for user-defined views into a type. Of course, the existence of a sound semantics does not guarantee a language we can actually implement. Checking whether a term is a definitional isomorphism is undecidable without being explicitly given its inverse. Moreover, depending on how extensional the model is, there may be terms that are isomorphisms in the model, but are not definitional isomorphisms in COVERTT.

5.2 Coproducts and Wadler Views

With variables as patterns, the next primitive patterns we need are constructors for a datatype. As we saw in Section 3.3, one way to model datatypes is with labelled variants. So if \mathcal{C} has coproducts, we can model datatype constructors as injections into a coproduct context, and we can amalgamate branches that match on all the constructors of a datatype using the universal property of a coproduct.

Labelled variants are only coproducts up to isomorphism, but we have seen that isomorphisms are always singleton covers, and below in Section 5.3 we see that composition preserves canonical covers. So it suffices to consider coproducts directly. Unfortunately, in an arbitrary category, coproduct injections are not guaranteed to form a canonical cover. Thm. 4 requires each representable to be a sheaf for the pullback of every cover, so we need coproducts to be stable under pullback, i.e. the pullback of a coproduct is the coproduct of pullbacks. Thankfully, pullback stability of coproducts holds if \mathcal{C} is **Set**, a presheaf category, a topos, or any other locally cartesian closed category (LCCC). We already want \mathcal{C} to be LCCC in order to support dependent functions.

Theorem 5. *Suppose \mathcal{C} has all pullbacks and that coproducts are disjoint and stable under pullback. Then $\{\text{inj}_j : \Delta_i \rightarrow \coprod_{i \in I} \Delta_i\}_j$ canonically cover $\coprod_{i \in I} \Delta_i$.*

The immediate result of this theorem is that for every inductive type, the constructors are covering for that type, so long as the inductive type is modelled

as the labelled variants of its constructor types. However, it is important to realize that this theorem applies for any decomposition of a type into the coproduct of other types, regardless of whether the injections correspond to constructors or not. Such a coverage introduces the possibility for views as introduced by Wadler [27], which act as first class pattern synonyms:

Corollary 2. *Consider a finite $I : \mathcal{U}$, a family $S : I \rightarrow \mathcal{U}$ and an indexed function $f : (i : I) \rightarrow S\ i \rightarrow T$ in COVERTT . Suppose we have a category \mathcal{C} with a CwF model of COVERTT supporting pattern matching as in Section 3.4 over a coverage $_ \triangleright _$. If $\llbracket T \rrbracket \cong \llbracket \Sigma(t : T)(i : I)(s : S\ i). t =_T f\ i\ s \rrbracket$, then there is also a model of COVERTT with coverage $(_ \triangleright _ \cup \{f\ i\}_i)$*

That is, if T is isomorphic to the sum of some types $S\ i$ over finite i , we can safely match on a value from T , where the i th pattern is the $S\ i$ value that it corresponds to, regardless of whether T is defined as an inductive type or the $S\ i$ are its constructor types.

5.3 Nesting and Composition

With variables and constructors as primitive covers, we now need a way to combine them to build more complex covers. Adding the composition of different coverages does not change their sheaves. Every cover has the same sheaves as the sieve it generates, i.e., the closure of the cover under precomposition with any arrow in \mathcal{C} [15, C2.1.3]. So canonical covers are closed under composition:

Theorem 6. *For a cover J of \mathcal{C} , If $\{f_i : \Delta_i \rightarrow \Delta\}_i \in J$, and for each i , we have $\{g_{ij} : \Delta_{ij} \rightarrow \Delta_i\}_{ij} \in J$, then the sheaves of (\mathcal{C}, J) are identical to the sheaves of $(\mathcal{C}, J \cup \{f_i \circ g_{ij} : \Delta_{ij} \rightarrow \Delta\}_i)$. So if $\{f_i : \Delta_i \rightarrow \Delta\}_i$ is canonical, and for each i , $\{g_{ij} : \Delta_{ij} \rightarrow \Delta_i\}_{ij}$ is canonical, then $\{f_i \circ g_{ij} : \Delta_{ij} \rightarrow \Delta\}_{ij}$ is canonical.*

This allows patterns to be nested: if $\{f_i\}_i$ are covering for Δ , and for each set of variables in those covering patterns, $\{g_{ij}\}_{ij}$ is covering, then we can case split each variable in the f_i into j cases corresponding to the g_{ij} patterns, and the entire resulting set can still be covering. This property gives semantic justification for the case-split operation of the Agda and Idris editor modes, where the programmer selects a variable in a pattern match, and the pattern containing the variable is replaced by the sequence of patterns that has each possible constructor application in place of the variable.

When we combine the closure of the canonical coverage under identity (isomorphisms), sum injections, and composition, we can recreate dependent pattern matching on non-indexed datatypes. However, we see in the next sections that the language of coverages also gives us the tools to handle indexing.

5.4 Pruning Absurd Contexts

If we allow ourselves some extensionality, then we can use sheaves to model absurd branches and empty cases. Suppose that \mathcal{C} has an initial object 0 , with

a unique arrow $0_\Gamma : \emptyset \rightarrow \Gamma$ for every Γ . The initial context denotes an empty or absurd context, since we can derive a term of any type from it. It turns out that we do not ever need to include branches for patterns whose contexts are empty. If we can amalgamate for a cover where one pattern has an empty context, we can amalgamate for the same cover with that arrow deleted, since any matching family for the smaller cover can be turned into one for the larger cover by adding the unique arrow out of the initial context.

Theorem 7. *Let c be a canonical cover with $\theta : \Delta \rightarrow \Gamma \in c$. If there exists an arrow $\iota : \Delta \rightarrow \emptyset$, then $c \setminus \{\theta\}$ is canonical.*

This property mirrors how Agda and Idris allow for the omission of empty cases. In some cases, these languages only allow branch right-hand sides to be removed after the programmer specifies an empty pattern, marking which part of the scrutinee has an impossible type. We view this empty pattern as a syntactic aid to tell the type checker when a context is isomorphic to the initial context, so we do not directly model the empty pattern.

Like our assumption about equality, the condition of having a (strong) initial object does not hold in the term model, since not all eliminations of the empty type are definitionally equal. So long as COVERTT does not contain any axioms that specifically distinguish empty eliminations, our model is still sound.

5.5 Propositional Equality

Since isomorphisms are always canonical singleton covers (Lem. 2), we can create a coverage for a sufficiently-extensional equality type.

Corollary 3. *If $\langle v, \langle v, \text{Refl}_A \rangle \rangle : \Gamma \triangleright A \rightarrow \Gamma \triangleright A \triangleright A\{p\} \triangleright \text{Id}(A\{p^2\}, v, v\{p\})$ is an isomorphism, then $\{\langle v, \langle v, \text{Refl}_A \rangle \rangle\}$ is canonical.*

In more readable, non CwF notation: if A is isomorphic to $\Sigma(x : A)(y : A). x =_A y$ in the model via the projections, then $\{(x, x, \text{refl}_x)\}$ can be a singleton cover for $\Sigma(x : A)(y : A). x =_A y$. Such an isomorphism holds if \mathcal{C} has extensional equality, since it asserts that there is a unique, internally constructible proof of equality between two equal terms.

For a dependent match targeting $(x : A), (y : A), (pf : x =_A y) \vdash P(x, y, pf) : \mathcal{U}$, the above cover only requires we provide a branch result with type $P(x, x, \text{Refl}_A)$. The variable y was replaced by x in the goal type. This captures “inaccessible” or “forced” patterns [22], known as dot-patterns in Agda and Idris. By matching on the propositional equality, we work with refined information about the context. Here, we match y against the variable x rather than a constructor. No branching or discrimination happening, since the cover is a singleton. Rather, x is the only possible value for y given the equality proof. Agda writes $.x$ to express this pattern. Section 5.6 extends this to equality proof between arbitrary terms, rather than variables.

5.6 Pullbacks and Unification

We have seen that the canonical coverage includes isomorphisms and injections and that it allows for composition. However, the definition of a sheaf-theoretic coverage enables stability under pullback: for any coverage, a sheaf for that coverage is still a sheaf if we add the pullback of any cover by any arrow. Closure under pullback is the key condition that separates a coverage from a set of arrows. The definition in Section 4.1 is presented in the style of Johnstone [15, C2], in a general way that does not assume the existence of pullbacks. However, when each arrow in a cover has a pullback along some morphism, we get the following saturation condition:

Theorem 8. *For a cover J of \mathcal{C} , if $\{\theta_i : \Delta_i \rightarrow \Delta\}_i \in J$, and $g : \Gamma \rightarrow \Delta$, where the pullback of each θ_i along g exists, then $J \cup \{\{g^*\theta_i : g^*\Delta_i \rightarrow \Gamma\}_i\}$ has the same sheaves as J . So if $\{\theta_i : \Delta_i \rightarrow \Delta\}_i$ is canonical, then so is $\{g^*\theta_i : g^*\Delta_i \rightarrow \Gamma\}_i$.*

This abstract property, known as **stability under base change**, can be exploited to build interesting covers.

Context Extension Base change lets us add new scrutinees to a pattern match. In any CwF, for $\theta : \Gamma \rightarrow \Delta$ and $T \in \text{Ty}(\Delta)$, pulling back by $\mathbf{p} : \Delta \triangleright T \rightarrow \Delta$ yields a morphism $\langle \theta \circ \mathbf{p}, \mathbf{v} \rangle : \Gamma \triangleright T \{\theta\} \rightarrow \Delta \triangleright T$ [14]. Combining this with Thm. 8 gives:

Theorem 9. *For canonical $\{\theta_i : \Delta_i \rightarrow \Delta\}_i$ and a type $T \in \text{Ty}(\Delta)$, there is also a canonical cover $\{\langle \theta_i \circ \mathbf{p}, \mathbf{v} \rangle : \Delta_i \triangleright T \{\theta_i\} \rightarrow \Delta \triangleright T\}_i$.*

So we can build a covering pattern for a context Δ and immediately obtain a covering context on $\Delta \triangleright T$ by appending a new variable $\mathbf{v} \in \text{Tm}(T\{\theta_i\})$ to each pattern in the cover. In a dependent match the new variable might have a different type in each branch: T may be indexed by variables in Δ , but each θ_i is a value for Δ with variables from Δ_i . Further case-splitting on the newly introduced variable can be achieved using composition à la Section 5.3.

Matching on Equality Suppose that for some $\Gamma \in \mathcal{C}$ and $T \in \text{Ty}(\Gamma)$, and that: $\{\text{Refl}_T : \Gamma \triangleright T \rightarrow \Gamma \triangleright T \triangleright T\{\mathbf{p}\} \triangleright \text{Id}(T\{\mathbf{p}^2\}, \mathbf{v}, \mathbf{v}\{\mathbf{p}\})\}$ is canonical. As in Section 5.5, such a property holds for an extensional model. With such a cover on equality, we can apply the base change theorem, the CwF laws, and the properties of equality (Section 3.3) to obtain a cover on contexts containing equalities by matching:

Theorem 10. *Suppose \mathcal{C} has all pullbacks. Consider $\Gamma \in \mathcal{C}$ with $T \in \text{Ty}(\Gamma)$, $t_1, t_2 \in \text{Tm}(T)$. Then pulling back Refl_T by $\langle \langle t_1\{\mathbf{p}\}, t_2\{\mathbf{p}\} \rangle, \mathbf{v} \rangle$ yield a context Δ and an arrow $\langle \theta, \text{Refl}_T(t_{12}) \rangle$, where $\theta : \Delta \rightarrow \Gamma$, $t_{12} \in \text{Tm}(T\{\theta\})$, and $t_1\{\theta\} = t_2\{\theta\} = t_{12}$. Moreover, if $\{\text{Refl}_T\}$ is canonical, then so is the cover $\{\langle \theta, \text{Refl}_T(t_{12}) \rangle : \Delta \rightarrow \Gamma \triangleright \text{Id}(T, t_1, t_2)\}$.*

For the intuition behind this, consider the pullback square to the right, translated to COVERTT-style notation for clarity. First, because the square commutes, we know that θ is a substitution that equates (t_1, t_2, pf) and

$$\begin{array}{ccc} \Delta & \xrightarrow{\theta \circ t_{12}} & \Gamma, (x : T) \\ \downarrow \theta \circ \text{refl}_{t_{12}} & & \downarrow (x \circ x \circ \text{refl}_x) \\ \Gamma, (pf : x =_T y) & \xrightarrow{(t_1 \circ t_2 \circ s)} & \Gamma, (x : T), (y : T), (x =_T y) \end{array}$$

(x, x, refl_x) , i.e., it is a **unifier** of t_1 and t_2 . Since a pullback is a limit, it is universal, so any other unifiers for t_1 and t_2 necessarily factor through θ . Thus, it is the **most general unifier** for t_1 and t_2 . This is precisely what the usual rule for pattern matching on equality uses: it unifies the two sides of the equality, treating syntactic variables as unification variables, and generates a substitution that is then applied to the goal. The context Δ consists of the variables that were in common between t_1 and t_2 which remain free in the unification t_{12} . In the case that t_1 and t_2 do not unify, then the pullback is an arrow out of an initial context, and the branch can be omitted completely (because absurd covers can be omitted, as in Section 5.4).

6 Example: Folding Without a Starting Value

We now have specified everything we need (sans recursion) for feature-parity with the original presentation of dependent pattern matching by Coquand [9]. Our sheaf-centric view generalizes the elaboration process of Goguen et al. [12], but directly within the model instead of as a syntactic elaboration. Constructors for a coproduct form a cover and variables form a singleton cover, acting as the basis from which other covers are generated. Covers can be composed, extended, refined by matching on an equality, or pared down by pruning absurd branches. Indexed data types can be handled using folding and matching on equality.

To see a non-trivial example of how to build a cover in the canonical coverage, consider the `foldr1` function found in the Agda standard library [10].

$\text{foldr}_1 : (A \rightarrow A \rightarrow A) \rightarrow \text{Vec } A \text{ (suc } n) \rightarrow A$

$\text{foldr}_1 f (\text{cons } x \text{ nil}) = x \mid \text{foldr}_1 f (\text{cons } x (\text{cons } y \text{ ys})) = f x (\text{foldr}_1 f (\text{cons } y \text{ ys}))$

Because the argument vector has length at least one, the case for `nil` can be omitted. The base case is then a vector of length one, and the inductive case is a vector of length two or more.

Assume an extensional CwF model of COVERTT in a LCCC \mathcal{C} where inductive types are labelled variants. We show how the patterns for `foldr1` are in the canonical coverage for \mathcal{C} , and hence `foldr1` can be modelled. Note that because labelled variants are defined in terms of isomorphism, we do not preclude initial algebra semantics for modelling the self-reference part of inductive types.

6.1 Translating to COVERTT

First, we translate the function to COVERTT-style by making the length argument explicit and replacing the indexed constructors with ones taking explicit

equality proofs. The datatype becomes:

```
data Vec (A :  $\mathcal{U}$ ) : (n :  $\mathbb{N}$ ) →  $\mathcal{U}$  where
  nil : (n = 0) → Vec A n
  cons : (m :  $\mathbb{N}$ ) → A → Vec A m → n = m + 1 → Vec A n.
```

We also abstract out the recursive calls, since we have not included them in COVERTT and have not required that our model category \mathcal{C} support them. Despite using recursion, foldr_1 is an ideal example because it is not contrived, and uses all the main saturation conditions we developed in Section 5. Since foldr_1 is decreasing in the length of the lists, its recursion can be modelled with well-known techniques orthogonal to our contribution.

```
foldr1 : (n :  $\mathbb{N}$ ) → (A → A → A) → Vec A (suc n)
          → (self : (A → A → A) → Vec A (suc n) → A) → A
foldr1 0 f (cons 0 x (nil Refl) Refl) self = x
foldr1 (m + 1) f (cons (m + 1) x (cons m y ys Refl) self)
              = f x (self f (cons y ys Refl) Refl)
```

6.2 Building the Coverage

Using COVERTT notation rather than CwF notation for clarity and space reasons, we now show how the rules of Section 5 can be used to build a cover:

$$\begin{aligned} & \{((m + 1) f (\text{cons } (m + 1) x (\text{nil Refl}) \text{Refl}) \text{self}), \\ & (m + 1) f (\text{cons } (m + 1) x (\text{cons } m y ys \text{Refl}) \text{Refl}) \text{self}\} \\ & \triangleright (n : \mathbb{N}), (f : A \rightarrow A \rightarrow A), (v : \text{Vec } A (\text{suc } n)), (\text{self} : \dots) \end{aligned}$$

- Identity (Cor. 1) has variables $(n)(f)(v)(\text{self})$ covering the scrutinee type;
- Coproduct (Cor. 2) has $\{(\text{nil } eq_{\text{nil}}), (\text{cons } m' x xs eq_{\text{cons}})\} \triangleright \text{Vec } A q$ for any q ;
- Composition (Thm. 6) allows us to construct the canonical cover $\{(n f (\text{nil } eq_{\text{nil}}) \text{self}), (n f (\text{cons } m' x xs eq_{\text{cons}}) \text{self})\}$;
- We have $eq_{\text{nil}} : n + 1 = 0$, but this type is empty, so applying the absurd rule (Thm. 7) gives a singleton cover $\{(n f (\text{cons } m' x xs eq_{\text{cons}}) \text{self})\}$;
- $eq_{\text{cons}} : n + 1 = m' + 1$, so we get a pullback substitution mapping m' to n and all other variables to themselves. Applying the Refl rule (Cor. 3), $n n \text{Refl}$ is a cover of $(m' : \mathbb{N}) \triangleright (n : \mathbb{N}) \triangleright (n + 1 = m' + 1)$. By composition, the scrutinee context has singleton cover $\{(n f (\text{cons } n x xs \text{Refl}) \text{self})\}$ in the canonical coverage;
- The coproduct property for Vec and composition yield a cover $\{(n f (\text{cons } n x (\text{nil } eq_{\text{nil}}) \text{Refl}) \text{self}), (n f (\text{cons } n x (\text{cons } m y ys eq_{\text{cons}}) \text{Refl}) \text{self})\}$;
- Finally, since $eq_{\text{nil}} : n = 0$ and $eq_{\text{cons}} : n = m + 1$, we apply the Refl rule for each proof, along with composition, to obtain the desired cover above.

Then, we can use the sheaf condition model how $f x (\text{self } f (\text{cons } y ys \text{Refl}) \text{Refl})$ and x are amalgamated into a denotation for the entire function.

7 Discussion

7.1 Related Work

Dependent pattern matching was first proposed by Coquand [9]. While this work contains no explicit mentions of sheaf theory, it originated the idea that patterns could be thought of in terms of coverings and partitions of a space, which greatly inspired our work. The theory and practice of both pattern matching and eliminators foundational developments in proof assistants: McBride [19] developed elimination for LEGO, and later EPIGRAM [20]. This was extended to views and with-clauses by McBride and McKinna [21].

Goguen et al. [12] show how pattern matching can be elaborated to primitive eliminators, and hence given semantics in any model that had semantics for eliminators. These are in turn given semantics using initial algebras [1, 2]. Cockx et al. [6] extend this to work with univalent theories. Elaboration is similar to amalgamation using the sheaf condition, but amalgamation occurs strictly in the model. Elaborating to eliminators also handles recursion, which is not yet explicitly included in our sheaf semantics.

To our knowledge, the first explicit connection between the sheaves and pattern matching was by Sherman et al. [25], which gave a framework for pattern-matching on real numbers using topological spaces. The thesis version of this work [24] generalizes the approach from topological spaces to Grothendieck topologies. Cockx et al. [7] give similar semantics to overlapping patterns by treating them as definitional equalities, using confluence rather than sheaves. Using equalizers or pullbacks to represent unification was originated by Rydeheard and Burstall [23], as well as Goguen [13].

7.2 Future Work

First-Class Pattern Synonyms An immediate application of this work would be to implement an enhanced version of pattern synonyms in a language like Agda. Currently, Agda lets the programmer declare pattern synonyms, but each name must map to a syntactic pattern i.e., a set of nested constructor applications. Agda checks if a definition is covering by elaborating to these patterns. Our framework could be used to build direct coverage checking for pattern synonyms, so the programmer could build their own alternate, extensible covers of a type, using sheaf theory to justify their coverage. This would provide direct semantics for the user-defined views of Wadler [27] and McBride and McKinna [21]. Further research is needed to extract a constructive procedure for amalgamating branches that can be implemented in practice.

Overlapping Pattern Matches Our semantics require non-overlapping, injective patterns, but our framework suggests a way to lift this restriction. Recall that the sheaf condition only requires that a matching family agrees on the overlap between covering patterns. This suggests two ways to give semantics to overlapping patterns: by ensuring that the right-hand sides of each pattern

match agree on the overlap of their left-hand sides, or by adding information to each pattern to ensure they are actually non-overlapping.

The latter approach matches current implementations: catch-all patterns are elaborated into multiple branches whose left-hand patterns are the constructors that have not yet been used. Unfortunately, to prove anything about a function defined this way, the programmer needs a proof case for each branch in the elaboration, even if they correspond to a single branch in the function as written. Our framework may support canonical covers which contain extra information preventing overlap, such as proofs that previous branches had not matched. These could be used to develop covers for matches with overlapping patterns that do not require creating additional cases during elaboration, enabling more succinct proofs about overlapping cases.

Inductive Datatypes and Termination Checking When patterns are not restricted to constructors, it is not immediately apparent which recursive pattern matching functions can be soundly modelled, since pattern variables may not be structurally smaller than the patterns in which they occur. Further study is needed to devise criteria for which recursive definitions are well founded with non-constructor patterns.

Beyond Top-Level Matches Our semantics only support top level pattern matches. Many of the results we used, such as the fundamental theorem of topos theory, are well suited to top-level matches but do not directly translate to terms in an arbitrary context. Additionally, if the scrutinee type of a pattern match is in a non-empty context, then matching affects not only the motive, but may refine the values or types of variables in the context on the left. These technical issues suggest a semantic theory of *telescopes*, which are objects representing extensions to a given context by some number of types, and *environments*, which extend substitutions by some number of terms.

With Clauses Our approach to matching on equality proofs gives an intuition for modelling Agda-style with-clauses and views [21], though a full account is beyond the scope of this paper. Suppose we are defining a pattern match with scrutinees of type Δ and result type $\Delta \vdash T : \mathcal{U}$, and we want to match on some intermediate expression $\Delta \vdash s : S$. There is an isomorphism $\iota : \Delta \cong \Delta, (x : S), (pf : x =_S s)$. So if we have a cover of $\{s_i : \Delta_i \rightarrow \Delta\}_i$ to match s against, we can use composition and extension to obtain a cover $\{s_i : \Delta, (pf : s_i = s) \rightarrow \Delta\}_i$. In the case that s_i and s unify, the Refl rule from above can be used to match on the equality, and in an extensional model, the goal type can be safely rewritten due to the existence of the equality proof.

Intensional Models Our current approach relies on extensional models, where equality proofs correspond with equality in the model. We can define models of COVERTT in terms of canonical coverages and non-syntactic equality, and we

can give a CwF term model for COVERTT because syntactic pattern matching fulfills the criteria of Section 3.4, but the term model for COVERTT cannot be described in terms of sheaf-theoretic coverages.

To show the issue, we show that $\{\text{true}, \text{false}\}$ is not a canonical cover of **Bool** for the CwF given by well-typed COVERTT terms quotiented by definitional equality. Consider a function $\text{haltsInN} : \text{SyntaxTree} \rightarrow \mathbb{N} \rightarrow \text{Bool}$, that looks at a syntax tree of an untyped lambda calculus term and checks whether it halts in n or fewer steps. Consider also $\Omega : \text{SyntaxTree}$, a representation of $(\lambda x. x \ x)(\lambda x. x \ x)$. If $\{\text{true}, \text{false}\}$ is a canonical cover, it is also a sieve in the canonical topology [15, C2.1.8], so pulling the sieve back by $\text{haltsInN} \ \Omega$ produces the set of arrows $\{h_j \circ !_{V_j} : V_j \rightarrow \mathbb{N} \mid h_j : \mathbb{1} \rightarrow \mathbb{N}\}$. This contains each arrow in $\mathbb{1} \rightarrow \mathbb{N}$, i.e., each natural number. For the cover to be canonical, for any type T there must be a way to amalgamate $\{t_j : T \mid h_j : \mathbb{1} \rightarrow \mathbb{N}\}$ into $\mathbb{N} \rightarrow T$. Then all set-theoretic infinite sequences of natural numbers could be amalgamated into type-theoretic functions $\mathbb{N} \rightarrow \mathbb{N}$, which is impossible.

The above example relies on the existence of an infinite cover. While infinite covers are allowed in sheaf theory, they do not correspond directly to pattern matches that a programmer can write down. So further exploration of finite and infinite covers may resolve the issue.

Another issue is that extensional models typically imply that all equality proofs of a given type are equal. As such our approach is incompatible with univalent theories like Homotopy or Cubical Type Theory [26, 8]. Both of these issues might be addressed by replacing sheaves with stacks or, even more generally, ∞ -stacks [16]. These replace the strict equality of the sheaf condition with higher structure. However, the technical and theoretical overhead of switching to stacks is considerable, and utilizing them for pattern matching will be a significant undertaking.

Toposes and Quasi-toposes Apart from pattern matching, the theory of sheaves plays a central role in categorical logic, since categories of sheaves over a coverage form *Grothendieck toposes*, which serve as models of constructive logic. Quasi-toposes relax the sheaf conditions to only require uniqueness of amalgamations. Future work should search for deeper connections to toposes or quasi-toposes. Two-level type theories [3] may yield some answers, since they describe the interactions between a model of a type theory and the category of presheaves over that model.

7.3 Conclusion

This work formalizes the connection between dependent pattern matching and the notion of sheaves over a site. We have provided a framework which is expressive enough to capture the semantics of current pattern matching implementations, while laying the groundwork for future enhancements. Our work demonstrates that elaboration to eliminators is not the only feasible semantics for dependent pattern matching, and that there is perspective to be gained from treating pattern matching as a core feature and using the lens of sheaf theory.

Bibliography

- [1] Abbott, M., Altenkirch, T., Ghani, N.: Containers: Constructing strictly positive types. *Theor. Comput. Sci.* **342**(1), 3–27 (2005), ISSN 0304-3975, <https://doi.org/10.1016/j.tcs.2005.06.002>
- [2] Altenkirch, T., Ghani, N., Hancock, P., McBride, C., Morris, P.: Indexed containers. *J. Funct. Program.* **25**, e5 (Jan 2015), ISSN 0956-7968, 1469-7653, <https://doi.org/10.1017/S095679681500009X>
- [3] Annenkov, D., Capriotti, P., Kraus, N., Sattler, C.: Two-level type theory and applications. *Math. Struct. Comput. Sci.* **33**(8), 688–743 (Sep 2023), ISSN 0960-1295, 1469-8072, <https://doi.org/10.1017/S0960129523000130>
- [4] Bertot, Y., Castéran, P.: *Interactive Theorem Proving and Program Development*. Springer-Verlag (2004)
- [5] Brady, E.: Idris 2: Quantitative Type Theory in Practice. In: 35th Eur. Conf. Object-Oriented Program. ECOOP 2021, Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2021), <https://doi.org/10.4230/LIPIcs.ECOOP.2021.9>
- [6] Cockx, J., Devriese, D., Piessens, F.: Pattern matching without K. In: Proc. 19th ACM SIGPLAN Int. Conf. Funct. Program., pp. 257–268, ICFP '14, ACM, New York, NY, USA (2014), ISBN 978-1-4503-2873-9, <https://doi.org/10.1145/2628136.2628139>
- [7] Cockx, J., Piessens, F., Devriese, D.: Overlapping and Order-Independent Patterns. In: Shao, Z. (ed.) *Program. Lang. Syst.*, pp. 87–106, Springer, Berlin, Heidelberg (2014), ISBN 978-3-642-54833-8, https://doi.org/10.1007/978-3-642-54833-8_6
- [8] Cohen, C., Coquand, T., Huber, S., Mörtberg, A.: Cubical type theory: A constructive interpretation of the univalence axiom. In: Uustalu, T. (ed.) 21st Int. Conf. Types Proofs Programs TYPES 2015, Leibniz International Proceedings in Informatics (LIPIcs), vol. 69, pp. 5:1–5:34, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2018), ISBN 978-3-95977-030-9, ISSN 1868-8969, <https://doi.org/10.4230/LIPIcs.TYPES.2015.5>
- [9] Coquand, T.: Pattern matching with dependent types. In: *Informal Proc. Log. Framew.*, vol. 92, pp. 66–79 (1992)
- [10] Documentation for Agda: Data.Vec.Base. <https://github.com/agda/agda-stdlib/blob/196766082e913de0d7cd98e3b672935a3b4528b8/src/Data/Vec/Base.agda> (2024)
- [11] Eremondi, J.: Joeyeremondi/lean-cwf: Lean proof for esop 2025 "coverage semantics for dependent pattern matching" (Jan 2025), <https://doi.org/10.5281/zenodo.14768609>, URL <https://doi.org/10.5281/zenodo.14768609>
- [12] Goguen, H., McBride, C., McKinna, J.: Eliminating dependent pattern matching. In: Futatsugi, K., Jouannaud, J.P., Meseguer, J. (eds.) *Algebra, Meaning, and Computation: Essays Dedicated to Joseph A. Goguen*

- on the Occasion of His 65th Birthday, pp. 521–540, Springer Berlin Heidelberg, Berlin, Heidelberg (2006), ISBN 978-3-540-35464-2, https://doi.org/10.1007/11780274_27
- [13] Goguen, J.A.: What is unification?: A categorical view of substitution, equation and solution. In: Algebraic Techniques, pp. 217–261, Elsevier (1989)
 - [14] Hofmann, M.: Syntax and Semantics of Dependent Types. In: Pitts, A.M., Dybjer, P. (eds.) Semantics and Logics of Computation, pp. 79–130, Publications of the Newton Institute, Cambridge University Press, Cambridge (1997), ISBN 978-0-521-58057-1, <https://doi.org/10.1017/CBO9780511526619.004>
 - [15] Johnstone, P.T.: Sketches of an Elephant: A Topos Theory Compendium. Oxford Logic Guides, Oxford University Press, Oxford, New York (Jul 2003), ISBN 978-0-19-852496-0
 - [16] Lurie, J.: Higher Topos Theory. No. no. 170 in Annals of Mathematics Studies, Princeton University Press, Princeton, N.J (2009), ISBN 978-0-691-14048-3 978-0-691-14049-0
 - [17] MacLane, S., Moerdijk, I.: Sheaves in Geometry and Logic: A First Introduction to Topos Theory. Springer, New York, NY, UNITED STATES (1992), ISBN 978-1-4612-0927-0
 - [18] McBride, C.: Dependently Typed Functional Programs and Their Proofs. Ph.D. thesis, University of Edinburgh, UK (2000)
 - [19] McBride, C.: Elimination with a Motive. In: Callaghan, P., Luo, Z., McKinna, J., Pollack, R., Pollack, R. (eds.) Types Proofs Programs, pp. 197–216, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg (2002), ISBN 978-3-540-45842-5, https://doi.org/10.1007/3-540-45842-5_13
 - [20] McBride, C.: Epigram: Practical Programming with Dependent Types. In: Vene, V., Uustalu, T. (eds.) Adv. Funct. Program., pp. 130–170, Springer, Berlin, Heidelberg (2005), ISBN 978-3-540-31872-9, https://doi.org/10.1007/11546382_3
 - [21] McBride, C., McKinna, J.: The view from the left. J. Funct. Prog. **14**(1), 69–111 (Jan 2004), ISSN 0956-7968, 1469-7653, <https://doi.org/10.1017/S0956796803004829>
 - [22] Norell, U.: Dependently typed programming in Agda. In: Proc. 4th Int. Workshop Types Lang. Des. Implement., pp. 1–2, TLDI '09, ACM, New York, NY, USA (2009), ISBN 978-1-60558-420-1, <https://doi.org/10.1145/1481861.1481862>
 - [23] Rydeheard, D.E., Burstall, R.M.: Computational category theory. Prentice Hall International (UK) Ltd., GBR (1988), ISBN 0131627368
 - [24] Sherman, B.: Making Discrete Decisions Based on Continuous Values. Master of Science, Massachusetts Institute of Technology (2017)
 - [25] Sherman, B., Sciarappa, L., Chlipala, A., Carbin, M.: Computable decision making on the reals and other spaces: Via partiality and nondeterminism. In: Proc. 33rd Annu. ACMIEEE Symp. Log. Comput. Sci., pp. 859–868, ACM, Oxford United Kingdom (Jul 2018), ISBN 978-1-4503-5583-4, <https://doi.org/10.1145/3209108.3209193>

- [26] Univalent Foundations Program, T.: Homotopy Type Theory: Univalent Foundations of Mathematics. <https://homotopytypetheory.org/book>, Institute for Advanced Study (2013)
- [27] Wadler, P.: Views: A way for pattern matching to cohabit with data abstraction. In: Proc. 14th ACM SIGACT-SIGPLAN Symp. Princ. Program. Lang., pp. 307–313, POPL '87, Association for Computing Machinery, New York, NY, USA (Oct 1987), ISBN 978-0-89791-215-0, <https://doi.org/10.1145/41625.41653>