
Monads for natural language semantics

CHUNG-CHIEH SHAN

Harvard University, 33 Oxford St, Cambridge, MA 02138, USA

ken@digitas.harvard.edu

ABSTRACT. Accounts of semantic phenomena often involve extending types of meanings and revising composition rules at the same time. The concept of *monads* allows many such accounts—for intensionality, variable binding, quantification and focus—to be stated uniformly and compositionally.

1 Introduction

The Montague grammar tradition formulates formal semantics for natural languages in terms of the λ -calculus. Each utterance is considered a tree in which each leaf node is a lexical item whose meaning is a (usually typed) value in the λ -calculus. The leaf node meanings then determine meanings for subtrees, through recursive application of one or more *composition rules*. A composition rule specifies the meaning of a tree in terms of the meanings of its sub-trees. One simple composition rule is function application:

$$[[x y]] = [[x]]([[y]]) : \beta \quad \text{where } [[x]] : \alpha \rightarrow \beta \text{ and } [[y]] : \alpha. \quad (24.1)$$

Here α and β are type variables, and we denote function types by \rightarrow .

To handle phenomena such as intensionality, variable binding, quantification and focus, we often introduce new types in which to embed existing aspects of meaning and accommodate additional ones. Having introduced new types, we then need to revise our composition rules to reimplement existing functionality. In this way, we often augment semantic theories by simultaneously extending the types of meanings and stipulating new composition rules. When we augment a grammar, its original lexical meanings and composition rules become invalid and require global renovation (typically described as “generalizing to the worst case” (Partee 1996)). Each time we consider a new aspect of meaning, all lexical meanings and composition rules have to be revised.

Over the past decade, the category-theoretic concept of *monads* has gained popularity in computer science as a tool to structure denotational semantics (Moggi 1990; Moggi 1991) and functional programs (Wadler 1992a;

Wadler 1992b). When used to structure computer programs, monads allow the substance of a computation to be defined separately from the plumbing that supports its execution, increasing modularity. Many accounts of phenomena in natural language semantics can also be phrased in terms of monads, thus clarifying the account and simplifying the presentation.

In this paper, I will present the concept of monads and show how they can be applied to natural language semantics. To illustrate the approach, I will use four monads to state analyses of well-known phenomena *uniformly* and *compositionally*. By “uniformly” I mean that, even though the analyses make use of a variety of monads, they all invoke monad primitives in the same way. By “compositionally” I mean that the analyses define composition rules in the spirit of Montague grammar. After presenting the monadic analyses, I will discuss combining monads to account for interaction between semantic phenomena.

2 Monadic analyses

Intuitively, a monad is a transformation on types equipped with a composition method for transformed values. Formally, a monad is a triple $(\mathbb{M}, \eta, \star)$, where \mathbb{M} is a type constructor (a map from each type α to a corresponding type $\mathbb{M}\alpha$), and η and \star are functions (pronounced “unit” and “bind”)

$$\eta : \alpha \rightarrow \mathbb{M}\alpha, \quad \star : \mathbb{M}\alpha \rightarrow (\alpha \rightarrow \mathbb{M}\beta) \rightarrow \mathbb{M}\beta. \quad (24.2)$$

These two functions are polymorphic in the sense that they must be defined for all types α and β . Roughly speaking, η specifies how ordinary values can be injected into the monad, and \star specifies how computations within the monad compose with each other.¹ Some concrete examples follow.

2.1 The powerset monad; interrogatives

As a first example, consider sets. Corresponding to each type α we have a type $\alpha \rightarrow t$, the type of subsets of α . We define²

$$\mathbb{M}\alpha = \alpha \rightarrow t \quad \forall \alpha, \quad (24.4a)$$

$$\eta(a) = \{a\} : \mathbb{M}\alpha \quad \forall a : \alpha, \quad (24.4b)$$

$$m \star k = \bigcup_{a \in m} k(a) : \mathbb{M}\beta \quad \forall m : \alpha \rightarrow t, k : \alpha \rightarrow \beta \rightarrow t. \quad (24.4c)$$

¹By definition, η and \star must satisfy left identity, right identity, and associativity:

$$\eta(a) \star k = k(a) \quad \forall a : \alpha, k : \alpha \rightarrow \mathbb{M}\beta, \quad (24.3a)$$

$$m \star \eta = m \quad \forall m : \mathbb{M}\alpha, \quad (24.3b)$$

$$(m \star k) \star l = m \star (\lambda v. k(v) \star l) \quad \forall m : \mathbb{M}\alpha, k : \alpha \rightarrow \mathbb{M}\beta, l : \beta \rightarrow \mathbb{M}\gamma. \quad (24.3c)$$

²In this section and the next, we treat types as sets in order to define the powerset and pointed powerset monads. These two monads do not exist in every model of the λ -calculus.

The powerset monad is a crude model of non-determinism. For example, the set of individuals m_1 defined by

$$m_1 = \{\text{John}, \text{Mary}\} : \mathbb{M}e$$

can be thought of as a non-deterministic individual—it is ambiguous between John and Mary. Similarly, the function k_1 defined by

$$\begin{aligned} k_1 &: e \rightarrow \mathbb{M}(e \rightarrow t) \\ k_1(a) &= \{\lambda x. \text{like}(a, x), \lambda x. \text{hate}(a, x)\} : \mathbb{M}(e \rightarrow t) \end{aligned}$$

maps each individual to a non-deterministic property. To apply the function k_1 to the individual m_1 , we compute

$$\begin{aligned} m_1 \star k_1 &= \bigcup_{a \in \{\text{John}, \text{Mary}\}} \{\lambda x. \text{like}(a, x), \lambda x. \text{hate}(a, x)\} \\ &= \{\lambda x. \text{like}(\text{John}, x), \lambda x. \text{hate}(\text{John}, x), \\ &\quad \lambda x. \text{like}(\text{Mary}, x), \lambda x. \text{hate}(\text{Mary}, x)\} : \mathbb{M}(e \rightarrow t). \end{aligned}$$

We see that the non-determinism in both m_1 and k_1 is carried through to produce a 4-way-ambiguous result.

Most words in natural language are not ambiguous in the way m_1 and k_1 are. To upgrade an ordinary (deterministic) value of any type α to the corresponding non-deterministic type $\mathbb{M}\alpha$, we can apply η to the ordinary value, say John:

$$\begin{aligned} \eta(\text{John}) &= \{\text{John}\} : \mathbb{M}e, \\ \{\text{John}\} \star k_1 &= \{\lambda x. \text{like}(\text{John}, x), \lambda x. \text{hate}(\text{John}, x)\} : \mathbb{M}(e \rightarrow t). \end{aligned} \tag{24.5}$$

Similarly, to convert an ordinary function to a non-deterministic function, we can apply η to the output of the ordinary function, say k_2 below:

$$\begin{aligned} k_2 &= \lambda a. \lambda x. \text{like}(a, x) : e \rightarrow e \rightarrow t, \\ \eta \circ k_2 &= \lambda a. \{\lambda x. \text{like}(a, x)\} : e \rightarrow \mathbb{M}(e \rightarrow t), \\ m_1 \star (\eta \circ k_2) &= \{\lambda x. \text{like}(\text{John}, x), \lambda x. \text{like}(\text{Mary}, x)\} : \mathbb{M}(e \rightarrow t). \end{aligned} \tag{24.6}$$

In both (24.5) and (24.6), an ordinary value is made to work with a non-deterministic value by upgrading it to the non-deterministic type.

Consider now the function application rule (24.1). We can regard it as a two-argument function, denoted A and defined by

$$\begin{aligned} A &: (\alpha \rightarrow \beta) \rightarrow \alpha \rightarrow \beta, \\ A(f)(x) &= f(x) : \beta \quad \forall f : \alpha \rightarrow \beta, x : \alpha. \end{aligned} \tag{24.7}$$

We can lift ordinary function application A to non-deterministic function application $A_{\mathbb{M}}$, defined by

$$\begin{aligned} A_{\mathbb{M}} &: \mathbb{M}(\alpha \rightarrow \beta) \rightarrow \mathbb{M}\alpha \rightarrow \mathbb{M}\beta, \\ A_{\mathbb{M}}(f)(x) &= f \star [\lambda a. x \star [\lambda b. \eta(a(b))]] : \mathbb{M}\beta \\ &\quad \forall f : \mathbb{M}(\alpha \rightarrow \beta), x : \mathbb{M}\alpha. \end{aligned} \tag{24.8}$$

Substituting (24.4) into (24.8), we get

$$A_{\mathbb{M}}(f)(x) = \{ a(b) \mid a \in f, b \in x \} \subseteq \beta \quad \forall f \subseteq \alpha \rightarrow \beta, x \subseteq \alpha. \quad (24.9)$$

Just as the definition of A in (24.7) gives rise to the original composition rule (24.1), that is,

$$\llbracket x y \rrbracket = A(\llbracket x \rrbracket)(\llbracket y \rrbracket), \quad (24.10)$$

the definition of $A_{\mathbb{M}}$ in (24.8) gives rise to the revised composition rule

$$\llbracket x y \rrbracket = A_{\mathbb{M}}(\llbracket x \rrbracket)(\llbracket y \rrbracket). \quad (24.11)$$

For the powerset monad, this revised rule is the set-tolerant composition rule in the alternative semantics analysis of interrogatives first proposed by Hamblin (1973). In Hamblin's analysis, the meaning of each interrogative constituent is a set of alternatives available as answers to the question; this corresponds to the definition of \mathbb{M} in (24.4). By contrast, the meaning of each non-interrogative constituent is a singleton set; this corresponds to the definition of η in (24.4).

To support question-taking verbs (such as *know* and *ask*), we (and Hamblin) need a secondary composition rule in which A is lifted with respect to the function f but not the argument x :

$$\begin{aligned} \llbracket x y \rrbracket &= A'_{\mathbb{M}}(\llbracket x \rrbracket)(\llbracket y \rrbracket) \quad \text{where} \\ A'_{\mathbb{M}} &: \mathbb{M}(\mathbb{M}\alpha \rightarrow \beta) \rightarrow \mathbb{M}\alpha \rightarrow \mathbb{M}\beta, \\ A'_{\mathbb{M}}(f)(x) &= f \star [\lambda a. \eta(a(x))] : \mathbb{M}\beta \quad \forall f : \mathbb{M}(\mathbb{M}\alpha \rightarrow \beta), x : \mathbb{M}\alpha. \end{aligned} \quad (24.12)$$

Substituting (24.4) into (24.12), we get

$$A'_{\mathbb{M}}(f)(x) = \{ a(x) \mid a \in f \} \subseteq \beta \quad \forall f \subseteq (\alpha \rightarrow t) \rightarrow \beta, x \subseteq \alpha, \quad (24.13)$$

Note that, for any given pair of types of $\llbracket x \rrbracket$ and $\llbracket y \rrbracket$, at most one of $A_{\mathbb{M}}$ (24.8) and $A'_{\mathbb{M}}$ (24.12) can apply. Thus the primary composition rule (24.11) and the secondary composition rule (24.12) never conflict.

2.2 The pointed powerset monad; focus

A variation on the powerset monad (24.4) is the *pointed powerset monad*; it is implicitly involved in Rooth's (1996) account of focus in natural language. A *pointed set* is a nonempty set with a distinguished member. In other words, a pointed set x is a pair $x = (x_0, x_1)$, such that x_0 is a member of the set x_1 . Define the pointed powerset monad by

$$\mathbb{M}\alpha = \{ (x_0, x_1) \mid x_0 \in x_1 \subseteq \alpha \} \quad \forall \alpha, \quad (24.14a)$$

$$\eta(a) = (a, \{a\}) : \mathbb{M}\alpha \quad \forall a : \alpha, \quad (24.14b)$$

$$m \star k = ([k(m_0)]_0, \bigcup_{a \in m_1} [k(a)]_1) : \mathbb{M}\beta \quad \forall m : \mathbb{M}\alpha, k : \alpha \rightarrow \mathbb{M}\beta. \quad (24.14c)$$

This definition captures the intuition that we want to keep track of both a set of non-deterministic alternatives and a particular alternative in the set. As with the powerset monad, the definition of $m \star k$ carries the non-determinism in both m and k through to the result.

Substituting our new monad definition (24.14) into the previously lifted application formula (24.8) gives

$$\begin{aligned} A_{\mathbb{M}}(f_0, f_1)(x_0, x_1) &= (f_0(x_0), \{a(b) \mid a \in f_1, b \in x_1\}) : \mathbb{M}\beta \\ &\quad \forall (f_0, f_1) : \mathbb{M}(\alpha \rightarrow \beta), (x_0, x_1) : \mathbb{M}\alpha. \end{aligned} \quad (24.15)$$

This formula, in conjunction with the primary composition rule (24.11), is equivalent to Rooth’s recursive definition of focus semantic values.

Crucially, even though the pointed powerset monad extends our meaning types to accommodate focus information, neither our definition of $A_{\mathbb{M}}$ in (24.8) nor our composition rule (24.11) needs to change from before. Moreover, the majority of our lexical meanings have nothing to do with focus and thus need not change either. For example, in the hypothetical lexicon entry $\llbracket \text{John} \rrbracket = \eta(\text{John})$, the upgrade from meaning type e to meaning type $\mathbb{M}e$ occurs automatically due to the redefinition of η .

2.3 The reader monad; intensionality and variable binding

Another monad often seen in computer science is the *reader monad*, also known as the *environment monad*. This monad encodes dependence of values on some given input. To define the reader monad, fix a type ρ —say the type s of possible worlds, or the type g of variable assignments—then let

$$\mathbb{M}\alpha = \rho \rightarrow \alpha \quad \forall \alpha, \quad (24.16a)$$

$$\eta(a) = \lambda w. a : \mathbb{M}\alpha \quad \forall a : \alpha, \quad (24.16b)$$

$$m \star k = \lambda w. k(m(w))(w) : \mathbb{M}\beta \quad \forall m : \mathbb{M}\alpha, k : \alpha \rightarrow \mathbb{M}\beta. \quad (24.16c)$$

Note how the definition of $m \star k$ threads the input w through both m and k to produce the result. To see this threading process in action, let us once again substitute our monad definition (24.16) into the definition of $A_{\mathbb{M}}$ in (24.8):

$$A_{\mathbb{M}}(f)(x) = \lambda w. f(w)(x(w)) : \mathbb{M}\beta \quad \forall f : \mathbb{M}(\alpha \rightarrow \beta), x : \mathbb{M}\alpha. \quad (24.17)$$

For $\rho = s$, we can think of \mathbb{M} as the intensionality monad, noting that (24.17) is exactly the usual extensional composition rule. While words such as *student* and *know* have meanings that depend on the possible world w , words such as *is* and *and* do not. We can upgrade the latter by applying η .

For $\rho = g$, we can think of \mathbb{M} as the variable binding monad, noting that (24.17) is the usual assignment-preserving composition rule. Except for pronominals, most word meanings do not refer to the variable assignment. Thus we can upgrade the majority of word meanings by applying η .

If we substitute the same monad definition (24.16) into the secondary composition rule (24.12), the result is

$$A'_{\mathbb{M}}(f)(x) = \lambda w. f(w)(x) : \mathbb{M}\beta \quad \forall f : \mathbb{M}(\mathbb{M}\alpha \rightarrow \beta), x : \mathbb{M}\alpha. \quad (24.18)$$

For $\rho = s$, this is the intensional composition rule; it handles sentence-taking verbs such as *know* and *believe* (of type $s \rightarrow (s \rightarrow t) \rightarrow e \rightarrow t$) by allowing them to take arguments of type $s \rightarrow t$ rather than type t . The monad laws, by the way, guarantee that $A'_{\mathbb{M}}(f)(x) = A_{\mathbb{M}}(f)(\eta(x))$ for all f and x ; the function η (in this case a map from $s \rightarrow t$ to $s \rightarrow s \rightarrow t$) is simply the intension (up) operator, usually written \wedge .

For $\rho = g$, the same formula (24.18) is often involved in accounts of quantification that assume quantifier raising at LF, such as that in Heim and Kratzer (1998). It handles raised quantifiers (of type $g \rightarrow (g \rightarrow t) \rightarrow t$) by allowing them to take arguments of type $g \rightarrow t$ rather than type t . The function η (in this case a map from $g \rightarrow t$ to $g \rightarrow g \rightarrow t$) is simply the variable abstraction operator.

2.4 The continuation monad; quantification

Barker (2000) proposed an analysis of quantification in terms of *continuations*. The basic idea is to *continuize* a grammar by replacing each meaning type α with its corresponding continuized type $(\alpha \rightarrow t) \rightarrow t$ throughout. As a special case, the meaning type of NPs is changed from e to $(e \rightarrow t) \rightarrow t$, matching the original treatment of English quantification by Montague (1974).

In general, for any fixed type ω (say t), we can define a *continuation monad* with *answer type* ω :

$$\mathbb{M}\alpha = (\alpha \rightarrow \omega) \rightarrow \omega \quad \forall \alpha, \quad (24.19a)$$

$$\eta(a) = \lambda c. c(a) : \mathbb{M}\alpha \quad \forall a : \alpha, \quad (24.19b)$$

$$m \star k = \lambda c. m(\lambda a. k(a)(c)) : \mathbb{M}\beta \quad \forall m : \mathbb{M}\alpha, k : \alpha \rightarrow \mathbb{M}\beta. \quad (24.19c)$$

The value c manipulated in these definitions is known as the *continuation*. Intuitively, “the continuation represents an entire (default) future for the computation” (Kelsey, Clinger, and Rees 1998). Each value of type $\mathbb{M}\alpha$ must turn a continuation (of type $\alpha \rightarrow \omega$) into an answer (of type ω). The most obvious way to do so, encoded in the definition of η above, is to feed the continuation a value of type α :

$$\llbracket \text{John} \rrbracket = \eta(\text{John}) = \lambda c. c(\text{John}) : \mathbb{M}e, \quad (24.20a)$$

$$\llbracket \text{smokes} \rrbracket = \eta(\text{smoke}) = \lambda c. c(\text{smoke}) : \mathbb{M}(e \rightarrow t). \quad (24.20b)$$

To compute the meaning of *John smokes*, we first substitute our monad definition (24.19) into the primary composition operation $A_{\mathbb{M}}$ (24.8):

$$A_{\mathbb{M}}(f)(x) = \lambda c. f(\lambda g. x(\lambda y. c(g(y)))) : \mathbb{M}\beta \quad (24.21)$$

$$\forall f : \mathbb{M}(\alpha \rightarrow \beta), x : \mathbb{M}\alpha.$$

Letting $f = \eta(\text{smoke})$ and $x = \eta(\text{John})$ then gives

$$\begin{aligned} \llbracket \text{John smokes} \rrbracket &= \lambda c. \eta(\text{smoke})(\lambda g. \eta(\text{John})(\lambda y. c(g(y)))) \\ &= \lambda c. \eta(\text{John})(\lambda y. c(\text{smoke}(y))) = \lambda c. c(\text{smoke}(\text{John})) : \mathbb{M}t. \end{aligned}$$

In the second step above, note how the term $\lambda y. c(\text{smoke}(y))$ represents the future for the computation of $\llbracket \text{John} \rrbracket$, namely to check whether he smokes, then pass the result to the context c containing the clause. If *John smokes* is the main clause, then the context c is simply the identity function id_ω . We define an evaluation operator $\varepsilon : \mathbb{M}\omega \rightarrow \omega$ by $\varepsilon(m) = m(\text{id}_\omega)$. Fixing $\omega = t$, we then have $\varepsilon(\llbracket \text{John smokes} \rrbracket) = \text{smoke}(\text{John})$, as desired.

Continuing to fix $\omega = t$, we can specify a meaning for *everyone*:

$$\llbracket \text{everyone} \rrbracket = \lambda c. \forall x. c(x) : \mathbb{M}e. \quad (24.22)$$

This formula is not of the form $\lambda c. c(\dots)$. In other words, the meaning of *everyone* non-trivially manipulates the continuation, and so cannot be obtained from applying η to an ordinary value. Using the continuized composition rule (24.21), we now compute a denotation for *everyone smokes*:

$$\begin{aligned} \llbracket \text{everyone smokes} \rrbracket &= \lambda c. \eta(\text{smoke})(\lambda g. \llbracket \text{everyone} \rrbracket(\lambda y. c(g(y)))) \\ &= \lambda c. \llbracket \text{everyone} \rrbracket(\lambda y. c(\text{smoke}(y))) = \lambda c. \forall x. c(\text{smoke}(x)) : \mathbb{M}t, \end{aligned}$$

giving $\varepsilon(\llbracket \text{everyone smokes} \rrbracket) = \forall x. \text{smoke}(x) : t$, as desired.

The main theoretical advantage of this analysis is that it is a compositional, in-situ analysis that does not invoke quantifier raising. Moreover, note that a grammar continuized is still a grammar—the continuized composition rule (24.21) is perfectly interpretable using the standard machinery of Montague grammar. In particular, we do not invoke any type ambiguity or flexibility as proposed by Partee and Rooth (1983) and Hendriks (1993); the interpretation mechanism performs no type-shifting at “run-time”.

This desirable property also holds of the other monadic analyses I have presented. For instance, in a grammar with intensionality, meanings that use intensionality (for example $\llbracket \text{student} \rrbracket$) are identical in type to meanings that do not (for example $\llbracket \text{is} \rrbracket$). The interpretation mechanism does not dynamically shift the type of *is* to match that of *student*.

It is worth relating the present analysis to the computer science literature. Danvy and Filinski (1990) studied *composable continuations*, which they manipulated using two operators “shift” and “reset”. We can define shift and reset for the continuation monad by (Wadler 1994)

$$\text{shift} = \lambda h. \lambda c. \varepsilon(h(\lambda a. \lambda c'. c'(c(a)))) : ((\alpha \rightarrow \mathbb{M}\omega) \rightarrow \mathbb{M}\omega) \rightarrow \mathbb{M}\alpha, \quad (24.23)$$

$$\text{reset} = \lambda m. \lambda c. c(\varepsilon(m)) : \mathbb{M}\omega \rightarrow \mathbb{M}\omega. \quad (24.24)$$

Assuming that the \forall operator is of type $(e \rightarrow t) \rightarrow t$, the meaning of *everyone* specified in (24.22) is simply $\text{shift}(\lambda c. (\eta \circ \forall)(\varepsilon \circ c))$. To encode scope islands, Barker implicitly used reset. Filinski (1999) proved that, in a certain sense, composable continuations can simulate monads.

3 Combining monads

Having placed various semantic phenomena in a monadic framework, we now ask a natural question: Can we somehow combine monads in a modular fashion to characterize interaction between semantic phenomena, for example between intensionality and quantification?

Unfortunately, there exists no general construction for composing two arbitrary monads, say $(\mathbb{M}_1, \eta_1, \star_1)$ and $(\mathbb{M}_2, \eta_2, \star_2)$, into a new monad of the form $(\mathbb{M}_3 = \mathbb{M}_1 \circ \mathbb{M}_2, \eta_3, \star_3)$ (King and Wadler 1993; Jones and Duponcheel 1993). One might still hope to specialize and combine monads with additional structure, to generalize and combine monads as instances of a broader concept, or even to find that obstacles in combining monads are reflected in semantic constraints in natural language.

Researchers in denotational semantics of programming languages have made several proposals towards combining monadic functionality, none of which are completely satisfactory (Moggi 1990; Steele 1994; Liang, Hudak, and Jones 1995; Espinosa 1995; Filinski 1999). In this section, I will relate one prominent approach to natural language semantics.

3.1 Monad morphisms

One approach to combining monads, taken by Moggi, Liang et al., and Filinski, is to compose *monad morphisms* instead of monads themselves. A monad morphism (also known as a *monad transformer* or a *monad layering*) is a map from monads to monads; it takes an arbitrary monad and transforms it into a new monad, presumably defined in terms of the old monad and supporting a new layer of functionality. For instance, given any monad $(\mathbb{M}_1, \eta_1, \star_1)$ and fixing a type ρ , the *reader monad morphism* constructs a new monad $(\mathbb{M}_2, \eta_2, \star_2)$, defined by

$$\mathbb{M}_2\alpha = \rho \rightarrow \mathbb{M}_1\alpha \quad \forall\alpha, \quad (24.25a)$$

$$\eta_2(a) = \lambda w. \eta_1(a) : \mathbb{M}_2\alpha \quad \forall a : \alpha, \quad (24.25b)$$

$$m \star_2 k = \lambda w. [m(w) \star_1 \lambda a. k(a)(w)] : \mathbb{M}_2\beta \quad \forall m : \mathbb{M}_2\alpha, \quad (24.25c)$$

$$k : \alpha \rightarrow \mathbb{M}_2\beta.$$

If we let the old monad $(\mathbb{M}_1, \eta_1, \star_1)$ be the *identity monad*, defined by $\mathbb{M}_1\alpha = \alpha$, $\eta_1(a) = a$, and $m \star_1 k = k(m)$, then the new monad (24.25) is just the reader monad (24.16). If we let the old monad be some other monad—even the reader monad itself—the new monad adds reader functionality.

By definition, each monad morphism must specify how to embed computations inside the old monad into the new monad. More precisely, each monad morphism must provide a function (pronounced “lift”)

$$\ell : \mathbb{M}_1\alpha \rightarrow \mathbb{M}_2\alpha, \quad (24.26)$$

polymorphic in α .³ For the reader monad morphism, ℓ is defined by

$$\ell(m) = \lambda w. m : \mathbb{M}_2\alpha \quad \forall m : \mathbb{M}_1\alpha. \quad (24.28)$$

The continuation monad also generalizes to a monad morphism. Fixing an answer type ω , the *continuation monad morphism* takes any monad $(\mathbb{M}_1, \eta_1, \star_1)$ to the monad $(\mathbb{M}_2, \eta_2, \star_2)$ defined by

$$\mathbb{M}_2\alpha = (\alpha \rightarrow \mathbb{M}_1\omega) \rightarrow \mathbb{M}_1\omega \quad \forall \alpha, \quad (24.29a)$$

$$\eta_2(a) = \lambda c. c(a) : \mathbb{M}_2\alpha \quad \forall a : \alpha, \quad (24.29b)$$

$$m \star_2 k = \lambda c. m(\lambda a. k(a)(c)) : \mathbb{M}_2\beta \quad \forall m : \mathbb{M}_2\alpha, k : \alpha \rightarrow \mathbb{M}_2\beta. \quad (24.29c)$$

The lifting function ℓ for the continuation monad morphism is defined by

$$\ell(m) = \lambda c. (m \star_1 c) : \mathbb{M}_2\alpha \quad \forall m : \mathbb{M}_1\alpha. \quad (24.30)$$

Monad morphisms can be freely composed with each other, though the order of composition is significant. Applying the continuation monad morphism to the reader monad is equivalent to applying to the identity monad the composition of the continuation monad morphism and the reader monad morphism, and yields a monad with type constructor $\mathbb{M}\alpha = (\alpha \rightarrow \rho \rightarrow \omega) \rightarrow \rho \rightarrow \omega$. Applying the reader monad morphism to the continuation monad is equivalent to applying to the identity monad the composition of the reader monad morphism and the continuation monad morphism, and yields a different monad, with type constructor $\mathbb{M}\alpha = \rho \rightarrow (\alpha \rightarrow \omega) \rightarrow \omega$.

3.2 Translating monads to monad morphisms

The monad morphisms (24.25) and (24.29) may appear mysterious, but we can in fact obtain them from their monad counterparts (24.16) and (24.19) via a mechanical translation. The translation takes a monad $(\mathbb{M}_0, \eta_0, \star_0)$ whose η_0 and \star_0 operations are λ -terms, and produces a morphism mapping any old monad $(\mathbb{M}_1, \eta_1, \star_1)$ to a new monad $(\mathbb{M}_2, \eta_2, \star_2)$. The translation is defined recursively on the structure of λ -types and λ -terms, as follows.⁴

Every type τ is either a function type or a base type. A function type has the form $\tau_1 \rightarrow \tau_2$, where τ_1 and τ_2 are types. A base type ι is a type fixed in \mathbb{M}_0 (ρ and ω in our cases), a polymorphic type variable (α and β

³By definition, ℓ must satisfy naturality:

$$\ell(\eta_1(a)) = \eta_2(a) \quad \forall a : \alpha, \quad (24.27a)$$

$$\ell(m \star_1 k) = \ell(m) \star_2 (\ell \circ k) \quad \forall m : \mathbb{M}_1\alpha, k : \alpha \rightarrow \mathbb{M}_1\beta. \quad (24.27b)$$

⁴Among other things, the translation requires \mathbb{M}_0 to be defined as a λ -type, and η_0 and \star_0 to be defined as λ -terms. Thus the translation cannot apply to the powerset and pointed powerset monads (see footnote 2). Nevertheless, any monad morphism (including ones produced by the translation) can be applied to any monad (including these two monads).

as appearing in η and \star (24.2) and ℓ (24.26)), or the terminal type $!$ (also known as the unit type or the void type). For every type τ , we recursively define its *computation translation* $\lceil \tau \rceil$ and its *value translation* $\lfloor \tau \rfloor$:

$$\lceil \iota \rceil = \mathbb{M}_1 \iota, \quad \lfloor \iota \rfloor = \iota, \quad \lceil \tau_1 \rightarrow \tau_2 \rceil = \lfloor \tau_1 \rightarrow \tau_2 \rfloor = \lfloor \tau_1 \rfloor \rightarrow \lceil \tau_2 \rceil, \quad (24.31)$$

where ι is any base type.

Each term $e : \tau$ is an application term, an abstraction term, a variable term, or the terminal term. An application term has the form $(e_1 : \tau_1 \rightarrow \tau_2) (e_2 : \tau_1) : \tau_2$, where e_1 and e_2 are terms. An abstraction term has the form $(\lambda x : \tau_1. e : \tau_2) : \tau_1 \rightarrow \tau_2$, where e is a term. A variable term has the form $x : \tau$, where x is the name of a variable of type τ . The terminal term is $! : !$ and represents the unique value of the terminal type $!$. For every term $e : \tau$, we recursively define its *term translation* $\lceil e \rceil : \lceil \tau \rceil$:

$$\lceil (e_1 : \iota \rightarrow \tau_1 \rightarrow \dots \rightarrow \tau_n \rightarrow \iota') (e_2 : \iota) \rceil = \lambda y_1 : \lceil \tau_1 \rceil. \dots \lambda y_n : \lceil \tau_n \rceil. \lceil [e_2] \star_1 (\lambda y_0 : \iota. [e_1] (y_0) \dots (y_n)) \rceil, \quad (24.32a)$$

$$\lceil (e_1 : (\tau_1 \rightarrow \tau_2) \rightarrow \tau_3) (e_2 : \tau_1 \rightarrow \tau_2) \rceil = \lceil e_1 \rceil (\lceil e_2 \rceil), \quad (24.32b)$$

$$\lceil \lambda x : \tau. e \rceil = \lambda x : \lceil \tau \rceil. \lceil e \rceil, \quad (24.32c)$$

$$\lceil x : \iota \rceil = \eta_1(x), \quad (24.32d)$$

$$\lceil x : \tau_1 \rightarrow \tau_2 \rceil = x, \quad (24.32e)$$

$$\lceil ! \rceil = \eta_1(!), \quad (24.32f)$$

where ι and ι' are any base types, and y_0, \dots, y_n are fresh variable names.

Finally, to construct the new monad $(\mathbb{M}_2, \eta_2, \star_2)$, we specify

$$\begin{aligned} \mathbb{M}_2 \alpha &= \lceil \mathbb{M}_0 \alpha \rceil, & \eta_2 &= \lceil \eta_0 \rceil, & \star_2 &= \lceil \star_0 \rceil, \\ \ell(m) &= \lceil \lambda f : ! \rightarrow \alpha. \eta_0(f(!)) \rceil (\lambda ! : !. m) \quad \forall m : \mathbb{M}_1 \alpha. \end{aligned} \quad (24.33)$$

To illustrate this translation, let us expand out \mathbb{M}_2 and \star_2 in the special case where $(\mathbb{M}_0, \eta_0, \star_0)$ is the reader monad (24.16). From the type translation rules (24.31) and the specification of \mathbb{M}_2 in (24.33), we have

$$\mathbb{M}_2 \alpha = \lceil \rho \rightarrow \alpha \rceil = \lfloor \rho \rfloor \rightarrow \lceil \alpha \rceil = \rho \rightarrow \mathbb{M}_1 \alpha,$$

matching (24.25a) as desired. From the term translation rules (24.32) and the specification of \star_2 in (24.33), we have

$$\begin{aligned} \star_2 &= \lceil \lambda m : \rho \rightarrow \alpha. \lambda k : \alpha \rightarrow \rho \rightarrow \beta. \lambda w : \rho. k(m(w))(w) \rceil \quad \text{by (24.16c)} \\ &= \lambda m : \rho \rightarrow \mathbb{M}_1 \alpha. \lambda k : \alpha \rightarrow \rho \rightarrow \mathbb{M}_1 \beta. \lambda w : \rho. \lceil k(m(w))(w) \rceil \quad \text{by (24.32c),} \end{aligned}$$

in which

$$\begin{aligned} \lceil k(m(w))(w) \rceil &= \eta_1(w) \star_1 \lambda y_0 : \rho. \lceil k(m(w)) \rceil (y_0) \quad \text{by (24.32a), (24.32d)} \\ &= \lceil k(m(w)) \rceil (w) \quad \text{by (24.3a)} \\ &= (\lceil w \rceil \star_1 \lambda y_0 : \rho. \lceil m \rceil (y_0)) \star_1 \lambda y_0 : \alpha. \lceil k \rceil (y_0)(w) \quad \text{by (24.32a)} \\ &= (\eta_1(w) \star_1 \lambda y_0 : \rho. m(y_0)) \star_1 \lambda y_0 : \alpha. k(y_0)(w) \quad \text{by (24.32d), (24.32e)} \\ &= m(w) \star_1 \lambda y_0 : \alpha. k(y_0)(w) \quad \text{by (24.3a),} \end{aligned}$$

matching (24.25c) as desired.

The intuition behind our translation is to treat the λ -calculus with which $(\mathbb{M}_0, \eta_0, \star_0)$ is defined as a programming language whose terms may have computational side effects. Our translation specifies a semantics for this programming language in terms of $(\mathbb{M}_1, \eta_1, \star_1)$ that is call-by-value and that allows side effects only at base types. That the semantics is call-by-value rather than call-by-name is reflected in the type translation rules (24.31), where we define $[\tau_1 \rightarrow \tau_2]$ to be $[\tau_1] \rightarrow [\tau_2]$ rather than $\lceil \tau_1 \rceil \rightarrow \lceil \tau_2 \rceil$. That side effects occur only at base types is also reflected in the rules, where we define $\lceil \tau_1 \rightarrow \tau_2 \rceil$ to be $[\tau_1 \rightarrow \tau_2]$ rather than $\mathbb{M}_1 [\tau_1 \rightarrow \tau_2]$. Overall, our translation is a hybrid between the call-by-name Algol translation (Benton, Hughes, and Moggi 2000, §3.1.2) and the standard call-by-value translation (Wadler 1992a, §8; Benton et al., §3.1.3).

3.3 A call-by-name translation of monads

Curiously, the semantic types generated by monad morphisms seem sometimes not powerful enough. As noted at the end of §3.1, the reader and continuation monad morphisms together give rise to two different monads, depending on the order in which we compose the monad morphisms. Fixing $\rho = s$ for the reader monad morphism and $\omega = t$ for the continuation monad morphism, the two combined monads have the type constructors

$$\mathbb{M}_{cr}\alpha = (\alpha \rightarrow s \rightarrow t) \rightarrow s \rightarrow t, \quad \mathbb{M}_{rc}\alpha = s \rightarrow (\alpha \rightarrow t) \rightarrow t. \quad (24.34)$$

Consider now sentences such as

$$\text{John wanted to date every professor (at the party).} \quad (24.35)$$

This sentence has a reading where every professor at the party is a person John wanted to date, but John may not be aware that they are professors. On this reading, note that the world where the property of professorship is evaluated is distinct from the world where the property of dating is evaluated. Therefore, assuming that *to date every professor* is a constituent, its semantic type should mention s in contravariant position at least twice. Unfortunately, the type constructors \mathbb{M}_{cr} and \mathbb{M}_{rc} (24.34) each have only one such occurrence, so neither $\mathbb{M}_{cr}t$ nor $\mathbb{M}_{rc}t$ can be the correct type.

Intuitively, the semantic type of *to date every professor* ought to be

$$((s \rightarrow t) \rightarrow s \rightarrow t) \rightarrow s \rightarrow t \quad (24.36)$$

or an even larger type. The type (24.36) is precisely equal to $\mathbb{M}_{cr}(s \rightarrow t)$, but simply assigning $\mathbb{M}_{cr}(s \rightarrow t)$ as the semantic type of *to date every professor* would be against our preference for the reader monad morphism to be the only component of the semantic system that knows about the type s . Instead, what we would like is to equip the transformation on types taking each α to $((s \rightarrow \alpha) \rightarrow s \rightarrow t) \rightarrow s \rightarrow t$ with a composition method for transformed values.

One tentative idea for synthesizing such a composition method is to replace the call-by-value translation described in §3.2 with a call-by-name translation, such as the Algol translation mentioned earlier (Benton et al., §3.1.2).⁵ For every type τ , this translation recursively defines a type $\llbracket \tau \rrbracket$:

$$\llbracket \iota \rrbracket = \mathbb{M}_1 \iota, \quad \llbracket \tau_1 \rightarrow \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \rightarrow \llbracket \tau_2 \rrbracket, \quad (24.37)$$

where ι is any base type. For every term $e : \tau$, this translation recursively defines a term $\llbracket e \rrbracket : \llbracket \tau \rrbracket$:

$$\begin{aligned} \llbracket e_1(e_2) \rrbracket &= \llbracket e_1 \rrbracket(\llbracket e_2 \rrbracket), & \llbracket x \rrbracket &= x, \\ \llbracket \lambda x : \tau. e \rrbracket &= \lambda x : \llbracket \tau \rrbracket. \llbracket e \rrbracket, & \llbracket ! \rrbracket &= \eta_1(!). \end{aligned} \quad (24.38)$$

Applying this translation to a monad $(\mathbb{M}_0, \eta_0, \star_0)$ gives the types

$$\llbracket \eta_0 \rrbracket : \mathbb{M}_1 \alpha \rightarrow \llbracket \mathbb{M}_0 \alpha \rrbracket, \quad (24.39a)$$

$$\llbracket \star_0 \rrbracket : \llbracket \mathbb{M}_0 \alpha \rrbracket \rightarrow (\mathbb{M}_1 \alpha \rightarrow \llbracket \mathbb{M}_0 \beta \rrbracket) \rightarrow \llbracket \mathbb{M}_0 \beta \rrbracket. \quad (24.39b)$$

If we let $(\mathbb{M}_0, \eta_0, \star_0)$ be the continuation monad (24.19) and $(\mathbb{M}_1, \eta_1, \star_1)$ be the reader monad (24.16), then the type t transformed is

$$\llbracket \mathbb{M}_0 t \rrbracket = (\mathbb{M}_1 t \rightarrow \mathbb{M}_1 t) \rightarrow \mathbb{M}_1 t = ((s \rightarrow t) \rightarrow s \rightarrow t) \rightarrow s \rightarrow t,$$

as desired. However, unless $(\mathbb{M}_1, \eta_1, \star_1)$ is the identity monad, the types in (24.39) do not match the definition of monads in (24.2). In other words, though our call-by-name translation does give the type transform we want as well as a composition method in some sense, its output is not a monad morphism.

4 Conclusion

In this paper, I used monads to characterize the similarity between several semantic accounts—for interrogatives, focus, intensionality, variable binding, and quantification.⁶ In each case, the same monadic composition rules and mostly the same lexicon were specialized to a different monad. The monad primitives η and \star recur in semantics with striking frequency.

It remains to be seen whether monads would provide the appropriate conceptual encapsulation for a semantic theory with broader coverage. In particular, for both natural and programming language semantics, combining monads—or perhaps monad-like objects—remains an open issue that promises additional insight.

Acknowledgments Thanks to Stuart Shieber, Dylan Thurston, Chris Barker, and the anonymous referees for helpful discussions and comments.

⁵Another possible translation is the standard (“Haskell”) call-by-name one (Wadler 1992a, §8; Benton et al., §3.1.1). It produces strictly larger types than the Algol call-by-name translation, for instance $s \rightarrow (s \rightarrow (s \rightarrow \alpha) \rightarrow s \rightarrow t) \rightarrow s \rightarrow t$.

⁶Other phenomena that may fall under the monadic umbrella include presuppositions (the error monad) and dynamic semantics (the state monad).

Bibliography

- Barker, C. (2000). Continuations and the nature of quantification. Manuscript, University of California, San Diego, 4 November 2000, <http://semanticsarchive.net/Archive/902ad5f7/>.
- Benton, N., J. Hughes, and E. Moggi (2000). Monads and effects. Lecture notes, International Summer School on Applied Semantics, 5 September 2000, <http://www.disi.unige.it/person/MoggiE/APPSEM00/>.
- Davy, O. and A. Filinski (1990). Abstracting control. In *Proceedings of the 1990 ACM Conference on Lisp and Functional Programming*, Nice, France, pp. 151–160. New York: ACM Press.
- Espinosa, D. A. (1995). *Semantic Lego*. Ph. D. thesis, Graduate School of Arts and Sciences, Columbia University.
- Filinski, A. (1999). Representing layered monads. In *POPL '99: Conference Record of the Annual ACM Symposium on Principles of Programming Languages*, San Antonio, TX, pp. 175–188. New York: ACM Press.
- Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language* 10, 41–53.
- Heim, I. and A. Kratzer (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.
- Hendriks, H. (1993). *Studied Flexibility: Categories and Types in Syntax and Semantics*. Ph. D. thesis, Institute for Logic, Language and Computation, Universiteit van Amsterdam.
- Jones, M. P. and L. Duponcheel (1993). Composing monads. Technical Report YALEU/DCS/RR-1004, Department of Computer Science, Yale University, New Haven.
- Kelsey, R., W. Clinger, and J. Rees (Eds.) (1998). Revised⁵ report on the algorithmic language Scheme. *Higher-Order and Symbolic Computation* 11(1), 7–105. Also in *ACM SIGPLAN Notices* 33(9), 26–76.
- King, D. J. and P. Wadler (1993). Combining monads. In J. Launchbury and P. M. Sansom (Eds.), *Functional Programming, Glasgow 1992: Proceedings of the 1992 Glasgow Workshop on Functional Programming*, Ayr, Scotland. Berlin: Springer-Verlag.

- Lappin, S. (Ed.) (1996). *The Handbook of Contemporary Semantic Theory*. Oxford: Blackwell.
- Liang, S., P. Hudak, and M. Jones (1995). Monad transformers and modular interpreters. In *POPL '95: Conference Record of the Annual ACM Symposium on Principles of Programming Languages*, San Francisco, CA, pp. 333–343. New York: ACM Press.
- Moggi, E. (1990). An abstract view of programming languages. Technical Report ECS-LFCS-90-113, Laboratory for Foundations of Computer Science, Department of Computer Science, University of Edinburgh, Edinburgh.
- Moggi, E. (1991). Notions of computation and monads. *Information and Computation* 93(1), 55–92.
- Montague, R. (1974). The proper treatment of quantification in ordinary English. In R. Thomason (Ed.), *Formal Philosophy: Selected Papers of Richard Montague*, pp. 247–270. New Haven: Yale University Press.
- Partee, B. (1996). The development of formal semantics. See Lappin (1996), pp. 11–38.
- Partee, B. and M. Rooth (1983). Generalized conjunction and type ambiguity. In R. Bauerle, C. Schwartze, and A. von Stechow (Eds.), *Meaning, Use and Interpretation of Language*, pp. 361–383. Berlin: De Gruyter.
- Rooth, M. (1996). Focus. See Lappin (1996), pp. 271–297.
- Steele, Jr., G. L. (1994). Building interpreters by composing monads. In *POPL '94: Conference Record of the Annual ACM Symposium on Principles of Programming Languages*, Portland, OR, pp. 472–492. New York: ACM Press.
- Wadler, P. (1992a). Comprehending monads. *Mathematical Structures in Computer Science* 2(4), 461–493.
- Wadler, P. (1992b). The essence of functional programming. In *POPL '92: Conference Record of the Annual ACM Symposium on Principles of Programming Languages*, Albuquerque, NM, pp. 1–14. New York: ACM Press.
- Wadler, P. (1994). Monads and composable continuations. *Lisp and Symbolic Computation* 7(1), 39–56.