MEANS-END RELATIONS AND ARTIFACTUAL FUNCTIONS: A SKETCH

JESSE HUGHES

Abstract. Artifactual functions are closely related to means-end relations. Whatever else "f is a function of o" might entail, it surely means that there is a means-end relation involving o in order to attain the goal involved in f. We take this observation as our starting point in creating a formal semantics for artifactual functions. Thus, we wish to first settle on formal semantics for means-end relations and then build on these semantics to formalize functions.

For the first part, we use a simplified Propositional Dynamic Logic (PDL) to give formal definitions of various kinds of means-end relations and try to sketch some practical consequences from such relations.

For the second part, we augment PDL with a set of objects (artifacts) and a simple type theory and give a formal representation of artifactual functions. Unlike the work in means-end relations, we do not define formally what an artifactual function is – functional ascriptions arise via through a complex combination of social and causal relations and the social relations cannot formalized away. Rather, we give a formal structure for what such ascriptions involve and initial hints at expected consequences of functional ascriptions.

"Buttons are to keep people warm."
Ruth Krauss, *A Hole is to Dig*

1. Introduction

Artifacts are function-bearing things: they have uses, whether by design, custom or accident, and these uses bring about related goals. In these terms, functions entail certain means-end relations: an artifact's use is a means to its functional goal. There is certainly more to be said about function than this, but this is a good starting point.

If we are to rely on the relation between functions and means-end relations in order to give a formal account of functions, we should start with a formal semantics for means-end relations. Indeed, much of our work to date has been on just this point: what is an appropriate formal semantics for means-end relations, especially as they pertain to artifactual function? The first section of the paper is a brief presentation of our answer to that question (see also [2, 3, 4] for more work on means-end semantics).

Roughly, we view a means-end relation as an assertion of causal sufficiency for a certain action (the *means*) to realize a certain condition (the *end*). There is a standard logical framework for reasoning about actions and the conditions they realize, namely Propositional Dynamic Logic (PDL) [1]. We take that as our foundation for a means-end semantics.

The bulk of our presentation is on artifactual functions, however. The work on functions is fundamentally different than our approach on means-end semantics. We take it that a means-end relation expresses a causal relation, but whether an artifact possesses a function or not depends on more than just causal relations: functions are *ascribed* to artifacts by someone. They are unlike means-end relations in that functional ascriptions include an irreducibly social component. This difference is emphasized by our terminology: means-end *relations* vs. functional *ascriptions.*

Thus, formal semantics cannot characterize functional ascriptions[1]. Instead, we intend to give a specification of functional ascription: a definition of the components that are required from any such ascription, whether implicitly or explicitly. In particular, we expect that a particular token may or

[1]Perhaps one could characterize ascriptions by adding agents capable of ascribing. This does not appear a very enlightening approach, however, and we do not intend to pursue it.

may not *fulfill* a given functional ascription. We want to characterize this notion of fulfillment and so we give a specification of ascriptions that includes the relevant features to define fulfillment.

Roughly, an artifact *fulfills* its function if the artifact fits into a causal structure so that the functional use is a means to the functional goal. Among the components needed to flesh out this intuition are: artifactual type, functional use and functional goal. In addition, other parameters relevant to the functional use and goal give appropriate *context* for the evaluation.

The reader may note that we have implicitly focused on a particular kind of functional ascriptions, namely those that involve active use. This is a bias in keeping with our interpretation of means-end relations in terms of dynamic logic. In both cases, we emphasize actions and their capacity to realize an end, but many artifacts have functions which require no apparent active use: a window may have the function of allowing sunlight into the room, but it is not obvious that this function has any associated use[2]. We accept that our presentation of artifactual function does not explain every functional ascription, but we believe that we have clarified some of the meaning of "active" functions, i.e. those functions involving active use. We postpone discussion of a broader class of function.

Our formal analysis may also obscure the real difficulties involved in analyzing functions. As we see it, fulfillment involves reasoning about what would be the case if one were to use a token in a particular context. Thus, our semantics involves counterfactual reasoning about causal relations. Such reasoning is notoriously difficult and problematic and we have sidestepped these difficulties by taking the causal relations as given. But our primary aim is a conceptual analysis of functional ascriptions. If our understanding of such ascriptions is correct, then they involve these problematic counterfactuals. Certainly, understanding how to evaluate the counterfactuals is an important question—where do our dynamic models come from? But we will assume at present that we can make sense of the counterfactuals in order to build dynamic models, so that our conceptual analysis can get off the ground. The difficulty is, from our perspective, already "in" the natural language functional ascriptions and so we expect it to be a feature of our semantics as well.

In summary: we present a brief introduction to our formal semantics for means-end relations and a longer application of these semantics to artifactual functions. For artifactual functions, we give a specification of some of the components in a functional ascription and a formal semantics for functional fulfillment (for both tokens and types). We give no account of the *origin* of such ascriptions, but rather analyze them as given. Our semantics for fulfillment involves reasoning about the causal structure that would result if a token was used in a given context, and thus is defined in terms of a counterfactual. For this, we assume that the causal structures are given and we ignore the deep and difficult problem of determining that structure.

2. Means-end relations

An end is a condition which some agent may desire. We take this in the broadest sense, so that any condition may be an end. Thus, it is reasonable to consider an end to be a formula in a formal language.

A means is a way to realize an end. Therefore, a means must be something one can do in order to change the world so that an end φ (which may currently be false) will become true. This suggests that means correspond to transitions between possible worlds. *Propositional Dynamic Logic* (PDL) is an appropriate language for modeling transitions between worlds via an agent's actions[3]. See [1] for an introduction to PDL. We will only sketch the semantics here.

2.1. **Propositional dynamic logic.**

The language of PDL is built from two non-empty disjoint atomic types: the set Φ_0 of atomic propositions and the set Π_0 of atomic actions. We use $P, Q, \ldots$

[2] Perhaps the installation of the window or its regular cleaning counts as the relevant use, but this seems a bit far removed from the ascription.

[3] It is common to refer to objects as means as well, which is opposed to our means-as-actions semantics. We hope to discuss how objects can be means in a later paper.

to range over Φ_0 and $m, n, \ldots$ to range over Π_0. The sets Φ of formulas and Π of actions are built via the following definitions, where $\varphi, \psi, \ldots$ range over Φ and $\alpha, \beta, \ldots$ range over Π.

$$\Phi = P \mid \top \mid \varphi \wedge \psi \mid \neg\varphi \mid [\alpha]\varphi$$
$$\Pi = m \mid \alpha;\beta \mid \alpha \cup \beta$$

We have omitted the iteration α^* and test $\varphi?$ actions from our logic, since these are not essential to our present purposes. The sentence $[\alpha]\varphi$ expresses that, if one does α, then φ will be realized. The construction $\alpha;\beta$ denotes sequential composition (first do α and then do β) and $\alpha \cup \beta$ denotes non-deterministic choice between α and β.

We introduce the connectives $\neg$, $\vee$ and $\rightarrow$ and the weak operator $\langle\alpha\rangle$ as usual.

A *PDL model* $\mathcal{M}$ for Π_0 consists of

- a set $\mathcal{W}$ of worlds (or states),
- an interpretation $v : \mathcal{W} \times \Phi_0 \rightarrow \{tt, \mathit{ff}\}$ assigning truth values to pairs of worlds and atomic propositions and
- a *dynamic interpretation* of actions. This dynamic interpretation consists of transitions between worlds, labeled by atomic actions. When an arrow $w \xrightarrow{m} w'$ exists, then w' is a possible outcome of doing m in world w.

The satisfaction relation $\models \ \subseteq \ \mathcal{W} \times \Phi$ is defined as usual for the boolean connectives. We write

$$w \models [m]\varphi \quad \text{iff} \quad \text{for every } w', \text{ if } w \xrightarrow{m} w', \text{ then } w' \models \varphi.$$

Consequently,

$$w \models \langle m\rangle\varphi \quad \text{iff} \quad \text{there is a } w' \text{ such that } w \xrightarrow{m} w' \text{ and } w' \models \varphi.$$

Thus, $w \models [m]\varphi$ just in case doing m ensures that φ will be true in whatever world results and $w \models \langle m\rangle\varphi$ just in case it is possible that φ will be true in the world that results from doing m.

For example, consider a world in which one may toss a coin. If we neglect all of the features but the coin toss, there are two possible outcomes: the coin may come up heads or it may come up tails. This is modeled by three worlds and two atomic propositions with the following dynamic structure, where the actual world is denoted by the filled circle[4].

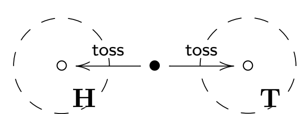

The actual world satisfies $\langle\mathsf{toss}\rangle\mathbf{H}$, but not $[\mathsf{toss}]\mathbf{H}$.

The action constructions $\alpha;\beta$ and $\alpha \cup \beta$ may be defined by the following axioms.

$$[\alpha;\beta]\varphi \leftrightarrow [\alpha][\beta]\varphi$$
$$[\alpha \cup \beta]\varphi \leftrightarrow [\alpha]\varphi \wedge [\beta]\varphi$$

The second axiom looks more natural in terms of the weak operator:

$$\langle\alpha \cup \beta\rangle\varphi \leftrightarrow \langle\alpha\rangle\varphi \vee \langle\beta\rangle\varphi.$$

Since both operations are clearly associative, we will drop parentheses indicating association hereafter.

We call an action α *prohibited in* w if there is no w' such that $w \xrightarrow{\alpha} w'$. Intuitively, such actions cannot be performed in w. If α is prohibited in w, then $w \models [\alpha]\varphi$ for any $\varphi \in \Phi$ (including $\bot$), but $w \not\models \langle\alpha\rangle\varphi$ for any $\varphi \in \Phi$ (not even $\top$).

[4]This example would be better handled by a semantics involving probabilities instead of non-determinism. That is the subject of current research on fuzzy set semantics for PDL, presented in [3] and applied to artifactual functions in [5].

2.2. **Means and ends in PDL.** In [2, 4], we presented a semantics for various means-end relations in PDL, including weakly and strongly sufficient and necessary means. We omit the discussion of necessary means here but briefly present the definitions for the two kinds of sufficient means.

Definition 2.1. Let $w \in \mathcal{W}$. We say that an action α is a *weakly sufficient means* to a formula φ in w if

$$w \models \langle\alpha\rangle\varphi,$$

i.e. , if there is a w' such that $w \xrightarrow{\alpha} w'$ and $w' \models \varphi$.

We say that α is a *strongly sufficient means* to φ in w if

$$w \models [\alpha]\varphi \wedge \langle\alpha\rangle\top,$$

i.e. , if for every w' such that $w \xrightarrow{\alpha} w'$ we have $w' \models \varphi$ and furthermore α is not prohibited in w.

While our agent has only a weak reason to do a particular weakly sufficient means, he is strongly motivated to do *some* weakly sufficient means. If he does not do any weakly sufficient means, then φ will *not* be realized. Perhaps, if we follow von Wright's analysis in [6], one may say that our agent must either do some weakly sufficient means or change his desires, on pain of practical irrationality.

We must be careful here to allow for the fact that sometimes, the world changes through no active effort on the part of our agent. If our agent desires light to read by and it is nearly dawn, then doing nothing is a means to his end. The action of doing nothing does not necessarily leave the world as it is. The world has a habit of changing on its own (and also due to the actions of other agents).

Intuitively, one should have greater motivation to perform strongly sufficient means for his end than to perform weakly sufficient means, but it is unclear how to express this intuition. Certainly, there are cases in which one prefers a weak means to a strong, because the strong means has undesirable side effects. The practical difference between strong and weak sufficiency is not clear to us at present.

3. Artifactual functions

One of the defining characteristics of artifacts is that they are function-bearing things. They acquire functions through various means, whether by design, custom acquired by repeated use or a one-time "accidental" use. A functional ascription expresses how an artifact may be used in order to achieve certain ends. In this respect, functions are closely related to means-end relations: a functional ascription expresses that an artifact plays a role in a certain means-end relation.

It is this aspect of artifactual function that motivates our treatment here. We focus on the relation between functions and means, to the neglect of diverting questions of how functions are ascribed. This latter topic is certainly worth investigation, but ascriptions are a fundamentally social act and thus we are not hopeful that formal semantics would provide much light on these issues.

In the following section, we identify certain characteristics of functional ascriptions that are relevant for our analysis.

3.1. **The primary characteristics of functions.** In this section, we discuss what we take to be the primary characteristics of artifactual functions with respect to our analysis. In choosing these characteristics, we have been guided by our desire to discuss the relation between functions and means-end relations and to give a formal definition of function fulfillment that is analogous to the corresponding natural language usage. Thus, we emphasize artifact use, context of use and the goals that distinguish success from failure, to the neglect of other characteristics of function equally worthy of study.

Functional ascriptions in natural language often take the form of "A gizmo is to ferfunkle," where *gizmo* is a type of artifact and *ferfunkle* is a verb. Whatever else such ascriptions convey, they assert that one can perform the act of ferfunkling with a gizmo. Functional ascriptions of this sort convey that a particular action can be done with the artifact. Certainly, not all functional ascriptions express active uses, either explicitly or implicitly, but we will be satisfied to give an account of those ascriptions which do and postpone a fuller analysis of artifactual functions.

Many such ascriptions include more than just the action: they include various other objects to which the action can be applied, such as "The gizmo is for ferfunkling *whatsits*" ("Staplers are for fastening papers together."). Other aspects of the context of use may be relevant as well, including parameters of use (how hard one pushes on the accelerator, for instance) or even the agent performing the use (whether a dart is likely to hit its target depends on who throws it). Thus, we presume that each functional ascription comes with a sequence of types, called the *context type*, listing the parameters relevant for determining the outcome of usage as well as evaluating its success.

Example 3.1. The basic action for using a thermostat to heat a room is *setting* it to a desired temperature. Thus, the context type includes the type $\mathcal{R}ange$ denoting a set of natural numbers representing the possible settings of the thermometer and the range of possible temperatures in the room.

The context type may include other types as well, including the heater to which the thermostat is attached, the space to be heated, etc. Alternatively, we may compare the adequacy of various thermostats when placed in a context fixed with respect to everything but temperature setting. We choose the latter, for simplicity's sake.

Ascriptions include more than an action and types of parameters. We take it that functions like this express that one can *do* a particular act with the artifact *in order* to realize some end. That is, functional ascriptions should express not only that an act can be done, but that doing so is likely to achieve an end. The end is implicit in many natural language expressions. "Staplers staple papers" appears to be a simple assertion that one can *do* the act of stapling *to* papers *with* a stapler[5]. But surely it includes an implicit goal, namely that the papers involved are fastened together after the act. Indeed, this goal seems implicit in the verb "staple".

That functions involve goals is more or less presupposed by some senses of malfunction. An artifact that, in normal circumstances, is incapable of reliably realizing its goal when compared to other artifacts of the same type is *malfunctioning*[6]. Such malfunctions assume that functions *have* associated goals and so we take these goals to be an essential part of the specification of a function.

Thus, functional ascriptions include the following components:

- the type of artifact;
- the action (use) involved;
- the goal to be achieved;
- the context type, a list of types for objects and parameters involved in the use and/or goal.

Example 3.2. Consider a thermostat in a classroom. It is part of system to regulate temperature in the room. It does so by comparing the ambient temperature t with the user-set desired temperature u. If $t < u$, then it activates the heater[7] (if off) and if $t \geq u$ then it deactivates the heater (if on).

Let $\mathcal{T}hermo$ be the type for thermostats and $\mathcal{R}ange$ a set of natural numbers[8] representing possible temperatures (say, in Celsius).

Then the function we've describe has the following components:

- The type of artifact is $\mathcal{T}hermo$.
- The action is adjust, the action of setting a thermostat to a particular temperature.
- The goal is that the temperature in the room reaches the thermostat setting or higher, denoted $\mathbf{Temp}_{\geq}$.
- The additional parameter is of type $\mathcal{R}ange$, representing the thermostat setting.

[5] It's not the first time that philosophers have tortured natural language.

[6] There are other kinds of malfunction, but this is certainly a way in which artifacts malfunction. We discuss this form of malfunction in [5].

[7] Real thermostats may do this only if t is sufficiently less than u, but we will assume a simpler model.

[8] We simplify matters somewhat by assuming a discrete set of temperatures, but this is inessential.

In the examples which follow, we are interested in the behavior of various thermostats functioning in the same physical setting, so we fix the room in which the thermostat is placed, as well as the heater to which it is attached, the current weather conditions and so on. The transition systems we consider, then, will represent what *would* be the causal structure[9] for each thermostat t if it were placed in this setting and if an agent set the thermostat to u.

3.2. **Functional ascriptions: the formal interpretations.** Propositional dynamic logic is a suitable setting for means-end relations, but in order to discuss artifactual functions, it is necessary to augment the model with a set of objects (including our artifacts). We will do so in the simplest manner we know, sacrificing logically plausible properties in favor of simplicity. We believe that this model is sufficient for conveying our basic definitions, whether or not in practice one would prefer a more robust model. In this section, we provide an overview of our models and leave the more technical details for the appendix.

Traditional approaches to adding objects to possible world semantics involve adding arities to propositions, amending the definition of valuation and so forth. This leads to a standard development of first-order dynamic logic, but it is rather more complicated than we need at present. Moreover, as we will see, we rely quite heavily on fictional "normal" tokens for each type and it is cumbersome to bend first order logic to include non-existent elements. Instead, we have chosen to introduce context via an indexed collection of models, as we will describe later. First, we discuss the rudimentary type theory we use and the specification of functional ascriptions.

Functions are ascribed to different artifactual types and types also determine which concrete contexts are relevant for the functional ascription. Thus, we add a rudimentary type theory to our formal models. We introduce a set $\mathbb{T}$ of types and a partial order $\leq$ representing the subtype relation.

Let $X^{<\omega}$ denote the set of finite sequences of elements of X. We write $\sigma, \tau, \ldots$ for sequences and we write $x * \sigma$ for the sequence obtained by prepending σ with the element x.

Definition 3.3. A *functional ascription* is a tuple

$$f = \langle T, \alpha, \varphi, \tau \rangle$$

where $T \in \mathbb{T}$ (the *artifact type*), $\alpha \in \Pi$ (the *function action*), $\varphi \in \Phi$ (the *function goal*) and $\tau \in \mathbb{T}^{<\omega}$ (the *context type*).

Definition 3.3 should be understood as giving a specification for functional ascriptions, not a condition for when a function should be ascribed to an artifactual type. We understand these ascriptions as asserting that, in contexts of the type τ, artifacts of type T have the function of realizing φ when α is done.

Example 3.4. For our thermostat, the relevant ascription is the tuple

$$\langle \mathcal{T}hermo, \mathsf{adjust}, \mathbf{Temp}_{\geq}, \mathcal{R}ange \rangle.$$

In other words, a function of thermostats is to effect that the room temperature is greater than n when one sets the thermostat to n.

Another ascription one may consider regarding thermostats is

$$\langle \mathcal{T}hermo, \mathsf{adjust}, \mathbf{IsSetTo}, \mathcal{R}ange \rangle.$$

Such an ascription would assert: it is a function of thermostats to be set to a particular n when the user sets it to n. This is linguistically odd, certainly, but it is true that one has an expectation that thermostats behave in this manner. It is perhaps odd to actually ascribe this function to thermostats, since one is unlikely to desire to realize **IsSetTo** for its own sake. Functional ascriptions are not determined by causal relations, but fulfillment of an artifactual function depends on such relations.

[9]We recognize the philosophical troubles we introduce here by applying counterfactuals to causal structures. For the purposes of our development, let us suppose that we are indeed capable of producing causal transition systems for counterfactual worlds.

Nonetheless, one certainly *could* claim that this is one of the functions of thermostats and indeed almost all thermostats could fulfill this function.

One may suppose that subtypes inherit functions, so that if we are given an ascription $\langle T, \alpha, \varphi, \tau\rangle$ and $T' \leq T$ then we may infer the ascription $\langle T', \alpha, \varphi, \tau\rangle$. But, on reflection, it does not seem that the subtype relation is so closely related to artifactual functions, since plausible counterexamples to this principle are easy to find. For instance, it seems reasonable to say that one of the functions of airplanes are to move people about, but it is not the case that remote-controlled surveillance drones have this function.

An artifact may fail to fulfill its function in a particular concrete context of appropriate type. That is, either due to malfunction, design, unforeseen or extreme circumstances or some other reason, doing α may fail to realize φ. We turn our attention to characterizing this notion of fulfilling one's function.

We let $\mathbb{O}$ denote a set "of objects" in a very broad sense: $\mathbb{O}$ consists of all the tokens for the various types[10] in $\mathbb{T}$. We assume that each $o \in \mathbb{O}$ has a primary (most specific) type, but this assumption is due to syntactic convenience, not deep ontological commitments. If o's primary type is T' and $T' \leq T$, then o is also an artifact of type T.

We refer to each sequence $\tau \in \mathbb{O}^{<\omega}$ as a *concrete context.* Let $\sigma \in \mathbb{T}^{<\omega}$ be a context type. If τ has the same length of σ and each τ_i has type σ_i, we say that τ is a *concrete context of type* σ.

A *PDL* $\mathbb{O}$*-model* is a set of PDL models $\mathcal{M}_\tau$ indexed by sequences $\tau \in \mathbb{O}^{<\omega}$, but with a common set $\mathcal{W}$ of worlds. Thus, each τ induces a dynamic transition structure $[\![-]\!]_\tau$ on $\mathcal{W}$ which represents the causal structure for the concrete context τ. Explicitly, the set $[\![\alpha]\!]_\tau(w)$ is the set of worlds which may result from performing α in world w and in context τ.

In addition, each model $\mathcal{M}_\tau$ comes with its distinct satisfaction relation $\models_\tau$. This allows the interpretation of atomic predicates P to depend on the context τ. Whether a world satisfies P may depend on the context in which we are interested. For instance, we may have among our propositional variables a relation F for "is fastened to" and interpret this relation so that $w \models_\tau F$ iff in world w, the objects $\tau_0, \tau_1, \ldots, \tau_{n-1}$ are somehow fastened to τ_n. (Presumably $w \not\models_\tau F$ if $\mathrm{len}(\tau) < 2$.)

We intend that each model $\mathcal{M}_\tau$ expresses the relevant causal and semantic structure for the concrete context τ. Thus, the set $\mathcal{W}$ of worlds are interpreted as worlds in concrete contexts, rather than some broader sense of possibility. Because our models describe what would be the case if an artifact were used in a particular way, independently of whether the artifact is so situated that it *could* be used thus, there is a strongly counterfactual component in our semantics. We accept this philosophically troublesome feature because we believe that the counterfactual component in our semantics is inherited from natural language ascriptions.

A given function $f = \langle T, \alpha, \varphi, \sigma\rangle$ defines a set of *relevant* contexts, where relevance is defined in relation to functional fulfillment. Specifically, a concrete context τ is *relevant to* f if τ is a context of type $T * \sigma$, the context type formed by prepending T to σ (see the appendix for details). Equivalently, τ is relevant just in case $\tau = o * \tau'$ where o is an object of type T and τ' is a concrete context of type σ. The ascription expresses an expectation regarding the outcome of α in just these concrete contexts and so it is these concrete contexts which are relevant to f.

Example 3.5. Let us consider a simple family of models for the thermostat example. We will take one action, adjust, representing the action of adjusting a thermostat, and two atomic propositions, **IsSetTo**, a relation expressing the temperature to which a thermostat is set, and $\mathbf{Temp}_\geq$, a relation expressing the ambient temperature surrounding the thermostat.

[10]We could view $\mathbb{O}$ as the most general type of all, but we do not pursue that approach. Instead, $\mathbb{O}$ is intended as a semantic domain for our types. Since our types may include mathematical structures, like the set of natural numbers, mathematical objects may be elements of $\mathbb{O}$, but one should not infer any ontological claims about the status of mathematical objects.

For our simple model, the only objects are thermostats and the natural numbers in $\mathcal{R}ange$. Thus, we have a collection of models indexed by sequences σ, where each σ_i is either a number or a thermostat.

A sequence σ is *relevant* to

$$f = \langle \mathcal{T}hermo, \mathsf{adjust}, \mathbf{Temp}_{\geq}, \mathcal{R}ange \rangle$$

if $\mathrm{len}(\sigma) = 2$, $\sigma_0 \in \mathcal{T}hermo$ and $\sigma_1 \in \mathcal{R}ange$. We are interested only in those models $\mathcal{M}_\sigma$ where σ is relevant.

Let $t \in |\mathcal{T}hermo|$ and $n \in |\mathcal{R}ange|$. The model $\mathcal{M}_{\langle t,\, n\rangle}$ represents the causal and semantic structure that would result if thermostat t is situated in our classroom (attached in the usual manner) and n is the parameter relevant for the adjust action. Thus, regarding the atomic propositions,

- $w \models_{\langle t,\, n\rangle} \mathbf{IsSetTo}$ iff in world w, the thermostat t is set to n.
- $w \models_{\langle t,\, n\rangle} \mathbf{Temp}_{\geq}$ iff the temperature of the air near the thermostat is *at least* n.

The transition system

$$[\![\mathsf{adjust}]\!]_{\langle t,\, n\rangle} : \mathcal{W} \to \mathcal{P}\mathcal{W}$$

represents the effects of setting thermostat t to n. These effects need not be immediate, and indeed our simple model has no explicit representation of time, but we intend that $w' \in [\![\mathsf{adjust}]\!]_{\langle t,\, n\rangle}(w)$ just in case world w' is an expected result of setting t to n in w. Unless t is *very* broken in w, we expect that $w \models_{\langle t,\, n\rangle} [\mathsf{adjust}]\mathbf{IsSetTo} \wedge \langle \mathsf{adjust}\rangle\top$, i.e. that for each thermostat-range pair, adjust is a strongly sufficient means to **IsSetTo**.

For normally functioning thermostats in normal situations, we expect that

$$w \models_{\langle t,\, n\rangle} [\mathsf{adjust}]\mathbf{Temp}_{\geq} \wedge \langle \mathsf{adjust}\rangle\top,$$

so that setting the thermostat is a strongly sufficient means to $\mathbf{Temp}_{\geq}$, too. This won't be the case for malfunctioning thermostats, or thermostats connected to faulty heaters. It won't be the case if the heater is too small to heat the area, either. But *normally* the condition should hold.

In Figure 1, we give four very simple models of possible worlds semantics for thermostats. In each case, we assume that $\mathcal{R}ange$ has three values $\{l, m, h\}$ (for *low*, *medium* and *high*). The transition systems show $\mathcal{M}_{\langle t,\, m\rangle}$ for some thermostat t. We assume that any temperature changes that occur are the result only of heater activity and not due to external influences.

Figure 1(a) depicts a normally functioning thermostat. For instance, the world marked $*$ in this figure is a world in which the thermostat is currently set to h but the ambient temperature is l (perhaps the thermostat has just been turned on). The effect of adjust is a world in which the thermostat is set to m and the ambient temperature is also m (because m is the temperature for our context). Thus, the effect of adjust is a world satisfying both **IsSetTo** and $\mathbf{Temp}_{\geq}$, i.e. adjust is a means to both **IsSetTo** and $\mathbf{Temp}_{\geq}$.

Figure 1(b) shows a thermostat that fails to raise the temperature of the room to m. This may be because the thermostat is broken, the heater is broken, the two are not designed to work together, the thermostat's thermometer is mis-calibrated so that it reports a higher temperature than is in the room, or many other reasons. (It may be possible to eliminate some of these reasons by looking at $\mathcal{M}_{\langle t,\, h\rangle}$, but we are not presenting a theory of diagnosis here.) The effect of adjust in world $*$ is a world in which the thermostat is set to m but the temperature remains unchanged. Thus, in $*$, adjust is a means to **IsSetTo** but not to $\mathbf{Temp}_{\geq}$.

Figure 1(c) shows the transition structure for a thermostat with a mis-calibrated thermometer. In this case, the thermometer registers a lower temperature than the actual ambient temperature. As a result, the thermostat activates the heater for a longer period than necessary, resulting in a temperature strictly higher than m. (Factors other than a mis-calibrated thermometer could produce the same causal structure, of course.) For instance, in $*$, the effect of adjust is a world in which the thermostat is set to m but the room has been overheated, so that the ambient temperature is h. But an overheated room still satisfies $\mathbf{Temp}_{\geq}$, so in this world, adjust is a sufficient means to both **IsSetTo** and $\mathbf{Temp}_{\geq}$.

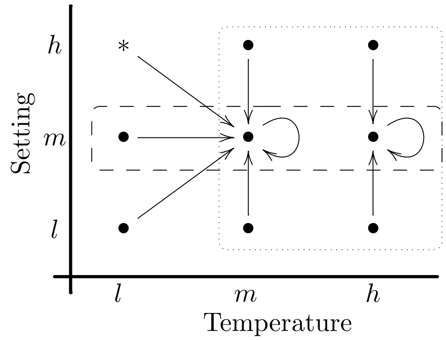

(a) A working thermostat.

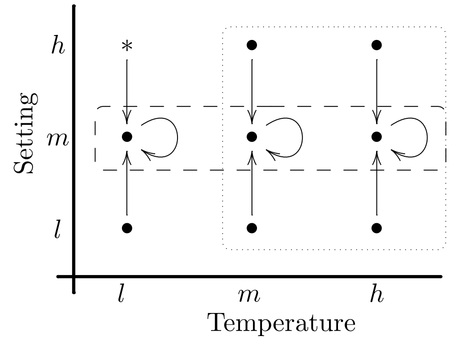

(b) A non-working thermostat.

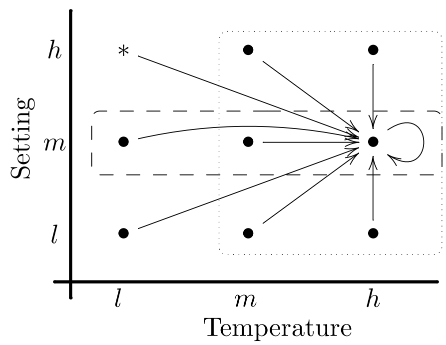

(c) A mis-calibrated thermostat.

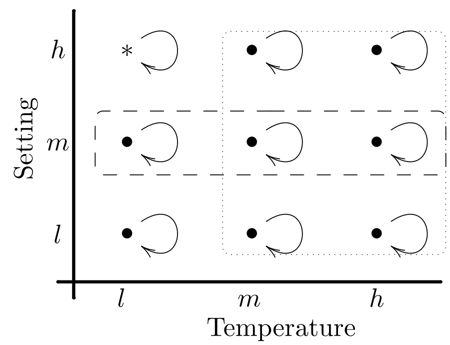

(d) A very broken thermostat.

FIGURE 1. Four models of thermostats. In each model $\mathcal{M}_{\langle t,\, m\rangle}$, the action adjust sets the thermostat to the middle of three possible values. The set of worlds bounded by the dashed line represents **IsSetTo**, and the dotted line represents $\mathbf{Temp}_{\geq}$. The world marked $*$ is a world in which the thermostat is currently set to h but the ambient temperature is l.

Figure 1(d) shows a thermostat that fails to even register the effect of adjust. Some digital thermostats are battery-powered, and if the battery is drained, then attempting to set the thermostat produces no change in the world at all (but we still suppose that the action adjust can be performed and that the thermostat *has* a current setting). Thus, the effect of adjust in $*$ is a world in which nothing has changed, i.e. the effect is again $*$. Clearly, adjust is neither a means to **IsSetTo** nor $\mathbf{Temp}_{\geq}$ in this world.

We have identified the causal structures relevant for a functional ascription and now we may give our definition of fulfillment.

Definition 3.6. Let $f = \langle T, \alpha, \varphi, \sigma\rangle$ be a functional ascription and let $o \in T$ and $\tau \in \mathbb{O}^{<\omega}$ a concrete context of type σ. We say that *o (weakly/strongly) fulfills f in concrete context τ* just in case the model $M_{o*\tau}$ satisfies: for every $w \in \mathcal{W}$, the action α is a (weak/strong) means to φ in w.

Thus,

- o weakly fulfills f in σ iff, in every world w,

$$w \models_{M_{o*\tau}} \langle\alpha\rangle\varphi;$$

- o strongly fulfills f in σ iff, in every world w,

$$w \models_{M_{o*\tau}} [\alpha]\varphi \wedge \langle\alpha\rangle\top.$$

Example 3.7. The thermostats featured in Figures 1(b) and 1(d) do *not* fulfill the function $\langle \mathcal{T}hermo, \mathsf{adjust}, \mathbf{Temp}_{\geq}, \mathcal{R}ange\rangle$ from Example 3.4. In each of these diagrams, there are worlds w such that

$$w \not\models_{\langle t,\, m\rangle} \langle \mathsf{adjust}\rangle \mathbf{Temp}_{\geq},$$

namely those worlds in which the ambient temperature has value l.

The thermostats in Figures 1(a) and 1(c) fulfill the function $\langle \mathcal{T}hermo, \mathsf{adjust}, \mathbf{Temp}_{\geq}, \mathcal{R}ange\rangle$. In every world, the world realized via adjust satisfies $\mathbf{Temp}_{\geq}$.

The fact that the thermostat in Figure 1(c) fulfills the function f may be surprising, since doing adjust *overheats* the room. But our function f only asserts that, if one does adjust, the ambient temperature in the resulting room will be at least m, and this is true for every world in Figure 1(c). One may certainly want to assert something stronger, such as: in every w, the action adjust realizes a temperature higher than m only if the temperature in w is higher than m. Such a functional ascription is not expressible in terms of Definition 3.3. For that, we must add conditional ascriptions, developing them along the lines of the conditional means-end relations introduced in [4].

An artifact o may fail to fulfill its function for many reasons, including malfunction, but also including that o is not designed to fulfill f in every concrete context or that o is functioning as designed but it is poorly designed. We do not think that an ascription f expresses the expectation that every o of appropriate type fulfills f in every appropriate concrete context. Rather, they express the expectation that "normally" artifacts of type T fulfill f in concrete contexts of type σ. We turn our attention to understanding this expectation.

3.3. **Universal and normal fulfillment.** The notion of fulfillment in Section 3.2 applies to tokens in concrete contexts, but very often one is interested in more general notions of fulfillment. For instance, given a function f for an artifact type T, we may be interested in comparing the performance of two subtypes T' and T'' of T. Or we may be interested in whether a particular token performs as well as can be expected for tokens of that type. For this, we need to introduce function fulfillments for types.

We do this in two steps. First, we introduce *partial contexts for context type σ.* These are sequences τ of objects and types such that each τ_i is either a token or subtype of type σ_i. Concrete contexts for σ are a special case of a partial context for σ, as is σ itself. Indeed, concrete contexts are the most specific example and σ itself the least specific, where specificity is a relation induced by the subtype relation.

Partial contexts induce another sense of relevance. A concrete context is *relevant* for a partial context τ just in case σ is more specific than τ. This gives a strong notion of artifactual fulfillment for partial contexts: A partial context τ *(weakly/strongly) fulfills* $f = \langle T, \alpha, \varphi, \sigma\rangle$ if

- τ is a partial context for $T * \sigma$ and
- each concrete context relevant to τ fulfills f.

In case the partial context τ is concrete, this definition coincides with the definition of fulfillment from the previous section. If τ is the broadest context induced by f (namely, $\tau = T * \sigma$), then τ fulfills f in this sense if and only if *every* concrete context fulfills f.

This is a very strong sense of fulfillment. If there is any broken token of type T, then $T * \sigma$ does *not* fulfill the function f in this sense. It is usually not the case that one who ascribes the function f expects $T * \sigma$ to fulfill its function in this sense. When saying that T-things have a certain function, we do not mean that every T-thing can be expected to fulfill its function.

Instead, we think that functional ascriptions imply expectations about *normal* tokens of type T in *normal* contexts of type σ. In order to express this in a formal setting, we assume that each type T comes with a set of *normal* tokens. This induces a concept of normal contexts, in the obvious way: a sequence τ is a *normal context for* σ if τ is a concrete context for σ and each τ_i is a normal token for type σ_i (or $\tau_i = \sigma_i$ when the latter is in $\mathbb{O}$).

We say that a partial context τ *normally (weakly/strongly) fulfills* f if the following hold.

- τ is a partial context for $T * \sigma$ and
- each *normal* context for τ (weakly/strongly) fulfills f.

We believe that when one ascribes the function f, he believes that $T * \sigma$ normally weakly fulfills f, that is, that it is possible to realize φ via α using normal tokens of type T and in normal contexts of type σ.

Example 3.8. In our thermostat example, a normal thermostat should be a working thermometer of an unexceptional quality, so that if t is a normal token of type $\mathcal{T}hermo$, then we expect $\mathcal{M}_{\langle t,\, m\rangle}$ to be the structure found in Figure 1(a) (assuming that the heater is appropriate for the task and so on).

Normal tokens of type $\mathcal{R}ange$ should probably include every element of $|\mathcal{R}ange|$, unless this set contains extreme values to which thermostats are rarely set.

The introduction of normal tokens raises many questions, of course. In particular, we must be careful how to interpret these normal tokens. First, we should be explicit that our idea of normality is about expectations regarding how a token ought to behave. It does not presuppose that any of the existing tokens actually *do* behave that way. If every token of a particular type happens to be broken, it does not follow that normal tokens of that type are broken.

Thus, our normal tokens are convenient fictions, and this is consistent with our formalism. Contexts determine the relevant causal structures. There is nothing about our formalism that requires each token in a concrete context literally exists. The models M_τ present the causal and semantics structures that are relevant in the hypothetical context τ. There is some real question how to make sense of these counterfactual structures, but the question does not seem much harder in the presence of fictional objects than otherwise.

But how does one come to expectations regarding fictional "normal" tokens of a particular type? We do not know the answer to this, but we nonetheless think that the concept of normal tokens is related to natural language functional ascriptions. We take it that functional ascriptions convey some expectations about realizing goals via uses and that these expectations do not hold for all artifacts of appropriate type, but for some kind of normal or typical artifacts. Formally, we represent this understanding by adding a set of normal tokens for each type, but we recognize that this move raises many new issues about the conceptual analysis. Nonetheless, we think that the normal tokens approximates natural language functional talk to such degree that we include them here, regardless of unsettled controversies.

Thus, we assume that one *does* have some intuitions about how "normal" tokens of an artifactual type would behave in various contexts and we use this assumption to define normal fulfillment. We claim that functional ascriptions convey an expectation of such normal fulfillments. Moreover, this gives an appropriate sense of fulfillment to compare the adequacy of various types of artifacts in achieving common functional ends. When one says that one artifact type is better suited to a task than another, it seems sensible to interpret this claim in terms of the adequacy of the normal tokens of each type (but we postpone an explicit discussion of this for another day).

4. Concluding remarks and further developments

Artifactual functions are closely related to means-end relations: a functional ascription implies that tokens of the appropriate types may be used to realize ends. This has been the motivation behind our work in formal semantics of artifactual functions (and was first suggested to the author by Sjoerd Zwart).

We have given a brief introduction to our semantics for means-end relations in Section 2, followed by a formal specification for functions and a semantics for fulfillment that borrows from our means-end semantics. Our interpretation includes philosophically suspect constructions, such as counterfactual causal structures (some involving fictional "normal" entities), but we stand by it as a close approximation of the natural language meaning of functional ascriptions.

PDL is a nice language for expressing the relation between functions and means-end talk, but it is a bit crude at distinguishing behavior. Like other modal logics, the semantics for PDL is directed towards distinguishing possible and necessary outcomes, without any hint of likelihoods. The next step in our development is to add probabilities to the picture so that we may more subtly discuss malfunction in cases in which the goal *might* be achieved, but is unlikely. This is a matter of adapting our development of *efficacy* for means-end relations, first presented in [3].

References

[1] David Harel. Dynamic logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume II, pages 497–604. D. Reidel Publishing Company, 1984.

[2] Jesse Hughes. "Means" means *What*? Submitted for APA Winter Meeting 2005, preprint available at `http://phiwumbda.org/~jesse/papers/index.html`, 2005.

[3] Jesse Hughes, Albert Esterline, and Bahram Kimiaghalam. Means-end semantics and a measure of efficacy. Submitted for Progic 2005, preprint available at `http://phiwumbda.org/~jesse/papers/index.html`, 2005.

[4] Jesse Hughes, Peter Kroes, and Sjoerd Zwart. A semantics for means-end relations. Accepted for presentation at SEP 2005, preprint available at `http://phiwumbda.org/~jesse/papers/index.html`, 2005.

[5] Jesse Hughes and Sjoerd Zwart. A semantics for functional efficacy. To be presented at SPT 2005, preprint available at `http://phiwumbda.org/~jesse/papers/index.html`, 2005.

[6] Georg Henrik von Wright. Practical inference. *The Philosophical Review*, 72(2):159–179, Apr. 1963.

Appendix A. Technical details

A.1. **Concrete contexts.** In this appendix, we present the technical details of our model of artifactual functions discussed in Section 3.

As we discussed there, our approach is the simplest means of adding objects and parameters to propositional dynamic logic we could find. We did not move immediately to a first-order dynamic logic as one might expect, because indexing our models by contexts suffices for our needs and allows a simpler presentation. Also, indexing the models by context stresses the counterfactual interpretation: the models express the causal and semantic structure that would result in the given context.

Let $\mathbb{O}$ be a set "of objects", in the loose sense described in Section 3 and let $\mathbb{T}$ be a set of types with the subtype relation denoted $\leq$. For $T \in \mathbb{T}$, let

$$\downarrow T = \{T' \in \mathbb{T} \ \mid \ T' \leq T\}.$$

We assume that each object in $\mathbb{O}$ has a primary or most specific type, and let $t : \mathbb{O} \to \mathbb{T}$ be the function taking an object to its type. Let the *support of* T be defined as

$$|T| = \{o \in \mathbb{O} \ \mid \ t(o) \in \downarrow T\}.$$

Clearly if $T \leq T'$ then $|T| \subseteq |T'|$.

Let $X^{<\omega}$ denote the set of sequences of X. If $\sigma \in X^{<\omega}$, then $\mathrm{len}(\sigma)$ denotes the length of σ. We write σ_i for the ith element of σ and begin indexing at 0, so that if $\mathrm{len}(\sigma) = n$, then $\sigma = \langle \sigma_0, \ldots, \sigma_{n-1}\rangle$. If $x \in X$, then $x * \sigma$ denotes the sequence $\langle x, \sigma_0, \ldots, \sigma_{\mathrm{len}(\sigma)-1}\rangle$.

We extend the support notation to sequences $\sigma \in \mathbb{T}^{<\omega}$ of types via the definition

$$|\sigma| = \left\{\tau \in \mathbb{O}^{<\omega} \ \mid \ \mathrm{len}(()\tau) = \mathrm{len}(()\sigma) \text{ and for all } i < \mathrm{len}(()\tau) \ . \ \tau_i \in |\sigma_i|\right\}.$$

Definition A.1. A *PDL* $\mathbb{O}$*-model* consists of a set $\mathcal{W}$ of worlds together with an $\mathbb{O}$-indexed set

$$\left\{\mathcal{M}_\tau \ \mid \ \tau \in \mathbb{O}^{<\omega}\right\}$$

of PDL models such that each $\mathcal{M}_\tau$ has $\mathcal{W}$ as its set of worlds.

We write $[\![-]\!]_\tau$ for the interpretation and valuation functions in the model $\mathcal{M}_\tau$. We write $\models_\tau$ for the corresponding satisfaction relation.

We think of τ as parameters to the formula φ (and to any dynamic operators α), so that we implicitly regard our set of propositional variables as (untyped) relations over $\mathbb{O}$. A propositional variable P in world w of $\mathcal{M}_\tau$ is true if P holds of τ in w. For example, the propositional variable **IsSetTo** from Example 3.5 is true in those worlds in model $\mathcal{M}_{\langle t,\, m\rangle}$ in which the thermostat is set to m. More precisely, **IsSetTo** is true in a world w in $\mathcal{M}_\tau$ just in case

- $\mathrm{len}(\tau) \geq 2$,
- $\tau_0 \in |\mathcal{T}hermo|$,
- $\tau_1 \in |\mathcal{R}ange|$ and
- in w, the thermostat τ_0 is set to τ_1.

An action α also depends on context. The action adjust from Example 3.5 only makes sense if the context includes both a thermostat and a value to be set, at least in our simple model. In other contexts, it is natural to interpret $[\![\mathsf{adjust}]\!](w)$ as the empty set, representing that adjust cannot be done in these contexts. For our thermostat models, we may choose the convention that $[\![\mathsf{adjust}]\!](w) = \emptyset$ *unless* the following hold:

- $\mathrm{len}(\tau) \geq 2$,
- $\tau_0 \in |\mathcal{T}hermo|$ and
- $\tau_1 \in |\mathcal{R}ange|$.

We repeat some of the definitions from Section 3 here.

Definition A.2. A *functional ascription* is a tuple

$$f = \langle T, \alpha, \varphi, \sigma\rangle$$

where $T \in \mathbb{T}$ (the artifact type), $\alpha \in \Pi$ (the function action), $\varphi \in \Phi$ (the function goal) and $\sigma \in \mathbb{T}^{<\omega}$ (the function context types).

Definition A.3. Let $f = \langle T, \alpha, \varphi, \sigma\rangle$ be a functional ascription. A *concrete context* $\tau \in \mathbb{O}^{<\omega}$ is *relevant to* f if $\tau \in |T * \sigma|$.

Definition A.4. Let τ be a concrete context relevant to $f = \langle T, \alpha, \varphi, \sigma\rangle$. We say that τ *weakly fulfills* f if for every $w \in \mathcal{W}$,

$$w \models_\tau \langle\alpha\rangle\varphi.$$

Similarly, τ *strongly fulfills* f if for every $w \in \mathcal{W}$,

$$w \models_\tau [\alpha]\varphi \wedge \langle\alpha\rangle\top.$$

In other words, weak (strong, resp.) fulfillment amounts to a weak (strong, resp.) means-end relation between function act and function goal in every world in $\mathcal{M}_\tau$. In the case that the transition system for $\mathcal{M}_\tau$ is deterministic (as in our examples in Figure 1), the two kinds of fulfillment are equivalent.

A.2. **Partial contexts.** The development in Section A.1 gives a specification for functional ascriptions and a semantics for token fulfillment of a function in context. Our functional ascriptions ascribe functions to a type, but our notion of fulfillment in Definition A.4 applies only to tokens. In this section, we introduce partial contexts, which are the first step to making sense of whether a type fulfills its function.

Let $||T|| = |T| \cup \downarrow T$ and for $\sigma \in \mathbb{T}^{<\omega}$, define

$$||\sigma|| = \left\{\tau \in (\mathbb{O} \cup \mathbb{T})^{<\omega} \;\mid\; \mathrm{len}(()\tau) = \mathrm{len}(()\sigma) \text{ and for all } i < \mathrm{len}(()\tau)\ .\ \tau_i \in ||\sigma_i||\right\}.$$

Definition A.5. Let $f = \langle T, \alpha, \varphi, \sigma\rangle$ be a functional ascription. A *partial context of* f is a sequence $\tau \in ||T * \sigma||$.

In other words, a partial context is a sequence $\tau \in (\mathbb{O} \cup \mathbb{T})^{<\omega}$ such that τ_0 is either an object of type T or a subtype of T and each τ_{i+1} is either an object of type σ_i or a subtype of σ_i.

Note that the concrete contexts from Section A.1 are examples of partial contexts.

Define a partial order $\leq$ on $||T||$ by $X \leq Y$ iff one of the following hold:

- X and Y are both types and $X \leq Y$ (as types).
- $X \in \mathbb{O}$ and $Y \in \mathbb{T}$ and $X \in |Y|$.
- $X, Y \in \mathbb{O}$ and $X = Y$.

This extends as a partial order on sequences $\tau \in (\mathbb{O} \cup \mathbb{T})^{<\omega}$ in a natural way, namely $\tau \leq \tau'$ iff $\mathrm{len}(\tau) = \mathrm{len}(\tau')$ and for each $i \leq \mathrm{len}(\tau)$, we have $\tau_i \leq \tau_i$.

Under this ordering, a concrete context for f is a minimal partial context. Also, the function ascription $f = \langle T, \alpha, \varphi, \sigma\rangle$ provides the greatest partial f-context, namely $T * \sigma$.

Example A.6. Suppose that we have two subtypes of thermostat, $\mathcal{D}igital$ and $\mathcal{A}nalog$, so-named for the kind of thermometer used in the thermostat. Then each of the following are partial contexts of the function f from Example 3.4.

- $\langle t,\, n\rangle$ for any $t \in \mathcal{T}hermo$ and $n \in \mathcal{R}ange$,
- $\langle \mathcal{D}igital,\, n\rangle$ and $\langle \mathcal{A}nalog,\, n\rangle$ for any $n \in \mathcal{R}ange$,
- $\langle t,\, \mathcal{R}ange\rangle$ for any $t \in \mathcal{T}hermo$,
- $\langle \mathcal{D}igital,\, \mathcal{R}ange\rangle$ and $\langle \mathcal{A}nalog,\, \mathcal{R}ange\rangle$,
- $\langle \mathcal{T}hermo,\, n\rangle$ for any $n \in \mathcal{R}ange$,
- $\langle \mathcal{T}hermo,\, \mathcal{R}ange\rangle$.

Each of these partial contexts specify a domain of interest with respect to fulfillment.

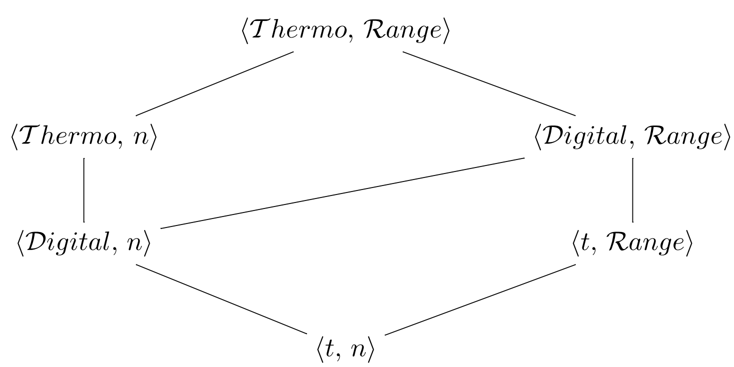

FIGURE 2. The partial order of partial f-contexts above $\langle t,\, n\rangle$.

Let us fix $t \in \mathcal{D}igital$ and $n \in \mathcal{R}ange$. The partial order of partial f-contexts above $\langle t,\, n\rangle$ is given in Figure 2.

We will use partial contexts to define sets of concrete contexts in which we are interested. Specifically, we will restrict our focus to concrete contexts which are consistent with the partial context in the sense described below.

Definition A.7. Let τ be a partial context of f. A concrete context τ' of f is a *τ-context* (or a *concrete context of f relative to τ*) iff $\tau' \leq \tau$.

The set of all I-contexts yields a natural definition of function fulfillment with respect to partial contexts.

Definition A.8. We say that a partial context τ of f *(weakly/strongly) fulfills* f if every I-context (weakly/strongly) fulfills f in the sense of Definition A.4.

We call this kind of partial-context fulfillment *universal*. In the case that I is a minimal (i.e. concrete) context, Definition A.8 is equivalent to Definition A.4.

Example A.9. The partial context $\langle \mathcal{D}igital,\, \mathcal{R}ange\rangle$ fulfills the function

$$f = \langle \mathcal{T}hermo,\, \mathsf{adjust},\, \mathbf{Temp}_{\geq},\, \mathcal{R}ange\rangle$$

just in case, for every $t \in |\mathcal{D}igital|$ and $n \in |\mathcal{R}ange|$, the concrete context $\langle t, n\rangle$ fulfills f, i.e. for every w,

$$w \models_{\langle t,\, n\rangle} \langle \mathsf{adjust}\rangle \mathbf{Temp}_{\geq}.$$

The partial context $\langle \mathcal{T}hermo, l\rangle$ fulfills f iff for every thermostat t and world w,

$$w \models_{\langle t,\, l\rangle} \langle \mathsf{adjust}\rangle \mathbf{Temp}_{\geq}.$$

But in our very simple models in Figure 1, this is trivially true (since we assumed that the temperature of the room was never lower than l).

A.3. **Normal contexts.** We expect that, when one makes a functional ascription

$$f = \langle T, \alpha, \varphi, \tau\rangle,$$

he intends that artifacts of type T are capable of fulfilling the ascription in context σ—at least weakly. But this is still ambiguous: does it mean that there is some concrete context which fulfills f? Or does it mean that $\langle T, \sigma\rangle$ universally fulfills f?

Surely neither of these is the correct interpretation. If one says that thermostats are for heating the room to a desired temperature, he means something more than: there is at least one thermostat that can perform this task. But just as surely, he does not mean *every* thermostat can fulfill this task. That claim would be false if there exists just one broken thermostat among the millions of working thermostats. Instead, we think that ascriptions convey the expectation that *normal* artifacts of type T will fulfill f in *normal* contexts of type σ.

Probably, what counts as a normal artifact of type T depends on the function f. For different functions, we may regard different tokens of T to be normal. But we will ignore this plausible dependency and assume that, for each type T, there is a set T^{normal} of "normal" T-tokens.

This induces a sense of normal contexts in the usual way, namely if $\sigma \in (\mathbb{O} \cup \mathbb{T})^{<\omega}$ then $\tau \in \mathbb{O}^{<\omega}$ is in σ^{normal} iff the following hold.

- $\mathrm{len}(\tau) = \mathrm{len}(\sigma)$;
- if $\sigma_i \in \mathbb{O}$ then $\tau_i = \sigma_i$;
- if $\sigma_i \in \mathbb{T}$ then $\tau_i \in {\sigma_i}^{\mathsf{normal}}$.

Definition A.10. Let $f = \langle T, \alpha, \varphi, \sigma\rangle$ be a functional ascription and τ a partial instantiation of f. Then τ *normally (weakly/strongly) fulfills* f if every $\tau' \in \tau^{\mathsf{normal}}$ (weakly/strongly) fulfills f in the sense of Definition 3.6.

(Jesse Hughes) Technical University of Eindhoven, Eindhoven, The Netherlands
E-mail address: J.Hughes@tm.tue.nl