A SEMANTICS FOR FUNCTIONAL EFFICACY

JESSE HUGHES AND SJOERD ZWART

Abstract. Artifactual functions are related to means-end relations. Whatever else a functional ascription conveys, it expresses the expectation that artifacts of the appropriate type may be used to achieve the functional goal. We use our existing work on means-end semantics to make this relationship precise and present a preliminary formal account of artifactual functions.

We adapt our work on probabilistic transition systems and fuzzy dynamic logic to add a measure of *efficacy* to both our means-end semantics and definition of function-fulfillment. Efficacy allows a rudimentary definition of one kind of malfunction, namely, a token which has much lower efficacy in achieving its functional goal than normal tokens of that type is *malfunctioning*.

"A whistle is to make people jump."
Ruth Krauss, *A Hole is to Dig*

1. Introduction

Functional ascriptions express expectations about how an artifact may be used. Such ascriptions include (implicitly or explicitly) a *function goal* and a *function action*. The former expresses the end of the function, the goal that can be attained by using the artifact in the right way. The latter specifies how to use the artifact to achieve this end[1].

The above quotation from Ruth Krauss's children book, *A Hole is to Dig*, clearly illustrates the kind of functional ascription we have in mind here. It ascribes a function to the artifact type $\mathcal{W}histle$. The function goal is to make people jump and the function action (implicit in the quotation) is to blow the whistle. In fact, we are so taken with how clearly this humorous ascription exemplifies the form of artifactual function ascriptions that we will use it as a running example throughout.

Functional ascriptions are closely related to means-end relations. Ascriptions express that *normally*, if one does the function action to an artifact of appropriate type, then he will realize the function goal. In other words, the function action is a *means* to the function end. We have presented a semantics for means-end relations via Propositional Dynamic Logic (PDL) in [7, 5] and we apply that work here to develop a formal semantics for artifactual functions (first presented in [4]).

[1]Not all functional ascriptions have obvious function actions. Ventilation holes in a cat carrier have a function, but it is not clear that there is any associated action. Our account here does not apply to such passive functions and we postpone consideration of these kinds of function for later work.

Artifacts do not always behave as expected. Poorly designed or poorly maintained artifacts may rarely realize their function goals, even when used as prescribed. We adapt our work on *efficacy* for means-end relations from [6] to represent this uncertainty by adding probabilities to the transition systems of PDL. This gives a preliminary definition of *malfunction*: a token is malfunctioning if it has considerably less efficacy to achieve its end than "normal" artifacts of the same narrow type.

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[2]. See [3] 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 $P, Q, \ldots$ 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

[2] 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.

- a set $\mathcal{W}$ of worlds (or states),
- an interpretation $v : \mathcal{W} \times \Phi_0 \to \{tt, f\!f\}$ 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 in w.

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[3].

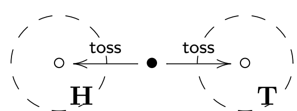

The actual world satisfies $\langle \mathsf{toss} \rangle \mathbf{H}$, but not $[\mathsf{toss}]\mathbf{H}$.

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.

[3] This example would be better handled by fuzzy set semantics involving probabilities instead of non-determinism. See Section 4.1.

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 $\perp$), but $w \not\models \langle\alpha\rangle\varphi$ for any $\varphi \in \Phi$ (not even $\top$).

2.2. **Means and ends in PDL.** In [5, 7], 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 (in w) an action α is a *weakly sufficient means* to a formula φ 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 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.

A weakly sufficient means may bring about one's end, but is not guaranteed to. A strongly sufficient means is certain to realize one's end. Nonetheless, there may be reasons to prefer a particular weakly sufficient means over a strongly sufficient means (including costs, undesirable side effects, etc.). If we follow von Wright's analysis [9], one may say that our agent must either do *some* weakly sufficient means to an end he desires 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).

3. Artifactual functions

Means-end relations are intuitively closely related to artifactual functions. When we say that a whistle is for making people jump, then we are asserting that there is a means involving the whistle (namely, blowing it) that will realize our end (making people jump). Perhaps not all artifactual functions involve explicit uses (means) like this, but at least many of them do. It is this kind of artifactual function we examine here.

In the remainder, we stress this particular characteristic of function, to the neglect of many other characteristics which are equally interesting but not equally easy to formalize or equally relevant to malfunction. In particular, we do not include any discussion of the origin of functional ascriptions. While means-end relations express only certain causal relations—certain propensities—functions are ascribed to an artifact by *somebody*, therefore the same function may be ascribed to artifacts with very different dissimilar underlying causal structures. The origins of artifactual functions (whether by design, repeated use or one-off accidental exploitation) include a social component that is not amenable to formal methods—but neither is it relevant to malfunction. We take the ascriptions as given and reason from there.

We presented an account of artifactual functions in [4], with technical details in the appendix. We summarize that account here, but we simplify it by omitting what we called *contexts* in ibid. Contexts give a fuller description of the expectations conveyed in functional ascriptions and also allow a more sophisticated representation of the causal structures involved: whether blowing a whistle makes people jump depends not only on the whistle, but also on how hard it is blown, the nervous qualities of the people that hear it and perhaps the skills of the blower. Nonetheless, we ignore these considerations for the sake of simplicity.

3.1. **Functional ascriptions.** In addition to the sets Φ_0 of atomic propositions and Π_0 of atomic actions, we add a set $\mathbb{T}$ of types and a set $\mathbb{O}$ of tokens. We assume that each token has a *primary type* and that the types are partially ordered by a subtype relation $\leq$. The primary type of a token is the minimal type containing it.

A *functional ascription* involves four components:

(1) a type T, called the *artifact type* of the ascription;

(2) an act α, called the *function action*;

(3) an end φ, called the *function goal*;

We typically write $f = (T, \alpha, \varphi)$ for a functional ascription.

The type T gives the artifact type the ascription is about. In our example ("A whistle is to make people jump."), the type is $\mathcal{W}histle$. The action is the way in which the type is to be used. In our case, this is implicit: to make people jump, one must blow the whistle. The third component is the end to be achieved: that people in our vicinity jump, a condition denoted **Jump**. Thus, we formalize the child's ascription as:

$$f = (\mathcal{W}histle, \mathsf{blow}, \mathbf{Jump})\,.$$

We assume that each atomic action involves some artifact—some token that one uses to do the action—so that they are appropriate to serve as function actions. Thus, rather that give a dynamic interpretation for each action $m \in \Pi_0$, we give dynamic interpretation for action-token pairs, (m, o). Thus, a model consists of the following:

(1) a set $\mathcal{W}$ of worlds;
(2) an interpretation $[\![P]\!] \subseteq \mathcal{W}$ of each $P \in \Phi_0$;
(3) a transition structure $[\![(m, o)]\!] : \mathcal{W} \to \mathcal{P}\mathcal{W}$ for each $m \in \Pi_0$ and $o \in \mathbb{O}$.

We assume that the actions involve some artifact, so that they are appropriate to serve as function actions. Thus, $[\![(m, o)]\!]$ is the dynamic structure representing "do m with o" and this extends to structures for (α, o) for any $\alpha \in \Pi$ in the obvious way. Thus, $[\![(\alpha, o)]\!](w)$ is the set of worlds which might result from doing α with o in w. We simplify our notation by omitting the parentheses for action-token pairs in dynamic operators, writing $\langle \alpha, o\rangle$ and $[\alpha, o]$ instead of $\langle(\alpha, o)\rangle$ and $[(\alpha, o)]$.

Action-token pairs (α, o) and (α, o') are intended to differ only in the object (o or o') used to do the action α. The remaining context of use should be the same in both cases. The dynamic structure $[\![(\alpha, o)]\!]$ expresses the counterfactual causal structures that would be observed if α was done with the object o. It may be the case that, in the actual world w, o is not at hand so we may not actually use o to do α, but this fact is not reflected in $[\![(\alpha, o)]\!](w)$. Instead, the transition structure for $[\![(\alpha, o)]\!](w)$ expresses what would happen if, contrary to facts, o was available and one used it to do α.

Of course, for very many action-token pairs, "do m with o" will not make much sense. We may blow whistles, horns, pinwheels and maybe even noses but if o is a piano or a pumpkin, then $[\![(\mathsf{blow}, o)]\!]$ makes little sense. But this is not a problem for our semantics: we simply define $[\![(\mathsf{blow}, o)]\!](w) = \emptyset$ for cases in which o is a piano or pumpkin.

3.2. **Token- and type-fulfillment.** A functional ascription $f = (T, \alpha, \varphi)$ asserts that tokens of type T can be used to realize φ, namely by doing α with them. This suggests the following sense of fulfillment, called *token-fulfillment*.

Definition 3.1. Let $f = (T, \alpha, \varphi)$ be a functional ascription and o a token of type T. We say that *o (weakly/strongly) fulfills f in w* if (α, o) is a (weakly/strongly) sufficient means to φ in w.

Definition 3.1 gives the clear relationship between functions and means-end relation. A functional ascription expresses the expectation that tokens of the right type fit into a particular means-end relation.

Of course, when one ascribes a function $f = (T, \alpha, \varphi)$, he does not believe that *every* token of type T fulfills f. Rather, *normal* tokens of type T express f. We believe that the natural semantics of functions involves reasoning about normal tokens of a given type. Thus, we assume that for each type T, there is a set T^{normal} of normal tokens of type T. This set of normal tokens is not intended as a set of actual, existing artifacts, but rather fictional tokens that specify the normal behavior of tokens of that type. In some cases, every existing token of a particular type may be broken, but normal tokens for the type should not be broken, so normal tokens must be taken as useful fictions.. The elements of T^{normal} allow us to specify normal behavior via their dynamic transition systems: each $o \in T^{\mathsf{normal}}$ induces a transition system for (α, o) (for each $\alpha \in \Pi$) and hence represents an expected or normal behavior of elements of type T. (Because there may be several radically different designs for elements of T, we allow that there is a set of normal behaviors.)

How does one come to expectations regarding the behavior of 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 a natural tool in interpreting natural language functional ascriptions. We recognize that adding fictional entities to an account already heavy with counterfactuals is a controversial move, but we wish to present a semantics closely related to the natural language semantics for functions and we believe that this involves beliefs about normal artifacts of a given type. Thus, we postpone the thorny questions of where such beliefs come from and simply assume them.

The concept of normal tokens gives a new sense of fulfillment, namely *type fulfillment*.

Definition 3.2. Let $f = (T, \alpha, \varphi)$ be a functional ascription and $T' \leq T$. We say that T' *normally (weakly/strongly) fulfills* f *in* w just in case every normal $o \in (T')^{\mathsf{normal}}$ (weakly/strongly) fulfills f in w in the sense of definition 3.1.

What are the commitments of a functional ascription? If an agent ascribes f, then what expectation is he expressing? Certainly, we do not think that he believes every token of type T fulfills f but he surely must expect the *normal* tokens to do so. In other words, he believes that T itself *normally* fulfills f, at least in the weak sense.

Should he expect that proper subtypes T' of T also fulfill f? Evidently not. If so, one could reasonably infer from f another functional ascription, namely (T', α, φ), but this is not plausible in general. Whistles may be for making people jump, but this is not a function of high-pitched dog whistles which emit a sound that cannot be heard by human ears. (Note: if $T' \leq T$ implies $(T')^{\mathsf{normal}} \subseteq T^{\mathsf{normal}}$, then the fact that T fulfills f would imply that T' does too. We do not assume that normal tokens of a subtype $T' \leq T$ are also normal for the type T.)

4. Efficacy and malfunction

Efficacy is the capacity of a means to realize its end. It is one of the primary characteristics we compare when selecting a suitable means to our end. In artifactual functions, efficacy allows one to distinguish well-designed types from poorly designed types and properly functioning tokens from malfunctioning tokens.

4.1. **Efficacy.**

Propositional dynamic logic is designed to distinguish possible and necessary outcomes of actions, but it is not well-suited for representing efficacy. For efficacy, one must be able to compare the likelihood of different outcomes—or of the same outcome via different actions. Instead of non-deterministic transition systems, it is natural to use probabilistic transition systems so that we include an explicit measure of the probability that a particular world w' is the outcome of doing α in w.

In [6], we presented an initial development in this direction. The fundamental idea is that we can use probabilities to construct fuzzy propositions, as discussed in [2, 1]. In particular, we interpret $\langle\alpha\rangle\varphi$ in each world w as the *fuzzy* proposition, "α reliably realizes φ in w."

Put differently, in Section 2.1, $[\![\langle\alpha\rangle\varphi]\!]$ was interpreted as the set of worlds in which doing α *might realize* φ. The definition of "might realize" was in terms of the transition structure for α—it was a graph theoretic property. Now, we amend the transition structure by putting probability weights on the transitions. This changes our interpretation of $\langle\alpha\rangle\varphi$: it allows one to take $[\![\langle\alpha\rangle\varphi]\!]$ to be a *fuzzy* set of worlds. Namely, $[\![\langle\alpha\rangle\varphi]\!]$ is the fuzzy set of worlds in which α reliably realizes φ. Explicitly, the degree to which $w \in [\![\langle\alpha\rangle\varphi]\!]$ is calculated as

$$[\![\langle\alpha\rangle\varphi]\!](w) = \sum_{w' \in \mathcal{W}} [\![\alpha]\!](w)(w') \cdot [\![\varphi]\!](w'), \tag{4.1}$$

where $[\![\alpha]\!](w)(w')$ is the probability that doing α in w results in w' and $[\![\varphi]\!](w')$ is the degree to which φ is true in w'. We interpret the other connectives in terms of the so-called standard interpretation of fuzzy sets. See [6].

In this setting, the distinction between weakly and strongly sufficient becomes rather strained and we drop the latter. Instead, we focus on the truth degree of the proposition that α reliably realizes φ in w. We call this truth degree the *efficacy of α as a means to φ in w.*

4.2. **Token/token comparisons.** Efficacy provides a measure of the suitability of tokens to a given task as well. In the previous section, we gave a definition of token-fulfillment that was boolean: either a token o fulfilled a given function in a given world or it did not. With probabilities, the situation is subtler. Tokens may fulfill their functions to different degrees.

Definition 4.1. Let $f = (T, \alpha, \varphi)$ be a functional ascription, o a token of type T and $w \in \mathcal{W}$. *The efficacy of o with respect to f in w* is the value $[\![\langle\alpha, o\rangle\varphi]\!](w)$. In other words, the efficacy of o with respect to f in w is the truth degree of the fuzzy proposition "(α, o) is a reliable means to φ in w."

Thus, we can compare the efficacy of distinct tokens of type T with in a given world. If the efficacy of o with respect to f is greater than that of o', then o is better suited for the task as far as reliability is concerned. We call this *token/token comparison.* Of course, efficacy is not the only feature one uses in choosing which artifact to use for a given job. In particular, one often wants not only to achieve a particular goal but to avoid undesirable side effects (including high cost, damage to the environment, long time required to realize the goal and so on). Nonetheless, we are interested here only in comparing tokens with respect to efficacy and omit discussion of these other characteristics.

Function goals may be fuzzy predicates, conditions which are true to greater or lesser degree rather than just true or false. For instance, the natural function goal for a room heater is a warm room, but the proposition that a room is warm is vaguely defined. It is natural to represent this proposition as a fuzzy proposition. In this case, a heater which has a low probability of making **warm** true to very high degree may have higher efficacy than a heater which has high probability of making **warm** to a smaller degree.

4.3. **Type/type comparisons.** We often wish to make *type/type comparisons* in choosing which type of artifact is suitable for realizing our ends. To do so, we must have a measure of efficacy for a type analogous to our definition of type fulfillment (Definition 3.2). If $f = (T, \alpha, \varphi)$ and $T' \leq T$, then how should one calculate the *efficacy of T' with respect to f in a world w*? As before, we turn to our set of normal tokens for T', but we this gives us a set of efficacy values. There are several options for computing a single value from this set. The most obvious and plausible options are to take the mean or to take the infimum. We prefer the latter, particular since taking the infimum amounts to a generalized conjunction in terms of our fuzzy logic operations.

Definition 4.2. Let $f = (T, \alpha, \varphi)$ be a functional ascription, $T' \leq T$ and $w \in \mathcal{W}$. *The efficacy of T with respect to f in w* is the value

$$\inf \left\{ [\![\langle \alpha, o \rangle \varphi]\!](w) \;\;|\;\; o \in T'^{\mathsf{normal}} \right\},$$

the greatest lower bound of the efficacies of the normal elements of T' with respect to f in w.

4.4. **Token/type comparisons.** The final kind of comparison we wish to make is a *token/type comparison.* That is, how does a particular token perform compared to normal tokens of the same type? If the efficacy of o with respect to f is much greater than the efficacy of T with respect to f, then o is a particularly reliable token of type T when it comes to doing f. If the efficacy of o is much less than the efficacy of T, then o is unreliable. A token may be unreliable for any number of reasons, including poor design (or not designed for the function f at all), operating outside the normal conditions for the token, bad maintenance and so on.

Some of these conditions are reasonably called malfunction and others not. In particular, a token that is unreliable because it is badly maintained is a malfunctioning token but a token that is unreliable due to bad design is not malfunctioning. It is functioning as it was designed to function and the result is unfortunate. We may distinguish these two cases by appealing to the token's *primary type* T'. If we assume that the primary type is narrowly defined, so that all tokens of common primary type share the same basic design, then we may distinguish tokens behaving as designed from those that do not. Thus, while it is true that a dog whistle is an unreliable means to make people jump (indeed, its efficacy should be 0 with respect to this function), it is not less reliable than other dog whistles.

Since we assume that the contexts of use of (α, o) and (α, o') are the same, both o and o' will be operating in the same conditions. If these conditions are normal for o and o' has the same primary type as o, then these conditions should also be normal for o' (alternatively, the conditions should be equally extreme for each). Thus the comparison of the two tokens is "fair" in this regard.

Therefore, we regard token/primary type comparisons to yield a reasonable definition of one important kind of malfunction. If o is much less reliable in fulfilling f than every normal artifact of type T' in the same circumstances, it is reasonable to say that o is *malfunctioning.*

Definition 4.3. Let $f = (T, \alpha, \varphi)$ be a functional ascription and o an artifact with primary type $T' \leq T$. Let $w \in \mathcal{W}$ and suppose that

$$[\![\langle\alpha, o\rangle\varphi]\!](w) \ll \inf\left\{[\![\langle\alpha, o\rangle\varphi]\!](w) \;\mid\; o \in {T'}^{\mathsf{normal}}\right\}$$

(where $\ll$ denotes "is much less than"). Then we say that *o is malfunctioning with respect to f in w.*

There are certainly other ways in which a token may malfunction. A television that emits deadly radiation may produce a perfectly sharp and beautiful picture (fulfilling its function), but it is nonetheless malfunctioning. However, it is not malfunctioning in the sense of Definition 4.3. But we make no claims that our definition covers every kind of malfunction. We are content if it captures an important and common kind of malfunction, and this is evidently the case.

5. Conclusion

This work is part of our continuing project to make clear the relationship between means-end relations and artifactual function. We have chosen to do so via formal semantics, since formalization provides a precision and clarity often lacking from informal analyses. Of course, formal semantics can also obscure the very difficult parts of the analysis, which, in our case, arises from the use of counterfactual reasoning, often about how fictional "normal" tokens ought to behave. There is some work to do in understanding these problematic kinds of reasoning, and this work won't be completed by playing with more formal models.

Nonetheless, if our models do presuppose some controversial reasoning, it is because this reasoning seems analogous to the natural language meanings we want to clarify. If our understanding regarding the natural meanings is correct, then we have provided an analogous formal semantics and with it, the usual advantages found in formalization (including precision, rigor and relatively clear consequences) as well as the usual disadvantages (including the temptation to oversimplify and the narrowing of focus to a small list of features of efficacy, function and means-end relations to the exclusion of other relevant and interesting features). We believe that the advantages of a formal semantics can outweigh the disadvantages when carefully applied and that the controversial features of our particular semantics are an unavoidable consequence of mirroring the natural language meanings.

Our main contribution in this paper is the application of fuzzy dynamic logic to the theory of artifactual functions we first presented in [4]. This provides a measure of efficacy for artifacts that allows token/token, type/type and token/type comparisons. That last allows a simple definition of a rudimentary but very common form of malfunction.

References

[1] Petr Hájek, Lluís Godo, and Francesc Esteva. Fuzzy logic and probability. In *Proceedings of the Eleventh Annual Conference on Uncertainty in Artificial Intelligence (UAI-95)*, Montreal, Quebec, Canada, August 1995.

[2] Petr Hájek, Lluís Godo, and Francesc Esteva. Reasoning about probability using fuzzy logic. *Neural Network World*, 10:811–824, 2000.

[3] 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.

[4] Jesse Hughes. Means-end relations and artifactual functions: a sketch. Presented at the *Norms, Reasoning and Knowledge in Technology* workshop, Boxmeer, the Netherlands., 2005.

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

[6] 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.

[7] 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.

[8] 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.

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

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

(Sjoerd Zwart) Technical University of Delft, Delft, The Netherlands *and* Technical University of Eindhoven, Eindhoven, The Netherlands
E-mail address: `S.D.Zwart@tbm.tudelft.nl`