

Verb Classes as Evaluativity Functor Classes

Pranav Anand and Kevin Reschke

University of California, Santa Cruz

Santa Cruz, CA, U.S.A.

{panand, kreschke}@ucsc.edu

1 Evaluation and Composition

In the past decade, the semantics of evaluative language – a hallmark of the generative semantics debates (see Fillmore (1985)) – has received renewed attention in both formal and computational linguistics, though in different guises: appraisal (Martin and White, 2005), expressive content (Potts, 2005), sentiment (Pang and Lee, 2008), and valuation (Jackendoff, 2007). This work has largely focused on evaluativity at the lexical level, i.e., the evaluative stance encoded in a particular word, or holistic judgments of phrasal level stance (e.g., what is the author’s sentiment in this sentence?). Largely absent from these discussions is the role of compositional interpretation in connecting lexical evaluativity to phrasal level stance. When compositionality has been investigated, it is in the explication of polarity preservers/inverters (e.g., copulas/negation) or the attempt to extend such operators to the event domain (Moilanen and Pulman (2007), Nasukawa and Yi (2003)).

While preservation or inversion may be valid for logical operators, we argue that evaluations of events are more complex. As the determination of truth conditions for event predicates relies on a knowledge of the event participants, so too evaluative stance towards an event is a product of the evaluative stances an assessor bears towards the predicate’s participants. We thus propose that, for the purpose of determining evaluative stance towards events, verbs and other predicates of events should be analyzed as functors (or, mappings) from n -ary evaluative tuples to an evaluative value. The task of determining event evaluativity thus reduces to the problem of determining the particular mapping from n -ary tuples to evaluative values that a given verb induces. This problem is complex (and potentially context-dependent), but we claim that verbal predicates fall into evaluative functor *classes* based on

their lexical entailments about their arguments. In particular, we identify three entailment types as prominent predictors of functor class: POSSESSION by a participant, EXISTENCE of a participant, and AFFECTEDNESS of a participant. We justify these claims by human annotation of a generated corpus as well as coding of naturally occurring text in the Gigaword corpus (Parker et al., 2009).

The paper is organized as follows. In section 2, we outline the theoretical machinery behind treating event-level evaluativity in terms of n -ary functors. Section 3 discusses how verbal predicates show commonalities based on the three entailment types above; it additionally demonstrates that the evaluative stance towards verbs of change of state is determinable in terms of the result state alone. Finally, section 4 presents our empirical assessment of the predictions of section 3. Section 5 concludes with a discussion about how to extend the machinery in section 2 to deal with more subtle differences in evaluative stance.

2 Arity of Evaluativity

Wilson (1975) discusses a contrast in *verbs of withholding* (1); this is the product of contradictory presuppositions regarding the desires of the protagonist (Gazdar, 1979). Events of *deprivation* differ from those of *sparing* in terms of whether the withholdee desired the outcome in question:

- (1) a. She deprived him of a day at the seaside.
- b. She spared him a day at the seaside.

But there is a further evaluative component distinguishing (1a) and (1b) – in a sense, (1b) is, if not infelicitous, somehow pragmatically marked, given the intuition that speakers have a tendency to describe events as instances of sparing insofar as they are positive outcomes from the speak-

ers’ perspectives.¹ And, indeed, in the human-annotated MPQA Subjectivity Lexicon (Wilson et al., 2005) *deprive* and *spare* are specified as strongly negative and positive, respectively. In contrast, the semi-supervised SentiWordnet lexicon (Esuli and Sebastiani, 2006) marks both terms as neutral (“objective”). This surprising disagreement, we argue, is the result of the fact that sparing/deprivation events are not *en masse* positive or negative; nor are they non-evaluative. Rather, particular events of sparing or deprivation are modulated by the author’s stances towards the event participants, as with the protagonists in (2). The deprivation of someone one feels positively towards (e.g., an ally) is typically evaluated as negative, and vice versa towards those one feels negatively associated with (e.g., an enemy):

- (2) a. My {ally, enemy} was deprived shelter.
 b. My {ally, enemy} was spared a dangerous mission.

A similar effect is noted by Nasukawa and Yi (2003) for predicates such as *have/lack* and Moilanen and Pulman (2007) for *fail*. The approaches in these papers consider only the direct object of a verb in determining overall evaluation, meaning that *lacking* and *failing* events are uniformly negative. Empirically, however, evaluation of an event of *lacking shelter* depends on the perceptions of the grammatical subject in an assessor’s mind: an enemy’s lack of shelter is arguably perceived as better than an ally’s. Thus, compositional computation of the evaluation of an event requires the evaluations of all arguments, not merely the internal ones. More precisely, if \mathcal{E} is the domain for evaluativity, an n -ary verb V induces an n -ary evaluativity functor:

$$(3) E_V : \mathcal{E}^n \rightarrow \mathcal{E}$$

Assuming that $\mathcal{E} = \{-, +\}$, the contrasts between withholding and possession verbs may be represented by the functions given in Table 1. As noted above, E_{have} and E_{lack} are opposites (equivalent to XNOR and XOR, respectively). However, $E_{deprive}$ and E_{spare} not. Rather, they are each partial functions of $E_{withhold}$, representing

¹The preferences of withholder give rise to an additional, distinct evaluative stance the withholder bears towards the event of withholding. The derivation of such event participant evaluativity is beyond the scope of this paper, though it too is expressible via the general machinery we introduce.

x	y	E_{hve}	E_{lck}	E_{wthld}	E_{dprv}	E_{spr}
+	+	+	-	-	-	#
+	-	-	+	+	#	+
-	+	-	+	+	+	#
-	-	+	-	-	#	-

x have/lack y
 a withhold/deprive/spare x of y

Table 1: Functors for verbs of possession and withholding.

the fact that *deprive* is an infelicitous description when event in question is assessed positively and *spare* infelicitous when the event is assessed negatively. Note additionally, that the agent for verbs of withholding does not affect event-level evaluativity (and thus $E_{withhold}$ and E_{lack} produce identical outcomes for the same x and y despite an arity difference), as one can determine the event-level evaluativity even in the absence of the assessor’s stance towards the agent. This is not a universal characteristic of verbal functors. For example, E_{visit} , $E_{entertain}$, and E_{meet} are all sensitive to agent evaluation: the identity of the one who visits, entertains, or meets is just as crucial as the identity of the one visited, entertained, or met for computing how one would feel about the event.

3 Evaluativity and Entailment Class

We propose that the contrast in sensitivity towards agentivity between $E_{withhold}$ and E_{visit} is not accidental. Rather, it, as well as the general character of a verb’s evaluativity may be understood on the basis of lexical entailments of the verb. Verbs of withholding, for example, are result verbs (Levin and Rappaport, 1995), entailing the existence of an unspecified causal event as well as a consequent lack of possession. Correspondingly, when felicitous, they are identical to E_{lack} . We propose that this insensitivity to the agent follows from a more general principle on changes of state (excepting AFFECTEDNESS, below):

- (4) CHANGE IRRELEVANCE: The evaluation of a change of state is equivalent to the evaluation of the end state.

CHANGE IRRELEVANCE thus correctly predicts that the evaluativity of verbs of profit and loss (e.g., *profit from*) are insensitive to the causing event.

x	E_{xst}	E_{nxt}	E_{pstv}	E_{ngtv}
+	+	-	+	-
-	-	+	-	+

x is existing/non-existent
 x has a positive/negative property

Table 2: Functors for the basic entailment states of EXISTENCE and AFFECTEDNESS.

A direct consequence of CHANGE IRRELEVANCE is that determining the evaluation of a complex event may be reduced to determining the evaluation of a state. While the factors affecting evaluation of a state are complex, we propose the following list of characteristic entailments:

- (5) a. EXISTENCE: A participant has/lacks existence.
- b. POSSESSION: One participant has/lacks possession of another.
- c. AFFECTEDNESS: A participant has a positive/negative property.

Thus, the evaluativity functor of an event whose result state is one involving POSSESSION will either be identical to E_{have} or its negation (E_{lack}). A similar situation holds for events resulting in states involving EXISTENCE and AFFECTEDNESS, whose corresponding functors are provided in Table 2. In particular, we predict that the change of state verb classes of creation/destruction, gain/lose possession, and benefit/injury will correspond to the basic functors in Tables 1 & 2 as follows: a) creation events result in states involving EXISTENCE, and hence their evaluation will be identical to $E_{existing}$; b) gain events result in states involving POSSESSION, and hence their functors will be determined by E_{have} ; finally, benefit events result in a positive property, and hence will follow $E_{positive}$. The antonymic classes of each of these will have functors determined by $E_{non-existent}$, E_{lack} , and $E_{negative}$, respectively. The predictions are summarized in Table 3.

Although there are class-based commonalities, it is important to acknowledge lexical idiosyncrasies that are directly evaluative. In (6), where *abuse* seems infelicitous regardless of the injury to a negatively evaluated individual, unlike *injure*. This contrast reflects the negative evaluation lexically encoded in *abuse*, which conflicts with the strongly positive evaluation of *defenseless child* –

VERB CLASS	RESULT STATE	FUNCTOR
creation	existence	$E_{existing}$
destruction	existence	$E_{non-existent}$
gain	possession	E_{have}
loss	possession	E_{lack}
benefit	affectedness	$E_{positive}$
injury	affectedness	$E_{negative}$

Table 3: Predicted functors for 6 change of state verb classes.

abuse events are inherently bad, and it is difficult to think of defenseless children committing them.

- (6) The defenseless child {#abused, injured, ?tortured} the monster.

We represent verbs like *abuse* as constant functors over their argument tuples. In large part because they are constant (and hence, most expressible in terms of simple lexical inferences), these functors have received the greatest attention in both the sentiment and expressive content literature. While that attention is deserving, the claim here is that they are part of a larger combinatoric system which may be obscured by focusing only on limiting cases.

4 Empirical Assessment

We conducted two empirical assessments of the predictions summarized in Table 3: an annotation study on constructed sentences and a corpus investigation of the arguments correlating with overt markers of evaluativity.

4.1 Annotation Study

In order to test the claims of the theory directly while controlling for variables of interest, we conducted an annotation study on 6480 constructed sentences. The sentences combined 48 predicates from the injury, benefit, destruction, creation, and transfer classes in FrameNet (Ruppenhofer et al., 2005) with 18 participants bearing canonical positive, neutral, and negative evaluations. Example sentences are provided in (7), and the complete list of verbs and nominals used are listed in Tables 5 and 6, respectively.

- (7) a. The {hero, man, villain} {comforted, assaulted} the {child, monster}.
- b. The {hero, man, villain} {assembled, defaced} the {cathedral, building, torture chamber}.

Twelve annotators were instructed to indicate overall evaluative stance (positive, negative, neutral) the author had towards the event described in each sentence. Each annotator had 10% overlap with each of two other annotators (i.e., each annotated 486 sentences uniquely, 54 with one annotator, and 54 with another).

An ordinal logistic regression was fit to the data, using the participant (agent, object, goal) polarities and verb classes, as well as random intercepts for annotator and verb. Significant interactions are as follows: Both positive affected arguments showed significant interaction with the injury/destruction classes ($p < 2.22 \times 10^{-16}$); inter-annotator agreement replicated this (Cohen’s $\kappa = 0.92$); that is, killing was judged more positive when the entity losing existence was an enemy and judged more negative when it was an ally. Neutral affected arguments showed behavior similar to positive arguments ($p < 2.22 \times 10^{-16}$), suggesting a principal of charity with respect to events of harm. Transfer verbs showed a weak sensitivity to possession. Table 4 illustrates the correlation between object evaluativity and annotated event evaluativity for dyadic verbs; the cells predicted by Table 3 are in bold. While gaining possession for positive possessors (e.g., a hero gaining a valuable watch) tracked the value of the object ($p < 1.34 \times 10^{-2}$, $\kappa = 0.89$), negative possessors showed less inter-annotator agreement ($p < 0.09$, $\kappa = 0.68$), with a tendency for neutral evaluation.

We also compared four putative constant negative functors: E_{abuse} , $E_{assault}$, E_{murder} , and E_{rape} . E_{murder} showed no difference from E_{kill} , but the remainder showed a strong negative evaluation in all +agent or +patient cases ($\kappa = 0.92$). However, in cases with a negative patient and a non-positive agent, there was less consensus ($\kappa = 0.68$), with a slight preference for a positive valuation. We suggest that the additional inference in these constant functors is negativity toward the agent. When the agent is otherwise positive, the negative response expresses disappointment in the agent, but not the outcome. When the agent’s behavior is deemed less relevant (e.g., cases where the agent is otherwise negatively judged), the basic characteristics of the injury class reappear.

4.2 Corpus Evaluation

To assess whether the predictions in Table 3 are attested in naturally occurring text, we searched the

OBJ	EVENT	BEN/CREAT	INJ/DESTR
-	-	449	125
-	n	128	74
-	+	73	449
n	-	17	563
n	n	161	68
n	+	470	19
+	-	28	588
+	n	105	16
+	+	518	44

Table 4: Dyadic Verb Annotations

one billion word Gigaword corpus (Parker et al., 2009) for ten target verbs. Event-level evaluativity was approximated by considering predicates extracted from three sentence frames which indicate stance directly: emotive factives (*X was happy/sad that ϕ*), polar adverbs (*Thankfully/Unfortunately, ϕ*), and promises/threats (*X promised/threatened to ϕ*). The target predicates were verbs of benefit (*help, cure, protect, reward*), injury (*kill, murder, execute, assault, injure*), and destruction (*destroy*). This procedure yielded approximately 6200 matches. Of these, we selected a sample ($n = 690$) for manual inspection.

In order to test the accuracy of our theory, we examined matches for the evaluativity of the verb object. According to the predictions of Table 3, object evaluativity should be deducible from the verb class and event-level evaluativity; thus, if the match in question involves a verb of destruction embedded in a positive frame (e.g., *Thankfully, they destroyed . . .*), we predict that the object should be perceived as negative by the author. Table 7 summarizes our judgments on object evaluativity (positive, or +obj and negative, or -obj) across the three verb classes and evaluative contexts (positive, or +c and negative, or -c); the cells predicted by Table 3 are in bold. Example positive, negative, and neutral objects are given in Table 8. Objects with no obvious polarity and non-entity objects (e.g. *nobody*) were marked ‘other.’

In general, the predictions of the theory are attested. The evaluativity of a benefit clause tends to match that the participant, while evaluativity of an injury clause tends to oppose it.² Likewise,

²The injury/+obj/+c counts are inflated. Half arose from context polarity misclassification; in remaining cases, context polarity was positive because it was contrastive – e.g., *Fortunately, the official was only injured in his hand.*

INJURY	abused, arrested, assaulted, beat up, executed, injured, insulted, killed, murdered, raped, scratched, undermined
BENEFIT	aided, comforted, cured, educated, helped, pardoned, protected, resuscitated, rewarded, strengthened, supported
DESTRUCTION	broke apart, crumbled, defaced, demolished, desecrated, destroyed, eliminated, fractured, obliterate, pulverized, shattered, wrecked
CREATION	assembled, brought forth, created, fabricated, faked, fashioned, generated, made, pieced together, produced, sculpted, synthesized
TRANSFER	accepted, acquired, bought, delivered, donated, gave, passed on, procured, sold, stole, surrendered, took

Table 5: Annotation Study Verbs

ANIMACY	POSITIVE	NEUTRAL	NEGATIVE
animate	child, hero, friend	co-worker, man, middle-aged individual	enemy, monster, villain
inanimate	cathedral, lovely silks, universal prosperity	building, cloth, general silence	torture-chamber, disgusting rags, universal poverty

Table 6: Annotation Study Nominals

	benefit		injury		destruction	
	+c	-c	+c	-c	+c	-c
+obj.	49	0	18	216	2	20
-obj.	1	8	15	2	39	1
other	27	10	150	115	7	10

+c: *thankfully; happy; promised*
+c: *unfortunately; sad; threatened*

Table 7: Token frequencies by verb-class, object, & context in Gigaword sample ($n = 690$).

the evaluativity of a destruction clause generally opposes the polarity of the the participant’s existence. For example, three-fourths of -obj/+c cases were instances of disarmament. The existence of weapons stockpiles is negative, thus their destruction (which results in loss of existence) is positive.

5 Further Directions

We have argued that event-level evaluativity should be considered in terms of verbal functors, and have shown that this allows us to capture entailment class based generalizations that appear distributionally valid. Given the sensitivities of the transfer annotations to not only polarity of evaluativity towards objects, but also the domain of evaluation, it would be wise to model the multiple di-

obj.	examples
+	(indexicals) <i>us</i> ; (concepts) <i>free speech</i> ; (allies) <i>French partisans</i> ; (valued objects) <i>treasures</i> ; (esteemed individuals) <i>champion</i> .
-	(concepts) <i>incompetence</i> ; (enemies) <i>bin Laden</i> ; (illegal possessions) <i>field of poppy</i> ; (weapons) <i>Iraqi Scuds</i> .

Table 8: Example judgments of evaluativity in verb objects

mensions of valuation that may be relevant within a judgment. Thus, the ontologies in both Jackendoff (2007) and Martin and White (2005) distinguish between ethical, teleological, and aesthetic evaluations. Hence, a gain of something aesthetically valuable may not give rise to the same evaluative intuition as something teleologically valuable. While it is straightforward to implement such distinctions via a multidimensional evaluative domain (i.e., $\mathcal{E} \cong \{-, 0, +\}^n$), the important and difficult work will be to determine how mappings across these finer-grained domains correlate with the semantic properties of the verbs they correspond to.

References

- Andrea Esuli and Fabrizio Sebastiani. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *Proceedings of LREC-06*.
- Charles Fillmore. 1985. Frames and the semantics of understanding. *Quaderni di Semantica*, 6(2):222–254.
- G. Gazdar. 1979. *Pragmatics: Implicature, Presupposition and Logical Form*. Academic Press.
- Ray Jackendoff. 2007. *Language, consciousness, culture*. MIT Press.
- B. Levin and M. Rappaport. 1995. *Unaccusativity*. MIT Press.
- J. R. Martin and P. R. R. White. 2005. *Language of Evaluation: Appraisal in English*. Palgrave Macmillan.
- Karo Moilanen and Stephen Pulman. 2007. Sentiment composition. In *Proceedings of RANLP 2007*.
- T. Nasukawa and J. Yi. 2003. Sentiment analysis: Capturing favorability using natural language processing. In *Proceedings of the 2nd international conference on Knowledge capture*.
- Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135.
- Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2009. English gigaword, 4th ed. Technical report, The LDC Corpus Catalog.
- Chris Potts. 2005. *The Logic of Conventional Implicature*. Oxford University Press.
- Josef Ruppenhofer, Michael Ellsworth, Miriam R. L. Petruck, and Christopher R. Johnson. 2005. Framenet ii: Extended theory and practice. Technical report, ICSI Technical Report.
- T. Wilson, J. Weibe, and P. Hoffman. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of EMNLP-05*.
- D. Wilson. 1975. *Presupposition and Non-Truth Conditional Semantics*. Academic Press.