Gen-Meta: Generating metaphors using a combination of AI reasoning and corpus-based modeling of formulaic expressions

Andrew Gargett
School of Computer Science
University of Birmingham
Birmingham, United Kingdom
A.D.Gargett@cs.bham.ac.uk

John Barnden
School of Computer Science
University of Birmingham
Birmingham, United Kingdom
J.A.Barnden@cs.bham.ac.uk

Abstract—Metaphor is important in all sorts of mundane discourse [19], [7]: ordinary conversation, news articles, popular novels, advertisements, etc. This presents a challenge to how Artificial Intelligence (AI) systems understand inter-human discourse (e.g. newspaper articles), or produce more natural-seeming language, as most AI research on metaphor has been about its understanding rather than its generation. To redress the balance towards generation of metaphor, we directly tackle the role of AI systems in communication, uniquely combining this with corpus-based results to guide output to more natural forms of expression.

Keywords-Artificial intelligence; Cognitive science; Natural language processing; Interactive system;

I. INTRODUCTION

Working out why a speaker might choose to use metaphor is very much an open question. By way of attempting to answer the more tractable question of why, after having decided to express things metaphorically, a speaker may choose one metaphorical expression over another, we have formulated a way of meeting the challenge of generating metaphor. We describe in this paper an approach to metaphor generation which uniquely combines reasoning with data-oriented techniques, potentially accounting not only for more conventional forms of metaphorical expression, but also novel extensions to established metaphor. Our approach, which we have dubbed "Gen-Meta",[1] provides a natural language generation (NLG) front-end for a state-of-the-art metaphor processing framework, ATT-Meta [4]. Based on patterns of metaphorical expressions mined via cutting-edge methods for discovering metaphor in natural language, our system links an AI-derived conceptual level to a corpus-derived linguistic level, thereby generating an appropriate expression for the target metaphorical meaning.

While still at the prototype stage, the system aims to coordinate in modular fashion the interaction between three existing frameworks: ATT-Meta, Embodied Construction Grammar [13], and Dynamic Syntax [23]. This paper reports on the development of this approach, as well as some initial findings.

[1]http://www.cs.bham.ac.uk/~gargetad/genmeta-index.html

II. GENERATING METAPHOR

A. Overview of Natural Language Generation (NLG)

Producing natural language utterances involves numerous choices about *what* to say and *how* to say it, the central problem of Natural Language Generation (NLG), the study of the use of computational techniques for adequately generating strings of natural language, from deciding the basic content of the utterance ("what"), through to determining how to resolve forms of reference, planning discourse structure, and realising appropriate words and their combinations [9]. Regarding approaches to modeling such decision-making in NLG, there seem to three broad classes:

1) Template-based: generating via predefined slots of a template
2) Pipeline-based: stepping through decisions about what to say and how to say it, like a sort of production line
3) Learning-based: adapting via a form of learning, to particular domains and/or users [21]

The first class is the most common, while the second and third are closest to our own, combining "knowledge intensive" approaches to metaphorical extension through inferential processing, with more data-oriented approaches crucial for modelling the wide variety of forms possible for expressing oneself metaphorically. While much (if not all) NLG takes actual usage into account, we directly incorporate patterns of metaphorical expression found in corpora, to produce texts more directly reflecting language use.[2]

B. Key issues in generating metaphor

Producing expressions that are in some sense "more natural" is a key aim in NLG, so that phenomenon as ubiquitous in everyday human communication as metaphor (e.g. [19], [7]) should be a priority within NLG, one would think. Yet, while there is a recognisable body of research

[2]We consider our approach to address the so-called "knowledge bottleneck" [20], by tackling the immense amount of (lexical, morphological, syntactical, etc) knowledge required to generate natural language. Of course, the challenge of building the resources required for work on large-scale corpora is not without problems (see e.g. [2]).

on the natural language understanding (NLU) of metaphor, much less research has been devoted to generating metaphor [15]. Both NLU and NLG face many of the same issues when modeling more general cognitive phenomena such as metaphor, which apparently require solutions to substantial parts of core artificial intelligence. While much NLG research has assumed content to be given, enabling a focus on *how* to realise such content in actual linguistics strings, generating metaphor is very much about modelling content, requiring new ways of thinking and new techniques, or at least fresh ways of using already established techniques.

Opting to metaphorically express an idea implies strategically choosing this form of expression over another, such as in emotionally charged encounters [7]. No current NLG system can generate metaphor in a way that is contextually appropriate, as humans do all the time when communicating with one another, yet there has been a variety of previous attempts at generating metaphor, some of which we consider below.

C. Inferential approaches to modelling metaphor

Past approaches to metaphor generation based on rule- or constraint-based methods include [22], [18] and [15]. We have chosen to compare two exemplary approaches.

1) MIDAS: The "computational theory of metaphor" proposed by [22] yielded the MIDAS system, having the capacity to both understand and generate metaphors in the narrow domain of the UNIX operating system, for example:

(a) How can I *get into* mail?
(b) How can I *get out of* emacs?
(c) How can I *kill* a file?

The italicised items are metaphorical, since a direct reading of verbs is non-sensical in these contexts (e.g. killing a file here cannot mean, directly, ending the life of something that is alive, but it can mean, less directly, ending a computer process, and even deleting some item of information stored on a computer). For [22], many such metaphors are largely conventional (see [19]), reflecting larger conceptual classes of which they are members (other examples being Argument-Is-War, Time-Is-Money). MIDAS stores such conventional metaphors in its lexicon, this being an instance of a knowledge-rich approach to metaphor processing, whereby understanding a particular (conventional) metaphor is largely a matter of being able to access the entry for that metaphor.

[12] points out that while MIDAS is apparently over-specialised to the domain, the coverage of MIDAS is certainly impressive. Since MIDAS, there have been few knowledge-rich approaches with a substantially greater coverage. Here, as elsewhere in NLG, how to model content adequately has been the chief obstacle to progress.

2) ATT-Meta: Our approach to metaphor employs Barnden's ATT-Meta system, a state-of-the-art AI system for modelling metaphor as reasoning-by-simulation [4], whereby those aspects of a metaphorical expression, like *How do I get out of emacs?*, which are clearly not about reality, on a par with *How do I get out of this house?*, are dealt with in a distinct mental space, a so-called metaphorical pretence cocoon, wherein reasoning about such propositions and inferences can be kept separate from propositions and reasoning about reality.

While ATT-Meta has until now been used for metaphor understanding, it turns out to be fairly straightforward to extend it to generation, due to a novel feature of the system, namely its ability to transfer information from target-to-source, as well as in the more usual source-to-target direction. The reversed transfer is held to be crucial for the understanding of some metaphor, but can be adapted also for generation. As we noted above, while in its day, MIDAS represented an innovation, it was somewhat specialised to the task it was built for, whereas ATT-Meta presents a number of interesting features allowing greater generalization, yet at the same time it retains a certain specificity in its operation which provides a basis for contextualised reasoning.

While ATT-Meta's reverse use of mappings can be readily deployed as a part of the process of generating metaphorical utterances, we need some way of causing a reverse use to happen, bearing in mind that ATT-Meta works entirely by backward-chaining reasoning, or goal-directed reasoning, a form of reasoning commonly used in rule-based systems. So we need either to add a forward-chaining capability to ATT-Meta (so that, given a reality-space representation, reasoning would step forwards into the pretences space across a mapping), or to emulate such forward chaining by constructing a certain type of rule of the following intuitive form:

(R1) IF reality situation X corresponds to pretence situation Y, and Y holds THEN can-state(Y).

where X and Y are variables. Here we are helped by a distinctive feature of ATT-Meta mappings, in that they have the form:

(R2) IF guard-condition G holds THEN real-U corresponds to pretend-V.

Thus, rule **(R2)** would only pick up those mappings whose guard conditions are satisfied. Then crucial to understanding the claim that Bill *has* a cold, is presuming a cold to be a physical, and hence possessable, object, permitting only mappings whose guards (antecedents in these conditional rule forms) are satisfied by that presumption.

Consider the following examples utterances expressing the metaphorical notion of a cold as a physical thing (including representations of these in ATT-Meta terms):

(1) *Bill has a cold* → the_episode(being_infected, bill, cold)
(2) *Bill gave Bob a cold* → the_episode(transfer, bill, bob, cold)

Here, "the_episode" refers to an instance of some general event, which is "being_infected" in (1). For generation we

focus on the right-hand side of (1) and (2), and assuming someone's cold can be regarded as a physical object, then only those mappings whose guards are satisfied by that condition will be picked up. Moreover, the satisfaction of the guard is relative to specific entities and facts, thereby causing, for example, some very specific instance of a mapping to hold, rather than having this hold of anyone's cold. So, John's having his cold is deemed to correspond to John's physically-possessing his cold (as in (1) above). Thus, rule **(R1)**, by the very fact of picking up on such specific mappings instances, will already at least partially instantiate Y to that specific situation.

Our work on incorporating rules such as **(R1)** into our system has revealed significant technical difficulties, resulting from the fact that ATT-Meta has a way of open-endedly generating variants of mappings (see discussion of so-called view neutral mappings in [3]). There is a danger, therefore, that **(R1)** would cause prolific over-generation, a problem we will address in future work.

D. Data-driven approaches to modelling metaphor

Mounting evidence suggests people frequently employ formulaic language to express figurative meanings such as metaphor (e.g. [11], [16]). The approaches of [10] and [7] are central to our approach to this. [10] prefers a corpus-driven approach to modelling metaphor, aiming to determine taxonomies directly from the corpora concerned. [7] presents evidence of metaphor "tuning" during reconciliation talks within the context of acts of terrorism, in particular the way in which someone can increase the impact of their contribution by employing metaphor to describe the effect on their lives of another's actions. Such evidence has inspired our own use of corpus studies to guide generation toward conventionalised forms of expressing metaphors (for further exemplary approaches, see also [1], [25])

III. GEN-META: COMBINING INFERENTIAL AND DATA-ORIENTED APPROACHES TO MODELLING METAPHOR

A. Overview

Gen-Meta attempts to explicitly combine inferential and data-oriented modelling of metaphor, by chaining together modules which, as individually validated approaches to modeling language, bring advantages to the resulting system as a whole. We have discussed ATT-Meta above, and the remaining modules are:

Embodied Construction Grammar (ECG): ECG is a language understanding (but not generation) system having aspects highly congenial to metaphor, and of interdisciplinary significance [13].[3] ECG models the links between the conceptual level, represented as interconnected schemas, and the linguistic level, represented as interconnected constructions. Schemas consist of "roles" together with constraints on these roles. Recalling examples (1) and (2): there is a schema for the concept of somebody realising a **transferer role** TRANSFERRING something to somebody else realising a **transferee role**. ECG's schemas are strongly geared towards conceptual representations, have broadly the same orientation as ATT-Meta's representations. Constructions will be familiar from work within Cognitive Linguistics [8], for example, the English verb *give* may employ a ditransitive construction having three constituents, **subject**, **direct object** and **indirect object** (see examples (1) & (2)) – ECG formally specifies the ordering constraints operating over such constituents, as well as the linkage between the form and the meaning of the construction.[4] Being more flexible than traditional grammatical representations, constructions can better model the diversity and flexibility of expression types, from single words to multi-word expressions (MWEs). Constructions also usefully represent conventional metaphorical form and meaning, with MWEs rather than individual words frequently having metaphorical meanings.[5]

Dynamic Syntax (DS): DS is an implemented computational approach to both generation and understanding that is specially geared to dialogue [23]. DS offers advantages over other approaches in that it is one of the few generation systems that are: (1) grammar-based, (2) fully incremental, and (3) context-dependence. DS provides a fully incremental parsing model, with update modelled as transitions between succeeding parse states, essentially, enrichment of partial tree structures. Parsing is then the sequence of pairings of natural language strings of terms s with the logical formula R representing the semantic structure of those terms:

$$(3) \quad \{\langle s(i), \text{R}_i \rangle, \langle s(i+1), \text{R}_{i+1} \rangle, \dots\}$$

Thus, R_i results from parsing $s(i)$. More generally, these successive parse states are modelled as triples of (i) (partial) tree structures, (ii) representations of partial tree structures in the formal language, and (iii) procedures for enacting transitions between pairs of partial trees. These trees represent semantic information, with syntax and lexical entries encoding instructions for building such trees.

Our approach develops a hybrid of DS and ECG, incorporating ECG-style form-meaning mapping, in particular between constructions and schemas, but extending ECG with parsing/generation techniques from DS.

B. Encoding the content: ATT-Meta logical forms, ECG schemas, DS Goal trees

Consider Figure (1), showing how the sentence "Bill gave Bob a cold" is differently represented by each module: starting with the Goal Logical Form, where ATT-Meta generates

---

[3]ECG has also recently been implemented [6].

[4]Meaning represented by an external ontology, see [13] for details.

[5]It should be noted that our presentation of the ECG formalism here is somewhat simplified, for ease of exposition.

GOAL LOGICAL FORM: $the\_episode(transfer, bill, bob, cold)$

GOAL SCHEMA:

Predication

scene:

Transfer
$agent$:
$theme$:
$recipient$:
$means$:

schema:

Give
$giver$
$givee$
$gift$

$agent$ **Referent**
resolved_reference: Bill'

$theme$ **Referent**
$recipient$ category: Cold'

Referent
resolved_reference: Bob'

GOAL TREE:

$Ty(t), Fo(Give'(Bob', Cold')(Bill'))$

$Ty(e),$
$Fo(Bill')$

$Ty(e \rightarrow t),$
$Fo(\lambda z(Give'(Bob', Cold')(z)))$

$Ty(e),$
$Fo(Cold')$

$Ty(e \rightarrow (e \rightarrow t)),$
$Fo(\lambda yz(Give'(Bob, y)(z)))$

$Ty(e),$
$Fo(Bob')$

$Ty(e \rightarrow (e \rightarrow (e \rightarrow t))),$
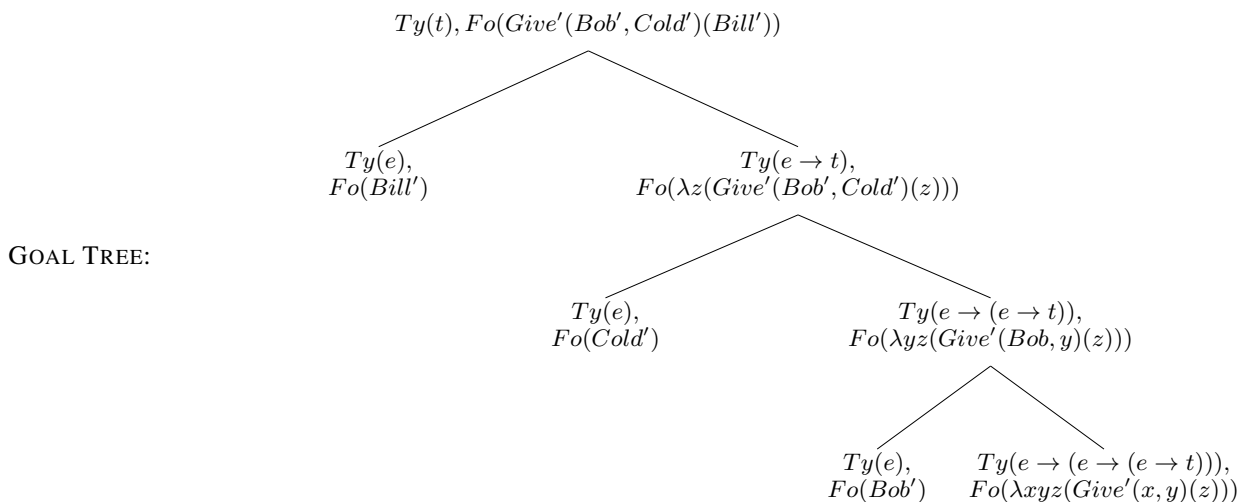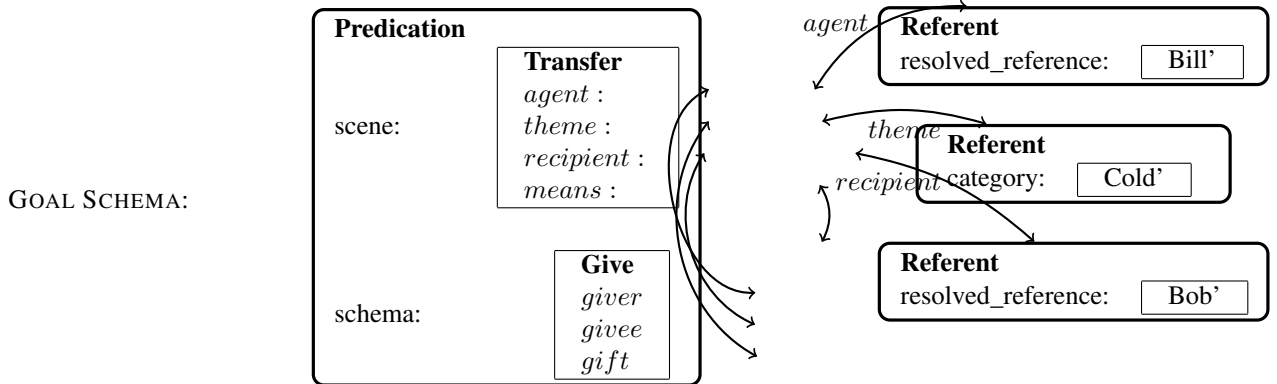$Fo(\lambda xyz(Give'(x, y)(z)))$

Figure 1: Overview of Gen-Meta: different representations of the goal content

reversal of the mapping ILLNESS-as-PHYSICAL OBJECT in the target-to-source direction (recall section (II-C2) above), this then being translated into the appropriate Goal ECG Schema. At this point, the missing ECG generation component presents a gap in our NLG pipeline, to be filled by the DS generation mechanism via a two-stage approach: stage 1 translates the Goal Schema into a Goal Tree and switches to the DS generation mechanism, thereby generating candidate DS trees to translate into a set of candidate constructions; stage 2 runs the mapping between conceptual and linguistic levels unique to ECG, reversing this in order to map the initial Goal Schema to the optimal construction from the set of candidates in stage 1.

For illustration, consider how to translate from Goal Logical Form to Goal ECG Schema. While the former is reducible to first-order predicate logic (FOL), the latter is reducible to (a subset of) feature logic (e.g. attribute-value matrices), and a well established result is that feature logics can be described in FOL [17]. However, note that an equally appropriate English utterance for the content expressed by "Bill gave Bob a cold" could be "Bill passed a cold to Bob", or even "Bill foisted a cold on Bob" (although the first is favored as a more conventionalized expression). So, Figure (1) in fact specifies content that can be realised by

a set of possible utterances, so that abstracting over the embedded schema component, here occupied by **Give**, yields a more general, albeit incompletely specified representation of this content. Other possible embedded schemas we could use here include (employing the notation of [17]):

(4) $\forall\ x,\ y,\ z\ pass'(x,y,z)\ \rightarrow\ x = person\ \&\ y = person$

In addition, the logical form of the more abstracted topmost **Predicate** schema (unspecified by information such as that in (4)) from Figure (1) is:

(5) $\forall\ x,\ y,\ z\ X,\ Y\ scene.X\ \&\ schema.Y\ \rightarrow X = transfer'(x,y,z)\ \&\ Y$

The logical form in (5) can be saturated by taking the logical form of (4) as argument.

Next, recall the subsumption relation between logical forms from feature logics: ϕ subsumes ψ iff every feature structure which satisfies ψ also satisfies ϕ (at least on one definition of this, see [17] for details), formally:

(6) ϕ subsumes ψ iff $\Lambda \models \phi \rightarrow \psi$.

Note that the logical form in (5) subsumes the Goal Logical Form in Figure (1), so that the Goal Schema in Figure (1) in fact adds information in the form of so-called "roles" (see [5] for details), and the delivery of such additional information to the ECG module by the ATT-Meta module needs to be included in the pipeline. Designing the set of such translation requirements enables us to formally specify the overall system by specifying the component modules (as per the original systems).

An advantage Gen-Meta has over other NLG approaches, is that reasoning done by the AI module increases overall system flexibility. Thus, if it turned out that, following example (1), Bill no longer has a cold, ATT-Meta's reasoning about such change in circumstances can be piped to the ECG module. This greater control over content specification also has the potential to directly address so-called *strategic generation*, a relatively under-researched area of NLG. Further, the data-driven aspect of Gen-Meta, means that candidate expressions are favored which more closely match formulaic expression of metaphor: so that the relatively formulaic "Bill gave a cold to Bill" would be favored over something as novel as "Bill foisted a cold on Bob".

C. Empirical results

Our combined AI/corpus-based approach enables fine-tuning of (tactical) generation by clothing AI-generated content in patterns of typical metaphorical expression, as determined via corpus-based discovery of conventional forms of expression. To this end, we have mined such web-based sources as online discussion forums, within our chosen domain of Discussion of Illness. Work on illness metaphors is long established, with [24] listing the following examples: AN ILLNESS IS A PUZZLE (e.g. "diabetes is a problem to solve"), A BODY IS A MACHINE (BODY_MACHINE) (e.g.

"your body repairs itself"), A DOCTOR IS A CONTROLLER (DOCTOR_CONTROLLER) (e.g. "my GP is trying to control my disease"), ILLNESS IS AN ATTACK (ILLNESS_ATTACK) (e.g. "an asthma attack").

We have collected a corpus of online discussion forums for illnesses of various kinds (e.g. diabetes, stress, infections, cancer), and we annotated illness metaphors in this corpus (see [14], for more details). We found metaphor types such as those reported in [24], but also novel modifications of these types, such as PATIENT IS A CONTROLLER (PATIENT_CONTROLLER) (e.g. "What most people do to control type 2 diabetes actually makes their blood sugar get worse!"),[6] as well as completely fresh metaphors, such as what might be dubbed "ILLNESS IS A RIDE" (e.g. "the diabetes rollercoaster").[7] Table I reports initial results for frequency of metaphor types for different illnesses. A Pearson chi-squared test on this data yields $\chi^2 = 125.8$ ($p < .005, df = 9$), suggesting metaphor is not independent of domain of illness.[8] These results are interesting in their own right (see [14]), yet can also be exploited to fine-tune Gen-Meta output. Roughly, relative size of the value in a cell in this table (indicated by standardized residuals) suggests relative contribution to the overall chi-squared value. For example, comparing standardized residuals for table cells, we could say that while we can be confident that a natural-seeming metaphor about stress is ILLNESS IS AN ATTACK (e.g. "stress attack" is quite common), this is not the case for diabetes (e.g. "diabetic attack" is far less common). Future work will refine how to best integrate such empirical findings for improving performance of our system.

Metaphor[9]	Diabetes	Infection	Cancer	Stress	ROWS
A	28(-.4)	9(2.3)	11(2.5)	5(-2.3)	53
B	3(.5)	0(-.6)	1(1.)	0(-1.)	4
C	117(3.8)	4(-2.2)	7(-1.9)	18(-3.1)	146
D	8(-5.2)	9(1.4)	8(.4)	47(6.7)	72
COLUMNS	156	22	27	70	**275**

Table I: Frequency of metaphor types for different illnesses (including standardised residuals in brackets), in online discussion forums (see text for details).

IV. CONCLUSION

Our generation approach combines AI techniques for producing metaphorical meanings, with corpus-based approaches for identifying conventionalised forms of metaphorical expressions. This enables three main advantages over existing approaches: (1) compared to other NLG

---

[6]Evoking the notion of "metaphor tuning" from [7].

[7]Metaphor labeling here is rather informal – see [14] for detailed discussion of issues regarding labeling.

[8]With H_0 "metaphor is independent of domain of illness."

[9]Key to types of metaphor: **A** = BODY_MACHINE, **B** = DOCTOR_CONTROLLER, **C** = PATIENT_CONTROLLER, **D** = ILLNESS_ATTACK.

approaches, Gen-Meta combines deep AI reasoning to increase flexibility in generating underlying content, (2) which together with data-driven techniques enables realization to favor formulaic expression of metaphor, and, finally, (3) the greater control over content specification which Gen-Meta affords suggests a new and exciting direction to follow in the under-researched area of *strategic generation*.

REFERENCES

[1] K. Abe, K. Sakamoto, and M. Nakagawa, "A computational model of the metaphor generation process," in *Cognitive Science Society*, 2006, pp. 937–942.

[2] S. Bangalore and O. Rambow, "Corpus-based lexical choice in natural language generation," in *38th Annual Meeting of the Association for Computational Linguistics, Hong Kong, China, October 1-8, 2000*, 2000.

[3] J. A. Barnden, "Metaphor and artificial intelligence: Why they matter to each other," in *The Cambridge Handbook of Metaphor and Thought*, ser. Cambridge Handbooks in Psychology, R. Gibbs, Ed. Cambridge: Cambridge University Press, 2008, pp. 311–338.

[4] ——, "Metaphor and context: A perspective from articial intelligence," in *Metaphor and Discourse*, A. Musolff and J. Zinken, Eds. Palgrave Macmillan, 2009, pp. 79–94.

[5] B. Bergen and N. Chang, "Embodied construction grammar in simulation-based language understanding," in *Construction Grammars: Cognitive grounding and theoretical extensions*, J.-O. Östman and M. Fried, Eds. John Benjamins, 2005.

[6] J. E. Bryant, "Best-fit constructional analysis," Ph.D. dissertation, EECS Department, University of California, Berkeley, Aug 2008.

[7] L. Cameron, "Metaphor and talk," in *The Cambridge Handbook of Metaphor and Thought*, ser. Cambridge Handbooks in Psychology, R. Gibbs, Ed. Cambridge: Cambridge University Press, 2008, pp. 197–211.

[8] W. Croft and D. A. Cruse, *Cognitive Linguistics*, ser. Cambridge Textbooks in Linguistics. Cambridge University Press, 2004.

[9] R. Dale and E. Reiter, *Building natural language generation systems*. Cambridge, UK: Cambridge University Press, 2000.

[10] A. Deignan, "Corpus linguistics and metaphor," in *The Cambridge Handbook of Metaphor and Thought*, ser. Cambridge Handbooks in Psychology, R. Gibbs, Ed. Cambridge: Cambridge University Press, 2008, pp. 280–294.

[11] D. Dobrovol'skij and E. Piirainen, *Figurative Language: Cross-cultural and Cross-linguistic Perspectives*, ser. Current Research in the Semantics / Pragmatics Interface. Elsevier Ltd, 2005.

[12] D. Fass, "Review of "a computational model of metaphor interpretation" by james h. martin. academic press 1990." *Comput. Linguist.*, vol. 17, no. 1, pp. 106–110, Mar. 1991.

[13] J. Feldman, "Embodied language, best-fit analysis, and formal compositionality," *Physics of Life Reviews*, vol. 7, no. 4, pp. 385–410, 2010.

[14] A. Gargett and J. Barnden, "Mining online discussion forums for illness metaphors," submitted.

[15] R. Hervás, R. P. Costa, H. Costa, P. Gervás, and F. C. Pereira, "Enrichment of automatically generated texts using metaphor," in *MICAI 2007: Advances in Artificial Intelligence*, ser. Lecture Notes in Computer Science, A. Gelbukh and n. Kuri Morales, Eds. Springer Berlin Heidelberg, 2007, vol. 4827, pp. 944–954.

[16] A. Idström and E. Piirainen, *Endangered Metaphors*, ser. Cognitive Linguistic Studies in Cultural Contexts (Book 2). John Benjamins, 2012, edited volume.

[17] M. Johnson, "Logic and feature structures," in *IJCAI], pages=992–996, year=1991*.

[18] M. A. Jones, "Generating a specific class of metaphors," in *Proceedings of the 30th annual meeting on Association for Computational Linguistics*, ser. ACL '92. Stroudsburg, PA, USA: Association for Computational Linguistics, 1992, pp. 321–323.

[19] G. Lakoff and M. Johnson, *Metaphors We Live By*. University of Chicago, 1980.

[20] I. Langkilde, "Forest-based statistical sentence generation," in *Proceedings of the 1st North American chapter of the Association for Computational Linguistics conference*, ser. NAACL 2000. Stroudsburg, PA, USA: Association for Computational Linguistics, 2000, pp. 170–177.

[21] O. Lemon, "Learning what to say and how to say it: Joint optimisation of spoken dialogue management and natural language generation," *Computer Speech & Language*, vol. 25, no. 2, pp. 210–221, 2011.

[22] J. H. Martin, "A computational theory of metaphor," University of California at Berkeley, Berkeley, CA, USA, Tech. Rep., 1988.

[23] M. Purver, R. Cann, and R. Kempson, "Grammars as parsers: Meeting the dialogue challenge," *Research on Language and Computation*, vol. 4, no. 2–3, pp. 289–326, 2006.

[24] J. R. Skelton, A. M. Wearn, and F. R. Hobbs, "A concordance-based study of metaphoric expressions used by general practitioners and patients in consultation," *The British Journal of General Practice*, vol. 52, no. 475, pp. 114–118, 2002.

[25] T. Veale and Y. Hao, "Comprehending and generating apt metaphors: a web-driven, case-based approach to figurative language," in *Proceedings of the 22nd national conference on Artificial intelligence - Volume 2*, ser. AAAI'07. AAAI Press, 2007, pp. 1471–1476.