

# Adaptation of a Lexical Organization for Social Engineering Detection and Response Generation

Archna Bhatia, Adam Dalton, Brodie Mather, Sashank Santhanam,  
 Samira Shaikh, Alan Zemel, Tomek Strzalkowski, Bonnie J. Dorr

The Florida Institute for Human and Machine Cognition, The University of North Carolina at Charlotte,  
 University of Albany NY, Rensselaer Polytechnic Institute NY  
 {abhatia,adalton,bmather,bdorr}@ihmc.us, {ssantha1,sshaikh2}@uncc.edu,  
 azemel@albany.edu, tomek@rpi.edu

## Abstract

We present a paradigm for extensible lexicon development based on Lexical Conceptual Structure to support social engineering detection and response generation. We leverage the central notions of *ask* (elicitation of behaviors such as providing access to money) and *framing* (risk/reward implied by the ask). We demonstrate improvements in ask/framing detection through refinements to our lexical organization and show that response generation qualitatively improves as ask/framing detection performance improves. The paradigm presents a systematic and efficient approach to resource adaptation for improved task-specific performance.

**Keywords:** resource adaptation, social engineering detection, response generation, NLP based bots for cyber defense

## 1. Introduction

Social engineering (SE) refers to sophisticated use of deception to manipulate individuals into divulging confidential or personal information for fraudulent purposes. Standard cybersecurity defenses are ineffective because attackers attempt to exploit humans rather than system vulnerabilities. Accordingly, we have built a *user alter-ego* application that detects and engages a potential attacker in ways that expose their identity and intentions.

Our system relies on a paradigm for extensible lexicon development that leverages the central notion of *ask*, i.e., elicitation of behaviors such as PERFORM (e.g., clicking a link) or GIVE (e.g., providing access to money). This paradigm also enables detection of risk/reward (or LOSE/GAIN) implied by an ask, which we call *framing* (e.g., *lose your job, get a raise*). These elements are used for countering attacks through bot-produced responses and actions. The system is tested in an email environment, but is applicable to other forms of online communications, e.g., SMS.

Email	Ask	Framing
(a) It is a pleasure to inform you that you have won 1.7Eu. Contact me. ( <a href="mailto:jw11@example.com">jw11@example.com</a> )	PERFORM contact ( <a href="mailto:jw11@...">jw11@...</a> )	GAIN won (1.7Eu)
(b) You won \$1K. Did you send money? Do that by 9pm or lose money. Respond asap.	GIVE send (money)	LOSE lose (money)
(c) Get 20% discount. Check <a href="#">eligibility</a> or paste this link: <a href="#">http....</a> Sign up for <a href="#">email alerts</a> .	PERFORM paste ( <a href="#">http...</a> )	GAIN get (20%)

Table 1: LCS+ Ask/Framing output for three SE emails

More formally, an *ask* is a statement that elicits a behavior from a potential victim, e.g., *please buy me a gift card*. Although asks are not always explicitly stated (Drew

and Couper-Kuhlen, 2014; Zemel, 2017), we discern these through navigation of semantically classified verbs. The task of ask detection specifically is targeted event detection based on parsing and/or Semantic Role Labeling (SRL), to identify semantic class triggers (Dorr et al., 2020). *Framing* sets the stage for the ask, i.e., the purported threat (LOSE) or benefit (GAIN) that the social engineer wants the potential victim to believe will obtain through compliance or lack thereof. It should be noted that there is no one-to-one ratio between ask and framing in the ask/framing detection output. Given the content, there may be none, one or more asks and/or framings in the output.

Our lexical organization is based on *Lexical Conceptual Structure* (LCS), a formalism that supports resource construction and extensions to new applications such as SE detection and response generation. Semantic classes of verbs with similar meanings (*give, donate*) are readily augmented through adoption of the STYLUS variant of LCS (Dorr and Voss, 2018) and (Dorr and Olsen, 2018). We derive LCS+ from asks/framings and employ CATVAR (Habash and Dorr, 2003) to relate word variants (e.g., *reference* and *refer*). Table 1 illustrates LCS+ Ask/Framing output for three (presumed) SE emails: two PERFORM asks and one GIVE ask.<sup>1</sup> Parentheses () refer to ask *arguments*, often a link that the potential victim might choose to click. Ask/framing outputs are provided to downstream response generation. For example, a possible response for Table 1(a) is *I will contact asap*.

A comparison of LCS+ to two related resources shows that our lexical organization supports refinements, improves ask/framing detection and top ask identification, and yields qualitative improvements in response generation. LCS+ is

<sup>1</sup>To view our system’s ask/framing outputs on a larger dataset (the same set of emails which were also used for ground truth (GT) creation described below), refer to <https://social-threats.github.io/panacea-ask-detection/data/case7LCS+AskDetectionOutput.txt>.

deployed in a SE detection and response generation system. Even though LCS+ is designed for the SE domain, the approach to development of LCS+ described in this paper serves as a guideline for developing similar lexica for other domains. Correspondingly, even though development of LCS+ is one of the contributions of this paper, the main contribution is not this resource but the systematic and efficient approach to resource adaptation for improved task-specific performance.

## 2. Method

In our experiments described in Section 3., we compare LCS+, our lexical resource we developed for the SE domain, against two strong baselines: STYLUS and Thesaurus.

**STYLUS baseline:** As one of the baselines for our experiments, we leverage a publicly available resource STYLUS that is based on Lexical Conceptual Structure (LCS) (Dorr and Voss, 2018) and (Dorr and Olsen, 2018). The LCS representation is an underlying representation of spatial and motion predicates (Jackendoff, 1983; Jackendoff, 1990; Dorr, 1993), such as *fill* and *go*, and their metaphorical extensions, e.g., temporal (the hour *flew* by) and possessional (he *sold* the book).<sup>2</sup> Prior work (Jackendoff, 1996; Levin, 1993; Olsen, 1994; Chang et al., 2007; Chang et al., 2010; Kipper et al., 2007; Palmer et al., 2017) has suggested that there is a close relation between underlying lexical-semantic structures of verbs and nominal predicates and their syntactic argument structure. We leverage this relationship to extend the existing STYLUS verb classes for the resource adaptation to SE domain through creation of LCS+ which is discussed below.

For our STYLUS verb list, we group verbs into four lists based on asks (PERFORM, GIVE) and framings (LOSE, GAIN). The STYLUS verb list can be accessed here: [https://social-threats.github.io/panacea-ask-detection/resources/original\\_lcs\\_classes\\_based\\_verbsList.txt](https://social-threats.github.io/panacea-ask-detection/resources/original_lcs_classes_based_verbsList.txt). Examples of this classification are shown below (with total verb count in parentheses):

- PERFORM (214): remove, redeem, refer
- GIVE (81): administer, contribute, donate
- LOSE (615): penalize, stick, punish, ruin
- GAIN (49): accept, earn, grab, win

Assignment of verbs to these four ask/framing categories is determined by a computational linguist, with approximately a person-day of human effort. Identification of genre-specific verbs is achieved through analysis of 46 emails (406 clauses) after parsing/POS/SRL is applied.

As an example, the verb *position* (Class 9.1) and the verb *delete* (Class 10.1) both have an underlying *placement* or *existence* component with an affected object (e.g., the cursor in *position your cursor* or the account in *delete your*

*account*), coupled with a location (e.g., *here* or *from the system*). Accordingly, *Put* verbs in Class 9.1 and *Remove* verbs in Class 10.1 are grouped together and aligned with a PERFORM ask (as are many other classes with similar properties: Banish, Steal, Cheat, Bring, Obtain, etc.). Analogously, verbs in the *Send* and *Give* classes are aligned with a GIVE ask, as all verbs in these two classes have a sender/giver and a recipient.

Lexical assignment of framings is handled similarly, i.e., verbs are aligned with LOSE and GAIN according to their argument structures and components of meaning. It is assumed that the potential victim of a SE attack serves to lose or gain something, depending on non-compliance or compliance with a social engineer’s ask. As an example, the framing associated with the verb *losing* (Class 10.5) in *Read carefully to avoid losing account access* indicates the risk of losing access to a service; Class 10.5 is thus aligned with LOSE. Analogously, the verb *win* (Class 13.5.1) in *You have won 1.7M Eu.* is an alluring statement with a purported gain to the potential victim; thus Class 13.5.1 is aligned with GAIN. In short, verbs in classes associated with LOSE imply negative consequences (Steal, Impact by Contact, Destroy, Leave) whereas verbs in classes associated with GAIN imply positive consequences (Get, Obtain).

Some classes are associated with more than one ask/framing category: Steal (Class 10.5) and Cheat (Class 10.6) are aligned with both PERFORM (*redeem, free*) and LOSE (*forfeit, deplete*). Such distinctions are not captured in the lexical resource, but are algorithmically resolved during ask/framing detection, where contextual clues provide disambiguation capability. For example, *Redeem coupon* is a directive with an implicit request to click a link, i.e., a PERFORM. By contrast, *Avoid losing account access* is a statement of risk, i.e., a LOSE. The focus here is not on the processes necessary for distinguishing between these contextually-determined senses, but on the organizing principles underlying both, in support of application-oriented resource construction.

**LCS+ resource for SE adapted from STYLUS:** Setting disambiguation aside, resource improvements are still necessary for the SE domain because, due to its size and coverage, STYLUS is likely to predict a large number of both true and false positives during ask/framing detection. To reduce false positives without taking a hit to true positives, we leverage an important property of the LCS paradigm: its extensible organizational structure wherein similar verbs are grouped together. With just one person-day of effort by two computational linguists (authors on the paper; the algorithm developer, also an author, was not involved in this process), a new lexical organization, referred to as “LCS+” is derived from STYLUS, taken together with asks/framings from a set of 46 malicious/legitimate emails.<sup>3</sup> These emails are a random

<sup>3</sup>It should be noted that this resource adaptation is based on an analysis of emails not related to, and without access to, the adjudicated ground truth described in section 3. That is, the 46 emails used for resource adaptation are distinct from the 20 emails used for creating adjudicated ground truth.

<sup>2</sup>LCS is publicly available at <https://github.com/ihmc/LCS>.

subset of 1000+ emails (69 malicious and 938 legitimate) sent from an external red team to five volunteers in a large government agency using social engineering tactics. Verbs from these emails are tied into particular LCS classes with matching semantic peers and argument structures. These emails are proprietary but the resulting lexicon is released here: [https://social-threats.github.io/panacea-ask-detection/resources/lcsPlus\\_classes\\_based\\_verbsList.txt](https://social-threats.github.io/panacea-ask-detection/resources/lcsPlus_classes_based_verbsList.txt). Two categories (PERFORM and LOSE) are modified from the adaptation of LCS+ beyond those in STYLUS:

- PERFORM (6 del, 44 added): copy, notify
- GIVE (no changes)
- LOSE (174 del, 11 added): forget, surrender
- GAIN (no changes)

Table 2 shows the refined lexical organization for LCS+ with ask categories (PERFORM, GIVE) and framing categories (GAIN, LOSE). Boldfaced class numbers indicate the STYLUS classes that were modified. The resulting LCS+ resource drives our SE detection/response system. Each class includes italicized examples with boldfaced triggers. The table details changes to PERFORM and LOSE categories. For PERFORM, there are 6 deleted verbs across 10.2 (Banish Verbs) and 30.2 (Sight Verbs) and also 44 new verbs added to 30.2. For LOSE, 7 classes are associated with additions and/or deletions, as detailed in the table.

**Thesaurus baseline:** The Thesaurus baseline is based on an expansion of simple forms of framings. Specifically, the verbs *gain*, *lose*, *give*, and *perform*, are used as search terms to find related verbs in a standard but robust resource [thesaurus.com](https://www.thesaurus.com) (referred to as “Thesaurus”). The verbs thus found are grouped into these same four categories:

- PERFORM (44): act, do, execute, perform
- GIVE (55): commit, donate, grant, provide
- LOSE (41): expend, forfeit, expend, squander
- GAIN (53): clean, get, obtain, profit, reap

The resulting Thesaurus verb list is publicly released here: [https://social-threats.github.io/panacea-ask-detection/resources/thesaurus\\_based\\_verbsList.txt](https://social-threats.github.io/panacea-ask-detection/resources/thesaurus_based_verbsList.txt).

We also adopt categorial variations through CATVAR (Habash and Dorr, 2003) to map between different parts of speech, e.g., *winner(N)*  $\rightarrow$  *win(V)*. STYLUS, LCS+ and Thesaurus contain verbs only, but asks/framings are often nominalized. For example, *you can reference your gift card* is an implicit ask to examine a gift card, yet without CATVAR this ask is potentially missed. CATVAR recognizes *reference* as a nominal form of *refer*, thus enabling the identification of this ask as a PERFORM.

### 3. Experiments and Results

*Intrinsic* evaluation of our resources is based on comparison of ask/framing detection to an adjudicated ground truth (henceforth, GT), a set of 472 clauses from system output on 20 unseen emails. These 20 emails are a random subset of 2600+ messages collected in an email account set up to receive messages from an internal red team as well as “legitimate” messages from corporate and academic mailing lists. As alluded to earlier, these 20 emails are distinct from the dataset used for resource adaptation to produce the task-related LCS+.

The GT is produced through human adjudication and correction by a computational linguist<sup>4</sup> of initial ask/framing labels automatically assigned by our system to the 472 clauses. System output also includes the identification of a “top ask” for each email, based on the degree to which ask argument positions are filled.<sup>5</sup> *Top asks* are adjudicated by the computational linguist once the ask/framing labels are adjudicated. The resulting GT is accessible here: <https://social-threats.github.io/panacea-ask-detection/data/>.

The GT is used to measure the precision/recall/F of three of three variants of ask detection output (Ask, Framing, and Top Ask) corresponding to our three lexica: Thesaurus, STYLUS, and LCS+. LCS+ is favored (with statistical significance) against the two very strong baselines, Thesaurus and STYLUS. Table 3 presents results: Recall for framings is highest for STYLUS, but at the cost of higher false positives (lower precision). F-scores increase for STYLUS over Thesaurus, and for LCS+ over STYLUS.

McNemar (McNemar, 1947) tests yield statistically significant differences for asks/framings at the 2% level between Thesaurus and LCS+ and between STYLUS and LCS+.<sup>6</sup> It should be noted that not all clauses in GT are ask or framing: vast majority (80%) are neither (i.e., they are true negatives).

We note that an alternative to the Thesaurus and LCS baselines would be a bag-of-words lexicon, with no organizational structure. However, the key contribution of this work is the ease of adaptation through classes, obviating the need for training data (which are exceedingly difficult to obtain). Classes enable extension of a small set of verbs to a larger range of options, e.g., if the human determines from a small set of task-related emails that *provide* is relevant, the task-adapted lexicon will include *administer*, *contribute*, and *donate* for free. If a class-based lexical organization is replaced by bag-of-words, we stand to lose efficient (1-person-day) resource adaptation and, moreover, training data would be needed.

A first step toward *extrinsic* evaluation is inspection of responses generated from each resource’s top ask/framing pairs. Table 1 (given earlier) shows LCS+ ask/framing pairs

<sup>4</sup>The adjudicator is an author but is not the algorithm developer, who is also an author.

<sup>5</sup>Argument positions express information such as the ask type (i.e. PERFORM), context to the ask (i.e. financial), and the ask target (e.g., “you” in “Did you send me the money?”).

<sup>6</sup>Tested values were TP+TN vs FP+FN, i.e., significance of change in total error rate.

**PERFORM:**

9.1 Put Verbs: *Position* your cursor here  
 10.1 Remove Verbs: *Delete* virus from machine  
 10.2 Banish Verbs→5 deleted (banish, deport, evacuate, extradite, recall): *Remove* fee from your account  
 10.5 Steal Verbs: *Redeem* coupon below  
 10.6 Cheat Verbs: *Free* yourself from debt  
 11.3 Bring and Take Verbs: *Bring* me a gift card  
 13.5.2 Obtain: *Purchase* two gift cards  
 30.2 Sight Verbs→1 deleted (regard), 44 added (e.g., check, eye, try, view, visit): *View* this website  
 37.1 Transfer of Message: *Ask* for a refund  
 37.2 Tell Verbs: *Tell* them \$50 per card  
 37.4 Communication: *Sign* the back of the card  
 42.1 Murder Verbs: *Eliminate* your debt here  
 44 Destroy Verbs: *Destroy* the card  
 54.4 Price Verbs: *Calculate* an amount here

**GIVE:**

11.1 Send Verbs: *Send* me the gift cards  
 13.1 Give Verbs: *Give* today  
 13.2 Contribute Verbs: *Donate*!  
 13.3 Future Having: *Advance* me \$100  
 13.4.1 Verbs of Fulfilling: *Credit* your account  
 32.1 Want Verbs: *I need* three gift cards

**LOSE:**

10.5 Steal Verbs→11 added (e.g., forfeit, lose, relinquish, sacrifice): *Don't forfeit* this chance!  
 10.6 Cheat Verbs: *Are* your funds **depleted**?  
 17.1 Throw Verbs: *Don't toss* out this coupon  
 17.2 Pelt Verbs: *Scams* **bombarding** you?  
 18.1 Hit Verbs: *Don't be beaten* by debt  
 18.2 Swat Verbs: *Sluggish* market getting you down?  
 18.3 Spank Verbs: *Clobbered* by fees?  
 18.4 Impact by Contact: *Avoid being hit* by malware  
 19 Poke Verbs: *Stuck* with debt?  
 29.2 Characterize Verbs→16 deleted (e.g., appreciate, envisage): *Repudiated* by creditors?  
 29.7 Orphan Verbs→5 deleted (apprentice, canonize, cuckold, knight, recruit): *Avoid crippling* debt  
 29.8 Captain Verbs→35 deleted (e.g., captain, coach, cox, escort): *Bullied* by bill collectors?  
 31.1 Amuse Verbs→91 deleted (e.g., amaze, amuse, gladden): *Don't be disarmed* by hackers  
 31.2 Admire Verbs→26 deleted (e.g., admire, exalt): *Are you lamenting* your credit score?  
 31.3 Marvel Verbs→1 deleted (feel): *Living in fear*?  
 33 Judgment Verbs: *Need to remove* **penalties**?  
 37.8 Complain Verbs: *Want your gripes* answered?  
 42.1 Murder Verbs: *Debt killing* your credit?  
 42.2 Poison Verbs: *Strangled* by debt?  
 44 Destroy Verbs: *PC destroyed* by malware?  
 48.2 Disappearance: *Your account will expire*  
 51.2 Leave Verbs: *Found your abandoned* prize

**GAIN:**

13.5.1 Get: *You are a winner* of 1M Eu.  
 13.5.2 Obtain: *You can recover* your credit rating

Table 2: Lexical organization of LCS+ relies on Ask Categories (PERFORM, GIVE) and Framing Categories (GIVE, LOSE). Italicized exemplars with boldfaced triggers illustrate usage for each class. Boldfaced class numbers indicate those STYLUS classes that were modified to yield the LCS+ resource.

Thesaurus	P	R	F
Ask:	0.273	0.042	0.072
Framing:	0.265	0.360	0.305
TopAsk:	0.273	0.057	0.094
STYLUS	P	R	F
Ask:	0.333	0.104	0.159
Framing:	0.298	<b>0.636</b>	0.406
TopAsk:	0.571	0.151	0.239
LCS+	P	R	F
Ask:	<b>0.667</b>	<b>0.411</b>	<b>0.508</b>
Framing:	<b>0.600</b>	0.600	<b>0.600</b>
TopAsk:	<b>0.692</b>	<b>0.340</b>	<b>0.456</b>

Table 3: Impact of lexical resources on ask/framing detection: Thesaurus, STYLUS, LCS+

whose corresponding (T)hesaurus and (S)TYLUS pairs are:

- (a) T: None, None  
 S: None, GAIN/won(1.7Eu)  
 (b) T: PERFORM/do(that), LOSE/lose(money)  
 S: GAIN/won(money), GIVE/send(money)  
 (c) T: None, GAIN/get(20%)  
 S: PERFORM/sign(http:...), GAIN/get(20%)

Below are corresponding examples of generated responses<sup>7</sup> for all 3 resources, based on a templatic approach that leverages ask/framing hierarchical structure and corresponding confidence scores. This module is part of a larger, separate publication.

- (a) T: How are you? Thanks.  
 S: ...too good to be true. What should I do?  
 L+: I will contact asap.  
 (b) T: Thanks for getting in touch, need more info.  
 S: Nervous about this. Your name?  
 L+: I would respond,<sup>8</sup> but I need more info.  
 (c) T: What should I do now?  
 S: Website doesn't open, is this the link?  
 L+: Thanks, need more info before I paste link

There are qualitative differences in these responses. For example, in (a) Thesaurus (T) yields no asks/framings; thus a canned response is generated. By contrast, the same email yields a more responsive output for STYLUS (S), and a more focused response for LCS+ (L). Similar distinctions are found for responses in (b) and (c). Note that in the LCS+ condition, if there is no match found using LCS+, downstream response generation prompts the attacker (e.g., "please clarify") until an interpretable ask or framing appears. In this SE task, not all responses move the conversation forward. A central goal of the SE task is to waste the attacker's time, play along, and possibly extract information that could unveil their identity.

#### 4. Related Work

LCS is used in interlingual machine translation (Voss and Dorr, 1995; Habash and Dorr, 2002), lexical acquisition

<sup>7</sup>For brevity, *excerpts* are shown in lieu of full emails.

<sup>8</sup>LCS+ detects both GIVE/send and PERFORM/respond.

(Habash et al., 2006), cross-language information retrieval (Levow et al., 2000), language generation (Traum and Habash, 2000), and intelligent language tutoring (Dorr, 1997). STYLUS (Dorr and Voss, 2018) and (Dorr and Olsen, 2018) systematizes LCS based on several studies (Levin and Rappaport Hovav, 1995; Rappaport Hovav and Levin, 1998), but to our knowledge our work is the first use of LCS in a conversational context, within a cyber domain. Our approach relates to work on conversational agents (CAs), where neural models automatically generate responses (Gao et al., 2019; Santhanam and Shaikh, 2019), topic models produce focused responses (Dziri et al., 2018), self-disclosure yields targeted responses (Ravichander and Black, 2018), and SE detection employs topic models (Bhakta and Harris, 2015) and NLP of conversations (Sawa et al., 2016). However, all such approaches are limited to a pre-defined set of topics, constrained by the training corpus. Other prior work focuses on persuasion detection/ prediction (Hidey and McKeown, 2018) by leveraging argument structure, but for the purpose of judging when a persuasive attempt might be successful in subreddit discussions dedicated to changing opinions (ChangeMyView). Our work aims to achieve effective dialogue for countering (rather than adopting) persuasive attempts.

Text-based semantic analysis for SE detection (Kim et al., 2018) is related to our work but differs in that our work focuses not just on *detecting* an attack, but on *engaging* with an attacker. Whereas a bot might be employed to warn a potential victim that an attack is underway, our bots are designed to communicate with a social engineer in ways that elicit identifying information.

## 5. Conclusions

Both STYLUS and LCS+ support ask/framing detection in service of bot-produced responses. Intrinsically, LCS+ is superior to both STYLUS and Thesaurus when measured against human-adjudicated output, verified for significance by McNemar tests at the 2% level. Extrinsically, STYLUS supports more responsive bot outputs and LCS+ supports more focused bot outputs.

A more general advantage of adapting LCS+ to the SE domain is that it can act as a guideline for developing similar resources for other domains which will similarly support focused outputs appropriate for particular domains. The main contribution of this paper is not development of a particular task-specific resource, nor to suggest that LCS+ is a generic resource for many tasks, but to present a systematic, efficient approach to resource adaptation technique that can generalize to other tasks for improved task-specific performance, e.g., understanding viewpoints in social media or detecting motives behind activities of political groups. We acknowledge that our extrinsic evaluation is limited. While we have demonstrated the efficacy of ask detection approaches on a set of representative emails, a quantitative evaluation is required to test the statistical significance of our extrinsic observations. Future work is planned to conduct experiments with crowd-sourced workers judging the efficacy and effectiveness of generated responses.

## Acknowledgments

This work was supported by DARPA through AFRL Contract FA8650-18-C-7881 and through Army Contract W31P4Q-17-C-0066. All statements of fact, opinion or conclusions contained herein are those of the authors and should not be construed as representing the official views or policies of DARPA, AFRL, Army, or the U.S. Government.

## 6. Bibliographical References

- Bhakta, R. and Harris, I. G. (2015). Semantic analysis of dialogs to detect social engineering attacks. *Proceedings of the 2015 IEEE 9th International Conference on Semantic Computing (IEEE ICSC 2015)*, pages 424–427.
- Chang, S. C., Shahani, R. C., Cipollone, D. J., Calcagno, M. V., Olsen, M. J. B., and Parkinson, D. J. (2007). Linguistic Object Model, January. 7,171,352.
- Chang, S. C., Shahani, R. C., Cipollone, D. J., Calcagno, M. V., Olsen, M. J. B., and Parkinson, D. J. (2010). Lexical Semantic Structure, March. 7,689,410.
- Dorr, B. J. and Olsen, M. B. (2018). Lexical conceptual structure of literal and metaphorical spatial language: A case study of push. In *Proceedings of the First International Workshop on Spatial Language Understanding*, pages 31–40.
- Dorr, B., Bhatia, A., Dalton, A., Mather, B., Hebenstreit, B., Santhanam, S., Cheng, Z., Zemel, S., and Strzalkowski, T. (2020). Detecting asks in social engineering attacks: Impact of linguistic and structural knowledge. In *Proceedings of Thirty-Fourth AAAI Conference on Artificial Intelligence 2020*.
- Dorr, B. J. (1993). *Machine Translation: A View from the Lexicon*. MIT Press, Cambridge, MA.
- Dorr, B. J. (1997). Large-Scale Dictionary Construction for Foreign Language Tutoring and Interlingual Machine Translation. *Machine Translation*, 12:271–322.
- Drew, P. and Couper-Kuhlen, E. (2014). *Requesting in social interaction*. John Benjamins Publishing Company.
- Dziri, N., Kamaloo, E., Mathewson, K. W., and Zaiane, O. (2018). Augmenting neural response generation with context-aware topical attention. *arXiv preprint arXiv:1811.01063*.
- Gao, J., Galley, M., Li, L., et al. (2019). Neural approaches to conversational ai. *Foundations and Trends® in Information Retrieval*, 13(2-3):127–298.
- Habash, N. and Dorr, B. J. (2002). Handling Translation Divergences: Combining Statistical and Symbolic Techniques in Generation-Heavy Machine Translation. In *Proceedings of the Fifth Conference of the Association for Machine Translation in the Americas*, pages 84–93, Tiburon, CA.
- Habash, N., Dorr, B. J., and Monz, C. (2006). Challenges in Building an Arabic GHMT system with SMT Components. In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas*, pages 56–65, Boston, MA, August.
- Hidey, C. and McKeown, K. (2018). Persuasive Influence Detection: The Role of Argument Sequencing. In *Proceedings of the Thirty-Second AAAI Conference on Arti-*

- ficial Intelligence*, pages 5173–5180, San Francisco, California, USA.
- Jackendoff, R. (1983). *Semantics and Cognition*. MIT Press, Cambridge, MA.
- Jackendoff, R. (1990). *Semantic Structures*. MIT Press, Cambridge, MA.
- Jackendoff, R. (1996). The Proper Treatment of Measuring Out, Telicity, and Perhaps Even Quantification in English. *Natural Language and Linguistic Theory*, 14:305–354.
- Kim, M., Song, C., Kim, H., Park, D., Kwon, Y., Namkung, E., Harris, I. G., and Carlsson, M. (2018). Catch me, yes we can!-pwning social engineers using natural language processing techniques in real-time.
- Kipper, K., Korhonen, A., Ryant, N., and Palmer, M. (2007). A Large-scale Classification of English Verbs. In *Language Resources and Evaluation*.
- Levin, B. and Rappaport Hovav, M. (1995). *Unaccusativity: At the Syntax-Lexical Semantics Interface, Linguistic Inquiry Monograph 26*. MIT Press, Cambridge, MA.
- Levin, B. (1993). *English Verb Classes and Alternations: A Preliminary Investigation*. The University of Chicago Press.
- Levow, G., Dorr, B. J., and Lin, D. (2000). Construction of Chinese-English Semantic Hierarchy for Cross-language Retrieval.
- Rappaport Hovav, M. and Levin, B. (1998). Building Verb Meanings. In M. Butt et al., editors, *The Projection of Arguments: Lexical and Compositional Factors*, pages 97–134. CSLI Publications, Stanford, CA.
- McNemar, Q. (1947). Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, jun.
- Olsen, M. B. (1994). The Semantics and Pragmatics of Lexical and Grammatical Aspect. *Studies in the Linguistic Sciences*, 24(1–2):361–375.
- Palmer, M., Bonial, C., and Hwang, J. D. (2017). VerbNet: Capturing English Verb behavior, Meaning and Usage.
- Ravichander, A. and Black, A. W. (2018). An empirical study of self-disclosure in spoken dialogue systems. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue, Melbourne, Australia, July 12-14, 2018*, pages 253–263.
- Santhanam, S. and Shaikh, S. (2019). A survey of natural language generation techniques with a focus on dialogue systems-past, present and future directions. *arXiv preprint arXiv:1906.00500*.
- Sawa, Y., Bhakta, R., Harris, I., and Hadnagy, C. (2016). Detection of social engineering attacks through natural language processing of conversations. In *Proceedings of the 2016 IEEE Tenth International Conference on Semantic Computing (ICSC)*, pages 262–265, 02.
- Traum, D. and Habash, N. (2000). Generation from Lexical Conceptual Structures. In *Proceedings of the Workshop on Applied Interlinguas, North American Association for Computational Linguistics / Applied NLP Conference*, pages 34–41.
- Voss, C. R. and Dorr, B. J. (1995). Toward a Lexicalized Grammar for Interlinguas. *J. of Machine Translation*, 10:143–184.
- Zemel, A. (2017). Texts as actions: Requests in online chats between reference librarians and library patrons. *Journal of the Association for Information Science and Technology*, 67(7):1687–1697.

## 7. Language Resource References

- Dorr, Bonnie and Voss, Clare. (2018). *STYLUS: A Resource for Systematically Derived Language Usage*. Association for Computational Linguistics.
- Nizar Habash and Bonnie J. Dorr. (2003). *A Categorical Variation Database for English*.