Verb-Particle Constructions and Lexical Resources

Aline Villavicencio

University of Cambridge Computer Laboratory,

William Gates Building, JJ Thomson Avenue,

Cambridge, CB3 0FD, UK

`Aline.Villavicencio@cl.cam.ac.uk`

Abstract

In this paper we investigate the phenomenon of verb-particle constructions, discussing their characteristics and their availability for use with NLP systems. We concentrate in particular on the coverage provided by some electronic resources. Given the constantly growing number of verb-particle combinations, possible ways of extending the coverage of the available resources are investigated, taking into account regular patterns found in some productive combinations of verbs and particles. We discuss, in particular, the use of Levin's (1993) classes of verbs as a means to obtain productive verb-particle constructions, and discuss the issues involved in adopting such an approach.

1 Introduction

In this paper we discuss verb-particle constructions (VPCs) in English and analyse some of the available sources of information about them for use in NLP systems. VPCs can range from idiosyncratic or semi-idiosyncratic combinations, such as *get on* (in e.g. *Bill got on well with his new colleagues*), to more regular ones, such as *tear up* (in e.g. *In a rage she tore up the letter Jack gave her*). However, examples of 'idiomatic' VPCs like *get on*, meaning *to be on friendly terms with someone*, where the meaning of the combination cannot be straightforwardly inferred from the meaning of the verb and the particle, fortunately seem to be a small minority (Side,

1990). Most cases seem to be more regular, with the particle compositionally adding a specific meaning to the construction and following a productive pattern (e.g. in *tear up*, *cut up* and *split up*, where the verbs are semantically related and *up* adds a sense of completion the action of these verbs).

VPCs have been the subject of a considerable amount of interest, and some investigation has been done on the subject of productive VPCs. For instance, even though the particle *up* occurs with a wide range of verbs, it only combines productively with some classes. Bame (1999) discusses two such cases: the resultative and the aspectual *up*. For example *Kim carried the television up* uses a resultative *up* and *Kim ate the sandwich up* an aspectual *up*. With the resultative *up*, the argument is affected (i.e., at the end of the action the television is *up*). In contrast, the aspectual or completive *up* suggests that the action is taken to some conclusion (i.e., the sandwich is totally consumed at the end of the action).

Fraser (1976) points out that semantic properties of verbs can affect their possibilities of combining with particles. Thus, semantic properties can influence the patterns of verb-particle combinations that verbs follow. For example, in the case of *bolt/cement/clam/glue/paste/nail* all are semantically similar verbs where the objects specified by the verbs are used to join material and they can all combine with *down*. There is clearly a common semantic thread running through this list, so that a new verb that is semantically similar to them can also be reasonably assumed to combine with *down*. Indeed Side notes that frequently new VPCs are formed by

analogy with existing ones, with often the verb being varied and the particle remaining (e.g. *hang on, hold on* and *wait on*).

By identifying classes of verbs that follow patterns such as these in VPCs, we are able to maximise the use of the information contained in lexical resources. For instance, we can make use of regular patterns to productively generate VPCs from verbs already listed in a lexical resource, according to their verbal classes (e.g. the resultative combinations *walk up/down/out/in/away/around/...* from *walk* and the directional/locative particles *up*, *down*, *out*, *in*, *away*, *around*, ...). We consider how we can use productive patterns to extend the coverage of current lexical resources, in the next sections. We start by characterising VPCs, and investigating the coverage provided by some available electronic dictionaries, in section 3. We also discuss the use of corpora to extend the coverage provided by these dictionaries. After that we investigate how more productive combinations can be generated from a semantic classification of verbs such as Levin's (1993).

2 Characterizing VPCs

VPCs are combinations of verbs and prepositional or adverbial particles, such as *break down* in *The old truck broke down*. In these constructions particles are characterised by containing features of motion-through-location and of completion or result in their core meaning (Bolinger, 1971). In syntactic terms in the example above we have an intransitive VPC, where no other verbal complement is required. Other VPCs may have further subcategorisation requirements, and in, for example, *They came across an old manuscript* we have a transitive VPC which has a further NP complement.

In this work we are looking exclusively at cases of VPCs, thus excluding prepositional verbs, where a verb subcategorises for a prepositional phrase (PP), such as *rely on*, in *He relies on his wife for everything*. Cases like this and others of adverbial modification need to be distinguished from VPCs. This difference may be quite subtle and, in order to distinguish VPCs from other constructions we use the following criteria:

- The particle may come either before or after the NP in transitive VPCs (e.g. *He backed up the team* vs *He backed the team up*). Whether a particle can be separated or not from the verb may depend on the degree of bondage of the particle with the verb, on the size of the NP, and on the kind of NP.

- In transitive VPCs unstressed personal pronouns must precede the particle (e.g. *They ate it up* but not **They ate up it*).

- The particle, in transitive VPCs, comes before a simple definite NP without taking it as its object (e.g. *He brought along his girlfriend* but not *It consists of two parts*).

- In VPCs subcategorising for other verbal complements, like PPs and sentential complements, the particle must come immediately after the verb.

- Verbs that subcategorise for an optional goal argument that is fullfilled by a locative or directional particle are considered to be VPCs with the particle further specifying the meaning of the verb (e.g. *walk up* in *Bill walked up the hill*).

As discussed by Bolinger (1971), many of the criteria proposed for diagnosing VPCs give different results for the same combination frequently including unwanted combinations and excluding genuine VPCs. Nonetheless, they provide us with at least the basis for this decision.

3 Dictionaries and VPCs

Dictionaries are a major source of information about VPCs. In Table 1 we can see the coverage of phrasal verbs (PVs) in several dictionaries and lexicons: Collins Cobuild Dictionary of Phrasal Verbs (Collins-PV), Cambridge International Dictionary of Phrasal Verbs (CIDE-PV), the electronic versions of the Alvey Natural Language Tools (ANLT) lexicon (Carroll and Grover, 1989) (which was derived from the Longman Dictionary of Contemporary English, LDOCE), the Comlex lexicon (Macleod and Grishman, 1998), and the LinGO English Resource Grammar (ERG) (Copestake and Flickinger, 2000) version of November 2001. This table shows in the second column the number of PV entries for each of

Figure 1: Top Ranked Particles I

these dictionaries, including not only VPCs but also other kinds of PV. The third column shows the number of VPC entries (available only for the electronic dictionaries).

Table 1: Phrasal Verb Entries in Dictionaries

Dictionary	PVs	VPCs
ANLT	6,439	2,906
CIDE-PV	over 4,500	-
Collins-PV	over 3,000	-
Comlex	12,564	4,039
ERG	533	337

As we can see from these numbers, each of these dictionaries has a considerable number of PV entries potentially providing us with a good starting point for handling VPCs. Table 2 shows some of the characteristics of each dictionary, in more detail, with respect to VPCs, where the seventh column shows the proportion of verbs used in VPCs (sixth column) from all verbs in a dictionary (second column).

Each of these dictionaries uses a different set of verbs and particles in its VPCs. However, with respect to the verbs listed in these dictionaries there is a high level of agreement among them with, for example, 93.26% of the verbs in Comlex being also listed in ANLT. In Table 2 we can see the increase in the number of verbs obtained by the union of the dictionaries, where A+C represents the union of ANLT and Comlex, A∩C their intersection and A+C+E the union of ANLT, Comlex and ERG. Because of the high level of agreement for their verbs, when joined together the contribution made by each dictionary is relatively small, so that the combination of the three (A+C+E) has only 7.3% more verbs than the ANLT

alone, for example.

In relation to VPCs, ANLT uses the largest number of particles, and with one exception all the particles contained in the ERG and Comlex are already contained in ANLT. When we rank the particles according to the frequency with which they occur in the VPCs, we get similar patterns for all of the dictionaries, as can be seen in Figure 1. This figure shows the 5 top ranked particles for each of the dictionaries, and for all of them *up* is the particle involved in the largest number of combinations. By analysing the VPCs in each of these dictionaries, we can also see that only a small proportion of the total number of verbs in a dictionary is used in its VPCs, Table 2. For example, only 20% of the verbs listed in ANLT form at least one VPC. For the other dictionaries this proportion is even lower. These tend to be very widely used and general verbs, such as *come, go, get, put, bring* and *take*. Whether the remaining verbs do not form valid VPCs or whether the combinations were simply omitted remains to be investigated.

Even though only a subset of verbs in dictionaries are used in VPCs, this subset generates a large number of combinations, as shown in Table 2. Each of these dictionaries specialises in a subset of VPCs. Because of this difference in coverage, when the dictionaries are combined, as each one is added it helps to significantly extend the coverage of VPCs. Although there is a significant number of entries[1] that are common among the different dictionaries, it seems to correspond only to a subset of the total number of entries each dictionary has. For instance, from the total number of entries obtained by combining ANLT and Comlex, Table 2, only 34% of the entries are listed in both dictionaries with the remaining 66% of the total number of entries being exclusive to one or the other of these dictionaries. Moreover, even with the large number of entries already obtained by combining these two dictionaries, a considerable proportion (16%) of the entries in the LinGO ERG lexicon are not listed in any of these two dictionaries (this proportion would increase if we took subcategorization etc into account).[2] Most

---

[1] These figures do not take into account subcategorisation information, where a given verb-particle construction can occur with more than one subcategorisation frame.

[2] The LinGO ERG lexicon was manually constructed with

Table 2: VPCs in Dictionaries

Dictionary	Verbs	VPC Entries	Distinct VPCs	Particles	Verbs in VPCs	Proportion of Verbs in VPCs
ANLT	5,667	2,906	2,250	44	1,135	20%
Comlex	5,577	4,039	1,909	23	990	17.75%
ERG	1,223	337	270	25	176	14.39%
A+C	6,043	-	3,111	44	1,394	23.07%
A∩C	5,201	-	1,052	23	731	14.05%
A+C+E	6,113	-	3,156	45	1,400	22.90%

of these are at least semi-compositional, e.g., *crisp up*, *come together*, *tie on*, and were probably omitted from the dictionaries for that reason,[3] though some others, such as *hack up*, are probably recent coinages. The coverage of these resources is quite limited and possible ways of extending it are a necessity for successful NLP systems.

4 VPCs in Corpora

The use of corpora to extract verb-particle combinations can contribute to extending the coverage of dictionaries. An investigation of the automatic extraction of VPCs from corpora is described in Baldwin and Villavicencio (2002). In this section we use VPCs extracted from the British National Corpus (BNC), comparing these VPCs with those contained in the combined A+C+E-VPCs, and discussing how the former can be used to complement the coverage provided by the latter.

The BNC is a 100 million word corpus containing samples of written text from a wide variety of sources, designed to represent as wide a range of modern British English as possible. Using the methods described in Baldwin and Villavicencio (2002), 8,751 VPC entries were extracted from the BNC. These entries are classified into intransitive and/or transitive VPCs, depending on their subcategorisation frame, and they result in 7,078 distinct VPCs. Some of these entries are not VPCs but rather noise, such as **** *off*, *'s down*, etc. After removing the most obvious cases of noise, there were 7,070 VPCs

left. These are formed by 2,542 verbs and 48 particles, as shown in Table 3.

Table 3: VPCs from Dictionaries and from BNC

Resources	VPCs	Verbs	Particles
A+C+E	3,156	1,400	45
BNC	7,070	2,542	48
A+C+E ∩ BNC	2,014	1,149	28
A+C+E - BNC	1,138	251	17
BNC - A+C+E	5,056	1,393	20
A+C+E+BNC	8,208	2,793	65

When comparing the VPCs in BNC (BNC-VPCs) with those in the combined dictionaries (A+C+E-VPCs) there are 1,149 verbs in common, corresponding to 82.1% of the verbs in the combined dictionaries. When these resources are joined together, there is a significant increase in the number of verbs and particles, with a total of 2,793 different verbs and 65 particles used in VPCs, Table 3. The verbs that appear in the largest number of VPCs are again general and widely used (e.g. *move, come, go, get* and *pull*). For these, the five particles that occur in the highest number of VPCs are shown in Figure 2, and they are basically the same as those in the dictionaries.

In terms of the VPCs, by joining A+C+E-VPCs with BNC-VPCs there is an increase of 160.30% in the number of VPCs. Among the extracted VPCs many form productive combinations: some containing a more informal or a recent use of verbs (e.g. *hop off*, *kangaroo down* and *skateboard away*). These VPCs provide a useful addition to those contained in the dictionaries. However, we are still able to ob-

---

most of the verb-particle entries being empirically motivated by the Verbmobil corpus. It is thus probably reasonably representative of a moderate-size domain-specific lexicon.

[3] The Cobuild Dictionary explicitly states that literal meanings and combinations are not given for all verbs.

Figure 2: Top Ranked Particles II

tain only a subset of the existing VPCs, and plausible combinations such as *hoover up* are not found in these combined resources. In the next section we discuss how to extend even further their coverage by making use of productive patterns found in classes of semantically related verbs.

5 VPC Patterns in Levin's Verb Classes

Fraser (1976) noted how semantic properties of verbs can affect their possibilities of combination with particles (e.g. *hunt/track/trail/follow down* and *bake/cook/fry/broil up*). Semantic properties of verbs can influence the patterns of combination that they follow (e.g. verbs of hunting and the resultative *down* and verbs of cooking and the aspectual *up*). By having a semantic classification of verbs we can determine how they combine with certain particles, and this can be used to extend the coverage of the available resources by productively generating VPCs from classes of related verbs according to the patterns that they follow. One such classification was proposed by Levin (1993). In Levin's classification, verbs are grouped into classes in terms of their syntactic and semantic properties. These classes were not developed specifically for VPCs, but it may be the case that some productive patterns of combinations correspond to certain classes of verbs. We investigated the possibility of using Levin's classes of verbs to generate a set of candidate VPCs, and in this section, we briefly discuss Levin's classes and describe how they can be used to predict productive verb-particle combinations.

There are 190 fine grained subclasses that capture 3,100 different verbs listed, resulting in 4,167 entries, since each verb can belong to more than one class. For example, the verb *to run* belongs

to classes 26.3 (Verbs of Preparing), 47.5.1 (Swarm Verbs), 47.7 (Meander Verbs) and 51.3.2 (Run Verbs). The number of elements in each class varies considerably, so that 60% of all of these classes have more than 10 elements, accounting for 88% of the verbs, while the other 40% of the classes have 10 or less elements, capturing the remaining 22% of the verbs. The 5 larger classes are shown in Table 4.

Table 4: Verb Entries in Levin's Classes

Class	Name	Entries
45.4	Other alternating verbs of change of state	257
31.1	Amuse	220
51.3.2	Run	124
43.2	Sound emission	119
9.9	Butter	109

It is possible that some productive patterns found in VPCs may be mapped onto the classes defined. In this case, some classes may be good predictors of productive VPCs, and to test this possibility we analysed the combinations generated by Levin's classes and a subset of four particles (*down, in, out, up*). To test the validity of a resulting combination, we searched for it first among the VPCs from the combined dictionaries, A+C+E-VPCs, and then among the much more numerous but potentially noisy A+C+E+BNC-VPCs.

All combinations of verbs in Levin's classes and these four particles were generated and tested for validity. We use the proportion of valid VPCs as a metric to determine the degree of productivity of a given class, so that the higher the proportion, the more productive the class, according to the combined resources. The classes are then ranked according to their productivity degree.

There are 16,668 possible combinations that can be generated, from the 4,167 entries in Levin's classes and four particles. However, from the 4,167 only 3,914 entries have verbs that are in A+C+E, so we will consider only 15,656 possible VPCs, when evaluating these results against the combined dictionaries.

When we compare the 15,656 possible VPCs with those in A+C+E, 2,456 were considered valid
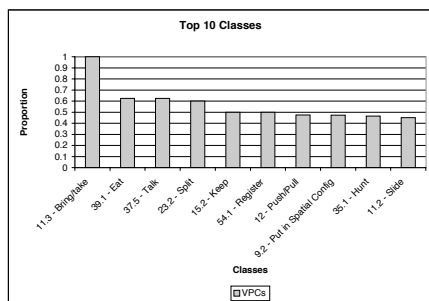
Figure 3: Levin's VPCs in Dictionaries

(15.69%). In Figure 3, we can see the degree of productivity of a class, for the 10 highest ranked classes, according to A+C+E-VPCs. From these classes, we can see two basic patterns:

- verbs that can form aspectual combinations, with the particle giving a sense of completion and/or increase/improvement to the action denoted by the verb, e.g. verbs of Eating (39.1) and Splitting (23.2),

- verbs that imply some motion or take a location, e.g. verbs of Bring and Take (11.3), Push and Pull (12) and Putting in spatial configuration (9.2), and can form resultative combinations.

However, apart from class 11.3, where all verbs form good combinations with all four particles, according to the dictionaries, the other classes have a lower proportion of valid combinations. As these results may be due to the coverage of the dictionaries, we compared these results with those obtained by also using BNC-VPCs to test the validity of a combination. In this case, from the 4,167 entries in Levin's classification, 3,925 have verbs that are in A+C+E+BNC-VPCs, generating 15,700 candidate VPCs, against which we perform the evaluation. Using this larger set of VPCs, further combinations are considered valid: 4,733 VPCs out of 15,700 candidates (30.15%). This time a considerable improvement in the results was verified, with a larger number of classes having the majority of its VPCs being considered valid. Figure 4 shows the ten top ranked classes found with A+C+E+BNC-VPCs. Confirming the trends suggested with the dictionaries, most of the top ranked classes have verbs implying some

kind of motion or taking a location (e.g. 11.3 - Bring and Take- and 53.2 - Rushing) forming resultative VPCs, or forming aspectual VPCs (e.g. 23.2 - Split).

All of the classes in Figure 4 have 70% or more of their verbs forming valid combinations, according to A+C+E+BNC-VPCs. For these classes a manual analysis of the VPCs generated was performed to test the predicted productivity of the class. All those combinations that were not attested were subject to human judgement. Cases of these are:

- *catapult down/up* - e.g. *More victories followed including a hard-fought points win over Lizo Matayi which should have catapulted him up for a national title challenge*,

- *split/tear in* - e.g. *The end of the square stick was then split in for a few inches*.

where all examples are from Google. This analysis revealed that all of the candidate VPCs in these classes are valid, which comes as a confirmation of the degree of productivity of these high ranked classes.

The classes that have a degree of productivity of 40% or more form 4,344 candidate VPCs, which when joined together with the combined resources obtain a total of 9,919 VPCs. This represents an increase of 20.74% in the coverage of A+C+E+BNC-VPCs, by making use of productive patterns found in VPCs.

As each of these particles occurs with a certain proportion of the verbs in a class, and this proportion varies considerably from class to class, and from particle to particle, further investigation was conducted to see the degree of productivity of individual class-particle pairs. The degree of productivity of each class-particle pair is determined by the proportion of verbs in that class that form valid combinations with that particle. Moreover, the larger the number of classes where the majority of verbs form valid VPCs with that particle, the more productive the particle is. Table 5 shows for each particle, the 5 classes that had the higher proportion of valid VPCs with that particle, according to A+C+E+BNC-VPCs. From these particles, the one that is involved in the larger number of combinations throughout more classes is *up*, which occurs with 40% or more of the verbs in a class for 54.7% of the
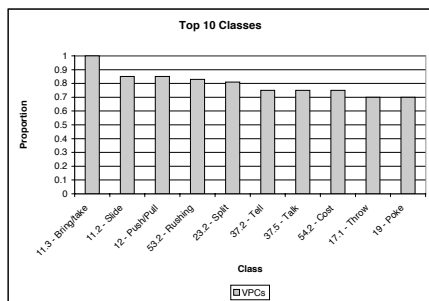
Figure 4: Levin's VPCs in Dictionaries + BNC

classes, and it is followed closely by *out*, as shown in Table 6. Thus *up* is the best predictor of valid verb-particle combinations, for most classes. On the other hand, the weakest predictor of valid combinations is *in*, which occurs in only a few classes, with 40% or more of the verbs. Class 11.3 is the best class predictor, allowing all verbs to combine with all particles.

The classes that have a degree of productivity of 40% or more with a given particle using this more specific measure, generate 4,719 VPCs, and these were used to extend the coverage of these resources obtaining a total of 9,896 VPCs. This represents an increase of 20.46% in the coverage of A+C+E+BNC-VPCs, by making use of productive patterns found in VPCs, and a very restricted set of particles.

Table 6: Classes with 40% or more of valid VPCs

Particle	Classes
up	54.7 %
out	53.7 %
down	46.7 %
in	9.9 %

These results suggest that patterns of productivity of VPCs can be mapped into Levin's classes. Whether choosing the more productive classes overall or the more productive class-particle pair the result is a significant increase in coverage of the lexical resources, when VPCs are generated from these classes. More investigation is needed to verify whether the unattested combinations, specially in the lower ranked classes are invalid or simply

did not occur in the dictionaries or in the corpus, because the problem of data sparseness is especially accute for VPCs. Moreover, it is also necessary to determine the precise semantics of these VPCs, even though we expect that the more productive classes generate VPCs compositionally, combining the semantics of the verb and particle together. Possible alternatives for dealing with this issue are discussed by both Bannard et al. (2003) and McCarthy et al. (2003). Furthermore, although there are some cases where it appears reasonable to treat VPCs as fully productive, there are also cases of semi-productivity (e.g. verbs denoting cooking processes and aspectual *up*: *boil up* and *heat up*, but not ?*sauté up*), as discussed by Villavicencio and Copestake (2002), so it is important to determine which classes are fully productive and which are not.

6 Discussion

We investigated the identification of regular patterns among verb-particle constructions using dictionaries, corpora and Levin's classes. These results suggest that Levin's classes provide us with productive patterns of VPCs. Candidate VPCs generated from these classes can help us improve the coverage of current lexical resources, as shown in this investigation. We used the available lexical resources and corpus data to give us an indication of class productivity, and we used this information to rank these classes. We took a sample of those classes that were considered to be good predictors of valid VPCs, and these were further investigated, through human judgements, confirming their correspondence with productive patterns in VPCs. Some of the patterns can also be applied to other related particles (e.g. the resultative pattern and locative/directional particles), but even using a small set of particles it was possible to considerably extended the coverage of these lexical resources.

More investigation into the productivity of the lower ranked classes is needed since the domain being considered was restricted to the combined resources, and we only considered a candidate VPC to be valid if it was listed in them. For instance, in a manual analysis of the combinations involving the class of *Roll* verbs (class 51.3.1, *bounce, drift,*

Table 5: Top 5 Classes for **up**,**down**, **out** and **in**

Class	UP	Class	DOWN	Class	OUT	Class	IN
11.2	100%	10.3	100%	10.3	100%	11.3	100%
11.3	100%	11.2	100%	11.2	100%	11.5	54%
12	100%	11.3	100%	11.3	100%	15.2	50%
19	100%	12	100%	10.9	100%	39.1	50%
37.5	100%	18.4	100%	37.2	100%	35.6	50%

drop, float, glide, move, roll, slide, swing) most of the verb-particles generated were considered acceptable.[4] In relation to the A+C+E-VPCs, we found that 64% of these combinations are not listed. The use of corpora significantly reduces this problem, so that when we also consider the BNC-VPCs, the results are much better, with 80.5% of the combinations being listed. But for some classes, such as those involving motion, not even the addition of corpus data helps, and a great proportion of the VPCs are not attested, even though most of the combinations are considered acceptable by native speakers. Thus, a more wide investigation using human judgement and a larger set of VPCs would be necessary, also using the World Wide Web as corpus. Nonetheless, these results are encouraging and confirm that these classes provide us with good predictors of VPC acceptability. Thus, the use of these classes to automatically generate verb-particle constructions, based on groups of verbs and particles presents a reasonable way of improving coverage of existing lexical resources.

Acknowledgments

References

Timothy Baldwin and Aline Villavicencio. 2002. Extracting the unextractable: A case study on verb-particles. In *Proc. of the 6th Conference on Natural Language Learning (CoNLL-2002)*, Taipei, Taiwan.

Ken Bame. 1999. Aspectual and resultative verb-particle constructions with Up. Handout for talk presented at the Ohio State University Linguistics Graduate Student Colloquium, May.

Colin Bannard, Timothy Baldwin, and Alex Lascarides. 2003. A statistical approach to the semantics of verb-particles. In *Proc. of the Workshop on Multiword Expressions: Analysis, Acquisition and Treatment*, Sapporo, Japan.

Dwight Bolinger. 1971. *The Phrasal Verb in English*. Harvard University Press, Harvard, USA.

John Carroll and Claire Grover. 1989. The derivation of a large computational lexicon of English from LDOCE. In B. Boguraev and E. Briscoe, editors, *Computational Lexicography for Natural Language Processing*. Longman.

Ann Copestake and Dan Flickinger. 2000. An open-source grammar development environment and broad-coverage English grammar using HPSG. In *Proc. of the 2nd International Conference on Language Resources and Evaluation (LREC 2000)*.

Bruce Fraser. 1976. *The verb-Particle Combination in English*. Academic Press, New York, USA.

Beth Levin. 1993. *English Verb Classes and Alternations - A Preliminary Investigation*. The University of Chicago Press.

Catherine Macleod and Ralph Grishman. 1998. Comlex syntax reference manual, Proteus Project. http://nlp.cs.nyu.edu/comlex.

Diana McCarthy, Bill Keller, and John Carroll. 2003. Detecting a continuum of compositionality in phrasal verbs. In *Proc. of the Workshop on Multiword Expressions: Analysis, Acquisition and Treatment*, Sapporo, Japan.

Richard Side. 1990. Phrasal verbs: sorting them out. *ELT Journal*, 44(2):144–52.

Aline Villavicencio and Ann Copestake. 2002. Verb-particle constructions in a computational grammar of English. In *Proc. of HPSG 2002*, Seoul, Korea.

[4]*drop up* is presumably disallowed because of contradictory directional properties.