

Quantifying Count/Mass Elasticity

Graham Katz and Roberto Zamparelli
Georgetown University and Università degli Studi di Trento

1. Introduction

The semantics of the mass/count distinction has long been a central issue in formal linguistics (Quine, 1960; Link, 1983), with much debate over the nature of these two classes of nouns (Doetjes, 1996; Chierchia, 2010; Rothstein, 2010; Pelletier, 2010). In the vast literature on the mass/count distinction the following criteria, listed by Chierchia (1998), have played a key role in distinguishing *mass uses* of nouns from *count uses*:

- (1) (i) determiners like *much, less, a little, a bit of, more* + N_{sing} are associated with mass uses;
- (ii) determiners like *every, a, each* are associated with count uses;
- (iii) ‘bare singulars’ arguments (*water/*table is wet, I touched water/*table*) are mass uses;
- (iv) plural uses are count uses;¹

A well-established semantic tradition holds that whether a noun can be used as a mass or a count depends on its denotation (Link, 1983). According to this view, mass nouns denote properties with *non-quantized* reference: p is non-quantized iff if $p(x)$ holds and y is part of x , then $p(y)$ holds. The intuition is that a (non-microscopic) subpart of a quantity of *water* still counts as *water* (Cheng, 1973). Count nouns, on the other hand, denote properties with *quantized* reference: p is quantized iff if $p(x)$ holds and y is part of x then $p(y)$ does not hold. A part of a *dog* is not a *dog*. Mass nouns, unlike singular count nouns, also denote cumulative properties (Quine, 1960): p is a mass property, then if $p(a)$ and $p(b)$ it follows that $p(a+b)$, where “ $a+b$ ” is the mereological sum of a and b . If a is quantity of *water* and b is a quantity of *water* then a and b together are *water*. The denotational view of the countability divide has long been known to have problems, one being the existence of many noun pairs, both within and across languages, which seem to denote the same objects but have different count/mass status. Examples are *shoes/footwear, coin/change* in English, pairs like Italian *capelli*_{count} vs. English *hair*_{mass}, etc., (Chierchia, 1998).

An alternative, popular among syntacticians and lexicographers alike (Chomsky 1965:82; Quirk et al. 1972:127), is the ‘lexicalist’ view, according to which nouns are lexically marked with a binary \pm COUNT feature which determines its syntactic countability status. The connection between this feature and the noun’s denotation is taken to be less direct. Nouns are assumed to bear this feature in much the way they bear grammatical gender or belong to a declension class. If this is the case, we should naturally expect, given the syntactic criteria in (1), that mass nouns should almost never occur in syntactically count contexts like (1-b) (e.g. **every oxygen*) and that count nouns should not be found in mass contexts like (1-a) (e.g. *much restaurant*). This position has been challenged by linguists and philosophers since at least the eighties (Allan, 1980; Borer, 2004; Pelletier & Schubert, 1989). According to Allan (1980), any noun can be in principle be ‘mass’ or ‘count’, and its countability status is fixed only at the level of the whole NP (in current terms, DP)—this is what Chierchia (2010) terms the *elasticity* of the mass/count distinction. Allan acknowledges that not all nouns are equally amenable to mass or count uses—nouns often prefer one or the other class. Since this preference cannot be entirely derived from their denotation (witness pairs like *coin/change, capelli/hair*), Allan proposes that nouns belong to various classes, with

*We would like to thank Marco Baroni, Egon Stemle and the audiences at the WCCFL 29 as well as QITL4 in Berlin for helpful suggestions. This research was carried out with the support of a DiSCoF fellowship from the University of Trento for the first author.

¹More exactly, plural uses which can be modified by a cardinal numeral; *plural mass nouns* such as *oats, belongings* do exist, but cannot be counted, see Acquaviva (2008).

an associated ‘massiness’ score. In the same spirit, Baldwin & Bond (2003) suggested a division of nouns into five major types, depending on their readiness to undergo countability shifts. Taken together, these results show that despite the existence of many intermediate or unclear cases, it still makes sense to talk about ‘mass noun’ or ‘count noun’, referring to those nouns with the largest/smallest propensity to be used as mass.

Lexicalists have pointed out that positing a binary lexical distinction does not imply that there is no way to move from ‘mass’ to ‘count’ or vice-versa. Indeed, three operations that shift the countability status of a noun by modifying its meaning have been widely discussed in the literature (Pelletier & Schubert 1989):

- (2) a. **kind-formation:** $N_{mass/count} \Rightarrow [type\ of\ N]_{count}$ (“Many wines are high in sulfites”)
 b. **packaging:** $N_{mass} \Rightarrow [standard\ container\ of\ N]_{count}$ (“I ordered three cokes”)
 c. **universal gridding:** $N_{count} \Rightarrow [material\ constituting\ N]_{mass}$ (“I put apple in the salad”)

What this implies is that while many and perhaps all nouns can be used in both mass and count contexts, their meaning will not be the same in the two contexts. Indeed, Barner & Snedeker (2005) give evidence that nouns are interpreted differently in syntactic contexts which favour a mass or a count reading. This is the starting point for the present study.

In much work on the mass/count distinction (such as Allan’s or Pelletier’s) hand-picked nouns considered representative of their lexical classes were used to determine the properties of these classes on the whole. One goal of this paper is to investigate this on a larger scale, using a corpus, and systematically verify some of the claims made in the theoretical literature. The central questions we propose to address in this paper are the following:

- (3) a. Is it true that the same nouns can appear both in ‘count’ and ‘mass’ contexts?
 b. Is it true that when a noun is used as a mass or a count its meaning shifts?

Answering the first question involves making use of reliable syntactic patterns that identify mass and count uses in a corpus. Answering the second one, however, requires the introduction of a measure of ‘meaning shift’. In this study we propose to apply the methods of *distributional* or *vector space* models (Landauer & Dumais, 1997; Sahlgren, 2006; Turney & Pantel, 2010) to measure meaning shift. The key idea is that if the plural use of a mass noun is always the result of the application of a shifting strategy (those in (2), or others), the singular/plural alternation should, in the case of a mass noun (“wine/wines”), be accompanied by a meaning shift which should be *smaller* than the shift one obtains in the singular/plural alternation of a plain count nouns (“dog/dogs”). Vector space models give us a way to quantify the degree of meaning shift in precise, quantitative terms, capitalizing on the idea that the meaning of content words can be captured by the pattern of co-occurrence with other content words within a large-scale corpus of natural language.

The structure of the paper is as follows. In the next section, we will address question (3-a), discussing two lexical patterns capable to identify count and mass *uses* of a noun. In Section 3 we describe in some detail the flavor of distributional semantic model we will need for the task. Section 4 addresses question (3-b): with a measure of semantic distance in place, we can order nouns in terms of how much their singular and plural forms differ and analyze those which are further apart. The issue is whether this shift correlates with their ability to appear in syntactically ‘mass’ contexts as defined in Section 2. Section 5 gives our interim conclusions and some hints to future work.

2. Corpus Study: Distinguishing mass and count via syntactic patterns

To study the elasticity of the the count/mass distinction we need reliable syntactic patterns that select each of these uses, and a large corpus. As data source we used UKWAC, a nearly 3-billion word corpus of English extracted from the WWW, POS-tagged and lemmatized (Baroni & Ueyama, 2006). Queries on this corpus were performed using the CQP tool, which is part of the Corpus Workbench software package.² Our preliminary goal was to determine the degree to which the syntactic criteria listed above distinguished a set of mass nouns from a set of count nouns.

²See <http://www.ims.uni-stuttgart.de/projekte/CorpusWorkbench/>

To design our query for extracting mass and count nouns, we started from the criteria in (1). We restricted the analysis to nouns that occurred more than 50 times in the corpus, so results would be reliable. To avoid misclassification due to NN-compounds (*much computer software*, *much computer inspired design*) we also excluded nouns followed by other nouns, participles and adjectives. Since UKWAC contains a large number of bare nouns which are parts of web pages titles or headlines, we started by using criteria (i) for mass and (ii) for count. To increase the recall, we used criterion (i) to extract some nouns, hand-selected those which seemed unquestionably mass (*water*, *patience*, etc.), and checked for modifiers which appeared with them but never with singular nouns which were ‘count’ according to criterion (ii). This gave us adjectives such as *considerable*, *boundless* and *ample*. The final pattern was:³

- (4) a. **Mass:** any singular noun not followed by another noun, participle or adjective and immediately preceded by *lots*, *plenty of*, *much*, *more*, *less enough*, *most*, *sufficient* *considerable*, *boundless*, *ample* or *limited* not preceded by *a(n)*
 b. **Count:** any singular noun not followed by another noun and immediately preceded by *a*, *an*, *one*, *every*, *first*, *each*, *another*.

We also extracted bare singular arguments by simply counting the number of times a non-compound noun appeared directly following a verb (as in *drink water*). At the end, the rate of occurrence as a bare singular argument turned out to be highly correlated (coefficient = 0.53) with the rate at which a noun occurred in mass contexts, but not correlated at all (0.08) with the rate in which a noun occurred in count contexts. This confirms that the various criteria are picking out the same set of nouns, by and large. Qualitative analysis shows that these distributional criteria select very much the kinds of nouns we would think of as lexically count, on the one hand, and lexically mass on the other. The 25 nouns used most frequently in **mass** contexts are listed in (5-a), while those used most frequently in **count** contexts are listed in (5-b)

- (5) a. **Mass contexts:** information, time, money, detail, space, fun, attention, info, part, work, interest, evidence, experience, energy, power, water, room, recipe, use, opportunity, effort, emphasis, support, research, trouble
 b. **Count contexts:** time, year, day, way, person, place, bit, week, man, opportunity, problem, lot, thing, role, company, basis, child, look, one, report, month, book, area, approach, hour


Notice that there are nouns that appear on both lists. While this does not necessarily mean that the countability status is a not a lexical distinction (words like *time* and *information* could simply be ambiguous), it turns out that nouns that appear in mass context often appear in count contexts, to varying degrees of frequency. This is illustrated by the frequency histogram in (6-a). Moreover, it is (perhaps surprisingly) **not** the case that the rate with which a noun is used as a mass expression is inversely proportional with the rate at which that noun is used as a count expression. Specifically, the rate of use with mass determiners is essentially uncorrelated (-0.028) with the rate of use with count determiners, on a per-noun basis (see Fig. (6-b)): Had there been an inverse correlation, we would have seen a band stretching obliquely from the y to the x-axis.

³The CQP queries were:

MASS = [word!="alan"] ("lots | ("plenty" "of") | "much" | "more" | "less" | "enough" | "most" | "sufficient" | "considerable" | "boundless" | "ample" | "limited") @[pos="NN"] [pos!="VVNIJJ.*IN.*"]

COUNT = [lemma="alanloneevery|firsteachlanother"] [pos="J.*"]* @[pos="NN"] [pos!="N.*"]

BAREARG = [pos="V.*"] @[pos="NN"] [pos!="N.*"]



While we were not surprised that some nouns appear in both mass and noun contexts—we were after all well-aware of the elasticity phenomenon—the wide range of mass uses for ‘obvious’ count nouns such as *car*, *ski* and *piano* does seem surprising:

- (7)
- a. How much car you can afford?
 - b. Phantom 87s were too much ski for me.
 - c. I always practice too much piano.

One thing these cases show is that the meaning shifts which seem to take place when a count form is used as a mass are a lot more diverse than the three simple rules in (2) suggest. For instance, (7-a) bears a resemblance with a universal gridding case, but not at the level of material substance (as if *car* was first reinterpreted as a more abstract concept, then measured; *much car* ends up referring to a scale of ‘car types’ whose unit of measure is cost or features); (7-b) might be a comment on the extent to which the length or rigidity of a pair of skis exceeded a comfortable level; (7-c) builds on the bare singular *practicing piano*, with a meaning akin to *I always do too much piano-practicing*. What is evident here and in many other examples produced by the experiment is the crucial role of *conceptual abstraction* in the noun’s reinterpretation. Indeed, a glance at the list in (5-a) shows a huge predominance of *abstract* nouns (the only real exception is *water*). The canonical mass nouns familiar from the linguistic and philosophical literature (e.g. *water*, *gold*, *furniture*) are nearly absent from the core data.


In the next section we investigate whether this meaning shifts can be effectively detected using the methodology of *vector space semantics*. While this method would not yet give us a way to identify the *type* of shift involved, it could be an important first step for any qualitative analysis of this data.

3. Vector Space Models: A brief overview

Vector space models of meaning (Landauer & Dumais, 1997; Sahlgren, 2006; Turney & Pantel, 2010) exploit the idea that a word’s meaning is reflected in its patterns of use: In large corpora, a word such as *cat* co-occurs frequently with such words as *food*, *house* and *mouse*, while *roof* co-occurs frequently with *house*, *wall* and *slate*. In vector space models, a word’s meaning is modeled as a vector that encodes a measure of how frequently this word co-occurs with other words, the “context” words, typically within a sentence or a small “window.” These patterns of co-occurrence can be used to derive measures of semantic similarity which closely model semantic similarity judgements elicited from speakers (Landauer & Dumais, 1997). To illustrate, in a small corpus, we might collect the co-occurrences of the words *roof*, *cat* and *dog* with the context words *house* and *food*. The counts displayed in (8-a) gives rise to this simple two dimensional vector space illustrated in (8-b).

(8) a.

	“food”	“house”
<i>roof</i>	1	4
<i>dog</i>	5	1
<i>cat</i>	4	2



In this space the distance between *dog* and *roof* is large, with *cat* somewhere between them. Typically semantic similarity between two words is computed on the basis of the cosine of the angle between these words, giving rise to a similarity metric normalized to the unit interval, with 1 corresponding to synonymy and 0 to unrelated. Some models are built so as to distinguish between unrelated words and those with antonymous meaning. Our model is not, since we were only concerned with degrees of similarity in the nominal domain, a domain for which antonymy is less well-defined.

Vector space models are distinguished in a number of ways (Turney & Pantel, 2010), including which context words are used, what the size of the co-occurrence window is, the way in which corpus-counts are used to produce a model of semantic similarity, as well as other factors. In the work described here we adopted the COALS vector space model (Rohde et al., 2005), a method that focusses on very local co-occurrence and which has been shown to be among the most accurate models of semantic similarity. In this model the 14,000 most frequent content words in the ukWaC corpus were chosen as the dimensions (we excluded 571 high-frequency words which are on a standard “stopword” list (Salton, 1971)). The model employs a 4 word ramped window, meaning that, for a given target word, only the four preceding and subsequent content words in the corpus were counted, and the co-occurrence was weighted by proximity. To normalize for frequency effects, these weighted frequency vectors are transformed into vectors of (positive) correlation coefficients for each word, giving a measure not just of how frequently two words co-occurred, but how informative that co-occurrence was. Although the COALS model is relatively sophisticated, there is no reason to believe that any other choices of context size, weighting, and normalization would significantly alter the results we report below.

An important question for any model concerns the identification of lexical items. In the earliest models (Landauer & Dumais, 1997; Lund & Burgess, 1996) a naive notion of lexical item was employed—words were taken to be sequences of characters separated by spaces. This, of course, is highly problematic, as it identifies homonyms—the noun *break* and the verb *break*, for example, are collapsed—but fails to combine information associated with wordforms that belong to the same lexeme—*break* and *broke*, for example, are treated as separate words. More recent models (Erk & Padó, 2008; Baroni & Lenci, 2010) have attempted to make use of more linguistic sophistication by computing vectors for lexemes, making use of part-of-speech tagging and lemmatization. Thus, different vectors are computed for the nominal uses of *break* on the one hand (combining counts for *break* and *breaks*) and the verbal use of *break*, on the other (combining *break*, *breaking*, *broke*, etc.).

Our model was more fine-grained, still. Making use of the part-of-speech annotation (based on the Penn Treebank tagset (Marcus et al., 1994) and lemmatization provided with the UKWAC corpus, generated using the TreeTagger (Schmid, 1999), we computed semantic vectors for words in a given part of speech—for example, *break* as a present tense verb, *break* as a progressive, *break* as a plural noun and *break* as a singular noun. The elements in our model, then, were lemma-POS pairs, such as **break-VBD** [*broke*], **break-NN** [(*a break*)], **break-NNS** [(*two breaks*)]. We made no effort to distinguish different senses of words (for example the physical and the temporal notion of *break*).

We collected co-occurrence statistics for the 20,000 most frequent lemma-POS pairs in the ukWaC corpus, making use of the implementation of the COALS model, made available by airhead software,⁴ to build a 14,000 dimension model. In the model thus generated the ten words most similar to *buses* (**bus-NNS**) were the following. Note that *bus* (**bus-NN**), is the fifth most-similar word:


⁴See <http://code.google.com/p/airhead-research/wiki/Coals>

		Similarity to bus-NNS
1	train-NNS	0.5442344251525053
2	taxi-NNS	0.4368242277700681
3	Metro-NP	0.4214840363161869
4	lorry-NNS	0.4186066577599463
(9)	5 coach-NNS	0.4084300268705176
6	bus-NN	0.374251933421374
7	Stagecoach-NP	0.33681906394181077
8	fleet-NN	0.33285009451198716
9	run-VVP	0.317420799073372
10	stop-VVP	0.3088151950886271

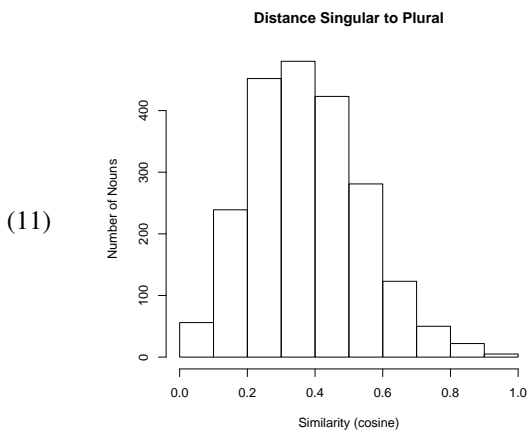
This gives a good sense of the kinds of semantic similarity that distributional techniques such as those described above tend to model.

4. Plural/Singular distance and the Count/Mass distinction

The question we had was whether this kind of model was fine grained enough to distinguish semantic meaning shifts of the sort presented by mass-count elasticity. Our expectation was that the semantic distance between the singular and plural form would be higher for nouns that have predominantly mass uses, since these nouns must undergo some sort of semantic coercion in order to be used as count nouns (and thus be used in the plural). If the model were sensitive to mass/count elasticity, *man* and *men* would have “meaning” vectors which were similar (or **near**) to one another while *water* and *waters* would have vectors that were somewhat less similar (or **far**) from one another, is illustrated in (10).



Of the 20,000 (lemma-POS) pairs in our model, there were 2114 nouns that appeared in both both singular (e.g., **apple-NN**) and (marked) plural (e.g., **apple-NNS**) forms.⁵ The average singular/plural similarity in our model was 0.378 with a standard deviation of 0.163. The distribution of nouns by singular/plural distance is illustrated in (11).



⁵We excluded from analysis 17 word forms such as *metadata* which appeared in the corpus tagged as both NN and NNS but had no marked plural, as the tagging for these items appeared to be quite unreliable.

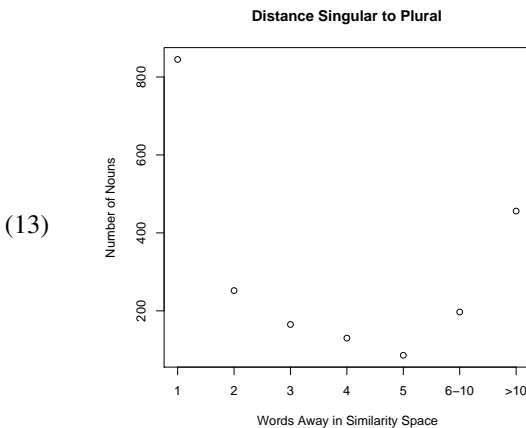
The twenty nouns whose singular/plural similarity were greatest are listed in (12-a), while those whose singular/plural semantic similarity was smallest are listed in (12-b)

- (12) a. **Near Nouns:** *minority, seeker, acid, surgeon, tale, examiner, beach, guitar, venture, bomber, cancer, sector, pool, locomotive, engine, cookie, speaker, century*
 b. **Far Nouns:** *leave, make, creator, con, humanity, extreme, good, disadvantage, toddler, strength, fortune, horizon, total, storey, hip, mouse, dozen, tip, monkey, security*

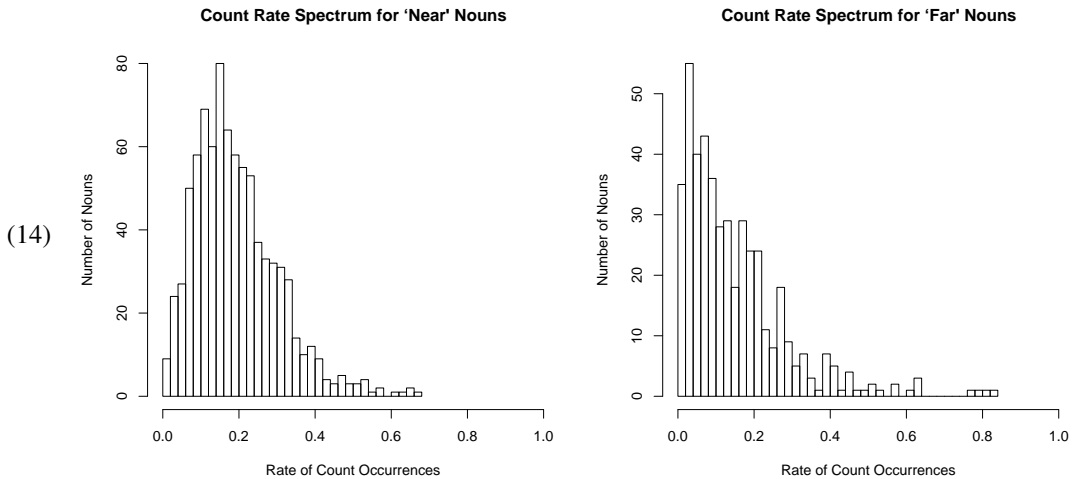
The intuitive contrast between the **near** and **far** nouns above is clear: All the **near** nouns are semantically similar for the same reason—they are nouns for which the singular and plural forms are directly semantically related (one *surgeon* many *surgeons*, for example)—while the **far** nouns differ semantically for a wide variety of reasons. These range from simple homography ((a) *con* (*job*) and (*the pros and*) *cons*) to lexical ambiguity (*security/securities* and *humanity/humanities*), to polysemy (*strength/strengths, fortune/fortunes*). In addition there are simple lemmatization error in the corpus; for example *leaves* was always treated as the plural of *leave* and never of *leaf*, so occurrences of *leaves* in arboreal contexts give rise to a large plural/singular distance.

In order to analyze the set of **near** and **far** nouns quantitatively, we introduce a slightly different measure of semantic similarity: the ordinal words-apart metric. As we saw above, the closest word to **bus-NNS** in the semantic space was **train-NNS**, with cosine similarity of just over 0.54. In contrast, the similarity between **accent-NN** and its nearest neighbor in the semantic space **dialect-NN** is just 0.34. The ordinal ranking metric measures semantic distance in terms of the number of words in the semantic space that separate two words. On this metric **train-NNS** and **bus-NNS** have a semantic distance of 1, as do **accent-NN** and **dialect-NN**.

As illustrated in (13), for 837 nouns in our model the singular (**NN**) form is the nearest neighbor of the plural (**NNS**) form, while for 453 of the nouns the singular wasn't among the ten words closest to the plural.



Because for nouns typically used predominantly in mass contexts, a plural use signals a semantic shift, we hypothesized these two sets of nouns would differ markedly in the degree to which the lexemes are used in count and in mass contexts. We expected the **near** nouns to be more predominantly count nouns and the **far** nouns to be more predominantly mass nouns. To investigate this, we used our count- and massiness measure from Section 2. Looking first at mass contexts we find that the average rate of mass contexts occurrence (by noun type) for the **far** nouns is significantly higher than that for **near** nouns (0.0028 vs. 0.0014, $p = 0.003041$). Even more dramatic is the difference in the distribution of count-context rates between the two groups:



The difference between group means (0.192 vs. 0.149) is highly significant ($p = 7.09e-09$). In short, the **near** group contains more predominantly count nouns than the **far** group, while the **far** group contains more predominantly mass nouns.

5. Conclusions and further work

What we have shown is that the class of nouns that undergo count/mass shifting is quite a bit wider than what is typically mentioned in the literature, and that the range of meaning shifts that this ‘elasticity’ entail appears to be wider than is typically accepted. Furthermore, we have dispelled the view that there is something like a *count/mass* spectrum, on which nouns fall: as we have seen, many nouns which appear frequently in count contexts also appear frequently in mass contexts.

In addition, we have seen that vector space models of meaning appear to be sensitive to the shift between count and mass nouns. Those nouns whose singular uses are semantically distinct from their plural uses are more likely to be those which are used more frequently in mass contexts. This indicates that the plural uses involve a meaning-shifting mass to count conversion. We have also uncovered a wide range of other cases in which the singular and plural forms of a noun have different distributions, itself an interesting result (and one frequently ignored by vector space models).

There is much that remains to be done. In this study, we did not directly distinguish count from mass uses, but instead used the plural/singular uses as proxy for this. This is problematic, since the same meaning shift we find in the plural might also have applied in the singular: plurals might indeed be all count, but singulars are unavoidably a mix of mass and count uses (consider *I like the wine/wines that grow(s) in this region* vs. *I just drank wine from a plastic cup*). Clearly, if it is the case that a given noun undergoes a semantic shift when its count or mass status changes, it would be better to try to detect this effect in syntactic environments which make the countability status as unambiguous as possible, e.g. *too much car* vs *a car*. In ongoing work, we are directly investigating the semantic distance between count singular and count mass occurrences. The next natural research question would be to establish if these methods can also be used to recognize meaning shifts which correlate with a mass/count shift from others, and ideally, discover shifts which are the hallmarks of a *specific* kind of mass/count alternation (say, kind formation). We also intend to pursue this issue in future work.

References

- Acquaviva, Paolo (2008). *Lexical plurals. A Morphosemantic Approach*. Oxford University Press, Oxford.
- Allan, Keith (1980). Nouns and countability. *Language* 56, 541–567.
- Baldwin, Timothy & Francis Bond (2003). Learning the countability of english nouns from corpus data. *ACL '03: Proceedings of the 41st Annual Meeting on Association for Computational Linguistics*, Association for Computational Linguistics, Morristown, NJ, USA, 463–470.
- Barner, David & Jesse Snedeker (2005). Quantity judgments and individuation: evidence that mass nouns count. *Cognition* 97, 41–66.

- Baroni, Marco & Alessandro Lenci (2010). Distributional memory: A general framework for corpus-based semantics. *Computational Linguistics* .
- Baroni, Marco & Motoko Ueyama (2006). Building general-and special-purpose corpora by web crawling. *Proceedings of the 13th NIJL International Symposium, Language Corpora: Their Compilation and Application*, Citeseer, 31–40.
- Borer, Hagit (2004). *Structuring sense: In name only*. Oxford University Press, Oxford.
- Cheng, C-Y. (1973). Response to Moravcsik. Hintikka, J., J.M.E. Moravczik & P. Suppes (eds.), *Approaches to Natural Language*, D.Reidel, Dordrecht, 286–288.
- Chierchia, Gennaro (1998). Plurality of mass nouns and the notion of “semantic parameter”. Rothstein, Susan (ed.), *Events and Grammar*, Kluwer, Dordrecht, 53–104.
- Chierchia, Gennaro (2010). Mass nouns, vagueness and semantic variation. *Synthese* 174:1, 99–149.
- Chomsky, Noam (1965). *Aspects of the Theory of Syntax*. MIT Press, Cambridge, Mass.
- Doetjes, Jenny (1996). Mass and count: syntax or semantics. Tech. rep., HIL/Leiden University. In: *Meaning on the HIL*. Occasional Papers in Linguistics.
- Erk, Katrin & Sebastian Padó (2008). A structured vector space model for word meaning in context. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, 897–906.
- Landauer, Thomas K. & Susan T. Dumais (1997). A solution to plato’s problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review* 104:2, 211–240.
- Link, Godehard (1983). The logical analysis of plurals and mass terms: a lattice-theoretical approach. Bäuerle, Rainer, Christoph Schwarze & Arnim von Stechow (eds.), *Meaning, Use, and the Interpretation of Language*, deGruyter, Berlin/New York, 302–323.
- Lund, Kevin & Curt Burgess (1996). Hyperspace analog to language (HAL): A general model of semantic representation. *Language and Cognitive Processes* .
- Marcus, Mitchell P., Beatrice Santorini & Mary A. Marcinkiewicz (1994). Building a large annotated corpus of english: The penn treebank. *Computational Linguistics* 19:2, 313–330.
- Pelletier, Francis Jeffrey (ed.) (2010). *Kinds, Things and Stuff*. New directions in cognitive science, Oxford University Press.
- Pelletier, Francis Jeffrey & Lenhart K. Schubert (1989). Mass expressions. Gabbay, D. & F. Guentner (eds.), *Handbook of Philosophical Logic*, Kluwer, Dordrecht, vol. 4, 327–407.
- Quine, Willard Van Orman (1960). *Word and Object*. MIT Press, Cambridge, Mass.
- Quirk, Randolph, Sidney Greenbaum, Geoffrey Leech & Jan Svartvik (1972). *A grammar of contemporary English*. Longman, London.
- Rohde, Douglas L. T., Laura M. Gonnerman & David C. Plaut (2005). An Improved Model of Semantic Similarity Based on Lexical Co-Occurrence. Ms. MIT.
- Rothstein, Susan (2010). Counting and the Mass/Count Distinction. *Journal of Semantics* 27:3, p. 343.
- Sahlgren, Magnus (2006). *The word-space model: Using distributional analysis to represent syntagmatic and paradigmatic relations between words in high-dimensional vector spaces*. Ph.D. thesis, Stockholm University, Stockholm, Sweden.
- Salton, Gerard (1971). *The SMART retrieval system—experiments in automatic document processing*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- Schmid, Helmut (1999). Improvements in part-of-speech tagging with an application to german. *Natural language processing using very large corpora* 11, 13–26.
- Turney, Peter & Patrick Pantel (2010). From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research* 37, 141–188.

Proceedings of the 29th West Coast Conference on Formal Linguistics

edited by Jaehoon Choi, E. Alan Hogue,
Jeffrey Punske, Deniz Tat,
Jessamyn Schertz, and Alex Trueman

Cascadilla Proceedings Project Somerville, MA 2012

Copyright information

Proceedings of the 29th West Coast Conference on Formal Linguistics
© 2012 Cascadilla Proceedings Project, Somerville, MA. All rights reserved

ISBN 978-1-57473-451-5 library binding

A copyright notice for each paper is located at the bottom of the first page of the paper.
Reprints for course packs can be authorized by Cascadilla Proceedings Project.

Ordering information

Orders for the library binding edition are handled by Cascadilla Press.
To place an order, go to www.lingref.com or contact:

Cascadilla Press, P.O. Box 440355, Somerville, MA 02144, USA
phone: 1-617-776-2370, fax: 1-617-776-2271, sales@cascadilla.com

Web access and citation information

This entire proceedings can also be viewed on the web at www.lingref.com. Each paper has a unique document # which can be added to citations to facilitate access. The document # should not replace the full citation.

This paper can be cited as:

Katz, Graham and Roberto Zamparelli. 2012. Quantifying Count/Mass Elasticity. In *Proceedings of the 29th West Coast Conference on Formal Linguistics*, ed. Jaehoon Choi et al., 371-379. Somerville, MA: Cascadilla Proceedings Project. www.lingref.com, document #2723.