Coverage-based Frequency Bands: A Proposal

Dale Brown

Kanazawa University

doi: <http://dx.doi.org/10.7820/vli.v06.2.Brown>

Abstract

Second language vocabulary research makes much use of word frequency lists and their division into bands. In recent years, bands of 1,000 items have become conventional. However, there does not seem to be any firm basis or rationale for this. Conventional banding may be questioned since the utility of words varies greatly depending on frequency, because there are enormous differences in frequency within higher bands, and because the reliability of the placement of words in bands becomes progressively poorer at lower frequency levels. This article suggests an alternative approach: basing bands on coverage levels. Because of the frequency distribution of words, this means the highest frequency bands would contain very few words, while lower frequency bands would contain a great many words. The article shows how such bands can be constructed and presents a re-analysis of the results of a vocabulary test designed with conventional bands in terms of coverage-based bands. This re-analysis produces a very different profile of learners' knowledge, and it is argued that the shape of this profile may be more useful in terms of guiding instruction in that it gives a clearer indication of which words should be targeted for a group of learners. It is further argued that the smaller number of words contained in coverage-based bands at higher frequency levels makes them a more feasible basis for instruction. The article thus concludes that coverage-based bands may be a fruitful avenue for researchers to explore.

1. Introduction

Much work on L2 vocabulary makes use of word frequency lists and in particular their division into bands. Bands of 1,000 items have become conventional through applications in vocabulary testing, such as Meara's yes/no tests (Meara & Milton, 2003; Meara & Miralpeix, 2006) and the vocabulary size test (Nation & Beglar, 2007), and the lexical profiling of texts (e.g., Nation, 2006; Webb & Rodgers, 2009) and of learners (Laufer & Nation, 1995; Milton, 2007).

There does not, however, appear to be any firm basis or an explicit rationale for the use of 1,000-item bands. Kremmel (2016) has recently criticized this convention since the utility of words varies greatly with frequency. This is clear in the coverage levels provided by successive bands. For example, Webb and Rodgers (2009) looked at coverage in a corpus of film scripts and found that while the first 1,000 word families covered 86.5% of the words in the corpus, the second 1,000 covered 4.2%, the third 1.7%, down to the tenth 1,000 which covered just 0.1% of the scripts. A second problem is that at higher frequency levels there is massive variation in the frequency of the items within a band. In Davies' COCA-based

(2008) lemmas list, for example, the most frequent lemma is 550 times more frequent than the 1,000th lemma, yet both, being in the first band, may be similarly regarded as high frequency items. By contrast, the 1,000th lemma is only 98 times as frequent as the 20,000th lemma. A third problem is that the reliability of the placement of words in particular bands is good at the highest frequency levels, but becomes progressively poorer for lower frequency bands. Comparing the first five 1,000-lemma bands from Davies with Kilgariff's (1995) similar BNC-based lemma bands, 83% of the lemmas overlap at the 1 K level, 61% at the 2 K level, 46% at the 3 K level, 34% at the 4 K level, and just 25% at the 5 K level. This problem is confirmed by Sorell's (2013) much more extensive data, while Miller and Biber (2015) raised concerns regarding the reliability of frequency rankings even when investigating a very narrow, focused corpus.

Kremmel's (2016) suggestion for reforming frequency banding was for bands of different widths at different frequency levels: 500-item bands at the highest frequency level, 1,000-item bands at mid-frequency levels, and 2,000-item bands at lower frequency levels. This article considers a more radical step in this direction and investigates whether the width of bands might be profitably based on coverage levels. Due to the frequency distribution of words, which is described by Zipf's (1935) law, this produces very small bands at the highest frequency levels and very wide bands at lower frequency levels. This is not an entirely new proposal: Edwards and Collins (2013), in developing a mathematical model of vocabulary acquisition, used bands based on the logarithmic function of a word frequency distribution, which has a similar outcome. There does not, however, seem to have been any previous discussion of how different approaches to banding may give us different perspectives on L2 vocabulary.

This article demonstrates how coverage-based bands may be constructed and considers how such bands may change our perceptions of vocabulary learning and teaching. This is achieved through the re-analysis of a data set from a yes/no test, which was constructed using conventional banding, on the basis of coverage-based bands.

2. Constructing Coverage-Based Frequency Bands

To construct coverage-based bands, Davies' COCA-based (2008) 20,000 lemmas list (which in fact includes a total of 20,200 lemmas) was used. The combined frequency of these lemmas is approximately 352 million, so, to create 10 bands, those at the top of the list were assigned to Band 1 until their combined frequency approximated 35.2 million. The next lemmas in the list were then assigned to Band 2 until approximately the same frequency was reached, and this process followed for each subsequent band. The result (Table 1) was that the highest bands are very narrow, while Band 10 is extremely wide. It was also not possible to divide the lemmas perfectly into 10 bands, Band 2 having a combined frequency somewhat larger than desired.

Next, a content-words-only approach was trialed. Gardner (2013) suggests that function words may be acquired differently from content words and notes they can distort frequency counts; Kremmel (2016) points out that most vocabulary tests do not include function words; while Milton (2009) suggests it can be useful at times for researchers to exclude them from frequency counts. Thus, all lemmas not listed

Table 1. Davies' Lemmas List Divided into 10 Bands Based on Coverage

Band	Lemmas included	Combined frequency of lemmas	No. of lemmas
1	1–2	34,584,731	2
2	3–6	38,225,595	4
3	7–15	35,052,835	9
4	16–35	35,165,890	20
5	36–83	35,017,770	48
6	84–210	34,979,487	127
7	211–514	35,034,325	304
8	515–1,243	35,010,971	729
9	1,244–3,315	35,001,760	2,072
10	3,316–20,200	33,922,071	16,885
Total		351,995,435	20,200

as nouns, verbs, adjectives, or adverbs were removed from the frequency list: a total of 446 (2.2%) of the lemmas. Ten bands were then formed with the 19,754 remaining lemmas, and this time the bands had very similar combined frequencies (Table 2).

These coverage-based bands directly deal with the first concern with conventional bands about differences in the utility of words of different frequency levels. They also address the second concern in that the items within each band are of more similar frequencies and there is a similar level of consistency across the bands. Finally, it is likely that coverage-based bands have greater reliability: with much wider bands at less frequent levels, about which there is less certainty of the frequency of items, it is more likely that most of the words would appear in the same band were bands constructed using another frequency list.

3. Using Coverage-Based Frequency Bands

One important use of frequency bands is in vocabulary testing, with a number of tests basing item selection on frequency bands. A test derived from

Table 2. Davies' Lemmas List, with Function Words Excluded, Divided into 10 Bands Based on Coverage

Band	Lemmas ranked	Combined frequency	No. of lemmas
1	1–3	19,423,367	3
2	4–27	19,305,146	24
3	28–88	19,583,571	61
4	89–207	19,561,491	119
5	208–412	19,590,257	205
6	413–767	19,557,796	355
7	768–1,375	19,564,642	608
8	1,376–2,525	19,559,456	1,150
9	2,526–5,275	19,555,439	2,750
10	5,276–19,754	19,652,327	14,479
Total		195,353,327	19,754

coverage-based bands could be designed in the same way. The vocabulary size test (Nation & Beglar, 2007), for example, uses a 1% sampling rate (i.e., one word on the test represents 100 words on the frequency list), while Meara's yes/no tests (Meara & Milton, 2003; Meara & Miralpeix, 2006) use a 2% sampling rate. Applying these rates to the coverage-based bands in Table 2, a test would require 201 items with a 1% sampling rate or 396 items with a 2% sampling rate (Table 3). These figures would, however, drop dramatically were Band 10 omitted. A test based on coverage-based bands is, therefore, quite feasible.

For the purposes of this article, it was unfortunately not possible to design and administer a test in this way. Instead, an existing data set was re-analyzed. These data were collected with a yes/no test, with items from Meara and Milton's (2003) X_Lex test and Meara and Miralpeix's (2006) Y_Lex test. The 240 items comprised 20 items each from 10 frequency bands, along with 40 pseudowords. Data were initially collected from 466 learners at four universities in Japan. Ninety-seven learners were excluded: 17 with an L1 other than Japanese, 10 since they had not fully completed the experimental tasks or were unwilling for their data to be used, and 70 with a false alarm rate over 25%. This left 369 learners in the study.

Figure 1 shows the mean percentage of words indicated as known in each of the 10 original frequency bands. This graph presents a familiar profile (e.g., Aizawa, 2006; Milton, 2007; Brown, 2012; Booth, 2013) with a gradual decline in the proportion of words known across the upper frequency bands. At lower frequency levels, there is more variability, which likely reflects the lower reliability of the placement of words in those bands.

The data were then re-analyzed by first finding the frequency ranking of each test item in Davies' lemmas list. The items were then assigned to the coverage-based bands in Table 2. Table 4 shows that the test contained no items at all from the first two bands, but had a reasonable sample of items from the other bands. Figure 2 thus shows the profile for the mean percentage of words indicated as known in the coverage-based bands. This shows high levels of knowledge (around ceiling level) up to Band 8, then a sharp drop at Band 9, and a further decline at Band 10.

Table 3. Sampling Coverage-Based Bands for a Vocabulary Test

Band	No. of lemmas	1% sample	Rounded (min. frequency = 1)	2% sample	Rounded (min. frequency = 1)
1	3	0.03	1	0.06	1
2	24	0.24	1	0.48	1
3	61	0.61	1	1.22	1
4	119	1.19	1	2.38	2
5	205	2.05	2	4.1	4
6	355	3.55	4	7.1	7
7	608	6.08	6	12.16	12
8	1,150	11.5	12	23	23
9	2,750	27.5	28	55	55
10	14,479	144.79	145	289.58	290
Total	19,754		201		396

Figure 1. Mean percentage of words known across the 10 original frequency bands.

Table 4. Sampling Rate of the Yes/No Test by Coverage-Based Frequency Bands

Band	No. of lemmas in band	No. of lemmas in yes/no test	Sampling rate (%)
1	3	0	0
2	24	0	0
3	61	3	4.9
4	119	4	3.4
5	205	10	4.9
6	355	9	2.5
7	608	19	3.1
8	1,150	17	1.5
9	2,750	35	1.3
10	14,479	91	0.6

A further analysis by deciles for both the 10 original bands and the 10 coverage-based bands is presented in Figures 3 and 4, respectively. Figure 3 shows that for each decile the approximately linear decline in the percentage of words known across the frequency bands is simply somewhat lower in each case. Figure 4 suggests that the pattern of ceiling levels of knowledge followed by a sharp drop seen in Figure 2 may begin at different points for different deciles.

These results suggest that coverage-based bands may be useful with regard to a second important application of word lists: guiding curriculum development and programs of instruction. The profiles in Figure 3 provide little guidance for instruction, it being unclear which of five or six bands should be targeted for any of the deciles. The profiles in Figure 4 seem somewhat more helpful. Although the results in this case are far from clear-cut, perhaps because many of the learners have already largely progressed through the majority of

Figure 2. Mean percentage of words known across 10 coverage-based frequency bands.

Figure 3. Mean percentage of words known across the 10 original frequency bands by deciles.

the bands and/or because the yes/no test was not designed specifically to investigate knowledge of these bands, a test derived from coverage-based bands may indicate more clearly the band at which learners' knowledge tails off. Coverage-based bands may, therefore, be of use diagnostically in identifying targets for instruction.

Figure 4. Mean percentage of words known across 10 coverage-based frequency bands by deciles.

Coverage-based bands may also be more practical with respect to instruction. Bands of 1,000-items are arguably too broad: providing instruction for 1,000 items is a considerable challenge within the typical horizon of a language course. Webb and Chang's (2012) study of Taiwanese learners over 5 years found typical growth of 200–300 word families a year, and they thus suggest a target of 400 families a year for a well-designed program. At this rate, it would take 2.5 years to progress through a single 1,000-family band. Coverage-based bands provide much smaller targets at the earlier stages of learning, which may well be more manageable and feasible (see also Dang and Webb [2016], who similarly argue that 1,000-item bands are impractical for pedagogical purposes and who divide their 800-item Essential Word List into 50-item bands to create readily attainable targets for teaching).

The corollary to small, manageable bands at high frequency levels is that at lower frequency levels the bands are much larger. These bands, of several thousand or many thousands of items, are clearly not manageable or feasible for explicit instruction. However, while there is debate about the exact point at which it should occur, it has been suggested that at some point in learners' development explicit teaching and learning should become secondary and implicit learning come to the fore (Nation, 2001; Schmitt, 2008). In addition, once learners have developed a good foundation of vocabulary, more specialized vocabulary appropriate to their learning purpose may be a more appropriate target for instruction (Nation, 2001). Thus, while the lowest bands may be infeasible for direct instruction, this may not be a major issue.

4. Conclusion

This article has proposed that dividing word frequency lists into bands based on coverage may have advantages over conventional 1,000-item bands. Conventional banding is questioned since words vary greatly in their utility depending on frequency, because such bands do not make the enormous differences in frequency within higher bands readily apparent, and because the reliability of the placement of words in bands becomes progressively poorer at lower frequency levels.

Basing bands on coverage produces upper bands that contain very few items, with the number of items in each band getting progressively bigger, until lower bands contain large numbers of items. Coverage-based bands deal directly with the issue of the variable utility of words, have smaller and more regular internal differences in frequency, and are likely to be more reliable in terms of the placement of words in bands.

A re-analysis of vocabulary test data in terms of coverage-based bands generated a very different profile of learners' knowledge, and it was argued that the shape of this profile may be more useful in guiding instruction. It was also argued that the smaller size of coverage-based bands at higher frequency levels makes them a more feasible basis for instruction.

Clearly the data presented here are very much preliminary. The bands constructed for this study, the method used, and the re-analysis of the vocabulary test data is intended to be illustrative. Word frequency lists and banding have multiple purposes and different approaches to banding may be appropriate for different purposes. Nevertheless, this article would like to argue that coverage-based bands are potentially informative and hopes that future research may make use of this proposal.

References

- Aizawa, K. (2006). Rethinking frequency markers for English-Japanese dictionaries. In S. Ishikawa, K. Minamide, M. Murata, & Y. Tono (Eds.), *English lexicography in Japan* (pp. 108–119). Tokyo: Taishukan Publishing Company.
- Booth, P. (2013). Vocabulary knowledge in relation to memory and analysis: An approximate replication of Milton's (2007) study on lexical profiles and learning style. *Language Teaching*, 46(3), 335–354. doi:10.1017/S0261444813000049
- Brown, D. (2012). The frequency model of vocabulary learning and Japanese learners. *Vocabulary Learning and Instruction*, 1(1), 20–28. doi: 10.7820/vli.v01.1.brown
- Dang, T.N.Y., & Webb, S. (2016). Making an essential word list for beginners. In I.S.P. Nation (Ed.), *Making and using word lists for language learning and testing* (pp. 153–167). Amsterdam: John Benjamins.
- Davies, M. (2008). *The Corpus of Contemporary American English (COCA): 400+ million words, 1990-present*. Retrieved from <http://www.americancorpus.org>

- Edwards, R., & Collins, L. (2013). Modelling L2 vocabulary learning. In S. Jarvis & M. Daller (Eds.), *Vocabulary knowledge: Human ratings and automated measures* (pp. 157–183). Amsterdam: John Benjamins.
- Gardner, D. (2013). *Exploring vocabulary*. London: Routledge.
- Kilgariff, A. (1995). *BNC database and word frequency lists*. Retrieved from <http://www.kilgariff.co.uk>
- Kremmel, B. (2016). Word families and frequency bands in vocabulary tests: Challenging conventions. *TESOL Quarterly*, 50(4), 976–987. doi: 10.1002/tesq.329
- Laufer, B., & Nation, P. (1995). Vocabulary size and use: Lexical richness in L2 written production. *Applied Linguistics*, 16(3), 307–322.
- Meara, P., & Milton, J. (2003). *X_Lex, The Swansea Levels Test*. Newbury: Express.
- Meara, P., & Miralpeix, I. (2006). *Y_Lex: The Swansea Advanced Vocabulary Test*. v2.05. Swansea: Lognostics.
- Miller, D., & Biber, D. (2015). Evaluating reliability in quantitative vocabulary studies: The influence of corpus design and composition. *International Journal of Corpus Linguistics*, 20(1), 30–53. doi:10.1075/ijcl.20.1.02mil
- Milton, J. (2007). Lexical profiles, learning styles and the construct validity of lexical size tests. In H. Daller, J. Milton, & J. Treffers-Daller (Eds.), *Modelling and assessing vocabulary knowledge* (pp. 47–58). Cambridge: Cambridge University Press.
- Milton, J. (2009). *Measuring second language vocabulary acquisition*. Bristol: Multilingual Matters.
- Nation, I.S.P. (2001). *Learning vocabulary in another language*. Cambridge: Cambridge University Press.
- Nation, I.S.P. (2006). How large a vocabulary is needed for reading and listening? *The Canadian Modern Language Review/La revue canadienne des langues vivantes*, 63(1), 59–81. doi: 10.3138/cmlr.63.1.59
- Nation, I.S.P., & Beglar, D. (2007). A vocabulary size test. *The Language Teacher*, 31(7), 9–13.
- Schmitt, N. (2008). Review article: Instructed second language vocabulary learning. *Language Teaching Research*, 12(3), 329–363. doi:10.1177/1362168808089921
- Sorell, C.J. (2013). *A study of issues and techniques for creating core vocabulary lists for English as an international language*. Doctoral dissertation, Victoria University of Wellington.
- Webb, S., & Rodgers, M.P.H. (2009). The lexical coverage of movies. *Applied Linguistics*, 30(3), 407–427. doi: 10.1093/applin/amp010
- Webb, S.A., & Chang, A.C.-S. (2012). Second language vocabulary growth. *RELC Journal*, 43(1), 113–126. doi: 10.1177/0033688212439367
- Zipf, G.K. (1935). *The psycho-biology of language*. Boston, MA: Houghton-Mifflin.