

# What words ought to exist?

Coining with coinduction

Dr. Tom Murphy VII Ph.D.\*

1 April 2011

## Abstract

This paper is an earnest attempt to answer the following question scientifically: What words ought to exist?

**Keywords:** computational cryptolexicography, n-Markov models, coinduction

## Introduction

During a recent high-stakes game of Scrabble-brand Crossword Puzzle<sup>1</sup> I had what could only be described as a killer bingo word (all 7 tiles) that, after careful study, I determined could not be placed anywhere on the board. Later in that same game, I had another sequence of letters that just totally seemed like it should be able to make some long-ass words, like for example “oilsoap” which turns out is not a legal Scrabble word.<sup>2</sup> This naturally made me frustrated and I wanted to do something about it. Why can’t “oilsoap” be a word? Or “loopsia”? Words are introduced into the lexicon all the time. My first reaction of course was to make an on-line version of Scrabble where all words are legal. This is called Scralbe (where they can *all be* words!<sup>3</sup>) This is available at <http://snoot.org/toys/scralbe>, and is pretty boring, I gotta be honest (Figure 1).

The thing is, it’s just more fun when some words

\*Copyright © 2011 the Regents of the Wikiplia Foundation. Appears in SIGBOVIK 2011 with the blessing of the Association for Computational Heresy; *IEEEEE!* press, Verlag-Verlag volume no. 0x40-2A. ¥0.00

<sup>1</sup>Scrabble is a registered trademark of Hasbro Inc./Milton Bradley, and Mattel/JW Spear & Sons plc.

<sup>2</sup>There are actually no 7-letter words that can be made from these letters. Don’t even bother. Even if playing off an existing letter on the board, the best we can do are the non-bingos “topsoil,” “topsail,” or “poloist” with an available *t*.

<sup>3</sup>As of 2011, the official Scrabble slogan is “every word’s a winner!” which is clearly false.




Figure 1: In-progress Scralbe game, 753 points.

aren’t words. Think about it: If all words were real, then you could never make a really devastatingly successful challenge in Scrabble that like, rocked the whole household and turned a formerly casual family games night into some kind of crying contest. Spelling bees could still exist, because while no matter what those kids spelled,<sup>4</sup> it would be a word, it would not necessarily be the *right* word, just like maybe a homophone. There would be fewer bar fights, but probably not that many fewer. Moreover, iuhwueg nznie a uaohahweih zmbgba bawuyg!

Clearly we need more words, but not all of them. So this raises the question: What words *ought* to exist? This paper explores several different approaches for scientifically answering this question, compares the results,

<sup>4</sup>Well, we have to consider the possibility that the kiddo would use a letter that doesn’t exist. In this particular fantasy, grant me also that every letter also exists, even  $\hat{\diamond}$ .

and proposes specific words that should be added, with their meanings.

Disclaimer possibly indicated for SIGBOVIK: The “research” contained herein is 100% legitimate.<sup>5</sup> I have attempted to present it in a tutorial style that assumes little mathematical or computer science background. I have also left off the last *S* for *Savings*.

## 1 First idea: Wishlist

My website “snoot.org” has a number of games on it, including a Scrabble clone called Scribble<sup>6</sup> and Boggle clone called Muddle.<sup>7</sup> This website has been running for almost ten years, comprising over 150,000 Scribble games totaling 3.8 million words placed and 628,000 Muddle games with over 10 million words found. During each game, players repeatedly attempt to play words that aren’t real. The computer rebukes them, but hope really springs eternal with these people. It’s like they truly deeply wish to break out of the shackles of the Official Scrabble Players Dictionary.<sup>8</sup> So the first approach to determining what words ought to exist is to analyze the words that people tried to play, in order to try to extract the essence of word-yearning.

This analysis is quite straightforward. I took the ten years of logs files and extracted each attempt to play a word in Scribble or Muddle. These log files are quite large, so the first step is just to get a count, for each alleged word, and store those in a more convenient format. There were 3,572,226 total words attempted<sup>9</sup> in Scribble and 13,727,511 in Muddle. The most frequent ones appear in Figure 2. Aside from the one-letter ones, the most frequent words are legitimate words, since players have a bias towards attempting words that will not be rebuked by the computer.

Seeing the words that people wish existed is a simple matter of filtering out the words that already exist, using the Scrabble dictionary. (I also filtered out

<sup>5</sup>Source code is available at <http://tom7misc.svn.sourceforge.net/viewvc/tom7misc/trunk/wishlist/>

<sup>6</sup><http://snoot.org/toys/scribble/>

<sup>7</sup><http://snoot.org/toys/muddle/>

<sup>8</sup>For the analyses in this section that depend on a list of legal words, I actually use a modified version of SOWPODS, which is the tournament list used in Australia and the UK, and significantly more permissive than the US Tournament Word List. Though the modified version is non-canonical, I stuck with it because it’s what’s been in use on the site for ten years.

<sup>9</sup>Here a word attempted is the major word of the play. This does not include incidental words (typically two-letter ones) formed in the perpendicular direction.

| Scribble |      | Muddle |      |
|----------|------|--------|------|
| Count    | Word | Count  | Word |
| 45,605   | a    | 20,412 | late |
| 42,315   | i    | 19,405 | rate |
| 32,499   | d*   | 19,276 | dear |
| 12,981   | in   | 19,049 | tear |
| 12,851   | oe   | 19,019 | date |
| 12,528   | s*   | 18,771 | lear |
| 12,207   | re   | 18,423 | deal |
| 11,159   | tv   | 18,231 | real |
| 10,720   | jo   | 18,138 | lead |
| 10,386   | it   | 18,076 | tale |
| 10,369   | et   | 17,969 | lane |
| 9,659    | qua  | 17,956 | sear |
| 9,218    | xi   | 17,570 | read |
| 9,099    | go   | 17,193 | teal |
| 9,052    | ow   | 17,170 | lean |
| 8,801    | qat  | 17,071 | dare |
| 8,602    | aa   | 16,923 | dale |
| 8,278    | un   | 16,892 | seal |
| 8,142    | en   | 16,806 | sale |
| 8,005    | or   | 16,465 | seat |

Figure 2: Most frequently attempted words in Scribble and Muddle. Asterisks indicate non-words.

one-letter “words”. It is easy to see that no one-letter words should exist, again because of ambiguities created in spelling bees. Not only when literally spelling “bees”, but according to the official Scripps National Spelling Bee rules, the speller may optionally pronounce the word to be spelled before and after spelling it. So if “s” were a word, then the following ridiculous exchange obtains: Judge: “S. The letter *s*. Etruscan origin.” Speller: “S. S. S.” and the judge cannot tell if the speller meant to state the word before and after, or thinks the word is spelled “sss”.) 22.3% of the words attempted in Scribble and 36.8% in Muddle were not real. The most frequent ones appear in Figure 3.

There’s a clear difference between these two lists. The Scribble list is dominated by words involving difficult-to-play letters like *v* (there are no legal two-letter *v* words). Most of the words would probably be acknowledged as real, just not legal in Scribble. The ones that don’t already have meanings, like “cho” and “int” and “que” seem to be pretty good candidates to exist. The Muddle list is all four-letter words (the minimum allowed length) using common letters. Other than the

| Scribble |        | Muddle |      |
|----------|--------|--------|------|
| Count    | Word   | Count  | Word |
| 11,159   | tv     | 16,251 | dane |
| 4,003    | ok     | 6,156  | rane |
| 2,862    | iraq   | 5,603  | sare |
| 2,725    | zen    | 5,576  | nate |
| 2,448    | cho    | 4,863  | mear |
| 1,538    | viz    | 4,750  | cale |
| 1,418    | sdasda | 4,616  | nees |
| 1,396    | von    | 4,568  | nale |
| 1,136    | etc    | 4,507  | fale |
| 878      | int    | 4,347  | deat |
| 829      | june   | 4,263  | tean |
| 745      | lp     | 4,251  | nile |
| 719      | zion   | 4,160  | mens |
| 665      | cia    | 4,087  | deel |
| 661      | jim    | 3,851  | deam |
| 651      | iraqi  | 3,828  | dana |
| 648      | ques   | 3,781  | beed |
| 542      | que    | 3,769  | lans |
| 502      | tim    | 3,725  | tade |

Figure 3: Most frequently attempted non-words in Scrabble and Muddle.

ones that are already words, like “dane” and “nile” and “mens” (as in “mens section” or “the powerfuel weapon kills hard so many mens”), these are all good candidates for words to exist. Probably if you were playing someone really intense in Scrabble, and he or she played one of these, and was super deadpan about it and maybe had caused some crying contests before, and a known sesquipedalianist, you would let these fly because they look like real words to me. A point in their favor is that they would be quite low-scoring words in Scrabble; not a *z* or *q* to be found. Even in the Scribble list there’s no “qzkwv” junk. The effect is probably due to a few factors: Players are less likely to attempt obvious non-words, common letters appear more often on the rack and on the board and so the opportunity to play words like in Figure 3 presents itself more frequently, and in Muddle, there is no advantage to using unusual letters, except the joy of being a weirdo. Nonetheless, these lists are surely biased by the specifics of Scribble and Muddle, and the question at hand is not just what words ought to exist for the purpose of internet word games, but for general purposes.

Another downside is that this method completely ig-




Figure 4: Cumulative distribution of word frequency. Approximately 25,000 different words (y axis) were issued 55 times or fewer (x axis). The “total” area does not appear much larger than its components because this is a log-log plot.

nores the many words that are attempted only once or a small number of times. Players are very creative; of the 564,610 unique words attempted, 501,939 of them aren’t real! The vast majority of words are attempted only a handful of times (Figure 4). Though those words individually are not good candidates to exist, like tiny stars wished upon in the night sky,<sup>10</sup> in aggregate they form a significant planetarium that may tell us what *kind* of words people wish existed. For example, if we saw that the words “sweeeeeeeet”, “sweeeeeeeeeeeet”, “sweeeet” and “sweeeeeeeeeeeet” occurred a few times each, we could infer that people wished that words like “sweet” with strictly more than two *es* were real words. They might even be indifferent to the absolute number of *es*, as long as there existed some legal variation with more than two *es*. (This appears to be borne out by data. According to Google’s estimates, the words “swe<sup>n</sup>t” for various medium-sized *n* (10–20) appear on the Internet with similar frequency. The only exception is “sweeeeeeeeeeeeeeeet”, with 19 *es*, which un-

<sup>10</sup>Astronomers now agree that stars do exist, by the way.

expectedly appears three times as often as 18 or 20 *es* does; see Figure 5.) In order to lance these two boils, in the next section I explore statistical methods for generalizing from lots of individual examples.

## 2 Statistical models

The reason that people are more likely to play words like “rane” is that the letters are common—they appear more often in words, and more often in the Scrabble bag. But it’s not simply a matter of the frequency of letters; if it were, we would expect to see words like “eee” dominating the list, since *e* is the most common letter in English.<sup>11</sup> People do not play such words often because they do not *seem* like real words. “oilsoap” seems more like a word than “ioaopsl” to most non-crazy people, even though they contain the same letters. This is because we have expectations on what letters are likely to appear next to one another in words. This section is about modeling expectations on what letters appear together, and then using that model to generate the most likely words that don’t yet exist.

**Markov chains.** This guy called Andrei Markov had an idea which is pretty obvious in retrospect, but he had it like a hundred years ago before any of us were born (probably; if not: you are old), which he didn’t call Markov chains but now they’re called Markov chains because I guess in the hopes that contemporary mathematicians will get stuff named after their dead selves if they keep the tradition of naming stuff after dead people alive. The idea is easiest to understand in the context of the current problem. Suppose we know that the words “hello”, “helpful” and “felafel” are the only real words. The following is a frequency table of how often each letter occurs.

|   |   |   |   |   |   |   |   |
|---|---|---|---|---|---|---|---|
| h | e | l | o | p | f | u | a |
| 2 | 4 | 6 | 1 | 1 | 3 | 1 | 1 |

This tells us that *l* is by far the most common letter, so the most likely word is probably “l” or “llllllll” or something. A Markov chain is like a frequency table, but instead of counting individual letters, we count how often one letter comes *after* another. Here is the Markov chain for those words.

<sup>11</sup>Tied for first place with *n*, *g*, *l*, *i*, *s*, and *h*.

|                      |     |                                     |
|----------------------|-----|-------------------------------------|
| 17,900,000           | 0   | swt                                 |
| 1,060,000            | 1   | swet                                |
| 580,000,000          | 2   | sweet                               |
| 1,310,000            | 3   | sweett                              |
| 806,000              | 4   | sweeeet                             |
| 509,000              | 5   | sweeeeeet                           |
| 283,000 <sup>1</sup> | 6   | sweeeeeet                           |
| 170,000              | 7   | sweeeeeeeet                         |
| 115,000              | 8   | sweeeeeeeeet                        |
| 75,200               | 9   | sweeeeeeeeeet                       |
| 94,300 <sup>2</sup>  | 10  | sweeeeeeeeeeeet                     |
| 51,700               | 11  | sweeeeeeeeeeeeet                    |
| 37,900               | 12  | sweeeeeeeeeeeeeet                   |
| 32,000               | 13  | sweeeeeeeeeeeeeeeet                 |
| 25,300               | 14  | sweeeeeeeeeeeeeeeeet                |
| 24,300               | 15  | sweeeeeeeeeeeeeeeeeet               |
| 41,000 <sup>3</sup>  | 16  | sweeeeeeeeeeeeeeeeeeeet             |
| 55,000               | 17  | sweeeeeeeeeeeeeeeeeeeeet            |
| 45,000               | 18  | sweeeeeeeeeeeeeeeeeeeeeet           |
| 133,000 <sup>4</sup> | 19  | sweeeeeeeeeeeeeeeeeeeeeeeet         |
| 34,800               | 20  | sweeeeeeeeeeeeeeeeeeeeeeeeet        |
| 16,100 <sup>5</sup>  | 25  | sweeeeeeeeeeeeeeeeeeeeeeeeeeeet     |
| 10,100               | 30  | sweeeeeeeeeeeeeeeeeeeeeeeeeeee...t  |
| 2,800                | 40  | sweeeeeeeeeeeeeeeeeeeeeeeeeeeee...t |
| 923                  | 50  | sweeeeeeeeeeeeeeeeeeeeeeeeeeeee...t |
| 118                  | 75  | sweeeeeeeeeeeeeeeeeeeeeeeeeeeee...t |
| 38                   | 100 | sweeeeeeeeeeeeeeeeeeeeeeeeeeeee...t |
| ? <sup>6</sup>       | 200 | sweeeeeeeeeeeeeeeeeeeeeeeeeeeee...t |

Figure 5: Frequency of “swe<sup>n</sup>t” on the internet for various *n*, estimated by Google. Notes: **(1)** Spell correction offered for “sweeeet”. **(2, 3, 4)** Spell corrections offered to e<sup>9</sup>, e<sup>14</sup> and e<sup>15</sup> respectively. **(5)** Spell correction offered for “weeeeeeeeeeeeeeeeeeeee t” (?) **(6)** With two hundred *es*, the word is too long for Google, which asks me to “try using a shorter word.” Thanks Google, but I already did try the shorter ones.

|          | <b>h</b> | <b>e</b> | <b>l</b> | <b>o</b> | <b>p</b> | <b>f</b> | <b>u</b> | <b>a</b> |
|----------|----------|----------|----------|----------|----------|----------|----------|----------|
| <b>h</b> | 0        | 0        | 0        | 0        | 0        | 0        | 0        | 0        |
| <b>e</b> | 2        | 0        | 0        | 0        | 0        | 2        | 0        | 0        |
| <b>l</b> | 0        | 4        | 1        | 0        | 0        | 0        | 1        | 0        |
| <b>o</b> | 0        | 0        | 1        | 0        | 0        | 0        | 0        | 0        |
| <b>p</b> | 0        | 0        | 1        | 0        | 0        | 0        | 0        | 0        |
| <b>f</b> | 0        | 0        | 0        | 0        | 1        | 0        | 0        | 1        |
| <b>u</b> | 0        | 0        | 0        | 0        | 0        | 1        | 0        | 0        |
| <b>a</b> | 0        | 0        | 1        | 0        | 0        | 0        | 0        | 0        |

The letters across the top are the “previous letter” and the ones across the left are the “next letter” and the box contains the corresponding count. For example, the pair “el” appears four times. (Pairs of letters are called “bigrams” by nerds, some nerd-poseurs, and Markov who I can’t tell if he was a nerd by his picture, because he does have a pretty austere beard, but also did a lot of math.) One of the useful things about a Markov chain is that it lets us predict the next letter that we might see. For example, if we see “half”, then the column labeled **f** above tells us that the next letter is twice as often an *e* than a *u*, and that no other letters ever occurred. Typically we think of these as being probabilities inferred from our observations, so we say there’s a 2/3 chance of *e* following *f* and a 1/3 chance of *u*. Now the word “lllll” isn’t so likely any more, because there’s only a 1/4 chance of the next letter being *l* once we see *l*.

Words are not just their interiors; it’s also important what letters tend to start and end words. We can do this by imagining that each word starts and ends with some fake letters, and include those in the Markov chain. Let’s use < for the start symbol and > for the end. So we pretend we observed “<hello>”, “<helpful>”, and “<felafel>”. Speaking of which, could you imagine if there were such a thing as a helpful felafel? Would you eat it? Because then it probably can’t help you any more, except to get fat.

|          | < | <b>h</b> | <b>e</b> | <b>l</b> | <b>o</b> | <b>p</b> | <b>f</b> | <b>u</b> | <b>a</b> |
|----------|---|----------|----------|----------|----------|----------|----------|----------|----------|
| <b>h</b> | 2 | 0        | 0        | 0        | 0        | 0        | 0        | 0        | 0        |
| <b>e</b> | 0 | 2        | 0        | 0        | 0        | 0        | 2        | 0        | 0        |
| <b>l</b> | 0 | 0        | 4        | 1        | 0        | 0        | 0        | 1        | 0        |
| <b>o</b> | 0 | 0        | 0        | 1        | 0        | 0        | 0        | 0        | 0        |
| <b>p</b> | 0 | 0        | 0        | 1        | 0        | 0        | 0        | 0        | 0        |
| <b>f</b> | 1 | 0        | 0        | 0        | 0        | 1        | 0        | 0        | 1        |
| <b>u</b> | 0 | 0        | 0        | 0        | 0        | 0        | 1        | 0        | 0        |
| <b>a</b> | 0 | 0        | 0        | 1        | 0        | 0        | 0        | 0        | 0        |
| >        | 0 | 0        | 0        | 2        | 1        | 0        | 0        | 0        | 0        |

We just added these like other letters, but since the beginning symbol < never occurs after other letters, we don’t need a row for it (it would be all zeroes), and similarly since no letters ever follow > we don’t need a column for it. Now the word “lllll” is impossible because no words start with *l*.

It basically makes sense to consider the probability of a whole word to be the chance of simultaneously seeing each pair of letters in it, which is just the product of all the probabilities. So the word “hel” is 2/3 (for <**h**) × 2/2 (for **he**) × 4/4 (for **el**) × 2/6 (for **l**>), which is 0.222. These are the most likely words overall (I discuss how to generate such lists in Section 2.2):

|       |      |      |          |
|-------|------|------|----------|
| 22.2% | hel  | 2.5% | helpfel  |
| 11.1% | helo | 2.5% | helafel  |
| 7.4%  | fel  | 1.9% | fulo     |
| 3.7%  | hell | 1.9% | hello    |
| 3.7%  | felo | 1.2% | fell     |
| 3.7%  | ful  | 1.2% | helafelo |

This is pretty good. These words resemble the ones we observed to build the Markov chain, but are novel. I think helafelo is a pretty rad word, right?

The next step is to build a Markov chain for a list of real words and see what results. I built one for the SOWPODS word list, which results in the table in Figure 6. These are the most likely words, with real words filtered out:

|       |    |       |    |
|-------|----|-------|----|
| 4.99% | s  | 0.17% | y  |
| 1.75% | d  | 0.17% | p  |
| 0.95% | g  | 0.16% | a  |
| 0.55% | c  | 0.16% | n  |
| 0.43% | r  | 0.15% | ps |
| 0.42% | t  | 0.13% | ms |
| 0.40% | e  | 0.13% | ts |
| 0.35% | m  | 0.13% | ds |
| 0.32% | ss | 0.11% | hy |
| 0.20% | rs | 0.11% | k  |
| 0.19% | h  | 0.11% | ng |
| 0.18% | l  | 0.11% | ly |

Ugh, poop city! Actually, it turns out that when you see enough words, you see enough pairs that all sorts of junk looks likely. For example, “ng” is easily explained by many words starting with *n*, *g* often following *n*, and many words ending with *g*. Even though each pair makes sense, the whole thing doesn’t look like a word, because we expect to at least see a vowel at some point, for one thing.




Figure 6: Markov chain for the SOWPODS word list, where darker squares indicate higher probability. The darkest is the transition from *q* to *u* (98%), which is not surprising.

There is a standard solution to this problem, which is to generalize the Markov chain to keep more than one letter of history. So instead of just tallying how often *g* follows *n*, we count how often *g* follows *in* (and any other pair of letters).<sup>12</sup> This makes the table pretty large, so you'll just have to look at Figure 6 again and imagine it being 28 times wider. But the good news is that it invents much better words:

<sup>12</sup>The details are straightforward, except possibly that we now imagine each word to start with two (or in general, *n*) copies of the start symbol, so that we see “<<helpful>”. The column corresponding to the history << tells us the frequency of letters that start words, and for example the column <h tells us the frequency of letters that follow *h* when it appears at the start of a word. We do not need to repeat the ending character > because once we see it, we never do anything but end the word.

|       |      |       |      |
|-------|------|-------|------|
| .709% | ing  | .110% | le   |
| .248% | ses  | .107% | der  |
| .169% | des  | .107% | ove  |
| .154% | nes  | .101% | gly  |
| .140% | sts  | .088% | hy   |
| .131% | se   | .085% | ung  |
| .128% | ings | .083% | cy   |
| .126% | ded  | .081% | pres |
| .117% | cal  | .080% | pers |

These are even, like, pronounceable. The best news is that they keep getting better the more history we keep:

|       |      |       |       |
|-------|------|-------|-------|
| .109% | des  | .038% | ent   |
| .078% | pers | .036% | dist  |
| .076% | cal  | .035% | ble   |
| .062% | pres | .035% | ches  |
| .045% | nons | .034% | gly   |
| .044% | ress | .034% | inted |
| .042% | ing  | .034% | dists |
| .040% | pred | .033% | lity  |

|       |        |       |        |
|-------|--------|-------|--------|
| .045% | unders | .017% | heters |
| .034% | dising | .016% | sters  |
| .029% | pers   | .015% | stic   |
| .028% | cally  | .014% | pering |
| .023% | inted  | .013% | dises  |
| .020% | heter  | .013% | ching  |
| .019% | tric   | .012% | shing  |
| .018% | ster   | .012% | dest   |
| .018% | hier   | .011% | teless |
| .018% | unded  | .011% | resis  |

|                                     |  |  |  |
|-------------------------------------|--|--|--|
| GetTempFileName failed with error 5 |  |  |  |
|-------------------------------------|--|--|--|

With four letters of history, the words produced are quite good! (The results at  $n = 5$  are somewhat disappointing since the program crashes from running out of memory. The table at  $n = 5$  would have over 481 million entries.) Many of these seem like real words. Some even suggest meaning because they contain common morphemes. To make the case that these are not just real-looking words but characteristic of the English language, compare the results of the same algorithm on the dictionary from the Italian language edition of Scrabble, which is probably called *Scrabblizzimo!* (Figure 8). Italian is lexicographically a more compact language than




Figure 7: Markov chain for the Italian language. Again darker cells indicate higher probability. Italian has more lexicographic structure recognizable from bigraphs than English does: Note that the extremely rare letters “j”, “k”, “q”, “w”, “x”, and “y” have almost empty rows. “z” very frequently follows “z”, as in *pizza*. Words almost always end in a vowel.

English (Figure 7); there are only 21 letters (outside of occasional interlopers in loan words like *jeans* and *taxi*). Moreover, even though the dictionary contains 585,000 words (twice as many as English), the probabilities of observing these non-words are much higher than the most likely English ones.

## 2.1 Usage-weighted methods

One criticism of this approach is that it considers every word in the word list to be equally important.<sup>13</sup> I object on the philosophical grounds that some words that already exist *ought to exist more* than other words that already exist. For example, *congenital* is a much nicer word than the plain ugly *congenial*, and is reflected by the fact that *congenital* is used five times more fre-

<sup>13</sup>In fact, the common part of words with many different conjugations is in essence counted many times. This means *ornithology* in its six different forms contributes six times as much to our model as the word *the!*

|       |         |       |          |
|-------|---------|-------|----------|
| .137% | ammo    | .026% | rino     |
| .071% | rice    | .025% | diste    |
| .061% | rico    | .024% | risti    |
| .055% | este    | .023% | disci    |
| .053% | scono   | .022% | riasse   |
| .049% | immo    | .022% | riassi   |
| .047% | assero  | .021% | cate     |
| .047% | scano   | .019% | rite     |
| .038% | rammo   | .019% | cando    |
| .034% | cata    | .018% | riassero |
| .034% | assimo  | .018% | riassimo |
| .032% | riate   | .018% | dete     |
| .032% | disce   | .018% | disca    |
| .030% | esti    | .017% | risca    |
| .029% | rica    | .017% | cente    |
| .028% | endo    | .016% | acci     |
| .027% | dissimo | .015% | centi    |
| .026% | rici    | .015% | girono   |

Figure 8: Most probable words induced by the Markov Markov chain for the Italian language ( $n = 4$ ).

quently than *congenial*.<sup>14</sup> In this section, we produce Markov models of words weighted by the frequency with which people tend to use them. This is just a simple matter of training the model on some natural language corpus (with many occurrences of each word, or no occurrences of unpopular words) rather than a flat list of all alleged words.

**Facebook.** Since the best research is intensely navel-gazing, I started by analyzing a corpus of my own writing, specifically my Facebook status updates since March 2006. There were 1,386 status updates containing such jems as “Tom Murphy VII thinks mathfrak is straight ballin” and “Tom Murphy VII global L\_50 reused for unused\_36166!!”. The most likely words with  $n = 4$ :

<sup>14</sup>14,200,000 times to 2,820,000, on the Internet, according to Google.

| My Facebook status updates, $n = 4$ . |            |       |         |
|---------------------------------------|------------|-------|---------|
| .252%                                 | pittsburgh | .097% | can't   |
| .209%                                 | steelers   | .083% | i'm     |
| .209%                                 | sfo        | .083% | icfp    |
| .195%                                 | it's       | .083% | app     |
| .125%                                 | bdl        | .069% | x       |
| .111%                                 | sigbovik   | .069% | drunj   |
| .109%                                 | facebook   | .069% | g       |
| .097%                                 | mic        | .061% | ther    |
| .097%                                 | s          | .055% | doesn't |

This is the worst. Not only does it contain loads of one-letter words that we have already determined are verboten,<sup>15</sup> but the rest are just non-words that I tend to use like the names of cities, prestigious conferences, or IATA airport codes. The main problem is that there is simply not enough data from which to generalize.

**Wikipedia.** I tried again, but with Wikipedia, using a snapshot of the English site from June 2009. This is 23 gigabytes of data, most of it expository text composed by native speakers, plus bathroom humor vandalism. The list produced by this analysis is much better, though it contains artifacts from non-English Wiki language used in articles. The unabridged list appears in the appendix; my hand-selected favorites:

| English Wikipedia, $n = 3$ . |         |         |           |
|------------------------------|---------|---------|-----------|
| .0287%                       | smally  | .00518% | reporth   |
| .0156%                       | websity | .00484% | delection |
| .0156%                       | stude   | .00459% | grounty   |
| .0124%                       | chool   | .00437% | betweek   |
| .0120%                       | fontry  | .00431% | fination  |
| .0102%                       | undex   | .00388% | manuary   |
| .0099%                       | octory  | .00360% | whicle    |
| .0096%                       | coibot  | .00262% | stategy   |
| .0084%                       | footnot |         |           |

Lots of these could be the names of tech startups or Pokémon.

## 2.2 Coining words with coinduction

In the earlier sections I blithely produced tables of the most probable words according to an  $n$ -Markov chain. It is not obvious how to do this (or that it is even possible), so I explain the algorithm in this section. It's safely skippable, I mean if you don't want to know about a

<sup>15</sup>Note that since  $n = 4$  these words have to actually appear in status updates to have nonzero probability for this list. "g" is explained by frequent occurrences of "e.g.", for example.

pretty cool algorithm that's not that complicated and might even be new, plus *dual math*.

Computing the probability of an individual word is easy. We prefix it with  $n$  copies of the start symbol  $<$ , suffix it with a single  $>$ , and then look up the probability of each symbol given its  $n$  preceding symbols in the table, and multiply those all together. We can compute the probability of any word this way. The problem with sorting all of the possible words by their probabilities is that there are an infinite number of them. We can't just look at short words first, either, because for example the word "thethethe" is many times more likely ( $p = 6.08 \times 10^{-11}$ ) than the shorter "qatzs" ( $9.07 \times 10^{-12}$ ).

The solution is to use coinduction. Most people remember induction from school, maybe, which is the one where you have some base case like "0 is even", and then you prove that all numbers are either even or odd by assuming " $n - 1$  is even or odd" and proving " $n$  is even or odd". From this we conclude that every number is either even or odd. The idea is the proof shows how to, for any given number  $m$ , count down to the base case "0 is even", and then repeatedly apply the  $n - 1$  step (inductive step) to get back up to  $m$ . This is a great way to prove facts about finite things like numbers. Think of induction as a way you prove a statement like "Good to the last drop," or "There's always room for Jello."

Coinduction is a good proof technique for infinite things, like a sorted infinite list of possible strings. The idea behind coinduction is kind of like, you prove something like "0 is a number" (the base case), then prove something like "if  $n$  is a number, then  $n + 1$  is a larger number", and then conclude that there exists an infinite series of numbers, each larger than the previous one. Think of coinduction as a way you prove a statement like "Once you pop, you can't stop," or "Never gonna give you up."

To sort the infinite list we don't actually use coinduction (we're not going to prove anything, just implement it), but its computational counterpart, corecursion. I just can't resist the "coin" pun.

What we do is define a function "most probable paths", which returns a (possibly infinite) stream of strings for a given starting state. Each string is finite and ends with the terminal symbol  $>$ , and they appear sorted by decreasing probability. (The most probable words overall will be just the first elements from the stream returned by this function when using a starting state like  $<<<$  for  $n = 3$ .) Since we don't want to explore all possible strings in order to produce this





```

fun most_probable_paths { lower_bound : real, state : state }
  : { string : symbol list, p : real } stream =
  let
    fun nexts i =
      case symbol_from_int i of
      NONE => nil
    | SOME sym =>
      let
        val p = (* probability of seeing sym in this state *)
      in
        if p < lower_bound
        then nexts (i + 1)
        else if sym = end_symbol
            then S.singleton { string = nil, p = p } :: nexts (i + 1)
            else
              let
                val lb' = lower_bound / p
                val tails =
                  most_probable_paths { lower_bound = lb',
                                         state = advance_state (state, sym) }
              in
                (* Now multiply through the probabilities and add the symbol
                 to the head of the strings. *)
                Stream.map (fn { string = t, p = p' } =>
                           { string = sym :: t, p = p * p' }) tails ::
                nexts (i + 1)
              end
            end
          end
      end
    (* Try all next symbols. *)
    val streams = nexts 0
  in
    S.merge_sorted bysecond_real_descending streams
  end

```

Figure 9: Pseudocode for `most_probable_paths`. `advance_state` gives a new state from a previous state and symbol observed, so that for example `advance_state(abc, z)` gives `bcz`. The pseudocode for `merge_sorted` is given in Figure 10.

```

fun merge_sorted cmp l =
  let
    fun ms nil () = Nil
      | ms (s :: t) () =
        case force s of
          Nil => ms t ()
        | Cons (v, ss) =>
          ms_insert v [ss] t
    and ms_insert bv sg nil =
      Cons (bv, delay (ms sg))
      | ms_insert bv sg (s :: t) =
        case force s of
          Nil => ms_insert bv sg t
        | Cons (v, ss) =>
          case cmp (bv, v) of
            GREATER =>
              ms_insert v (singleton bv :: ss :: sg) t
          | _ => ms_insert bv (s :: sg) t
  in
    delay (ms l)
  end

```

Figure 10: Pseudocode for `merge_sorted`. `ms` merges a sorted list, and `ms_insert` is a helper where we have a candidate best value `bv` which will either be the one we return at the head of the stream, or we'll replace it and then stick `bv` somewhere to be returned later. (This algorithm can be improved by making a data structure like a (co)heap; this is just a simple first pass.)

|   |             |          |       |
|---|-------------|----------|-------|
| observation (the word <i>mezzotinting</i> ); there are actually   | 1.000000000 | -worms   | 69/69 |
| hundreds of such unique prefixes and suffixes. We need            | 1.000000000 | -worm    | 69/69 |
| a better list. <sup>16</sup> Removing the ones that appear just a | 1.000000000 | -working | 21/21 |
| single time doesn't really help that much:                        |             |          |       |

|             |          |     |
|-------------|----------|-----|
| 1.000000000 | -zzazzes | 3/3 |
| 1.000000000 | -zzazz   | 3/3 |
| 1.000000000 | -zzans   | 3/3 |
| 1.000000000 | -zzanim  | 2/2 |
| 1.000000000 | -zzan    | 3/3 |
| 1.000000000 | -zygotic | 3/3 |

Still bad. Let's turn up the juice to prefixes and suffixes that appear at least 10 times.

|             |           |       |
|-------------|-----------|-------|
| 1.000000000 | -wrought  | 10/10 |
| 1.000000000 | -writings | 12/12 |
| 1.000000000 | -wraps    | 10/10 |
| 1.000000000 | -wrap     | 11/11 |

Much better! But the next step is going to be to try removing these prefixes and suffixes from words that have them, to find new words. Since these have modularity of 100%, we already know that every time we apply them, the result will already be a word. So they are useless for our analysis. Here are the most modular prefixes and suffixes with modularity *strictly less than* 1.

|             |         |         |
|-------------|---------|---------|
| 0.985714286 | -makers | 69/70   |
| 0.985714286 | -maker  | 69/70   |
| 0.983606557 | -wood   | 120/122 |
| 0.983471074 | -woods  | 119/121 |
| 0.982758621 | -down   | 57/58   |
| 0.982658960 | -works  | 170/173 |
| 0.981818182 | -houses | 108/110 |
| 0.981818182 | -house  | 108/110 |
| 0.981132075 | kilo-   | 52/53   |

<sup>16</sup>The right thing to do here is probably to use binomial likelihood rather than the scale-independent fraction. But simpler approaches produce pretty good lists.

|             |             |           |
|-------------|-------------|-----------|
| 0.980752406 | -less       | 1121/1143 |
| 0.980743395 | over-       | 2190/2233 |
| 0.980000000 | -books      | 49/50     |
| 0.980000000 | -book       | 49/50     |
| 0.979591837 | -proof      | 48/49     |
| 0.979310345 | -lessnesses | 142/145   |
| 0.979069767 | -ships      | 421/430   |
| 0.978723404 | -lessness   | 184/188   |
| 0.978723404 | -board      | 138/141   |
| 0.978494624 | -woman      | 91/93     |
| 0.978021978 | -women      | 89/91     |
| 0.977528090 | -ship       | 435/445   |
| 0.977272727 | -manship    | 43/44     |
| 0.976744186 | -weeds      | 84/86     |
| 0.976470588 | after-      | 83/85     |
| 0.976190476 | -manships   | 41/42     |
| 0.976190476 | -making     | 41/42     |
| 0.976190476 | -craft      | 41/42     |
| 0.976190476 | -boats      | 41/42     |
| 0.976190476 | -boat       | 41/42     |

Wow, now we're talking! The single word that cannot have "-maker" removed is *comaker*, suggesting that *co* should be word (noun: "What a comaker makes.").

Given this list, the next step is to identify potential words that can be backformed by removing prefixes or adding suffixes from existing words. Such a string can often be found via multiple prefixes and suffixes. For example, *twing* can be formed by removing "-ing" from *twinging* (a false positive, since the root word is actually *twinge* in this case) as well as by removing the prefix "lef-", which has modularity of 20% (including splits such as "lef/tie"). Maybe not good justification, but *twing* is a pretty good word anyway.

We define the probability of a word as its Markov probability (with  $n = 4$ , as this seems to produce the best results), times the probability that at least one of the potential backformation rules applies.<sup>17</sup> Here are the most likely words by backformation:

| word  | prob  | most likely backformation rules                          |
|-------|-------|--|
| dises | .023% | para- (0.42) fluori- (0.39)<br>melo- (0.35) bran- (0.31) |
| tring | .020% | hams- (0.36) scep- (0.35)<br>bows- (0.33) hearts- (0.29) |
| disms | .017% | triba- (0.31) drui- (0.30)<br>bar- (0.27) invali- (0.27) |

<sup>17</sup>As above we only allow backformation rules that have at least 10 occurrences, to prevent degeneracy.

|       |       |  |
|-------|-------|--|
| ching | .017% | day- (0.86) hot- (0.69)<br>star- (0.51) guillo- (0.50) |
| sking | .017% | dama- (0.24) imbo- (0.18)<br>fri- (0.18) atta- (0.17)  |
| cally | .015% | anti- (0.78) specifi- (0.61)<br>magnifi- (0.55) phoni- |
| pring | .015% | days- (0.67) heads- (0.62)<br>outs- (0.54) ups- (0.51) |

I think that this approach shows promise, but there appear to be a few problems: Many of these "rules" can be explained by bad segmentation ("heads-" appearing to be modular, for example, is really just "head-" plus "s" being a common letter.) Second, I believe the disjunctive probability of any rule applying is too naive for determining the score. For example, *tions* has almost a thousand different prefixes that could apply to it; the chance of *any one* of them applying is very nearly 1. But this is actually because "tions" is just a common way for a word to end. Legitimate root words to which many good prefixes are applied cannot be easily distinguished from common suffixes by this symmetric algorithm. More work is called for here.

## 5 Survey

On occasion I have been accused of "overthinking" problems, whatever that means. So to compare, I next hazarded a tried and true technique from grade school, the survey.

I asked a few people who happened to be around, "What word ought to exist?" Most people did not know what to make of this question, and also, because people seem to revel in the opportunity to get (well deserved) revenge on me by being disruptive trolls, many of the answers were designed to be unusable. In order to not reprint everyone's bullshit—but not introduce bias by selectively removing data—I discarded random subsets of the data until it did not contain bullshit any more.

|          |                 |
|----------|-----------------|
| Rob:     | etsy, nuuog     |
| Chris:   | nurm            |
| David:   | wafflucinations |
| Lea:     | hnfff           |
| Reed:    | pansepticon     |
| Jessica: | gruntle         |

From this we can conclude that 16% of people wish *nurm* were a word, and so on. These words did not come with definitions, except for *gruntle*, which Jessica gives

as “the opposite of disgruntle”. This is actually already a word, but it was the inspiration for Section 4. *etsy* is the name of a popular on-line crafts community so I don’t know why Rob would suggest that. The meaning of *wafflucinations* is clear from morphological analysis.

## 6 Conclusion

In this paper I investigated several different ways of answering the question: What words ought to exist? Each method produces different words, and some don’t work that well, but nonetheless we have several rich sources of words, each time scientifically justified. I conclude with a section of recommendations for words that ought to exist, along with definitions.

### 6.1 Recommendations

*Sweeeeeeeeeeeeeeeeeet* with 19 *es* is the clear favorite based on analysis of usage, so this one should be introduced. It means “Really sweet.”

*Rane* sounds too much like *rain*, but *sare* has a unique pronunciation and many people seem to think it’s already a word. I propose that *sare* be introduced as a noun meaning, “a word that sounds real but isn’t.”

*Cho* is similarly easy to pronounce and spell. I propose that it be defined as “A kind of cheese,” so that we can really nail the new triple entendre on the classic joke. *Chomaker* is someone who makes that kind of cheese.

*Unders* was one of the most frequently occurring words towards the top of many analyses. This word should be a colloquialism for underwear, which would probably already be understood from context.

*Dise* is suggested by both the Markov model (as *dises*, *dising*) and backformation (as *dises*). I like thinking of it as being the root of *paradise*, where *para-* means something like “along side of” or “resembling”. So *dise* is the place you’re really looking for when you get to paradise and realize it’s just a mediocre country club.

*Helafelo* is one hell of a fellow.

**Addendum.** During the preparation of this paper, the Scralbe game has converged on a culture where the words played are real-seeming, with creative definitions. Examples: *frodeo* (“Gandalf is the clown.”) *pridefax* (“An unproven treatment for telephone anxiety.”) *eeeeee* (“eeeeee”) *orzigato* (“Move mr. roboto. for great justice.”) *stovebed* (“Now you don’t have to get out from under the covers to make breakfast.”) *ovawiki* (“The

free egg cell that anyone can edit.”) *gaptave* (“Two discontinuous musical intervals.”) *achoolane* (“Nostril (colloq.)”) *gplerious* (“Completely overcome by software licenses.”) *bestcano* (“Some eruptions are better than others.”) Thanks to the players for their contributions, especially Ben Blum, Chrisamaphone, Rob Simons, and Flammin Fingers.

## Appendix

Here are the most likely words induced by the English Wikipedia, with  $n = 3$ . I have left off the probabilities; they can be reproduced by downloading Wikipedia and running the software yourself, which only takes like 9 hours.

s ther w en t c d f m b ser e diff u sup p n de tal  
othe th wikipedia rd con id r ame edia cal uk wiki und  
als st l km wp g x ext j des td ter h co k sout don o  
cound v ver wer med res ind la em re ins afd ands div  
htm infor rom wher ple mi acces ii ent unive ff refere  
ger bution lish ince wikiped auth dist oth smally rea al  
y est cle fr lation tral gov nd mily isbn cond bel be  
for hould ber pers inder les gree mation lar chan dia  
ove stary compan mon val san cate nothe pdf tes ster  
nort pril hist sq clast ful ment q aft tr nown lign del  
z ave dify stor sity artic pring rese mas li ord ca thes  
oldid lin es ress flage blace pland mar mes conce ex  
ter wate mp jame lity ft cology spany mic ve han thist  
spam cand sher smal subdiv tem jan ture inded bothe  
ared cour rect thers ope mm dom der mand tt dr phy  
cd websity stude lor bbc ru mer flor succes ence tely  
aust stan clas et outh pl ints vfd fc wikipedit sus for  
man contry au edu inding el lic bord mr ston whic ide  
hous mich auser google pt le doi engle ed gen used  
col num nate cer withis clude shous inted wikipedian  
thered andex chool ning attle chael sta plation ampion  
crea fontry lege prese grouh iii mader augh ll ching fb  
furt yearch examp octor nove furth bein tra nover tex  
ter mont sted gener quary includ ander sco ast gr wome  
ality lon hough ano exas polid sease dec sor ga nov unty  
shor imdb offic partic oved forn lan fm undex rus arted  
alia cong vill eason gue shough cc vil tember octory jr  
gove writy thand anding rev colog gover edian publist  
mics comple syster olding coor useum spect expan pu  
ble bate tribut oned coibot betwo conter cles leason fina  
freen eng fa que fied tre hamp guage wered nic awar yout  
usic cgi andy nrhp clar arting bc cated rober feb ames  
ven jun thould defere artics sen frand cred ada cana rec  
tion modiff sal los ret ched decial collow notalk nl yor

etele sain ance cens chand apr delect dea stant bruary  
commer righ ch tor peral andia arth roup founy hor  
scies alth pete footnot html ath tribs ac chas eith pa tl  
refor bes bor cently fonter neer tration reat hi awa chris  
da wikipedits atter pp ents dvd arge wan sch sr fere inds  
norts tha nament anot oper aspx rence website minal  
brity dts apring goalso alber rown graphy anded min  
alian ength mility abbr talign bec com authe hite pic  
pation jone ords demy alt wated alife totalk lond tage  
ad ana evers amg gened colo engton mina lates porth fff  
lifer footballly stry shing hase lis eletion rol mally inton  
ded doesn ny mus pro du pre somes dired na becamp  
ire inces cologo mong swith acts aree ang els noth unic  
mous dest ali deat smale poing dc ission iv oct positivity  
pred poss kore anda ja ms gard coi frome afc nhl pub  
lice mothe jers maded qual geory roy vid pur studing  
ip ign fath cally coundex col abouth twork threen ing  
ess lishe lable reland cing chi thern dete ar ving di eur  
wers il dio hl ando seast sation orge ble sep db classion  
cm las conver evel ing popula uns polic vs frans ker ann  
plich ovember alish cher therence prock ife feate user  
vice indid unded powere hs phot assion von dece theight  
proview ama prote kg birthe specian userve se nois cfm  
ma len ause gle jp flagic differe counder rance swed lee  
tel ont sould whics artion rans ric parly overy ves sa grea  
ans une tary ss jul ne fami dever ange hms oright da  
tion cription prothe univer lat tota un adming aaa am  
pan perfor athe seased pov becaust alson thed methe aff  
spand inc sis stral refered endex xfd texamp valign sha  
dists fron heade lood thest unived italk brary colled uc  
first ite gan ep rary oney wikipedition ra ex scord citation  
hout clos bon lp imple acade flagico bost isn spection  
samp ral tead sant roces yorks ories bords unition che  
dra epison eff stal about ach earch basic reate conves ses  
hims mone fran carly nated lished fra pm oh att delet  
unior reporth daving emple tworld novel trand ross ie  
vall orld vi becauser caust dard exand feder ches yound  
gence bol tribut rement gf sugges tity forder aug systed  
mathe websited abover ltd remember colle wils mos sman  
lete dow eles throup ba unt expr formany df ments tect  
pape hool exper butions hited delection deces counng gon  
gers stors abc partics writion furthe rade mult bb lism  
afries frence bish arence tems musion colore ethe bute  
vie ia hing oute hom clain nom missue hus ns chich ox  
for contribut vict autom boile novers aus airst jos proble  
werence georgan rel eas opedia rs ned aut iss pg ct othis  
indo shi dst alf ian grounty hile mor jor grap parth elet  
regory them ell septed pc url gnf unds appy studed  
offere gb timent sed dary eved hally deate hr taly nom-

ing duri cented thead aliam mored pmid hered hong  
recore vious wates fathe goved squary soment mout rela  
eview unives shere trate betweek gia mory cas dal fi  
enge stribs colspan vol facted amer reser loc fination  
deven andid leade tand protes gar dely til alled sary  
tro alist ng cha charly ab hel begion mana sk powed  
ted calign wikipe dar nm churce mil ren dise rever sel  
howere nexter commen sonal oly subser coiboth shool  
hal froman eu waship pr rease ph rical op assed bility  
ording conces mal stitle iucn sri os mond lations mance  
artment rigin rt brite ffff jsp lb dd wikipeding lant sett  
arter te yource sh contribs footbal zh aution acter ad  
dit governs coln useries ract thich tain inten sec aa blo  
janual hon nowing letion tition pute goldid communit  
wors fontribut areat titute rected summent station al hu  
counded sp eld ples yeare nating fing foll reque fiel pol  
josed atland phile manuary civision ok feded ult dj sas  
alle cg buildid kee redia ap unk ko cr inver nomy sund  
toon reg aaaa cath dll lc ters communive bal prope refef  
ania mrs hased worder pria arry centing spon ade vand  
dor souch tooks seph capan pbb colong sql petition shors  
fift prily prial brition fl seption nz fonto cbs fally whill  
losed cluded mations monal sland pict russion profilm  
coundo reture harly soute sc rr titution therent shouser  
pager oute quent spons fn ited proces bable atte whicle  
sile si rece unitor yan grougt cos pian comput trans  
jource arthe tury specia artice provid raction deleter  
alo fored ti eague wiking conth ree musly ars mf ups  
school rey donal sected writed fret formate cre co  
lution opers chen bd hp sember folled pagest smaless  
wen socies cs engin aften chur delp bia ce nors cource  
ussion fontribs summan arties jon dism lang skinng au  
gus theren xs seve deball er uning truction nomic creat  
gian acco sv pera alls rict relete rects coord stre ences  
aread spane meth brea ener thernal thire whical rouny  
thu cl pola una thround uservation volved hund aga  
comes menter froma uset publis frey libe posite northe  
ska appedia recial ro fant hought mination geo presult  
andon occup ala nas thists ew maring secore nee dian  
inders inves infore hower sume scal kans indust eding ka  
aux brid var orded gived backgrouh hugh ado prive eb  
rading joing hown euro lim kan chus alifor wille deced  
conven blogy nbc quall coller partion folle eleter reven  
reques lition humber moder backgroup tribution flages  
lears forch reight pault evely mayork sition feath recon  
guit metion hony devieu nume confere fle unitory rican  
ima accore deside ris pedia lisher acture aland userview  
thate polor kar nel sents pany cluding vard bing woul  
havel greer thang genry effere locall desh thors empt

sult quence mir blook porthe tran rew bording shought  
kely neight lam stion publishe hors wis apan nony sourn  
nage orce profest jonal fonth younty deta pd wa lo peries  
flign fource ros rfc mely mt smalled ge retion relect vern  
alid aligh userves speer willa infl undon dels aftern gres  
namedia advan distor oria subdiving lx presign blocal  
edway def scho fic couns belf sily metry cread pi novery  
oces cribut holor reces inte howed cout chis ency ac-  
tory lange cf nc bornia adven revie sames newspanshan  
orial userry suff eth moto showere hunty nucle witz ider  
mune flagicont aire shtm fri guary rigion mome afted  
heate mo formedia oble cnn nea ha canage borning tol  
navbox ava octors md frent offical wal arent anon artish  
creation rell dir auding ira butional labout bott southe  
dence sporth formall ps intry thamp stube unus user-  
ing ron usert conse brote lib oces bg proad tas wer-  
ent these lyn atta pris oile drage unre whenry ement  
stada chara raphy tecton rica reletion foothe titler las-  
sion areen them conship formand ution unio ided guid  
riging brited buch septem eleted conds preser ta ort feld  
ine actional alize tex fer aused engly eque sevel atted  
smalle thome alty reath gation prefero hd nr peration  
covember govember livision fren dq aren deve va ffc  
jo prian afl gion medy stategory nomich tured mel ene  
appa licy linal rought areate bethe othey eral oport  
ara kw whild cance communitid varing brothe pia mu  
sched socian poin vis obs squal paking sma convers no-  
ing maken cathe ni solor alry flagicons aaaaa weren arti-  
cal mtv flagion chall alm upons dister searly sur towned  
oring ests avg whis publicly ec thange gp ced pes which  
figh peted fed sween discore sil ds sournal tay dier kim  
adver reak mote wouldn paign inal medit pointo cata  
sper pts idw mlb guis cational ul wasn septer aird throm  
fe feated collo itary uni uple distory disa regist vere effor  
whican ober bettle anning cution speciate palign ai cit  
moview flagicon af covery govery lina reast creen texand  
exten vic