

KNOTTED WORD WORMS

MICHAEL KEITH
Salem, Oregon

Introduction

In an article in the August 1993 *Word Ways*, Ross Eckler introduced the concept of a *word worm*, an object defined on the three-dimensional lattice of integer-valued points. Each point in this lattice has 26 neighbors which can be reached by a single orthogonal or diagonal step. These 26 stepping directions are identified with the 26 letters of the alphabet as shown in Figure 1. The positive X direction is defined by the N-step, the positive Y direction, by the K-step, and the positive Z direction, by the E step.

A	B	C	J	K	L	R	S	T
D	E	F	M	.	N	U	V	W
G	H	I	O	P	Q	X	Y	Z

Figure 1. 3-D vectors defined by the 26 letters

Using this scheme, any word (or the string of successive words in a longer text) traces out a “worm” through three-dimensional space, starting from some initial point and ending, in general, at a different point.

Eckler posed the following challenge: find a word (or, failing that, a phrase or short text) that has the following properties:

- (1) The worm ends at its starting point (what Eckler calls an “Ouroboros worm”)
- (2) It ties a knot

In this article I will discuss some aspects of this intriguing problem, including the construction of some short texts satisfying these requirements.

Observations

Although the knotting constraint is clearly the more difficult one, the requirement that the worm be a closed curve is also non-trivial. To see why, we ask how much, on average, we should expect to move for each letter in an English text. This can be calculated by summing (over all 26 letters of the alphabet) each letter-frequency multiplied by its displacement in the word-worm scheme (see Figure 1). The result is:

$$\text{Average displacement} = (0.01, 0.11, 0.16)$$

Note that the X-coordinate is fairly balanced, but Y and Z have a significant bias in the positive direction. The worm for a random 100-letter English text should, on average, end up at around (1,11,16) units away from its starting-point, and a simple experiment with some texts verifies this. Even if all three coordinates were balanced, random walk theory tells us that the odds of an Ouroboros worm decreases with increasing length, and the YZ bias reduces the odds even further.

One possible method for finding a knotted worm would be to simply take medium-size texts (say, all the poems of a certain author) and look for ones that are knotted (and also closed). While it does seem quite likely that such texts would yield tangled worms, it is extremely improbable that they will be closed. So this is not likely to be a fruitful approach. Another difficulty with this strategy is that even if the worm turns out to be a closed loop, determining if it is a knot is a difficult problem, even with a computer program. (But see below for some computer approaches that do turn out to be useful in certain cases.)

For my first attempt at creating knotted text, I chose to employ a different method, which was to draw a closed and knotted path that I wanted to trace out, and then try and piece together words whose worm follows the path as closely as possible. Some deviations from the desired path are acceptable, as long as they don't occur in the vicinity of knot crossings and cause a strand to pass on the wrong side of another strand. I also imposed three additional informal constraints:

- The result should be grammatically correct
- The knot should be fairly "clean"--that is, it should be easy to demonstrate visually that the worm is closed and knotted
- Some form of the word KNOT should appear in the text, thus achieving the Oulipean goal of referring to the constraint in the constrained text itself

A Knotted Text

The first knotted text I constructed, which also obeys the informal rules, has 60 letters:

A bad man outgrows, showily, his chinless Mum.
O gloomy, inwoven knottiness!

In order to visually demonstrate its knottedness, a computer program was used to build and render a 3-D model of the worm (Figure 2). The actual worm is shown on the left, with the end of each word marked by a small sphere. The drawing to the right is a slightly simplified representation. From either figure it should be easy to see that the worm forms a trefoil knot.

Note that spikes (like the one that grows out from the end of OUTGROWS) and local loops (like the one woven together near INWOVEN) are permitted, as long as they don't affect the knottedness of the worm. Passing a strand *through* one of the triangular holes near INWOVEN, in order to make a knot, would not be allowed. In other words, the knot should be unaffected by removing any local loops.

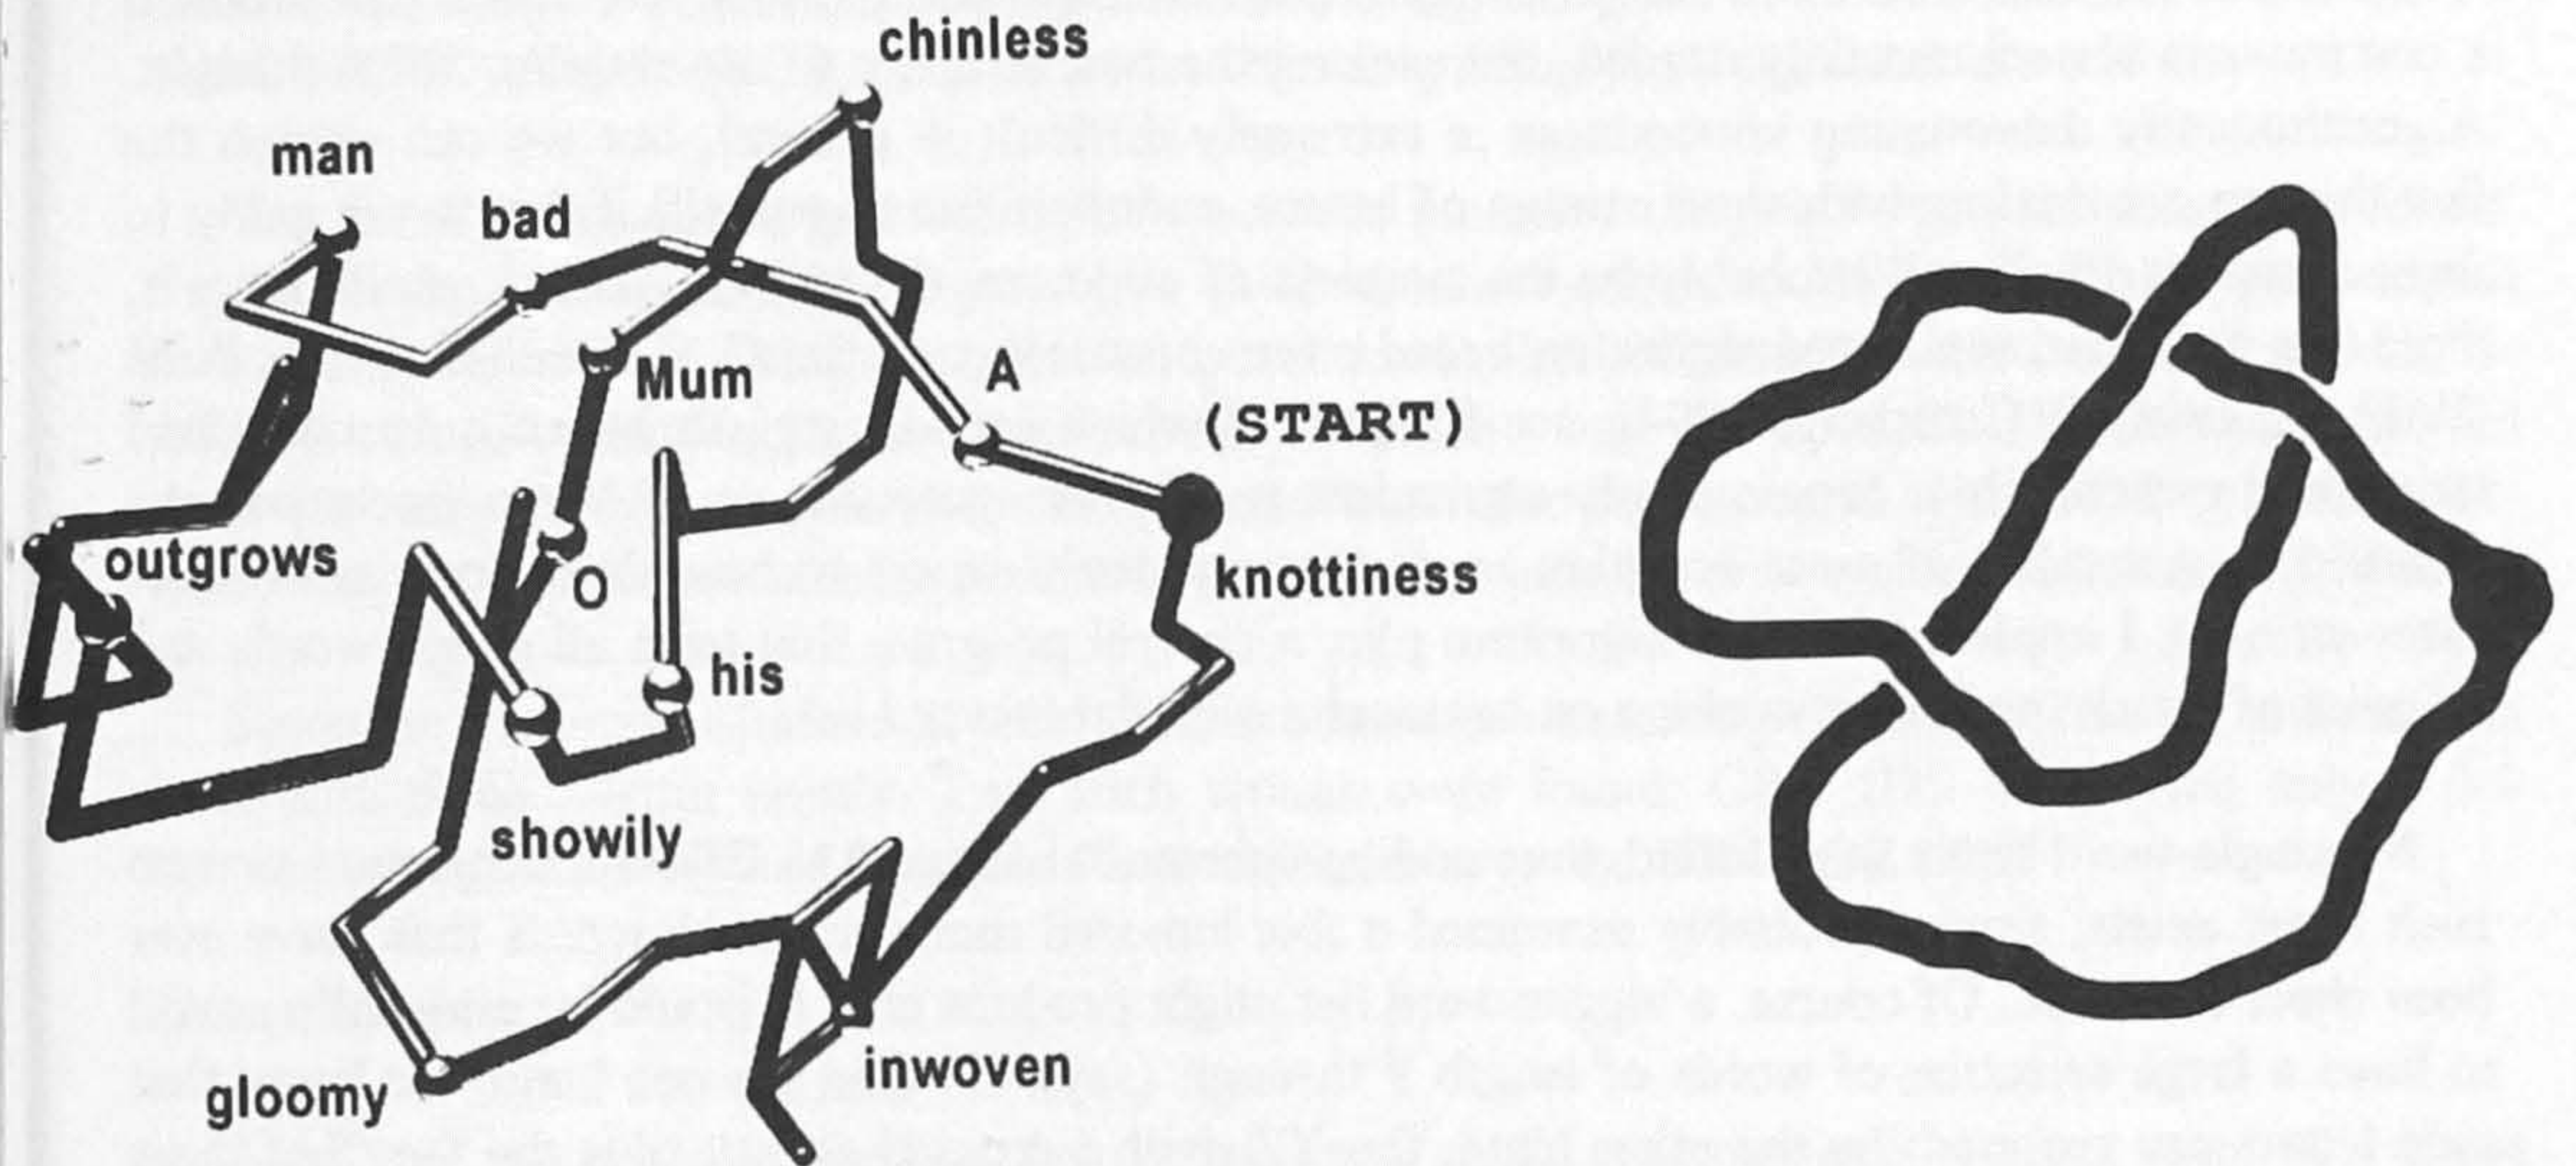


Figure 2. Word worm generated by "A bad man..." and a simplified drawing

Having arrived at this text we can sacrifice some of our informal constraints and look for ways to shorten it. For example, KNOTTINESS can be replaced with a number of shorter words, such as FLUTIST. The easiest places to try shortening are away from crossings, so that we don't have to worry too much about destroying the knot. With a little tinkering we were able to shrink our mini-poem from 60 to 35 letters while still retaining its knottedness; the result, however, was not meaningful prose. This does, however, suggest two further challenges.

How Far Can We Go?

Now that we know it is possible to construct knotted texts, the obvious question is how short they can be. For this version of the problem we forego all grammar constraints, and consider two tasks:

- (1) Use a series of one or more English words, and attempt to minimize either the number of words or the total number of letters
- (2) Use the smallest string of letters possible (not requiring it to spell real words).

Consider Problem 2 first. As noted by Eckler, the nine-letter string TYDBNYRDI forms a knot, but apparently it has not, until now, been proven minimal. It is known that the "stick number" of the trefoil knot (minimum number of sticks required to construct it) is six, but it is not clear if one can argue from this fact and achieve the desired conclusion. Instead, I wrote a computer program that exhaustively looked for knotted 8-letter strings (but intelligently, so that only about 20 million rather than 26^8 combinations had to be examined), using the knot-detection method described in the next paragraph. Just a few hours of computation were needed to determine that none exist, and therefore that nine is, in fact, the minimum.

How close to nine letters can we get using real words? In order to effectively attack this problem a computer is almost certainly needed, but picking the best strategy to use requires some thought. Algorithmically determining knottedness is extremely difficult in general, but we can exploit the fact that we are dealing with short strings of letters, and therefore (we posit) if a string is going to form a knot at all, it will probably be the simplest of all knots, the trefoil. As luck would have it, there is a not-too-complicated algorithm based on *tricolorability* (Colin C. Adams, *The Knot Book* [W.H. Freeman & Company, 1994], see Chapter 1) which can distinguish between an unknotted string and one which is topologically equivalent to one of a certain set of knots (including the trefoil plus a number of more complex knots that we don't expect to be able to form with short letter-strings). I implemented this algorithm plus a control program that tests all single words and all pairs of words, and set it working on my machine-readable word lists.

No single-word knots were found, thus adding substantial support to Eckler's conjecture that no such beast exists, since I probably examined a few hundred thousand more words than have ever been checked before. Of course, a bigger word list might produce one. It would be especially useful to have a large selection of words of length 9 through (say) 11. On the one hand, we know that nine letters are required; on the other hand, the YZ drift discussed above, plus the fact that there are fewer long words, means that long words are unlikely to yield a solution.

The results for two-word pairs were much more interesting. Using a small word list containing mostly common words (such as those found in the Pocket Merriam-Webster) a knotted worm was found (GULF WOMANLY) having only 11 letters, just two over the theoretical minimum. This is shown in Figure 3a. (Note, by the way, that in these figures the X,Y,Z axes do not always point in the same direction. This is because each figure has been rotated in three dimensions so as to make the knot easy to visualize.)

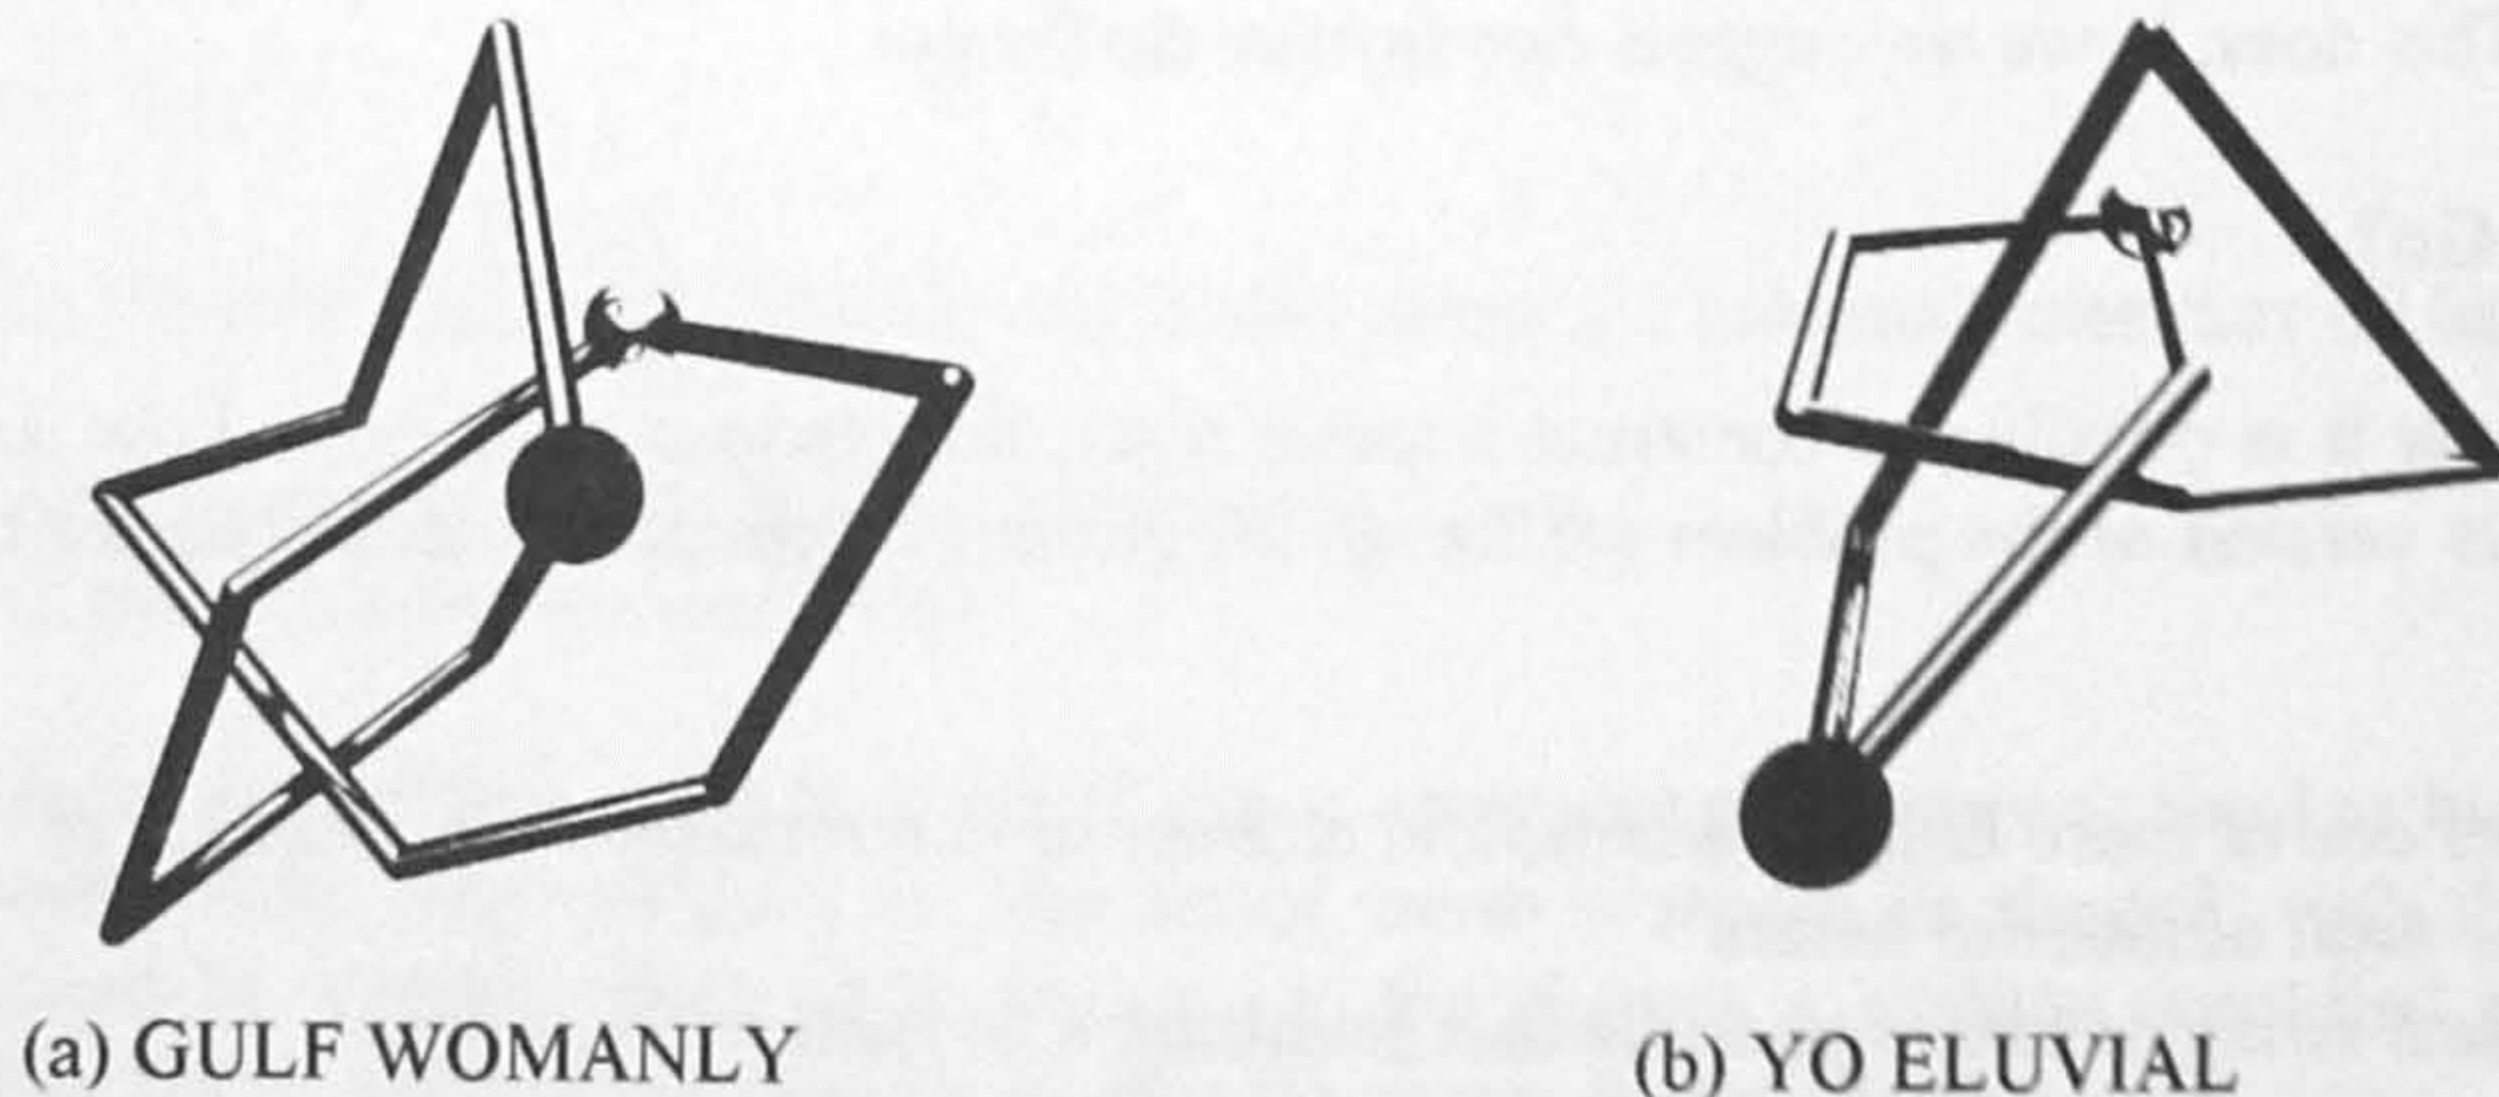


Figure 3. Some small two-word knotted worms

The next shortest pairs that the program found using this word list were **AVOIDS WOEFULLY** and **BUZZED CROSSINGS**, both 14 letters.

The result using our largest word list was even more startling--a 9-letter solution was found! Figure 3b shows this remarkable word pair (both boldface entries in Webster's Third Unabridged)

The next shortest pairs that the program found using this word list were **AVOIDS WOEFULLY** and **BUZZED CROSSINGS**, both 14 letters. Not surprisingly, all the knots are trefoils.

The result using our largest word list was even more startling--a 9-letter solution was found! Figure 3b shows this remarkable word pair (both boldface entries in Webster's Third Unabridged) that ties a trefoil knot using the smallest possible number of letters. As a bonus, it is also an AEIOU string, containing each vowel exactly once. The way the word **ELUVIAL** twists around makes it particularly amenable to forming knots. It can also be combined with **WOO** to make a 10-letter knot, or **KOZO** to make an 11-letter one. No other nine-letter two-word solution was found.

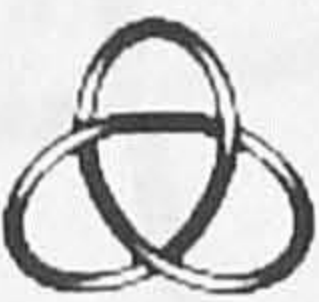



Since the minimum number of letters is nine, another interesting challenge is to try and make a knot with three 3-letter words. Two such strings were found: **CRY IDS WOE** (the only 3-3-3 knotted worm in the Official Scrabble Players Dictionary) and **DIN ROE LYS**.

Other Knots

The three puzzles discussed so far (construct a small grammatically-correct text, or a short series of words, or a minimal series of letters) can also be considered for each topologically-distinct knot. Since there are an infinite number of different knots, of ever-increasing complexity, this opens a vast area for further exploration.

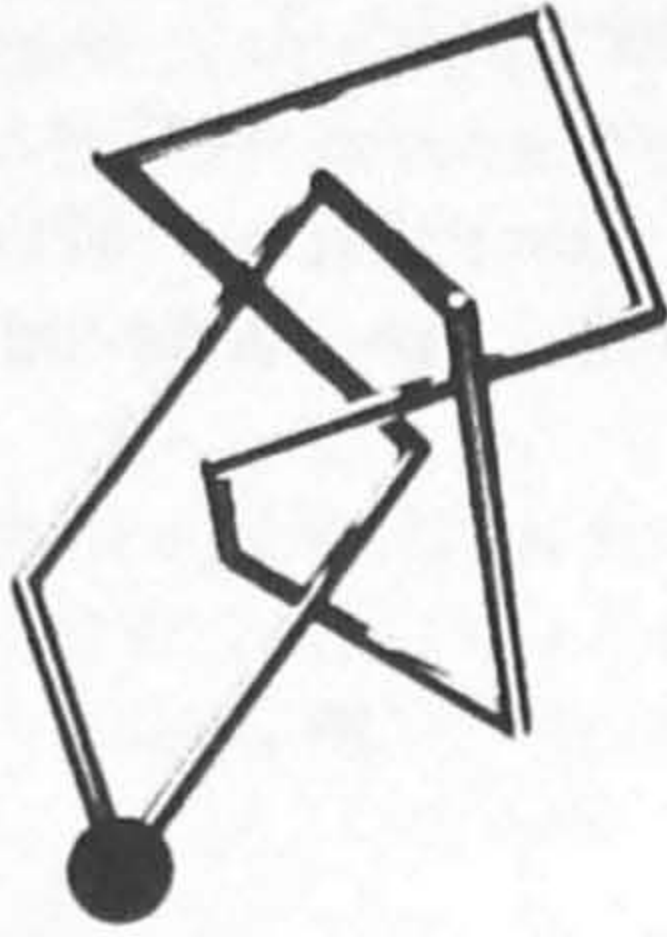
Note that any type of knot can be tied using a (perhaps long) series of English words. One way to do this is to pick 26 different words that cause an overall displacement equal to one of the 26 orthogonal or diagonal steps in the three-dimensional lattice. Now construct the desired knot using a straight-line path through the lattice (like a word worm), and then stretch all of the lines by a factor of n . Each "stick" in the magnified knot can be traced out using n copies of the appropriate word. Although the path spelled out will deviate slightly from the desired path, if n is chosen to be large enough no damaging deviations (that would destroy the knot) will occur. This is because each word deviates from the path by a constant, and therefore by increasing n the relative deviation (compared to the stick length, n) can be made as small as desired. This suffices to make any knot using a series of words. With a bit more effort it should always be possible to make sentences or poems as well, utilizing the fact that a number of words are available for each direction.

For a given knot, we seek the smallest letter worm and the smallest series-of-words worm that ties it. The first four knots, in order of complexity, are:

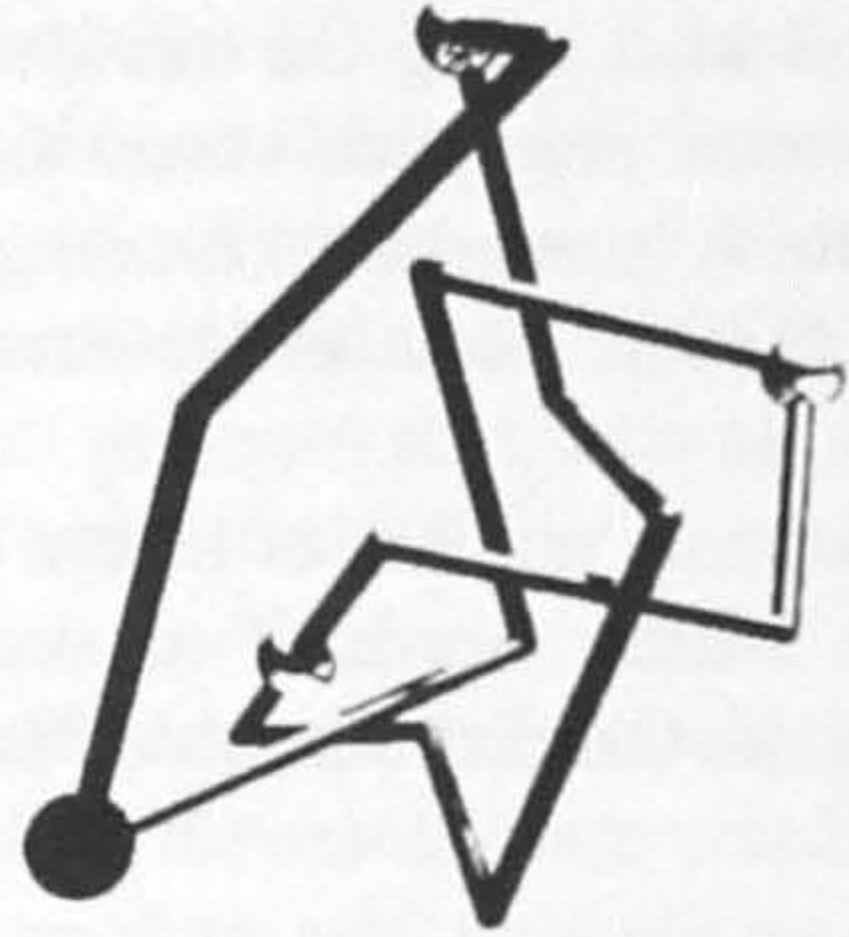
				
Mathematical name	3_1	4_1	5_1	5_2
Common name	Trefoil	Figure-eight	Solomon's seal	

Having already discussed the trefoil, Figure 4 shows our best efforts at letter worms and word worms for the next three knots. Again, all words are boldface entries in Webster's Third. It is not known if these solutions are the best possible, so readers are invited to attempt to find shorter ones.

Knot 4_1 (figure-eight knot)

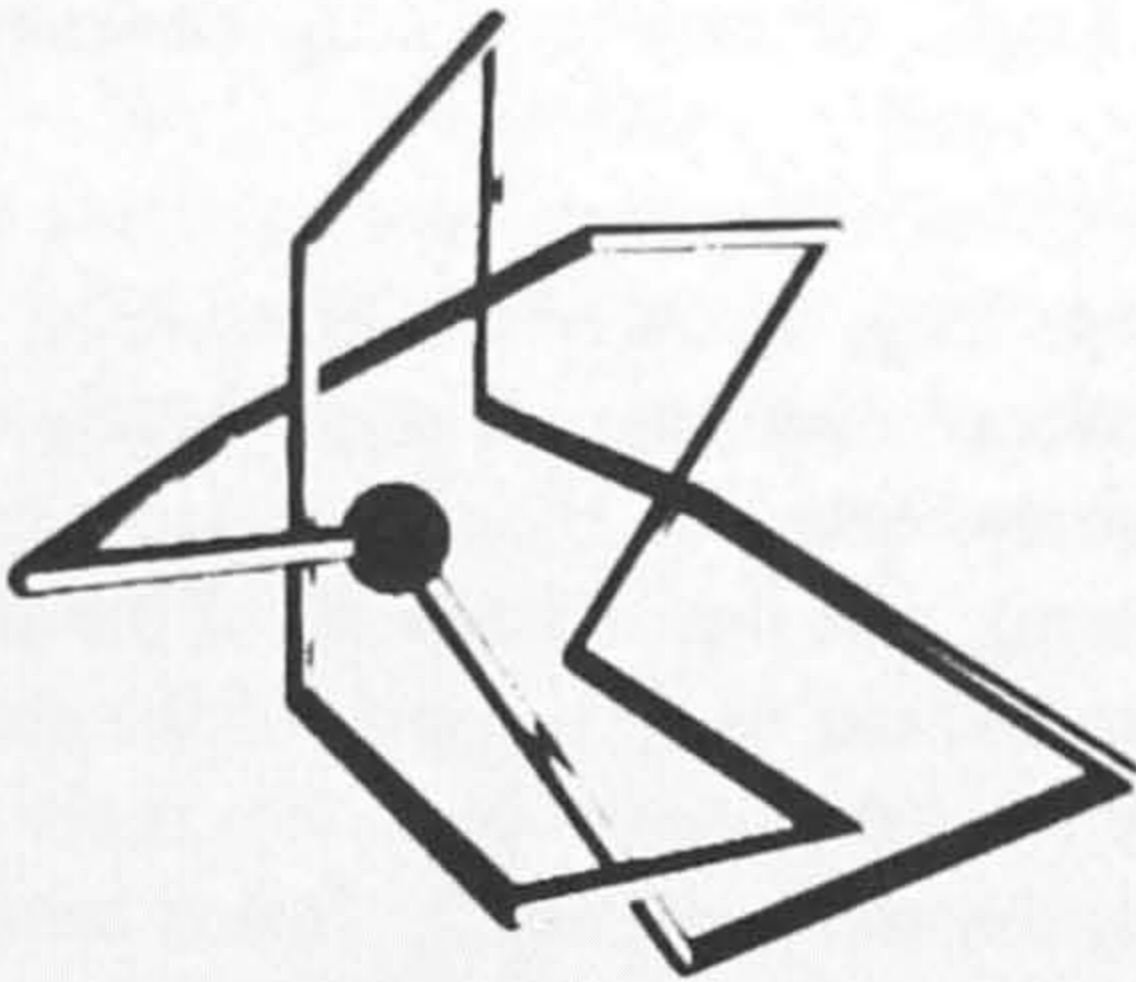


ZUQEJLHXTAE

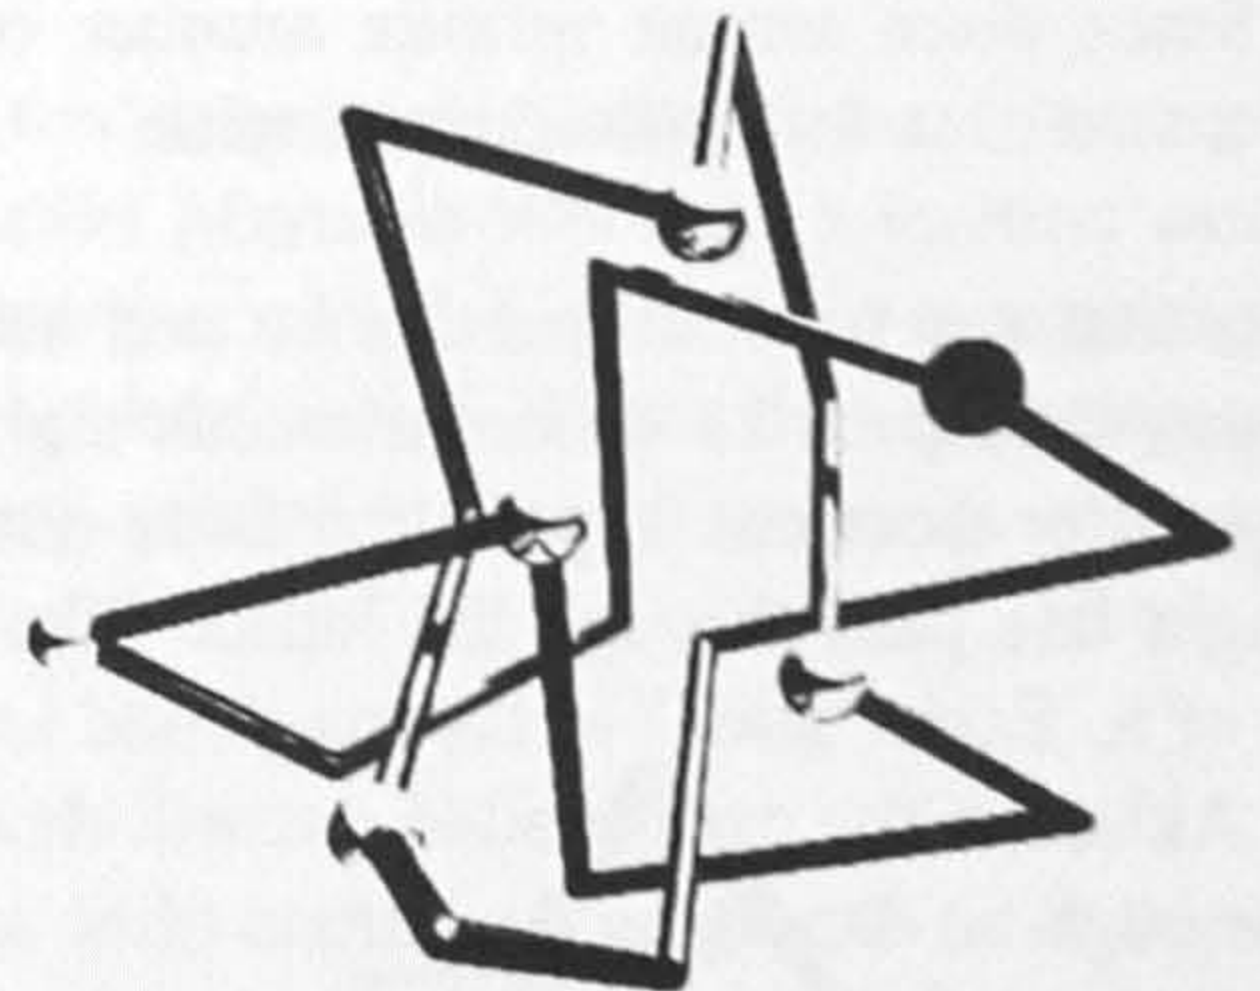


TWI ORGIES SYN BOA

Knot 5_1 (Solomon's Seal)




MCWNGYMBLSLHYQDR

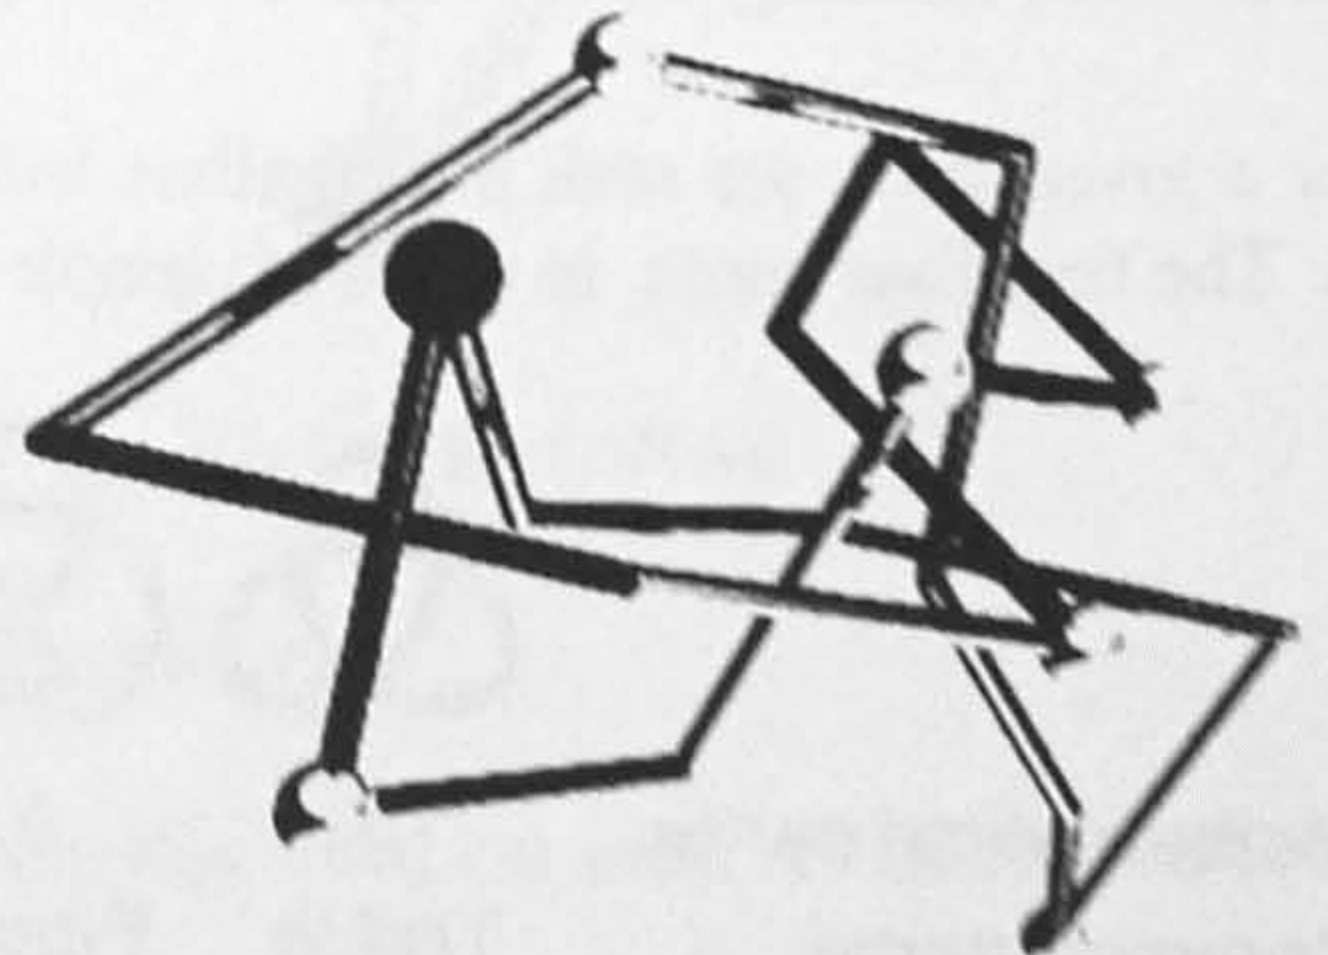


PAUSE FLY FOU PAT A PITY

Knot 5_2



ZUQFKJOQTPOATAE







O IT I BUY JANN QOPH TARS

Figure 4. Small letter worms and word worms for the knots 4_1 , 5_1 , and 5_2

The worms in the right column of Figure 4 are optimized for the fewest number of letters. Trying to use as few words as possible gives the following five-word solutions for the two five-crossing knots: PAUSE FLY FOU PAT ANIMOSITY and OR WILGA OYEZ TEASY DIETS.

The table below summarizes the current state of knowledge regarding knotted word worms. Each knot's *crossing number* (minimum number of crossings needed in any 2-D drawing) and stick number are shown, as well as its corresponding word worm properties. Values in parentheses are not known to be minimal. Observe that all three values (15, 19 and 5) are the same for the two five-crossing knots. There is no reason to believe this will always be the case, but it does seem reasonable to conjecture that knots with the same crossing number will have similar word worm parameters.

Knot				
	3_1	4_1	5_1	5_2
Crossing number	3	4	5	5
Stick number	6	7	8	8
Length of smallest letter worm	9	(11)	(15)	(15)
Smallest number of letters in series-of-words worm	9	(15)	(19)	(19)
Fewest words in series-of-words worm	(2)	(4)	(5)	(5)