

XXWord v2.0

The Crossword Solver's Apprentice

by Graham Wheeler
(c) 1994--1997 All Rights Reserved

Contents

[Welcome](#)

[Setting up a Crossword Grid](#)

[Giving XWord Additional Clue Information](#)

[Using the Mouse to define the Grid](#)

[Solving the Crossword](#)

[The Menus](#)

[Dictionary Consultations](#)

[Concept Keys](#)

[Crossword Save File Format](#)

Welcome

Welcome to XWord!

XWord is a crossword solver's helper. It uses some efficient dictionary searching algorithms for computing the set of possible solutions to a partially complete crossword. It can then show you which possible letters can be placed in each square, and which possible words can go through each square. The idea is that you use XWord as you solve the crossword, entering letters as you determine them, and using the feedback provided by XWord as additional clues. XWord can also sometimes surprise you by completing a half to two-thirds finished crossword correctly (although this depends on giving XWord some additional information).

XWord is distributed as part of WordsWorth. As such, it uses the same dictionary as WordsWorth, which does not contain proper nouns. This restricts its usefulness somewhat, but the problem may be addressed soon (Pentiums help!).

DISCLAIMER - This software is distributed without a warranty of any sort. You use this software at entirely your own risk. I will not be responsible for any loss, damage, injury, death, disaster, plagues, famine, divorces, disagreements, etc that may result from the use of this software.

Setting up a Crossword Grid

XWord is quite a powerful tool with a number of features. The best way to learn how to use it is to get a crossword and start using it. This manual gives a brief description of the process that you should follow to use XWord to help you solve the crossword.

The first task is to set up a grid. Choose 'New' from the File Menu. Enter the size of the grid in the dialog box (you can use the sliders to change the values). The grid will then be shown with all squares set to black. You now have to enter the layout of the white square. In the top left-hand corner is a red rectangle - this is the board cursor position. You can move it around the board with the cursor keys, or by clicking the mouse on a part of the grid. To mark a square as empty (white), put the cursor on it and type '?' (space will also do). To reset it to black, put the cursor on it and enter ". To enter a known letter at a square, put the cursor on it and type the letter.

As you enter square contents, the cursor moves automatically, generally following the direction you establish initially by pressing the right or down cursor keys. An exception is when entering a letter, where the movement follows the direction of empty white squares.

It is a good idea to set up the board with just " and '?' before entering any letters. Boards generally have some symmetry - this is exploited by having some commands to flip and reflect the grid in various ways. These are available from the 'Grid' menu. You can also set the default symmetry to use 'on the fly' as you enter the grid, from the 'Options' menu. As most crosswords are reflected and inverted about the center, it is a good idea to use this as the default symmetry (called 'reflect and flip' in the menu).

After entering the grid, it is a good idea to prevent accidental changes by 'locking' it (selected from the 'Grid' menu). You should then enter any letters that you already know. Once you have done that, you need to give XWord additional information about the clues that you still have to solve. This is described in the next section.

Giving XWord Additional Clue Information

XWord uses its dictionary to work out all possible solutions to the crossword. However, if there is anything wrong with the board layout constraints, the results will be unpredictable, and may include XWord concluding that the crossword has no solution at all.

There are two ways these constraints can be incorrect: you may have entered letters on the board that are wrong, or XWord is being led astray by the limitations of its dictionary. There is not much that can be done about the former, but there is some hope in the latter case!

Firstly, you need to mark any multi-word clues as being multiple words. You can do this using the mouse, as is described in the next section, or using the keyboard, which we describe here.

Move on to the letters that start new words (other than the first word) and use one (or both) of the following keys:

- ^ Mark square as start of a new down word (brown);
- > Mark square as start of a new across word (magenta)

The square will be displayed with a red line at the left edge and/or the top edge, indicating the start of the new word. It is a good idea to also mark squares which follow hyphens, as WordsWorth's dictionary does not include any hyphenated words.

If you suspect that there are words that are not in the dictionary (such as proper nouns), you need to tell XWord to not consult the dictionary for that word. Move on to each letter of the word, and mark it with one of the following:

- | Don't limit current square to admissible down words
- Don't limit current square to admissible across words
- + Don't limit current square by dictionary at all (a combination of the above two)

These attributes will appear as horizontal and/or vertical lines through the center of the square.

XWord will automatically add such override attributes to any word in the grid that is longer than the longest word in the XWord dictionary (13 letters at the time of writing).

Using the Mouse to define the Grid

To set up the grid with the mouse, you can click on a square to select it (the cursor will move to the square). If the square is already selected, then clicking with the *right* mouse button will bring up a dialog box which will allow you to set the options for the square.

You can also specify an allowed subset of letters that can be used for the square. This is useful in the case when you know that a clue is an anagram of a certain set of letters, but you don't know which anagram (you would presumably have done a dictionary consultation to find the anagram, so this is most likely to occur when the consult failed to find the anagram, usually because it is a proper noun). You can then at least tell XWord that it should only allow the known letters to be used for each square in that word (unfortunately you need to do this for each square in the word; on the other hand this gives a finer level of control). Actually, the new Anagram attribute mentioned below makes this feature largely redundant, but it has been left in in case it proves useful.

Clicking on the selected square with with the *left* mouse button brings up a dialog box that lets you specify settings for a whole clue (down and/or across). You can specify the length of the word(s) down and/or across, and whether the dictionary should be overridden down and/or across. You can also specify multi-word clues by separating the word lengths with commas. For example, a clue whose answer consists of a two letter word followed by a three letter word should be entered as 2,3.

Two additional attributes that can be set are the 'concept keys' for the word, and, if you know the word is an anagram of some other word, you can enter that other word here too. Using these attributes slows down the computation of constraints by XWord quite considerably, but they can be extremely useful. For more information about concept keys, see Section .

You can use any combination of mouse and keyboard actions when setting up the board. XWord will handle these consistently. The fastest way to set up a board is to use the left mouse button to enter all the clues which have squares in the top left quadrant (if the grid has two axes of symmetry) or in the top half of the grid (if the board has one axis of symmetry), save the result, use the Grid menu to complete the other three-quarters or half of the grid, and then enter any dictionary overrides and multi-word clue information in the rest of the grid, either with the keyboard or the mouse.

Solving the Crossword

After entering the additional info, select `Enable Constraints` from the `Solve` menu and then select `Compute Constraints` from that menu (or click on the white area surrounding the grid; this has the same effect). XWord will attempt to solve the crossword, or at least find the possible words that can be put in the grid, based on its dictionary, and the constraints you have entered on the grid. If you haven't filled in many letters yet, this may take a few minutes.

You can then move the cursor around to squares for which you do not know the letter (and which XWord didn't fill in when you asked it to solve the crossword). As you move about, you will see two things changing. At the bottom of the screen XWord displays all characters that it believes can be entered in the current square. If the square is empty, on the right hand side of the screen you will see up to 8 down words and 8 across words that can go through the square. The total number of possible words across and down is also shown.

As you move around, XWord works out all possible words that can go down and across through that square, and hence what letters can be placed on the square. When the board is nearly empty, the number of possible words shown on the right may be in the thousands. When it is a lot less than that, but still more than eight, you can view the entire list by selecting `View Across` or `View Down` from the `Solve` menu.

You now proceed by using the lists of possible words and letters that XWord displays, together with the crossword clues, to work out more letters. Every time you enter another letter XWord can narrow down the possibilities further. If XWord finds that a square can have only one possible letter, it will display the letter in red, without actually committing the letter to the board permanently. Solving the crossword then becomes a collaborative process between you and XWord.

If the number of possible words is ever shown as zero, then you have a problem. Somewhere you have made an error in setting up the grid, or (often the case) there is a word that is not in the dictionary and for which you have not told XWord to override the dictionary (or you have specified concept key restrictions that are not matching the right word, or incorrect anagram restrictions). As a result, XWord has concluded that the crossword has no solution. You will have to try to identify the problem and correct it (this can be tricky - a good approach is to move around the grid, discarding the constraints (from the 'Solve' menu) before each move, and look for places on the grid that still show zero possibilities - these are the potential problem areas).

HOWEVER, THERE IS ONE CASE IN WHICH SEEING ZERO POSSIBILITIES IS OK! This is if the word in the crossword is longer than the longest word in the dictionary. If this is the case, XWord will automatically override the resulting artificial constraint in any perpendicular cross words. You will notice this happening often in the demo version, where the dictionary is limited to words of six letters or less; any words of more than six letters will show up with zero choices, but the words that cross it will be unaffected.

The Menus

The File Menu

The Solver Menu

The Grid Menu

The Options Menu

Dictionary Consultations

The dialog box for dictionary consultations lets you generate lists of anagrams and words, given a set of letters. The set of letters you specify can include blank tiles, as well as things like 'any vowel' or 'any non-vowel'.

The method used is fairly complex but worth understanding if you want to make full use of it. Underlying each consultation are two pieces of information:

- o a word template which specifies which letters are allowed in which positions in matching words;
- o a letter pool, which specifies which letters are available for use in unconstrained entries in the word template.

A few examples can be used to illustrate the idea:

- o for the literal word ROPE the word template is "ROPE" and the letter pool is empty;
- o for anagrams of ROPE the word template is unconstrained (we will write this as "####") and the letter pool consists of the letters R, O, P and E;
- o for anagrams of ROPE that begin with R and have P in the third place the word template is "R#P#" and the letter pool is "O,E".

Given a word template and letter pool, WordsWorth tries to find matches by plugging different letters from the letter pool in to '#' positions in the word template (actually this is completely untrue - WordsWorth proceeds through the whole dictionary, and then for each word tries to find a way to construct that word given the template and letter pool. The effect is the same, however; it is just easier to prevent duplicate matches by stepping through the dictionary, which is sufficient reason to do it this way even though it is considerably slower).

In the consult dialog, the first choice you have is the type of match to use. There are four possibilities:

Use all letters Matches must match the whole word template;

Match prefixes Matches may match any leading subsection of the word template;

Match suffixes Matches may match any trailing subsection of the word template;

Multi-word anagrams This is a special case. The whole word template must be matched but the match may consist of multiple words.

The first and last case should be clear; the prefix and suffix matches can be illustrated by continuing with our last example. Using that example and prefix matches, WordsWorth would find words that match any of the following situations:

Word Template	Letter Pool
=====	=====
R	OE
R#	OE
R#P	OE
R#P#	OE

Using suffix matching, WordsWorth would instead use:

Word Template	Letter Pool
=====	=====
#	OE
P#	OE
#P#	OE
R#P#	OE

Thus in the first case, WordsWorth will see if any of the following are valid words:

R RO RE ROP REP ROPE REPO

In the suffix case, WordsWorth will check the following:

O E PO PE OPE EPO ROPE REPO

Note that in the case that the word template is unconstrained (that is, all 's), then prefix and suffix matching are identical, both having the meaning "words consisting of one or more letters from...". Reread the last sentence if you didn't understand it.

When making consultations, you don't explicitly specify the word template and letter pool. Rather, you specify a single pattern. In the simplest case a pattern is a sequence of letters (whitespace may be present but is ignored). The letters may be either upper case or lower case, the former meaning that the letter must occur in that position (i.e. is a word template constraint), while lower case means the letter need not be in that position (i.e. it is in the letter pool). Thus our example above would be specified as "RoPe" (or equivalently as "RePo"). A few more examples should help here:

Pattern	Word Template	Letter Pool
informatIOn	#####ION	informat
MISinformed	MIS#####	informed
Treated	T#####D	reate

The first two are typical examples of patterns that would be used for suffix and prefix matching, respectively.

Patterns can also have some special characters. Each has an 'upper case' and a 'lower case' form, analogous to letters. The special characters are:

Meaning	Upper Case	Lower Case
Any letter	:	.
Any vowel	*	+
Any consonant	=	-

Patterns can also have 'letter ranges' enclosed in square brackets, such as '[A-Z]' (any letter), '[AEIOU]' (any one vowel) or '[A-MZ]' (any one of A through M, or Z). To specify the opposite of a range, start the range with an exclamation mark or caret. For example, '[!AEIOU]' matches any non-vowel. The letter case of the first letter in the range determines whether the range will be used as a word template constraint or as a letter pool entry; other than this letter case is ignored. Thus, for example, "[a-z]" is legitimate and equivalent to "[a-z]". Similarly "[A-z]" is equivalent to "[A-Z]". You cannot use the other special characters within a range (thus for example, "[*y]" is not valid).

Some dictionary consultations can take a long time. You can interrupt a consultation by clicking on 'Cancel' in the dialog box. The matches WordsWorth has found up to that point will then be displayed. Quite often there will be thousands of possible anagrams and it becomes impossible to search out the good ones. There is a technique to help in this case - look through some of the possibilities, and select a word you think is appropriate in the final anagram. Then subtract this word from the input, and see what words can be made from the remaining letters. For example, consider VEGETARIAN. We see that some anagrams have the word RAVE, and decide to use this. This means we must find anagrams of the remaining letters, namely GETIAN. We find the word EATING. Thus VEGETARIAN=EATING RAVE. You can specify minimum and maximum word lengths for prefix, suffix and multi-word anagram matches. For the latter you can also specify the minimum and maximum number of words allowed per match. If these are set to zero (the default) they will have no effect. They can be useful when generating multi-word anagrams with patterns that contain several A's or I's, as WordsWorth will likely churn out thousands of matches in these cases starting with the words 'A', 'AA', and/or 'I'. Setting the minimum length to three (for example) can help prevent this.

Concept Keys

A new feature in version 2.0 is the use of 'concept keys'. In WordsWorth you can only use these when doing dictionary consultations; in XWord you can also associate concept keys with words in the crossword grid.

Concept keys are essentially semantic categorisations of words, somewhat like a thesaurus. The file containing the keys is named "thes.dat" for this reason. This file contains semantic associations between nearly 25,000 words. When you specify a concept key, the dictionary will be constrained to match only those words with possible semantic associations to the key. These semantic associations can be quite loose, as they are based on transitive relations. That is, a word A is considered to be potentially related to a concept key C if there is a word B such that A and B are related in meaning, and so are B and C. Furthermore, the way that concept key restrictions are implemented in WordsWorth/XWord is not infallible - extra words can be admitted that are not actually related, although the actual number of erroneous matches should be a very small percentage. This fallibility is a tradeoff; the algorithm used sacrifices complete accuracy for speed. Concept key restrictions have a very low impact on the performance of dictionary consultations (although somewhat more of an impact in XWord grids).

You can specify more than one key, separated by commas. Keys must be single words. If more than one key is listed, the intersection is used, but you can change this to a union by starting the list with "+". If there are no words satisfying the concept key(s), they are ignored. This is enforced more strongly in the XWord grid, where the word lengths are also used; i.e. if you specify concept keys for a six letter word in the grid, and there are no six letter words matching the keys, then the keys are ignored.

To illustrate the use of concept keys, here is an example. Say you have a six-letter word for which the clue is "A swamp". You know that the second letter is "O". Then, either by entering it on the grid via the word settings dialog, or in dictionary consultations, you can specify the pattern ":O::::" and the concept key "swamp" to restrict the search. This will result in the following matches:

BOLDLY
FOXILY
MORASS
NORMAL
TOPPLE

If, on the other hand, you had specified the concept key "bog", then the following words would be matched:

COARSE
COBWEB
HOOKER
HOOPLA
LOTTED
MODERN
MOLDER
MORASS
MOTHER
ROTTEN
SOILED
TOILET

The combination of the two ("swamp,bog") gives the best results, namely the intersection of these two lists which is the single (correct) word "MORASS".

The obvious place to use concept keys with XWord is when you have clues consisting of a single word. You can then enter that word as a concept key. It is sometimes a good idea to try a dictionary consultation with the concept key first - you may find that there are only a few matches, none of which is the word that you are looking for, in which case using the concept key in the grid would be a bad idea.

Crossword Save File Format

For those who are interested in such things, this section describes the format of saved crossword files used by XWord.

The file begins with the dimensions of the board, and the basic grid layout. Here is an example:

```
13 13
```

```
CHAFF#S.....E
A#.#OMAHA#.#R
N....#GENUINE
A#.#..O..#.#C
SANTA##TEAPOT
TA.....#.#.#
AR..#.#.#...P
#O#.#.#...H...E
.N..K.##A...R
.#.#O.D.R#.#V
....A.O#ACCRA
.#.#L.V.R#.#D
RAMPAGE#ELOPE
```

Because we only have a single letter to represent each square, if a square contains a letter and a flag or more than one flag, not all of the information can be represented. However, in this case the board can be followed by additional lines of the form:

```
<row> <column> <'command' character>
```

or:

```
<row> <column> <'command character'><argument>
```

to specify this additional information. The full set of command characters that can be used is:

.	Empty square
?	Empty square
#	Black square
^	Start of new down word in multi-word clue
>	Start of new across word in multi-word clue
+	Do not restrict this square at all
	Do not restrict this square in a downward direction
-	Do not restrict this square in an across direction
A-Z	Put this letter in the square
=	Concept keys across
:	Concept keys down
&	Anagram across
*	Anagram down

The last four are followed by a string argument containing the concept key list or anagram letters. Here is an example of such a file as saved by XWord:

```
15 15
```

```
; Created by XWord
```

```
#.....#.....#
#.#.#.#.#.#.#.#
....#.....
#.#.#.#.#.#.#.#
.....#.....
```

```
#.###.#.#.#.#.#
.....#.....#
###.###.###.###
#.....#.....
#.#.#.#.#.###.#
.....#.....
#.#.#.#.#.#.#.#
.....#.....
#.#.#.#.#.#.#.#
#.....#.....#
```

```
2 11 >
6 7 ^
8 10 >
8 12 >
8 13 |
9 13 |
10 13 ^
11 7 ^
12 6 >
```

Note that saved board files can have comment lines starting with semicolons. However, these CANNOT appear before the board dimensions!

The File Menu

The file menu in Xword mode lets you save and load your crosswords and grids, start new grids, or quit XWord.

The Solver Menu

Enable Constraints - Enabling constraints turns on XWord's word matching as you move about the grid. Usually you will not do this until you have finished entering the grid, as it serves no purpose at that stage. When constraints are enabled, moving around the grid with the keyboard is much slower.

Disable Constraints - If you have enabled constraints but find that you need to move around the grid quickly with the keyboard for some reason, you can use this to disable constraint computation.

Compute Constraints - XWord will work out the current constraints on the entire grid. The more empty squares on the grid, the longer this will take. Nonetheless, it is usually acceptable to do this each time you solve a clue and enter it in the grid. A fast way to do this is to click on the white area surrounding the grid.

Discard Constraints - XWord discards the constraints it has worked out. XWord will do this automatically if you change the grid in any way other than entering a letter, but just in case you can use this menu option to force a discard.

View Across this pops up a window with a list of all the words that can pass through the current square in a horizontal direction.

View Down this pops up a window with a list of all the words that can pass through the current square in a vertical direction.

Consult Dictionary this is identical to dictionary consultations in WordsWorth. It is discussed further below.

The Grid Menu

Flip Grid Mirror-reflect top half of board setup left-to-right;

Reflect Horizontally Copy left half board setup to right half;

Reflect Vertically Copy top half board setup to bottom half.

Set Grid Symmetry This lets you set the 'on-the-fly' symmetry to use while the grid is being entered. Initially the default symmetry (set from the 'Options' menu) is used; you can use this option to change to a symmetry other than the default;

Lock/Unlock Grid Locking the grid prevents any further alteration of the basic black square/white square layout. This is useful once you have entered the grid, to prevent it from being accidentally modified, which can happen quite easily if you enter a word with automoving on, and the current automove direction is not set correctly. Unlocking the grid will allow you to edit it again.

Enable/Disable Automove This lets you turn cursor automoving on or off.

The Options Menu

Configure Colours This lets you choose the colours that XWord should use when drawing the grid;

Configure Options This lets you set the default crossword size and symmetry, and select the dictionary file that should be used.

