Directions for Scripting Analysis in Arena 3.0, For Conversion to AIF (Analysis Interchange Format), by Kenneth W. Regan, revised 2/23/13.

1. Download Arena 3.0 from http://www.playwitharena.com/?Download:Arena_3.0 (http://www.playwitharena.com/) You will get a Zip file, or WinRAR file if you have that installed. Extract to a folder named "Arena3.0" in a convenient location. The program itself comes as a ready-to-run .exe file---no installer needs to be run.

2. Download the engine if you do not already have it. It does not matter where you place it---it need not be in the Arena "Engines" folder. For example, you can download the Stockfish zip file from http://www.stockfishchess.com, extract the distribution folder, find the executable engine(s) appropriate to your system and place them anywhere---though I do have mine in a sub-folder of Arena3.0\Engines\. Later on as Stockfish is updated, I just add the new .exe files to the same folder.

3. Open Arena 3.0. There are several configuration choices to make, but let's install the engine first. Select Engines→Install New Engine, and navigate to where the desired .exe file is located. If there are both 1-cpu and multi-CPU ("Deep") versions of the engine, select the 1-cpu version. Besides 32-bit and 64-bit versions, you may have choices compiled for particular CPU architectures, or particular optimized versions such as "-ja-" for compiles by engine-source expert Jim Ablett. Select what seems appropriate, and if there is no 1-cpu version already, do step (4a) below. Select the UCI option, and answer yes to the dialog to start the engine.

4. The wide bottom band of the Arena window (above the narrow gutter at the very bottom which may say "Tournament Game in 5 Minutes") is where engine-setting info and analysis appear. Right-click in that area, and select "Configure Stockfish-231-64-ja" (or whatever engine name) from the top of that menu. Ctrl-1 also selects that option without the need to right-click.

4a. If the engine has a box saying how many "Threads" or "CPUs", and you find it set to a number greater than 1, select it to 1. Then click "Save as new engine." Add "1-cpu" somewhere to the name of the engine---e.g. replace the "[001]" affixed by Arena by "1-cpu".

5. **Multi-PV And Other Settings.** The project is based on the theory that one needs to have in-depth analysis of all reasonable moves to judge how a player decides based on his/her options, and one also needs an accurate appraisal of blunders. All positions have a point where the Nth-best move is so bad that there is no need to distinguish moves as worse-than-N. Many positions have 20+ reasonable moves, however, so I believe a setting cutting off analysis after N moves needs N to be at least 30. Very few have more than 50 legal moves, and even though some engines allow N up to 99, the cost of extra blank lines showing up in Arena's output files isn't worth it. Here are recommended settings with particular engines:

i. Rybka 3.0 and higher: Set MultiPV = 50, and importantly, also set the "MultiPV_cp" box to 400. The latter "cap" feature cuts off analysis of variations that are over 4.00 pawns worse than the best move, which is a considerable time-saver with Rybka versions. [Rybka has a bunch of other settings, especially v4.1: Set "Always score main move" ON, and "Display Upperbounds" ON. Leave "Random" at 0, leave Nalimov frequency at the default (36), leave PreserveAnalysis OFF, and leave Ponder OFF. Leave every setting in the right-hand column as-is. I don't have an opinion about UseLargePages or MainProcessPriority or the use of a hash file, nor the Log or Logx option. On a laptop computer, set CPU Usage to 95% or 90% or thereabouts.]

ii. Houdini 3: Set MultiPV to its max of 32. Leave MultiPV_cp at 0---unlike with Rybka it seems not to save much time by pruning, and possibly it affects only the display. Leave other defaults as-is: Ponder OFF, TB paths and hash and learning files empty.

iii. Stockfish 2.3.1: Set MultiPV to 32 or 50---your choice. There is no MultiPV_cp option. Leave other defaults as-is (especially OwnBook off), but can turn Ponder off.

Optionally, you can make and save a separate Multi-PV version of the 1-cpu engine. Make sure it is active (Engines→Manage→Details→highlight it and click the "Start this engine right now!" button), right-click to configure it again, set Multi-PV as above, click "Save as new engine", and add something like "MPV" to the name.

Now we proceed to settings that affect the particular quality of analysis and its output format. Some of these, especially tablebases and hashtable/caches, are managed by Arena as "Common" settings across all engines. Others, most importantly the search depth (13 for Rybka, 17 for Houdini, 19 for Stockfish) are selected when starting the analysis itself.

Configuring Scripted Analysis in Arena, For Conversion to "Analysis Interchange Format" (.aif)

We will eventually broaden our scripts for AIF conversion to make some of the following steps unnecessary, but they are helpful for now. Some steps can be varied, such as hash table size.

6. Go into Engines→Manage and click the Details tab (if it is not selected already).

6a. In the tabs at right, leave everything under "General" as-is, but click "Special". Then if the box "Values absolute (always from white's side)" is not grayed out, make it checked---note also step 8 below. Under the "Books" tab, make sure there is no attached book (i.e. leave path empty and leave "Use Arena mainbooks" OFF). If there was not already an option to drop the CPU usage to 90—95% in step 5., you can select "Strength(%)" = 90% or 95% here.

6b. Now select the "Options" master tab. Set "Engine names replace player names" OFF, leave the rest as-is.

6c. Now select the UCI tab. I think now it does not matter whether "Common own-book setting" and "Engines use their own book" are left checked ON; I've had them on for Rybka 3. For the hash table you have two options:

i. Leave "Common hashtable size" checked, and enter 256 MB, or 512 MB if you have fewer than 4 cores or more than 2GB of RAM.

ii. Uncheck it. Go back to step 4, i.e. hit Ctrl+1 or right-click in the analysis pane (below the chessboard) and select "Configure <name-of-engine>)." Now you should see a new box (usually near the top) for setting the hash size---choose 256 or 512 there.

6d. Regarding endgame tablebases, you have 3 choices, of which I've used the first two:

i. Put "Common endgame tablebases" and the Nalimov cache unchecked. (Then go back to the engine configuration and make sure anything with "Nalimov" or "Gaviota" or "Scorpio" is off.)

ii. Download and install just the 3-piece and 4-piece Nalimov tablebases, e.g. from http://kirill-kryukov.com/chess/tablebases-online/. Set the TB path accordingly and cache = 4MB.

iii. Download and install the 5-piece set, but remove the "kqpkq" files. The endgame of K+Q+P versus K+Q has cases of mate-in-50+ moves that are beyond the ability of mortals to comprehend, and not only cause a mismatch to human perception, but also derange the program's values of other positions. If you do this, use the default 32MB cache.

7. Either from the Engines menu again, or by right-click as in step 4., select "Configure Engine Analysis Output." There are 3 main tabs.

7a. Under "Design of the Info-Boxes", I've changed "Nodes total" and "Nodes per second" OFF, and "Tablebase usage" ON even with TB's off. (I don't remember whether this matters.)

7b. Under "Mainlines", leave all six "Visible Elements" ON. Below that, *important*: select "Move number of the current game". I don't think "Show ponder analyses" and "Show personality" matter---I have them ON. Under "Notation of moves" (currently impt), select "With hyphen", and THEN select "Short notation in the mainlines" which grays it out. Ignore the Winboard engines group, and below it, leave "Show selective search depth" and "No newline at same principal variations" ON. It is OK to leave "Show times with milliseconds" off.

7c. Under "Options", IMPT (for now): Change "Insert mainlines at the top" TO "append mainlines at the bottom." And under "Alignment of the analysis elements", select "Left-aligned". Everything else can be left as-is, except that the new "knodes/sec" option in Arena 3.0 should probably be OFF.

Click OK to save. There are still more Options to configure in the GUI itself…

8. Select Options→Appearance. At left click "Movelist", and make sure "Short Notation" is selected at upper right. Under PGN, change "Go to move 1 after loading" to "0". Then click "Other Settings" and **IMPT**: set "Values always from White's point of view" ON. Leave "Check 50 Move Rule" set to "Always", and other things I think don't matter. Click OK to finish.

9. I believe Options→"Configure Tablebases" is redundant, but it can't hurt to make it agree with what was chosen under step 6.

10. *Finally,* for long scripted analysis, we want to **shut off** the automated logging of individual UCI commands and feedback. To do this in Arena 3.0, go Engines→Log Window… Under the "Settings" tab (which should be on by default), find the "Logging" pulldown, and set it to "nothing" or "a few." (Failing to do this seems to cause analyses to "bog down" after a few hours, with engines like Stockfish or Houdini whose output is more verbose than Rybka's.)

This completes the configuration of Arena itself; last we configure Automatic Analysis.

Scripting Analysis in Arena. [This section is now fairly complete but may need revisions.]

11. Do Shift-Ctrl-A, or Engines→Automatic Analysis. The window has 4 tabs. Some settings need to be re-selected at each invocation, while others are preserved. In reverse order:

11a. "Options" tab: set "Minimum search depth" = 0, leave other things unchecked.

11b. "Output" tab: make sure "Write in protocol file" is checked and "Append" is set. Choose a path to where you will collect analysis files, and see below for naming the file itself. Leave everything else unchecked---the "Comprehesive report" is much less informative.

11c. "Engines" tab: use "Loaded engines" as IMHO the multi-engine option brings unhelpful complication. For "Level" select "Use plydepth" and depending on the engine, enter:

i. Rybka versions: 13

ii. Houdini versions: 17

iii. Stockfish versions: 19

This seems like a weird difference, but my tests show these strengths to be approximately equal, about 2700 overall---which may be 2900 in early middlegame sliding down to 2400 in endgames. Rybka groups the bottom 4 nodes of its search, and some also allege "obfuscation" as masking a "true depth" of 3 or 4 ply higher (that I sometimes see "Depth: -2" show up in the Arena display makes me suspect a simple subtraction of 3 or 4). They also take roughly the same time to complete. Critter and Komodo versions also have rough parity at depth 17. There is a technical reason to prefer an odd number.

11d. "Source" tab: Use "forward" direction, and under "Colours" check both White and Black moves. (Even when a run has crashed after a White move, I avoid re-starting on a Black move because one may forget to re-check the "White moves" box, and also the next time White will be checked off by default which is easy to miss.) The rest depends on using PGN vs. EPD.

PGN: Load one or more PGN files in the middle window. Uncheck "All moves" and enter 9 to 250 (the 250 is a default). Check "All games" to analyze all games (see re-starting note below), or uncheck it to enter a range of games. Unlike the current (buggy?) behavior with EPD ranges, Arena will reflect changes to the range made while analyzing---which is especially helpful as it allows changing the lower range to 1 so as to pick up the next file at the start.

EPD: Owing to particularities and what I regard as bugs in Arena, I find it confusing to load more than one EPD file at a time. The main reason to work with EPD at all is that unlike with PGN, Arena issues a "ucinewgame" command after each move, which most engines implement as a command to clear the hash tables. This promotes reproducibility, and promises better reflection of my theory about predictions for moves whose values swing up or down at higher depths. Later we may also save time by preparing the EPD files to avoid re-analyzing theoretical positions.

Although Arena shows "999,999" as the default length for EPD runs, it writes only the first three digits of the position number, and I limit individual runs to 1,000—1,500 positions with Rybka and Houdini, 500—750 positions with Stockfish. Hence if the file itself has more, we must select a range. I prefer to end a range on the last move of a game, rather than at a round number, but this requires loading the EPD file from the main menus first in order to look. The "gotcha" is that currently *Arena ignores changes to EPD ranges made while analysis is going.*

12. Naming files: My convention is to use the name of the PGN-or-EPD file minus the extension, then (optionally) an underscore, then an abbreviation for the engine, depth, "e" if using EPD, a Roman numeral for the number of PVs, lowercase letters b,c,d… for separate output files from the same event, and finally ".log" as the extension. Here are some engine and settings identifiers:

i. R3d13L.log: Rybka 3 to depth 13 on a PGN file in 50-PV mode.

ii. R41d13eLb.log: Rybka 4.1 to depth 13 on an EPD file in 50-PV mode, second range in the file.

iii. Hou3d17eV.log: Houdini 3 to depth 17 on an EPD file in 32-PV mode ("V" since 32 = 2^5).

iv. SF231d19.log: Stockfish 2.3.1 to depth 19 on a PGN file in Single-PV mode.

I think of 512MB hash as the default for MultiPV, 256MB as the default for Single-PV runs as those are for quick-checks and comparison. Changes from this can be noted in the file or in the name. I myself use a modification to the Stockfish source code (GrainSize = 1 in "evaluate.cpp"; "is_draw" in "position.cpp" revised to return "true" only on the third occurrence of a position; make sure "uci.cpp" has "TT.clear()" as the action for the "ucinewgame" token---I am not sure how much this matters for fixed-depth analysis) and call it "SFG1" rather than "SF231".

13. Aborting and re-starting an analysis. The cleanest way to abort an analysis is just to X the entire Arena window, but sometimes (especially with Houdini and with Rybka within the first 15 or so seconds after a move) this does not complete, and the engine process and/or Arena itself needs to be killed via TaskManager. The surest way is to go to the front tab ("Source") of the "Automatic Analysis" window itself and click "Cancel", but then Arena writes partial output for the current move (sometimes twice), which should then be deleted by editing the output file manually. You may also wish to delete the last White move even if it completed.

It is OK to write to the same output file after re-starting, **but make sure "Append" is selected** (I wish the "Overwrite" option weren't there at all, for safety). Even though re-starting writes the game-header info again, it is fine to leave that in the output file.

To re-start at a particular move of a particular game in a PGN file, start "Games (PGN)" from the number of that game, and start "Moves" from the number of the move to re-start at. (One can also un-check "White" under "Colours" to re-start at a Black move, then quickly re-check it, but per above I avoid this.) After re-starting, re-set "From 9" under "Moves" (or whatever move you start game analysis from), and re-check "All games" under "Games" (or adjust the starting game number) if there are more PGN files to come in the batch.

Re-starting at a particular position in an EPD file is similar, but note that both "Moves" and "Colours" are inapplicable, and there is no point in adjusting the range of positions after re-starting since on-the-fly changes to it are not picked up.

14. Communicating analysis files. Even Single-PV analysis files of 45-game tournaments can be too big for e-mail attachments. I am able to post mine on the Web for others to get. Please contact me for other particular details. The author of BitTorrent was the first non-family member to learn of the birth of my son; perhaps that counts for something…