

Eclipse BETA v10.0.0.17

INPUT

New stroke data format to include shadows and confidence scores

Steno strokes throughout the program (except in dictionaries) now contain additional data that is stored, displayed, and used for processing. Historically, they contained data about the keys pressed, the keyboard layout type, the absolute time of day, and the relative time offset from the beginning of the recording.

Two additional data sets have been added:

Shadows will be stored for each key when using a Passport writer, so you can see the same indication of weakly-hit or partially-hit keys both in the note file and in the steno tracking window.

Confidence scores are stored for each stroke based on the frequency of use of the dictionary entry and the complexity of the stroke pattern (or, for ASR or voicewriting, based on the confidence level returned by the speech recognition algorithm.)

For now, you can see confidence scores in the tracking window during editing if you also turn on relative timecodes. Further refinement is expected for the display, including possibly a separate option.

User settings/Input USB realtime setup button opens Device Manager

If you have COM port selected as the input device, it opens the port setup, or for voicewriting, it opens the speech settings. With USB, there are no settings in Eclipse, so it opens the device manager as a convenience so that you can quickly check to see if the device is present.

Expression note reading should emulate Diamante to read SgStn

The ProCAT Expression emulates a Diamante with its note file format, so using the "read notes" function now takes that into account and reads the notes in Diamante format.

Make Dragon "recognition history" available

For voicewriters, the vox utilities function inside Eclipse now includes access to Dragon's recognition history dialog, which should be accessed while a session is still active.

List dragon profile name in job info in file manager

It is automatically stored as one of the job variables, called DRAGONPROFILE=Your Profile Name, so you can see it in the info window in the file manager along with the other form field variables such as the case number, witness name, etc. And similarly can also be searched for.

TRANSLATION

ASR Assistant

Automatic Speech Recognition technology is not sufficiently powerful to produce an accurate transcript without significant editing, and likely never will be. It cannot reliably punctuate, identify speakers, resolve subtle ambiguities, or editorialize obvious verbal errors such as stumbling over words, mumbling, repeating, or overtalk. They also make frequent errors on acoustically similar very short words such as "an" and "in." Human reporters do all of these things easily.

However, under ideal conditions, ASR is quite good at difficult vocabulary or fast speaking, two things that human reporters sometimes struggle with.

The ASR assistant is found under User settings/Translation/Translation Magic, and has an activity slider much like the Translation Magic setting (these two features are not mutually exclusive and can be used together.) This allows you to decide how much leeway the ASR system has to help.

In order for this feature to work, you must make sure that channel 0 is an audio recording set to PCM, 16000hz, 16-bit.

You must have an Internet connection to use this service, and while the beta period is free, the final release will include some free minutes and a per-minute charge after that.

When writing with the ASR assistant, the software compares what you write with what the ASR results are, concentrating only on the words. If they are the same, it does nothing. If they are different, it compares the confidence scores for the steno with the confidence scores for the ASR. If the ASR result is better, it will be applied to the text as an edit. The assistant creates a jobname_ASR.not file containing the voice notes from the recognition process. This allows the file to be re-translated with assistance without re-processing the audio.

These edits will appear as User text style 8. The recommendation is to set the background color to something obvious, and to set the User settings/Edit/Typed in text so that User 8 is no longer considered typed-in; this will allow you to easily make globals to words that have had ASR corrections made.

The corrections are done as minimal edits, so if the original steno was "climb" and the ASR result was "climbed" it will simply add the "ed" in User 8 color.

The algorithm has dozens of hard-line rules about what not to do. It will never remove or insert punctuation, or paragraphs, or Q&A or speaker designations because ASR is terrible at those things. It will avoid duplicating words, inserting verbal "noise" like "you know" if the steno reporter has deliberately omitted it, and many other rules.

If you find the ASR too often editing correct steno results and replacing them with errors, there are several things you can do:

1. Move the slider more towards steno to prioritize the steno results more. This boosts the perceived confidence of the steno and will tend to leave steno results alone even if they're wrong.
2. Improve the ASR results with better audio and more initial information. If you build a job dictionary ahead of time, the text of the job dictionary entries will be given to the ASR, which can include both vocabulary that the ASR does not know, and also uncommon words that are phonetically ambiguous to give the ASR a hint about which one is likely to be correct, such as "Shirley" vs. "surely."
3. Use the User settings/Programming/ASR Exceptions list to add words which should disqualify the results if the ASR hears them. This could include everything ranging from profane or embarrassing words to extremely common homonyms or other words which the ASR gets wrong more often than it gets right.

Note that drops or untranslates on the part of the steno reporter are still registered as zero confidence and can still be replaced by ASR results regardless of these settings changes. To avoid those, don't drop and don't write untranslates. You're welcome.

REALTIME

CART Window Speaker Color Assignments

If you have assigned individual speaker colors in the speaker table, these colors will now appear on the CART window as well.

CART Window musical note

The paragraph symbol used to indicate a musical note on the closed captioning display will now also appear as a musical note on the CART Window.

CART Window should populate from top when clear

Rather than rolling up from the bottom, clearing the CART window (or starting a translation) will cause the new text to start at the top and will only start scrolling when the window is full.

Changing Output Formats in Realtime

In order to limit potential crashes and undefined behavior, some additional limitations have been placed on what can be modified in realtime. It is no longer possible to use the "Delete" button to remove an output type after realtime has started, and the output settings details now restrict a number of the settings so that they're disabled during a realtime job.

Show untranslates in pending translation display phonetically

If the User settings/Translate/Phonetics option is on, the pending translation will also be phonetic. Note that this is independent of the Translation Magic setting.

Edit word buffer should support suspend output

After considering carefully how this use-case would work, it was decided that when using Edit Word Buffer, it's the editor who has control of the output and needs to decide exactly when to suspend output or not. To that end, it suspends immediately, regardless of what has been written, so when the editor gets to the point where the highlighted text (which has not been broadcast yet) is the portion that should not go out, the editor should suspend the output, move the cursor past the text that should be skipped, and then resume the output, leaving it entirely independent of the steno.

Also, note that skipped output does not prevent the text from being re-broadcast with the "apply edits" feature. The assumption is that the "apply edits" feature is disabled, since that feature is essentially mutually exclusive with the edit word buffer in the first place; you would never use both at the same time.

ASR preview

At the bottom of the steno window, in the pane where the dictionary entries usually appear, an ASR preview will show up if using the assistant in realtime. This shows the intermediate results as they're heard before it generates the final results that will be applied to the text around 30 seconds later.

AUTO-BRIEF

ASR feed into auto-brief

In addition to applying corrections to the document, if the ASR assistant hears a word that you do not have a way to write in your currently selected dictionaries, it will immediately generate a brief for it.

Allow users to customize order of brief types displayed

Under User settings/Realtime/Autobrief, there are now separate ordering indicators so that you specify which category is displayed first, second, third, etc.

Refine pop-up brief window auto hide feature

The window should not auto-hide unless it's placed such that it would be obscuring text that you're attempting to edit. Some users have screens large enough that it can be placed in the gap between the right margin and the edge of the screen, which is unobtrusive and should no longer auto-hide. Users that have to place it somewhere over the text will still see it auto-hide when the cursor moves up high enough on the screen.

Show short dictionary entry comments on auto-brief display

If any dictionary entry has a comment that is two characters or less, that comment will appear on the auto-brief display. This will allow the MK (Mark Kislingbury) AN (Anissa Nierenberger) and EV (Ed Varallo) expert hints to appear with the initials designated whose brief it is, but you can also use two-letter codes in your own dictionary entries if you wish for them to appear in the reminders for any reason.

Allow emergency auto briefs when AB Steno Theory is too strict

If the steno theory limitations make it completely impossible for the auto-brief feature to come up with a brief that fits the text, it will make a second pass with reduced rules. This will generally result in a double-stroke brief, which is better than none at all.

Auto-brief reminder dictionary should reference original source dictionary

If you see a reminder that is actually a mistake or slop stroke or old theory or other type of dictionary entry that you would like to remove, open the REMIND dictionary file and if an entry appears with no comment, it's a copy of an entry in your main dictionary. If it came from a job or special dictionary, that dictionary name will appear as the comment for that entry, making it easy to find and remove or repair.

AUDIO

Opus compression support

Opus is magic. It compresses extremely high quality 48000hz, 16-bit audio, which would normally be about 350mb per hour, down to about 25mb per hour with virtually no audible quality loss. It makes GSM sound like a gramophone.

Opus is completely incompatible with .wav files, which equates to some compatibility issues with other software, but now that Stenograph has jumped on the bandwagon, this is less of a concern. Eclipse can play back jobname.opus files recorded on Case Catalyst or the Luminex II writer. Opus files recorded in Eclipse should theoretically work on Stenograph products, as well. No conversion necessary.

Note that an opus recording will create an .opus file, not a .wav file. It's an entirely different file format not just a different compression mechanism.

The speed controls and volume threshold features work fine with the Opus codec since Eclipse extracts the data and processes it as raw waveform data before sending it to the playback device.

Due to its small size and high quality, it makes an excellent solution for Team Editing and ASR applications since both require sending audio through cloud-based services. It has been tested with IBM Watson and Google Speech API.

ASIO device support for multi-channel audio

ASIO is an industry standard created by Steinberg for multi-channel audio recording devices for the music industry. As such, it has broad support across the audio interface device world, making it possible to use an extremely large variety of audio interfaces.

This required adding an "Advanced" audio input mechanism, and all devices supported by this mechanism will start with A:devicename in the list of input devices. When it's available, make sure to select the one with ASIO in the name.

This interface type perfectly synchronizes the samples in each channel so that they can be played simultaneously without echo, and also ensuring the integrity of the mix. It's especially useful when recording large numbers of channels such as 12 or 16.

Identify audio devices by name instead of number

When several audio devices used on a system are plugged into USB ports, some of them could be plugged in in a different order from one day to the next, or one day a device could be plugged in and another day it could be omitted. Because Eclipse stores the current audio device by device number from 1 to the total number of devices, this can cause the device setting in Eclipse to be wrong if the devices weren't plugged in exactly as they were when they were originally selected.

In order to eliminate this problem, Eclipse now stores the device name separately. When loading the user settings, it checks the list of available devices to see if the previously selected device has changed to a different number and will move to that device number if it has. That ensures that the correct device will always be selected if it is present when you run Eclipse. (If it is not present, no change is made.)

Speaker designations shown on waveform display

If you are recording one microphone per speaker, it may be helpful to see the speaker names by the waveforms. After the "Mix" channel, the rest of the channels should show the current speaker name.

Channel ignore feature

If you left-click on the waveform of one channel in a multi-channel recording, it will toggle on and off an indication that reads *IGNORED*. This channel will be ignored for purposes of transcription. That means that it will be omitted from the mix, which also has the result that it will not be sent to team editors or be analyzed by the ASR assistant. Helpfully, however, the jobname_channelnumber audio file WILL still have a recording of what the person said while they were being ignored in case you do need to go back and review it and possibly add it to the transcript.

This ignore feature is especially helpful if you have situations where many people are individually miked and have frequent side-conversations with the person next to them when they're not supposed to be on the record. While many microphones have a mute function to prevent this, the attendees may not know this or may not use it, or may not cover the mic, etc. This not only gets chatter into the mix, but also dumps data into the ASR features such as the assistant.

It can also be helpful during breaks so that the assistant/record are not recording off-the-record chatter. Note that if you "ignore" the MIX channel, this merely prevents it from being sent to the ASR, but does not stop it from recording.

Store separate timecode offset and ms/sec for each audio channel in each job

The "Timecodes" dialog, found under User settings/Document/Timecodes or Tools/Multimedia/Audio/Timecodes now has a "channel" selector.

This is because the timecode offset and wav ms/sec settings could theoretically be required to be different for different .wav files if they were recorded on different devices. For example, a backup audio .mp3 recorder could have been started at a different time from the original .wav file, and it could "drift" away from the original PC recording by a small percentage.

For that reason, you can now set these items differently for each channel.

Allow multiple sync to cursor measurements to automatically determine ms/sec

The "sync to cursor" button is very good at easily synchronizing a .wav file that was recorded outside of Eclipse, such as a backup recorder or videographer, especially now that it works when the audio playback is paused, allowing you to find the location of the synchronization more more precisely.

The wav ms/sec setting, however, is a great deal more difficult to find the right value for. Historically, it has required guesswork, and slowly bumping it up or down to find the proportion by which the backup recording drifts away from realtime timecodes.

This version completely automates the process. If you use the "sync to cursor" function in the document in one location, it sets the offset. If you then go to a different location (at least ten minutes away) and use the "sync to cursor" function, the software now has two

data points to look at, which is enough information to calculate the ms/sec proportion automatically. You can repeat this process multiple times if you like; each time it is only looking at the last two measurements, and measurements farther away from each other will produce more accurate results.

Audio auto stop/play/both should add the new "pause" as an option

If you select "pause" then any time a global or "type in text" command is started or stopped, the audio will be paused and unpaused instead of stopped and played from the cursor position.

Wave form display should appear in realtime stats window

If you have selected "waveform" that will now appear in place of the "audio 1" bar. Note that the "audio 2" bar, which is historically for stereo recordings and for voicewriting, will stay the same, because stereo is now disabled, and Dragon does not provide waveform data, just a level.

"Block write" filenames with underscores should not partition .wav file

Using block/write to write out an excerpt of a file would, in version 9, write out an excerpt of the .wav audio data to go with it. This will still work if using a typical filename syntax.

However, the jobname_subpart syntax is used to have multiple files associated with the same master, and as such, a jobname_subpart.ecl will still play the jobname.wav file, a feature that dates back many years. For this reason, if the block/write function is performed with an underscore in the filename, it will no longer make an excerpt of the audio and will no longer modify the timecodes of the excerpted text.

Expand audio DSP range

A number of additional sampling rates and bit rates have been added to the list of rates eligible for the DSP (Digital Signal Processing) features in Eclipse. That means that playing back an audio file with a different speed or volume threshold should theoretically work for just about any codec/settings selection.

EDITING

Proofreader special error tags

Several additional special tags have been added to the proofreader for paragraph-based structural errors that are technically not part of the text, but which require the proofreader to look at several paragraphs to make a determination that there's a problem.

In order to make it possible to detect these, it adds some text to the beginning of the proofread paragraph content text. Just before the paragraph designation which looks like <<Q>> it will add >>EEEE with an error code in place of the four E's, as follows:

>>NOBY – Indicates a missing by line

>>DBLQ – Double question.

>>DBLA – Double answer

>>DBLS – Double speaker (with the same name)

>>MTPG – Empty page (caused by a page break/new page command after a blank line at the end of a form page that has just barely stretched over to a new page. Easy to miss, and really annoying when printing or creating a PDF and proofreading it.)

Proofreading rules can easily be added to the existing list just by adding these codes, like so:

Double Question=>>DBLQ

because all the proofreader has to do is see that this sequence of text is present to know that there's an error.

Block mark for globalling Typed-in-text

If block marking text before hitting Ctrl+G, much older versions of the software used to assume that was a text global. However, too many users were accidentally doing text globals instead of steno globals without intending to, so that was changed to do a normal steno global. It was made mandatory to check the "text" box in order to force a text global instead.

However, if block marking a sequence of text that is typed-in, where no steno tracking is available, and then hitting the global function, there's no downside to automatically detecting that the likely intention is a text global, so the text box will be checked in that circumstance.

Have special character window default to the last font selected

This only applies for the current editing session, but should still be convenient for anyone inserting numerous special characters from a particular font.

List file "Edit" button should jump to item

Highlight the list file item first and the "edit" button will both open the list file and perform a search for the currently highlighted item. Why didn't it do that before? Good question. I didn't think of it, and no one asked until now. A good example of a ten-minute feature. Keep the feedback coming.

Add PGH properties to context menu

When you right-click on plain text within a document, there is now both a "properties" selection for the text, and under the "Paragraphs" submenu, a "properties" selection for the paragraph properties, allowing direct access to the paragraph status, type, font, label and spacing.

Typing a speaker from scratch should populate both master and current

When translating, if you write STPHAO as {S:MR. ROBERTS} it probably replaces a {S:MR. LEFT} master entry, or a {S:SPEAKER01} master entry. If you write ROB ROB, on the other hand, there's probably no master, so it is added to the speaker table with MR. ROBERTS both in the master and current columns.

Typing a speaker from scratch in the document should be consistent with that, so if you're not replacing or defining a particular master speaker to be changed, it should enter the text in both the master and current column.

Allow speaker table to swap speaker "current" assignments while leaving "master" the same

Tricky. If you use the "move up" and "move down" buttons on the speaker table, what are you moving? Go back far enough, and it would move everything, including the shortcut, which was probably not what you wanted (why would you want speaker 3 above speaker 2?)

But what to do with the master/current designations? If you're setting up the speaker table and realize you've changed your mind and you want to re-assign your numbered master designations to different people, there's no easy way to do that. However, there's no guarantee that that's what you're doing.

In order to attempt to disambiguate this without adding extra buttons or extra options, the software attempts to determine this automatically. If the speakers are the same (MR. ROBERTS, MR. ROBERTS) then BOTH will be moved, because clearly those should stay together. If the "current" is EMPTY and the master is populated, then the MASTER will be moved, because clearly you're re-arranging your master speaker table (perhaps you want SKWRAO as 1 and STPHAO as 2 instead of as 2 and 1.)

However, if the master and current are BOTH populated, and DIFFERENT, then ONLY the current speaker name will move, assuming that what you're doing is re-arranging which name is assigned to which standard speaker designation.

Insert time/date should use cursor position instead of current time if cursor is not at end of document

Even in realtime, people do go back several pages and insert time/date commands into form pages during a break, such as recesses, titles, etc., and in this case, it should use the timestamp on the current stroke in the document, not the current time, to automatically determine the timecode you're likely to be inserting.

TEAM EDITING

Team editing host option for "LAN only"

Occasionally, a reporter has access to a scopist or team of scopists in the same building, on the same local network. In this case, it makes the team editing much faster and more reliable if they connect through the LAN. However, if the reporter has an Internet connection, the default behavior is to start the team editing session on the Internet.

This version has an explicit "LAN only" checkbox on the session dialog, making it possible to force the feature to work on the LAN without having to disable the Internet connection first.

Note that you cannot mix and match LAN and Internet editors in the same session. If you need both, everyone must connect through the Internet.

Save team editing chat

The chat log will now be saved under sessionname_chatlog.txt (regardless of the document name) in each user's jobs folder when the session is closed.

Remember Persistent Session name

If you use a persistent session name, that same name and password will appear in the dialog the next time you start a session. Remember to clear it or change it if you want a whole new session. A new document will always be started, so this could be a convenient way to dramatically speed up the connection process for throw-downs on a day-to-day basis.

Paragraph status auto-assignment should trigger on exit

The automatic assignment of paragraph status for things like "edited" and "proofed" happened as soon as the cursor touched the paragraph in previous versions. This version will temper that and will only change the assignment when the cursor leaves the paragraph, much like the process for deciding when to transmit the paragraph to the Bridge protocol for the "apply edits" feature.

Minimize paragraph status colors on Display

The block background color turned out to be a bit jarring even with subtle colors. Paragraph statuses will now be limited to small stripes on the left side of the paragraph.

Make "reverted" status optional

If you remove or rename the "reverted" description from the paragraph statuses, the document reconciliation process will no longer attempt to flag paragraphs which have been re-transmitted from the reporter during a reconciliation.

Team editing background color warning options

The User Settings/Display color options now include four separate Team Editing warning colors that can be modified to present different background color warnings for editing status:

Reconciliation – indicates that a discrepancy has been detected somewhere in the document and the document is reconciling with the reporter.

Recon pgh – indicates that this specific paragraph is different from the reporter's version and needs to be retransmitted from the host.

Last pgh – indicates the last paragraph in the document (the one that the reporter is currently writing into) which should generally be avoided by editors unless working on a LAN or fast, reliable Internet.

No server – indicates that the Internet connection has dropped and is not currently available and continuing to edit will likely result in the edits you make will be reverted. (A few edits will be kept in a queue, but it's rolling the dice whether they would conflict with other editors you cannot see due to the connection outage.)

PRODUCTION

Allow additional nesting of index items

It is possible to create additional index categories for PDF files for multi-part indexes by placing descriptive index lines within the created indexes. This version allows you to also specify an additional header parent with the HEADER//item text syntax, such as follows:

```
ix:Witnesses//Joe Smith
```

And any index lines for direct/cross/redirect/etc., instead of being inside just a Joe Smith index heading, will be placed inside a Witnesses heading, inside a Joe Smith subheading.

Bundles should allow extra PDF files that aren't transcripts (billing pages, etc.)

There is now a third type of PDF bundle available in the delivery feature. In addition to "multi-type" and "multi-volume" there is now a "custom bundle" type which allows you to specify both the contents and the name of the bundle extremely specifically.

After selecting "custom bundle" as the type, you are given a field to type in a template for what you want in the bundle. This can include PDF files, ASCII files, or any other file you like, and can include specific file names, or file names including the current job being delivered, and can allow you to specify the name for the bundle itself. The only two special commands are as follows:

%J – replace with the job name

%B – replace with the job name, and use this for the bundle name

Here is an example custom bundle template.

Final-%B.PDF,%J_MULTI-IX.PDF,%J.TXT,%J-Billing.PDF,OrderForm.png

With this example, if you delivered a Smith.ecl job, you would get the following:

Smith_MULTI-IX.PDF (presumably, the multi-page with index)

Smith.txt (the ASCII)

Smith-Billing.PDF (A custom billing PDF you created using whatever software.)

OrderForm.png (A generic paper order form image for future orders.)

all of the above files will be placed in a bundle named

Final-Smith.PDF

Delivery firm name, e-mail, etc., options for separate bundles

The delivery e-mail feature does allow you to select other file extensions for different file types to e-mail to clients, and the destination allows you to specify a number of different e-mails addresses to receive the delivery.

But what if different clients want different things? Perhaps one doesn't want a multi-page, or one wants an HTML version, or whatever.

In the "Destination" options, you can now specify specifically which files get transmitted to each separate e-mail address, and once again, you can use the template tokens described above.

For example, here is what your "destination" list might look like for two different clients:

bob@lawfirm.com,%J_Bundle.pdf,%J_invoice.txt,businesscard.jpg

joe@freelance.net,%J_MULTI-IX.pdf,joeinvoice.pdf,businesscard.jpg

PDF: "Use Display Colors" Enhancement

The "use display colors" option for the PDF output will now also use the background color if it differs from the page background color.

Line Number Formatting

You can now add a line number formatting option in the page number format separated by a pipe sign. For example {P####}|{L<0###} shows up to a four-digit page number and will pad all line numbers to three digits with zeros, such as 005, 012, etc.

Print more dictionary data

The dictionary print will now contain the "used" instances field in conjunctionp with the created/modified/used dates.

Output to PDF while in a note or dix file should indicate completion

When creating a PDF of a note file or dictionary, the software will now open the PDF file once it's complete so that you can review it.

ANALYSIS

Analyze documents briefs and frequency improvements

The briefs included in the document analysis will now include many additional suggestions which could now include items that you may want to add to your main dictionary, including common word phrases that you do not already have dictionary entries for. Previous versions limited the phrases to capitalized words only.

In order to make this, and the other analysis information, more useful, the other improvement is that the analysis report will include the usage score visible in the text, such as "Consolidated Widgets [14]" which tells you that this phrase appeared 14 times in the document. The items are still sorted by usage with the most frequent ones first (after the supercategory of capitalized vs. uncapitalized words.)

In order to make that information as useful as possible, note that it is NOT necessary to remove the bracketed frequency score if using this file directly in the "build dictionary" command. If you run "build dictionary" on a text file and it contains a line like the above, it will still only ask you for an entry for "Consolidated Widgets" without the frequency text.

Note file statistics including delete stroke analysis

If you are in a note file and you run the Tools/Job report function, you will get statistics just based on the notes. This includes the total number of strokes, the total number of strokes deleted, and then a breakdown of the number of times each key was pressed, the number of times each key was deleted, and a proportional comparison of the two accurate to five decimal places (since you probably care more about how often you delete a particular key compared to other keys, rather than the pure number of times you deleted it.)

Method to clear drag/drop history

In the User settings/Translate/Translation Magic History dialog, you can now right-click on any key to clear the statistics for that key, or right-click on the heading at the top to clear all of them.

Font Size Adjustment for Lesson Player

The lesson player now adjusts the font size in two separate ways: First, it follows the Windows scaling setting, making it work better for high-resolution displays, and secondly, it follows the zoom level last set for the lesson player dialog globally (Note that this adjustment is not made dynamically since the lesson player does not use the standard user interface font, so it may be necessary to zoom the lesson player and then close and re-open the dialog for the font to change.)

Complex constructions in SpeechText.dix

There's a subtle problem with the SpeechText.dix, in that it tries to figure out what the ideal original voisteno would be for a particular sequence of text, based on entries in the dictionary.

Sometimes, that's easy. If you have defairs = Department of Internal Affairs in your SpeechText.dix, and it sees "Department of Internal Affairs" in your document (whether you typed it, wrote it out the long way, or whatever) it will turn it to "defairs" in the raw

jobname.txt file that you're using for Dragon food. Easy.

Other combinations containing number triggers like volume = Vol.{#R} or paragraphs like khaki = {A}yes{Q} are MUCH harder to reverse engineer from the text.

In order to address every possible combination of things like this, speech users with advanced shorthand (shortvoice?) can now add "fixers" to their SpeechText.dix that is a catch-all for any possible combination of things (number triggers, paragraphs, conflicts, whatever) that cause the resulting output not to correctly interpret a sequence as the original voisteno shortcut and instead write it out as the long-form individual strokes.

The catchall is this: Use an exclamation point to indicate retroactive briefs to spell out the shortcut. You can make entries such as

kakapo = !kak okay peerko

That means if the software sees "kak okay peerko" in the voisteno it's about to send to the jobname.txt file, it will change it to kakapo (because you have a kakapo = {Q}okay{.} dictionary entry and the speechevent parser didn't figure out that particular combination of things in reverse from the text in the document.)

USER INTERFACE

Show keystrokes assigned to macros on macro list

Under User settings/Edit/Macros you can now see the currently relevant keystrokes next to the macros (if you're in Hyperkeys, you'll see hyperkey assignments, otherwise you'll see standard keyboard assignments.)

Preserve keyboard customization editing position

When modifying a keyboard setting, the User settings/Edit/Keyboard dialog will now re-select the setting you previously had selected.

Allow user to type icon name in customer toolbars add/remove columns

Under Window/Customize toolbars, you can now press the first letter of the command name to navigate through the list and jump to the next item that starts with that letter.

Allow text changes for "Macro 1-15" toolbar items

You can extend the macro names for toolbar-activated macros with the pipe sign, such as Macro1|Prospective Juror and that text will now be used on the toolbar.

However, note that due to size limitations, it will only use the first ten letters of the description since every button has to be the same width, and a single long description would spread out all of the buttons. So the above macro would appear as "Prospectiv" on the toolbar. Abbreviations such as "P. Juror" are recommended.

Macro list should have a duplicate checker

When exiting from the macro list, if you have modified anything, you will be warned of a duplicate macro name.

Store user selection dialog attributes universally

If you change the size, zoom, etc., of the "pick a user" dialog that appears when you first enter Eclipse, that information goes into your user.ini file, which means that the software doesn't have access to it when the program first opens.

In order to allow customization of that dialog to be useful, it now stores the data (for that dialog only) in the Eclipse.ini file so that it's available immediately when the program is opened.

Allow recent file menu to show only ECLs

It's rare that users want to go back into recently opened dictionaries or note files directly, since those are more likely opened while working on a document.

File Manager Scaling/Item Spacing

Adjustments have been made to the file manager scaling and item spacing to account for somewhat crowded displays on high-resolution monitors.

"Open File" dialog should default to empty

Previously, it would often default to containing the currently job name, which only makes sense under rare circumstances, such as when you're in a job file and you hit "open notes" and you're likely to want to look at the note file for the current job. It now starts empty most of the time instead.

Spanish added to Primary Transcription Language list

This allows for any additional Spanish-language-specific adjustments to the translation algorithms and editor that cannot be accomplished purely by modifying settings.

CLOSED CAPTIONING

Conditionally create paragraph breaks for each script line when sending/timecoding

If the current script paragraph doesn't have steno associated with each line, then in order to timecode each line in the realtime document, the send script line command **MUST** create a new paragraph for each line.

The software checks the available steno for the script paragraphs and should make the determination automatically about whether this is necessary.

SRT paragraph formatting

Most .SRT captioning exports, as it turns out, would be better if each individual line of the paragraph were exported as its own caption, rather than exporting each paragraph as a caption.

"Force Margins" prohibited except under special circumstances

Turning on force margins when it's not necessary is such a common mistake that it's now simply prohibited except under a very specific circumstance: If outputting multiple formats, where the primary format follows the paragraph margin settings, and the secondary format has to be different. Otherwise the option is simply disabled.

BUG FIXES

Captive Portals in Networks and Eclipse Startup

Connecting to Eclipse while on a hotel network or other network requiring a login was getting Eclipse stuck in a server connection loop if no access was available. This version should resume normal offline activity if the Internet connection is "working" but only showing a login screen.

Ctrl-I Command in Note Files

The Index command was timecoding the current stroke.

Typeover Tracking and audio playback

Playing back only the currently marked block is a function that should not trigger if the block is marked because the user is in the middle of a typeover tracking suggestion.

Single Character entries/Period Punctuation Toggling

A single digit followed by a period is no longer considered an abbreviation and adding a comma after it will replace the period.

Stenolink and relative timecodes

The stenolink feature was sending timecode values of zero. It should now contain valid timecodes.

"Are you sure you wish to stop recording" yes/not box escape key

The escape key now means "continue" recording and not "stop" recording.

Fill in description for missing WAV file codecs

The uLaw and A-Law codec descriptions have been added. Others can also be added if necessary.

Automatic Host Transfer failing

The "do you wish to transfer host status" message that appears when quitting from the host should now word correctly.

Copy/paste or RTF conversion doing some unicode incorrectly

All text movement from Eclipse to other software should handle this properly now.

ASIO recording devices not properly resuming recording after pause

The pause/resume is now universal.

ASIO Skipping speech in recordings

Short duplicated sections of audio are no longer a problem when recording with the ASIO multichannel interface specification.

ASIO Buffer underflows

ASIO multi-channel recording was running into multi-threading issues which have been solved.

Redacted audio not playing with correct offset and ms setting

When the multi-channel timecode offset and ms/sec options were created, no allowance was made for the fact that the jobname_redacted.wav file looks like a whole different channel. It now uses the same offset/ms as channel 0.

GSM 6.10 44.100khz producing static

Correction made for high-resolution GSM files

"Ignore Written numbers" nullifying {DECIMAL}

The ignore feature should only apply to "point" not to {DECIMAL}

Seating table did not get Channel 0 update

Seating table will correctly number microphones. However, it will not add a microphone for channel 0, assuming it's the mix of the other individual microphones. All multi-channel recording should be using a mix for the sake of compatibility and easy of processing. It's only an "optional" feature because of the fact that some multi-channel recordings (such as voicewriters and room audio) are not intended to be a mix of participant recordings.

Formfields and Autoreplacements

Auto-replacement input was spanning form fields.

Crash on Matching Russian entries in Phonetics Table

Fixed.

New Page Number Format Character Wrapping

The new page number format print command can now word-wrap much like index lines, comment lines, etc.

Soft comma followed by real comma doubles

A real comma should have permission to delete a soft comma in the metadictionary.

Don't allow empty text globals in dictionary

The dictionary will no longer process empty text globals.

Inconsistent mark->stitch behavior

Note that the options are force lower, leave alone, force upper. The center option does not force mixed case, it leaves the text as-is.

Global Window - Soft punctuation Special Entries Hotkey

The hotkeys for soft punctuation were incorrectly coded and should now work correctly.

Translation Magic history cutting off large number text

It should be able to handle three-digit numbers.

SRT/SCC Exporting w/o Captioning Outputs

These export types require a captioning output for reference about what the exported file should look like.

CART Window "Spacing %"

The spacing percentage was not working for the new CART window that allowed edits to be applied.

Automagic Suggestion - "Ecl Unknown"

Some automagic suggestions were causing errors when selected.

Default Footer forcing double spaces between version

Moving between versions was adding and/or removing carriage-return linefeed pairs due to inconsistencies between versions of Windows and historical versions of Eclipse. Duplicate blank lines are now suppressed.

Alternate PDF output location disables post-generation automatic preview

The PDF preview now changes to the appropriate directory before attempting to open the file from Windows.

"VERIFY 14" proofreading error on "i3" not ignoring all when hitting <esc>

The proofreading ignore was disabled for two-character errors, but has been dialed back to only be disabled for one-character errors.

"Archive Suggested Briefs"

A jobname_AB.dix file should not be created if the user has the brief suggestions turned off.

Regex Find/Replace

Errors in constructing regular expressions will now explicitly report the syntax errors rather

than crashing.

HYP T no longer works if it stops on a non resolved immediately following labels

Re-scanning will now advance the cursor correctly.

Full Screen Mode in Windows 10

Full screen mode was doing some incorrect math when storing/retrieving the original screen size and position.

File Manager: Double-Clicking Columns to Sort files opens First Item

It will now just sort.

Formfields at the beginning of Page 1

Form field to start a document can now be filled in.

Stentura Protege Translate window Text glitch

Protege description had bad unicode characters for accented letters.

Duplicate bad file locations can cause a lot of folders to be made

This would happen if the folder was a relative location syntax rather than an absolute one. For example, a location name of archive/2019 would mean to put the file in a 2019 folder, inside an archive folder, inside whatever the CURRENT folder was! Ack! So if you were in a block file when you needed that location, it would create a blocks/archive/2019 folder, and if you were on your desktop, you would get a desktop/archive/2019...

This version now requires all file locations to be absolute, which means they must start with a drive letter OR a {PARENT} designation of some kind, which all have to point back to an absolute path eventually ({JOB}blocks points back to JOB={MAIN}Name which points back to MAIN={DOC}Eclipse which points back to your Documents folder, which is an absolute location.)

Automagic spelling suggestions and cursor position

If the cursor is in the middle of a word, the spelling suggestion now replaces the whole word instead of just from the cursor position forward.

Last pgh of RTF export shows as single spaced while the rest don't when viewed in Wordpad.

An extra \par tag was necessary.

Tools->CM->Send Message

Empty messages are now ignored to prevent resting hands key repeats (or cat feet) spamming the chat.

Initial/Final block name display anomaly if very long

These text indicators now use the ellipses for text that's too long.

Inconsistent spellcheck underline

Spellcheck underscores now have consistent colors.

Zoom in CM chat window should change text size

The chat text now follows the zoom window size.

Pasting text containing vertical tabs causes access violation

Vertical tabs now filtered out of clipboard.

Spacing is not being inserted after periods if `{.}{N}` is reglobaled as `{.}`

Spacing adjusted.

Remove Watson profanity filter

People swear, and it needs to be in the record.

Dictionary block/separate leaves the cursor too low

Cursor position is now correctly recalculated.

Auto Magic popup box can hide the cursor if it hits the top of your screen

The automagic box now jumps below the cursor if it has no other choice.

di and mi words triggering number conversion suggestions in infobar

Roman numerals. Number suggestions will no longer appear if the letters in question spell out a correctly spelled English word. Scrabble players will note some lost crossovers such as "xi" but that's a small loss to avoid words like "did" coming up with number replacement suggestions.

Dictionary regex find/replace producing odd results

Regex processing repaired.

Spellcheck and "Change all" with hyphenated words

Oh, you hyphens. Are you part of a word or not? Get in line with the apostrophe and we'll talk.

Adjustment made. This will now work fine.

Toolbars in 1080p and DPI Scaling

Toolbar buttons are now scaled with a smoother algorithm that should make them look less

pixelated.

The things I thought I'd never have to go back to when I stopped working on videogame programming...

Auto-indexing Preview Dialog

The preview dialog could turn to non-modal in strange situations and allow editing while it was still open, which is bad. It will now behave correctly.

Stuck Illuminating "Pause" button while recording

Pause button should now behave.

CART Window Refactor and label spacing

Label spacing should be consistent with paragraph and document settings.

Dictionary PDF creation - DPI Scaling

The PDF printout of dictionaries should not be larger just because users set their font size larger on the display for a high DPI monitor. It will now scale proportionally.

CART Window Refactor and Long Dictionary Entries

Long dictionary entries will now correctly arbitrarily break at the CART window margin despite no logical word-wrap position such as a space.

Vox-specific: Local Replacement Redundancy

Performing vox local replacements on words that are in the dictionary as plain words could cause spacing to be incorrect.

Dictionary "Find/Replace All" when sorting by Text

Since changing the text changed the indexing, this could get caught in some bad situations and crash. Should work fine now.

"Encountered improper Argument" - Integral Pre/Suffix Setup

Integral prefix suffix modify dialog can now deal with improperly formatted entries without crashing.

Number being translated into Digits= format if followed by a decimal + more numbers

A decimal point should imply a quantity.

CART window Output Suspend/Resume issues

Suspending and resuming the new CART window was creating a whole new window.

Wav form display tied to "Translation stats"

It was not updating frequently enough, changing only every few seconds. It will now update continuously as it should.

Job Report - Access Violation

Closing a job should close the report dialog automatically so it doesn't crash attempting to access a job that's no longer there.

FM File icons disappear if DIX= exists and you double click a file to open it

File icons should stay stable now.

"Auto-rename Jobs to:" Datetime Option

The datetime option should never save as the actual filename template.

Clipboard path bug

Copy/paste from an empty paragraph no longer shows the clipboard path in the visible text.

Scrolling in File Manager

Scrolling is inconsistent and follows the highlight even when not necessary.

Speaker table entries only saving if master is not empty

Speaker table entries should save in the job dictionary under all circumstances in order to populate future jobs, if desired.

"Read along" Spellcheck feature Disables "Pause" functionality if Enabled

Pause functionality now works again.

Multi-Selecting and Updating blocks in File Manager

Opening block files in the file manager should open the correct files now.

Reverse scan to proofreader highlights not working

Technically an enhancement. Now works fine.

CART Window if always on top is unchecked the CART Window will slide behind Eclipse when dragged around`

CART Window now stays in front when dragging.

Briefs with * not displaying phonetically if phonetic briefs are selected

The asterisk will now show phonetically in briefs regardless of phonetic setting that would normally clear it (in order to resolve ambiguity between strokes with and without the asterisk.)

12-16-2019

Dictionary print/PDF needs a progress display

Large dictionaries could look like the software was locked up.

Passport shadows missing from realtime

The read-in notes were working, but not the realtime input.

Single-digit confidence in note files display anomaly

Fully confident strokes were showing as 0 instead of 00 and actually should have been showing as 99.

Translator not releasing jobname_ASR after translation

The note file would remain inaccessible until Eclipse was entirely closed.

ASR improvements

Dozens of changes were made to reduce false positives and change confidence grading in order to achieve the most reliable result possible. This includes the addition of an "ASR Exceptions" category in the programming tab to add words that should always be considered extremely unlikely (very low confidence) if the ASR hears them.

1-9-2020

Waveform Display: Change *MUTE* to *IGNORE*

This is more indicative of what's happening, since the channel is still being recorded, just not transcribed.

Steno conflicts should keep entry confidence if selections have same text

This has the most bearing on the ASR assistant when it's comparing steno confidence vs. ASR confidence. A conflict between "\Todd\to do" should be downgraded so that the ASR can jump in and fix it if it heard something clearly different from what the AI selected. However, an \exhibit\Exhibit{#N} conflict or an \on the job\on-the-job conflict should not have its confidence downgraded because the ASR won't hear a difference and shouldn't offer a second opinion.

Multi-channel Recording Devices and Opus - Relative Timecode issues

Relative timecodes were not coming up correctly in Opus file when recorded using a

stereo device with channel splitting.

ASIO Recording drivers + OPUS causes C++ Runtime error and MMPE

Opus now works with ASIO devices.

Opus attribute Description currently displaying PCM 48khz

Opus files will show the correct description in the infobar now.

"Contacting License Server" Loop

The timeout for contacting the server has been decreased to 7 seconds.

Microphones in Seating Chart - "Auto" in Playback control panel

The "auto" function in the audio playback feature was not taking the mix channel handling and the new channel numbering scheme into account correctly.

"Current" Speakers - Waveform Display

The speaker names in the waveform display were always being taken from the user settings instead of the current job.

Left Notebar Border extended over scroll-bar

Notebar drawing will now stay inside the border.

ASR assistant should not look at print commands

Text in print commands is now not examined or replaced.

ASR Assistance: Opus support

The ASR assistant can now process .opus files in realtime. It was previously only working with pre-recorded files.

ASR assistant should treat words with ^ed ending the same as without.

There are too many phrases such as "stake down" and "staked down" which are pronounced identically (or with extremely subtle differences) making it a frequent ASR error. Words with and without an "ed" ending will now be considered "equivalent" for the purposes of ASR comparison. The ASR assistant will no longer jump in with a change if it only contains this difference. It will still correctly add or replace words with and without "ed" if they are missing or incorrect in the original steno.