

ReadScoreLib Optical Music Recognition Library

(documentation version 4.00)

ReadScoreLib from Dolphin Computing is a comprehensive OMR system for converting printed/engraved music notation to MIDI and MusicXML. The library and output is fully compatible with Dolphin's SeeScore SDK which provides rendering and MusicXML query capabilities.

The performance of the library varies according to platform but on Windows a page of music takes about 3 seconds to translate. ReadScoreLib can cope with photographically distorted images and can process pages photographed with a device camera. ReadScoreLib can compensate for the missing rests and tuplet marks which are common in printed music, and can deal with variable size systems, missing time signatures and transposing instruments. The MIDI files generated by ReadScoreLib play naturally and continuously and without the jerks and jumps sometimes seen in OMR output. ReadScoreLib is also the only OMR system to support the varied tremolo notations often found in printed music.

Platform APIs

This documentation describes the C language interface to ReadScoreLib. The ReadScoreLib SDKs for iOS/macOS and Android provide APIs for Swift, and Java respectively. These APIs are equivalent to the C API described here. Their use will be apparent from the documentation below.

Recognition from a memory resident bitmap image

rscore_convert accepts a page of music in memory and generates corresponding MIDI and MusicXML files subject to selected options

<code>rscore*</code>	<code>rscore_convert(</code>		<code>read an image (32bpp RGBA) and convert to MIDI and/or XML files</code>
	<code>const unsigned*</code>	<code>data</code>	array of 32bit pixels in (32bpp RGBA) format with the alpha channel (bits 24 - 31) should be zero
	<code>int</code>	<code>imageWidth</code>	the pixel width of the image
	<code>int</code>	<code>imageHeight</code>	The pixel height of the image
	<code>int</code>	<code>rowbytes</code>	the byte offset from one row to the next and must be a multiple of 4
	<code>enum rscore_imageorientation</code>	<code>orientation</code>	the orientation of the image. See <code>rscore_imageorientation</code>
	<code>rscore_progress_callback</code>	<code>cb</code>	user-supplied callback called

		periodically to report progress and allow the process to be abandoned pass NULL if no callback desired
Void*	arg	the context argument to be passed to cb
const char*	midifilepath	if non-null the MIDI file is written to the full path midifilepath
const char*	xmlfilepath	if non-null the MusicXML file is written to the full path xmlfilepath
const rscore_options*	options	Information including bits specifying selected options (see rscore_optionsflags)
Unsigned*	Image_out	for use only when rscore_optLens selected: rectilinearised image written to image_out. Set NULL otherwise
rscore_errorinfo*	err	if non-null points to a struct which receives any error information
Return		handle to rscore instance. handle should be deleted on completion with rscore_delete

Image dimensions

The input image is laid out as a series of concatenated rows each pixel of which occupies one 32-bit word. *image* points to the top left word. The layout is RGB with the remaining (ms) byte set to zero.

The width and height arguments are in pixel units. The *rowbytes* argument allows a section of a larger image to be processed. This is done by setting *image* to the top left word of the subimage, width and height to those of the *subimage* and *rowbytes* the width of the full image.

Music segments

ReadScoreLib can recognise multiple areas of musical notation as for example might be found on a page of text containing musical examples. The positioning and size of these areas may be varied but will be read as a vertical sequence. Multiple columns of music are not currently supported.

Where there are multiple areas, or where a page containing multiple short pieces, or where one piece ends and another begins, the output MIDI and MusicXML will be continuous (as anything else would alter the music). However, by using *rscore_getbarinfo* and in particular the *rscore_firstOfSection* flag it is possible for an application to identify boundaries. And by using coordinate information (see *rscore_optXmlDefaultCoords* and *optXmlObjectBounds*) a rich graphical application is possible.

Preprocess only mode (see *rscore_optLens*)

For some applications access may be needed to the preprocessed image (see table below). This can be useful where images are distorted or poorly defined. The preprocessed image in black and white with distortions removed is often easier for a human to read and may be more amenable to further graphical processing.

When the *rscore_optLens* *ptions* flag is set only preprocessing takes place with the resulting image passed back through *the image_out* argument. The user supplied memory bitmap specified by *image_out* has exactly the same format as image except that *rowbytes* will be ignored and the image will be written to a contiguous block of memory of size $(width * height) * 4$.

NB The ReadScoreLib *rscore_optXmlObjectBounds* feature provides bounding box information for musical objects relative to the original input image; that is before preprocessing. If an application requires bounding box information relative to the preprocessed image it can be obtained by first preprocessing the image using the *rscore_optXmlObjectBounds* option and then submitting the preprocessed image for normal recognition without the flag.

NB Lens processing can only be performed using *rscore_convert*, one image at a time.

Progress callback

The *cb* argument allows the caller to specify a function to be called periodically to report progress as a percentage. *cb* should conform to the prototype

```
typedef bool (*rscore_progress_callback)(const rscore_convert_progressinfo *info, void *arg);
```

for example

```
bool cb(const rscore_convert_progressinfo *info, void *arg) {
    printf("%d% complete ", info->progress_percent);
    return true;
}
```

cb will be called by ReadScoreLib with the *progress_percent* member of the passed in *rscore_convert_progressinfo* struct giving the percentage of the job completed.

cb should normally return true. A false return will cause the job to be abandoned and control from *rscore_convert* to return.

Recognition from multiple files

rscore_bconvert set of functions allow multiple memory-resident page images to be built into single MIDI and MusicXML files subject to selected options

rscore*	rscore_bconvert_begin(read one or more page image files and output corresponding MIDI and MusicXML files
	rscore_progress_callback	cb	user-supplied callback called periodically during the build process to report progress and allow the process to be abandoned pass NULL if no callback desired
	void*	arg	the context argument to be passed to cb
	rscore_options*	options	Information including bits specifying selected options (see <i>rscore_optionsflags</i>)
Return	rscore_errorinfo		handle to rscore instance

rscore*	rscore_bconvert(read one or more page image files and output corresponding MIDI and MusicXML files
	rscore*	rsc	The rscore instance handle returned from <i>rscore_bconvert_begin</i>
	const unsigned*	data	array of 32bit pixels in (32bpp RGBA) format with the alpha channel (bits 24 - 31) should be zero
	int	imageWidth	the pixel width of the image
	int	imageHeight	The pixel height of the image
	int	rowbytes	the byte offset from one row to the next and must be a multiple of 4
	enum rscore_imageorientation	orientation	the orientation of the image. See <i>rscore_imageorientation</i>
Return	rscore_errorinfo		error information

rscore*	rscore_bconvert_end(read one or more page image files and output corresponding MIDI and MusicXML files
	rscore*	rsc	The rscore instance handle returned from <i>rscore_bconvert_begin</i>
	const char*	boundsFilePath	If non-NULL the full path to which object bound information will be written (see <i>rscore_optXmlObjectBounds</i>)
	const char*	midiFilePath	if non-null the full path to which the MIDI file will be written
	const char*	xmlFilePath	if non-null the full path to which the MusicXML file will be written
Return	rscore_errorinfo		error information

The three `rscore_bconvert` functions allow multi-page music to be built into MIDI and MusicXML files. These functions can also be used to create a text file specifying the coordinates of objects within the images.

When a series of images is built using the `rscore_bconvert` functions they are built as continuous music rather than as a series of pages. As new files are submitted with `rscore_bconvert` a parse tree is built for each page. When `rscore_bconvert_end` is called the sequence is converted to a linear relational format. This representation is continuous over pages so that cross-page features such as part bars, slurs and ties can be translated faithfully. At this point also, statistical information is extracted and applied globally to improve general recognition accuracy. Finally MusxicXML is generated from the internal relational representation.

Since any sequence of pages submitted between `rscore_bconvert_begin` and `rscore_bconvert_end` will be built into a continuous MusicXML stream, it is best to break the music up into individual movements and to submit these separately. For each sequence, `rscore_bconvert_begin` is called first. This establishes any progress/cancellation callback and returns a handle to the ReadScoreLib for the present session.

The `rscore` handle is then used to submit pages of music in sequence via `rscore_bconvert`. This function takes the same image parameters as the single-page `rscore_convert` API described above. See description above for details. The `rscore_error` return may be checked for status information relating to the image.

When all pages have been submitted through `rscore_bconvert`, `rscore_bconvert_end` should be called to complete the build. This function takes the paths to desired MIDI, MusicXML and bounding-box-ID file.

Progress callback

ReadScoreLib processes music in two stages. Rather than treat each page as a separate unit, `rscore_bconvert` creates a syntax tree for the whole input, page by page as described above.

Because of this organisation the progress callback functions for `rscore_convert` and `rscore_bconvert` work differently. Whereas `rscore_convert` calls the progress function in a single sequence with `progress_percent` advancing from 0 to 100, `rscore_bconvert` does this for each file and then one additional time when `rscore_bconvert_end` is called for the final stage of processing. `progress_percent` takes the value zero exactly once for each file allowing the transition from one to the next to be detected. When every page has been processed `progress_percent` once again begins at zero, reaching a maximum close to 100 when the whole task is complete. If the callback should return false at any point the whole build is abandoned.

This arrangement allows the client application to implement a progress indicator for the whole process, incrementing only on the zero transitions, or perhaps a few times during each `rscore_convert` call. Since the number of pages will be known the progress indicator can be arranged to advance smoothly from 0 to 100. If desired a per-page progress indicator is also possible.

General information

Input images

ReadScoreLib accepts a wide range of input quality, resolution, exposure and geometric distortion. For good results the scale of the submitted bitmaps should be such that the distance in pixels between individual staff lines is at least 10. For best results this distance should be between 15 and 24. Larger scales are possible but performance deteriorates as the square of linear distance. If the image is clear a metric of 16 will often give good results at two or three seconds per page on an iPhone 6.

The table below gives suitable ranges for image capture from music using different capture methods. For best results for a particular type of music, experiment with any parameters under your control such as light and scale.

Device	Scale	Comments
Scanner	Scan at 280 – 310DPI	
PDF to JPG converters	Set converter to 300 – 500DPI	It is worth trying a wide range of conversion DPI, especially where the image quality appears poor
Photography	Aim for a staff metric between 14 and 18	Take image in good light and aim for a minimum of spherical and other photographic distortion.
From a book		Ensure that the page lies as flat as possible. The left side especially, where the system line and the signatures are should be clear
PDF and other images created by score writers such as Sibelius, Finale, Dorico etc	Convert from PDF at 150 – 300DPI	Ensure that the image is dark enough to show features such as staff lines and stems clearly. These images can sometimes be bright and have thin lines. This can cause poor results

The image may be in colour, greyscale or monochrome. Greyscale is usually better as ReadScoreLib can choose its own threshold at different areas of the image. See below for more information on the preprocessing transformations applied

Image preprocessing

Before any music recognition takes place ReadScoreLib prepares each image as follows:

Preprocessing transformations	
Thresholding	Image is converted from colour, greyscale or monochrome to binary. The techniques used attempt to allow for variations in colour and shading
clipping	Material surrounding one or more areas of musical (for example in a device photograph) notation is stripped away
Linearisation	The image squared up: Distortions such as photographic distortions are reversed leaving staff lines straight and level
Segmented rotation	In some circumstances the image is broken into sections that are independently deskewed and dovetailed

The preprocessed image can be useful for some applications and is available using the *rscore_optLens* option.

When enabled the *rscore_optDeskewMonochrome* option causes monochrome pages to be subjected to segmented rotation instead of thresholding and linearisation

Image scale

For good accuracy ReadScoreLib requires an image where the vertical distance between staff lines is at least 9 pixels. The ideal range is 15 - 25

Operation abort

The calling process can cancel processing by returning false from the progress callback

Options

The options argument allows the caller to specify certain processing options as well as the application name and version as required for the XML identification element. Music parsing and MusicXML generation options are selected through bits set in the *rscore_options::flags* word (see below)

Certain flags such as *rscore_optParseTremolo* can be used to reduce the risk of recognition false positives by suppressing parsing for particular constructs. Normally these flags should be set unless the music is known to be free of a particular construct.

Option flags	
<i>rscore_optExcludeUpperStaffs</i>	Ignore all staves above the lower two, taken system by system
<i>rscore_optParseTremolo</i>	recognise tremolo notations
<i>rscore_optXmlDefaultCoords</i>	Generate MusicXML Default x, y attributes
<i>rscore_optXmlObjectBounds</i>	Generate a text file listing the bounds of objects. The text file has the same name as the XML file but with the TXT extension
<i>rscore_optTransposingInstruments</i>	Guess transposing instrument transpositions where possible from key signature
<i>rscore_optFrenchTimeSignatures</i>	Recognise French time signatures (those engraved in parentheses above the barline)

Tremolo

ReadScoreLib supports the following tremolo notations. Between one and four strokes are recognised.



Transposing instruments

Normally ReadScoreLib assumes all instruments untransposing and makes use of all key signature information in establishing the key. If the `rscore_optTransposingInstruments` flag is set this assumption is no longer made and transposing instruments are where possible identified from their key signatures. In these cases the appropriate transpositions are written into the MusicXML file and the correct key signature specified. However note that not all transposing instruments can be recognised in this way. These are of two types:

- 1) Octave transpositions such as the piccolo and the double bass.
- 2) Instruments such as horn and trumpet. These are normally written without key signature and cannot be deduced.

ReadScoreLib will shortly be providing support for OCR text (see below). This will include instrument names, thereby allowing more comprehensive support for transposing instruments.

Establishing the location of objects

ReadScoreLib provides several features that allow an application to obtain the identity and location of recognised score objects. The requirements of such a system go beyond the scope of MusicXML. MusicXML does support a *Default x/y* attribute type but these are unsuitable because they apply to some objects only and supply only a single pair of coordinates, rather

than full bounding box information. Moreover *Default x/y* is a hint facility and if used rigidly to specify the coordinates of objects on the original page unpredictable behaviour is likely when the music is rendered. ReadScoreLib can generate *Default x/y* but the XML produced is intended for use as a key into the text file (*rscore_optXmlObjectBounds*) and not as an aid to rendering.

As MusicXML cannot itself contain the information required to locate and bound objects in the original image, ReadScoreLib writes this information to a separate file. The pathname to the file is specified in the *rscore_bconvert_end* API. In addition, the same information is added as comments to the MusicXML file.

The present version (3.1) of MusicXML lacks a comprehensive system of unique identifiers for objects. ReadScoreLib therefore provides an alternative mechanism based on the *Default x/y* attribute, used as a key into the text file. Alternatively if XML comments can be read by the XML client, bounding box information can be accessed directly.

When *rscore_optXmlObjectBounds* is set ReadScoreLib generates bounding box information in both forms, as comments in the XML file and as a separate text file. If the XML client cannot read these comments in the course of reading the object itself, the *rscore_optXmlDefaultCoord* flag can be used to select *Default x/y* attribute generation. The application software can then use this as a key into the corresponding bounding box information in the text file. As it reads an object (eg a rest), the MusicXML client (the XML parsing software) will read the *Default x/y* attribute. The coordinate information can then be used to look up the bounding box information in the text file. To make key duplication unlikely *Default x/y* attribute is given a five digit floating point precision. The last two digits after the decimal point are randomised slightly to provide a unique key without affecting the positioning of objects.

The bounding-box-ID text file has a simple format consisting of one line of text for each object.

<barnumber>_<ref-x>_<ref-y>: <object-name> (<pagenumber>, <bounds>)

For example the following represents a clef found on page 5 (numbered from zero) enclosed within a box whose bottom-left and top-right corners lie at (584,2735) and (612,2797)

2_51.001_406:clef (5,584,2735,612,2797)

Some objects specify coordinates adapted to their positioning as musical objects. A note head for example specifies its centre only, as other dimensions can be deduced from the staff gauge.

<ref-x> and <ref-y> are not intended to specify a location in space, but to provide a unique link between MusicXML and The bounding-box-ID text file as described above.

Barline information

For applications that require only barline information (eg for score following) ReadScoreLib can be queried directly for the position and height of every barline in the score without recourse to the MusicXML or the text file. The *getbarcount* and *rscore_barinfo* APIs are used to obtain the positions and heights of barlines in the coordinates of the original page. If ReadScoreLib has transformed the geometry of the page (see *rscore_optLens*), then the transformations are reversed before being passed back.

int	<i>rscore_getbarcount</i>	Return the total number of bars in the input
	<i>rscore*</i>	handle to <i>rscore</i> instance

<i>rscore_barinfo</i>	<i>rscore_getbarinfo</i> (
	<i>rscore*</i>	<i>rsc</i>	handle to <i>rscore</i> instance
	int	<i>barindex</i>	the index of the bar [0.. <i>rscore_getbarcount</i> -1]

struct	<i>rscore_barinfo</i>	Barline information	
	<i>rscore_barline</i>	<i>opening_barline</i>	position of opening barline
	<i>rscore_barline</i>	<i>closing_barline</i>	position of closing barline
	int	<i>startbeat</i>	the starting beat number of this bar
	int	<i>beatcount</i>	the number of beats in this bar
	Unsigned	<i>flags</i>	see <i>rscore_barflags</i>
	int	<i>first_beat_offset</i>	**percentage** of bar width to the left of the first note or rest

struct	<i>rscore_barline</i>	Barline dimensions	
	<i>rscore_point</i>	<i>base</i>	the barline lower point
	int	<i>height</i>	The height in pixels

enum	<i>rscore_barflags</i>	Barline flags	
	<i>rscore_firstOfSystem</i>	marks the leftmost bar in the system	
	<i>rscore_eoscore</i>	marks the last bar in the score	
	<i>rscore_splitbarL</i>	box encloses the part of the bar at the end of one system	
	<i>rscore_splitbarR</i>	box encloses the part of the bar at the beginning of the next system	
	<i>rscore_firstOfSection</i>	marks the first bar of a section, for example where one movement ends and another begins on the same page	

Errors

Error status is normally reported through a passed-in *rscore_errorinfo* struct giving details of the error (see *rscore.h* for details).

Version information can be retrieved through the *rscore_version* struct.

Supported platforms

ReadScoreLib libraries are available for Windows, Mac, Android and iOS.

New Features

Two new features are expected to become available during 2018

- 1) support for text such as lyrics, instrument names, metronome marks and directions
- 2) A mask word input to convert functions specifying which staves from the image(s) are built into the MIDI file. This allows any combination of staves to be played.