

1
2
3
4

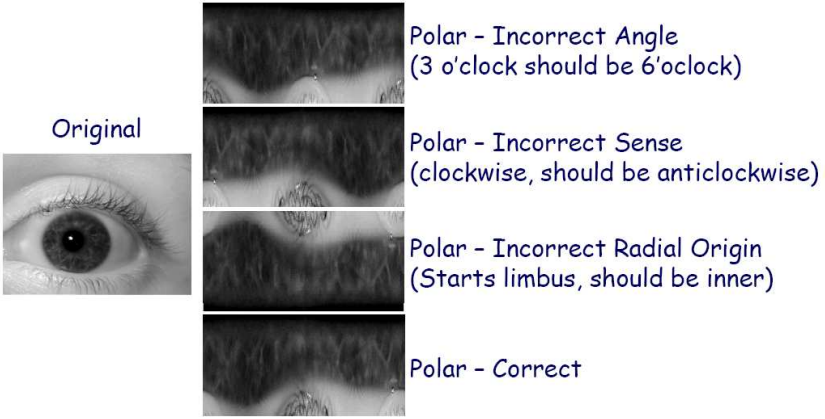
IREX Frequently Asked Questions

This FAQ was last updated on September 15, 2008.
The IREX Test Plan was last updated on September 15, 2008.

Date	Question	Answer
6/13/08	It is stated in section 5.14('phased testing') that "NIST will ... conduct this test in just a single phase. ... the results of testing will be published sometime after the implementations are received, without interim disclosure of results to supplier." Meanwhile, section A7.1 ('Reports') talks about 'phase I results' and 'phase II' results and how they will be handled differently.	Talk of "Phase I and Phase II" is an error. IREX 08 only has time for Phase I before the public report is due per ISO deadlines. The Test Plan was modified to fix this error June 20, 2008.
6/13/08	Can a company withdraw after July 2 but before SDK submission deadline September 5.	Yes, but NIST wants to give ISO SC 37 accurate information on expected participation. In any case we do not intend to report the names of companies that withdraw. The Test Plan was clarified June 20, 2008.
6/21/08	What about in-plane rotation of the eye? Should our software rotate the image when it prepares the Table 6 record?	No. The Table 6 record includes a field for "Rotation" which inherits the semantics of ISO/IEC 19794-6:2005. This field should be set with a rotation estimate. Thereafter a feature extraction or matching algorithm will want to use this value. See June 25 clarification text in 7.2.4.
6/23/08	The function <code>convert_raster_to_cropped_rectilinear(...)</code> provides clear data on size of the input image, but not on the output image. Is the size of the output image assumed to be 320x320? Or am I missing something?	Your software would estimate the diameter of the iris and cut out a square or rectangular region about M pixels bigger in each direction - see the bullets in 6.3 of the IREX Test Plan. The output image size does NOT need to be 320 x 320. For example, for an iris of diameter 180 pixels, the "M = max(60,0.4R)" suggestion implies a width of 180 + 2 (left and right sides) * max(60, 0.4*90) = 300.
6/23/08	Memory for the various output images is allocated by the NIST client of but our function does not have info on its dimension. How is this to be handled?	NIST will allocate more than enough memory for your output image (for example, the size of record header plus the size of the original uncropped input, which would be a maximum size.
6/23/08	NIST says it will survey over the cropping size, but the various function calls to not take crop size input parameters. So how will NIST do this?	NIST will re-crop the images at varying sizes, from large to small. We have added two output parameters to the API: the top left coordinates (x, y) of the bounding box in the parent image. The Test Plan was clarified June 25, 2008.
6/25/08	NIST says it will allocate memory for the proprietary template, but how much.	NIST will allocate 8KB for this buffer. If you'll need a larger amount email us in advance. The Test Plan was clarified June 25, 2008.
7/27/08	How many organizations indicated their intent to participate?	Thirteen.
7/27/08	In the "IREX Test Plan and API (version 5)", new "bbox_*" parameters were introduced in function " <code>convert_raster_to_cropped_rectilinear()</code> ".	Yes, intentionally NIST did not add the <code>bbox_*</code> output parameters to the <code>crop-and-mask</code>

	However, they were not added to the "convert_raster_to_cropped_and_masked_rectilinear()" function. Were they intentionally skipped?	function. Reason: There's no need. NIST will retain the bounding box coordinates computed in calls to the function "convert_raster_to_cropped_rectilinear" and use these coordinates for re-cropping of the masked image.
7/27/08	If we send multiple files as part of an SDK, how should these be packaged?	"zip" or "tar.gz" or "tar.bz2" archives are acceptable. These MUST be encrypted. Unencrypted SDKs will be deleted immediately.
7/27/08	You say (in section 5.10) "Organizations may enter two SDKs per class". Does it mean that maximum 6 SDKs can be submitted from each organization?	Yes, but NIST requires you to nominate a primary and secondary SDKs. ISO/IEC deadlines for contribution of content toward the revised ISO/IEC 19794-6 standard are in early December 2008 and this may require NIST to prioritize testing: This means SDKs labeled "primary" will be tested first; those labeled secondary will be tested later. The naming convention is as follows. No spaces in filenames: Please use underscores "_" instead. The class "X" or "Y" or "Z" should appear in the filename If two SDKs are submitted the these should be labeled using something like this: "companyName_primary_X.zip" and "companyName_secondary_X.zip"
7/27/08	If SDKs are submitted for X, Y and Z classes from the same organization (all of them are able to generate cropped rectilinear image but perhaps in slightly different manner), will this ability be tested for each SDK separately? I mean, that would give the maximum 6 different implementations for the same task.	While NIST does not encourage needless/minor variants, vendors are free to do this. But please note that the SDKs labeled "primary" will be tested before those labeled "secondary" issue, and this is important vis-a-vis ISO/IEC JTC 1 SC37 deadlines.
7/27/08	The IREX record allows implementers of the unsegmented polar and cropped-ROI-masked formats to supplement the mandatory fields with either or both of the boundary encodings ..": Q1 Will the same boundary fields be available for the compressed images?	An SDK may chose to compute and store the boundary in the output record. The input data to this computation will be the original uncompressed image. Later, NIST will apply compression to the image data. NIST will never adjust the boundary data, which therefore will be available to all feature extractors receiving the record. This mimics the operational situation where some segmentation is done at the point of capture before compression and transmission. This is not the only possible operational use-case of course.
	Q2 Do only "unsegmented polar" and "cropped-ROI-masked" formats can be supplemented with boundaries data (but "cropped" format cannot)?	All of the formats may be generated with the optional boundary contours. The crop-only format is a simple format that will serve as a baseline with against which we can compare performance of the two "major" formats. If boundary information is included NIST will evaluate matching performance with it and without it (by stripping it out).
	Q3 Are boundaries recorded in respect to the original image or the image saved in the record?	The boundaries are assigned with respect to the image stored in the record NOT the parent image. Reason: So that a receiver can process the record without need for external information.
	Q4 How can boundaries be stored for the unsegmented polar format?	Recall that when an unsegmented polar image is stored, the following is also present:

		<ul style="list-style-type: none"> - the width and height of the parent image (lines 16-17 of Table 6) - the (x,y) coordinates of the circle center (lines 35-36 of Table 6) with respect to the parent image <p>These allow reconstruction of the iris region vis reverse polar transformation.</p> <p>The Table 6 record can optionally include "best fitting ellipse" and "Freeman chain" boundaries for the "unsegmented polar" format. These should be stored with respect to the parent image.</p>
7/27/08	It is said (in section 7.2.7) "The SDK does not have to support Kind = 16 because NIST will execute any needed reverse polar transforms using NIST code to make Kind = 48 instances." But section 5.9 states that classes Y and Z should be able "to extract features from an [IRI07] polar image and match." Which one is correct?	<p>NIST considers the actual implementation of reverse-polar transformation to of secondary importance to whether an SDK can extract features from the rectilinear form produced by reverse-polar transformation.</p> <p>So NIST will release code to execute the forward and reverse-polar transformations. Implementers can then a) use the NIST code, b) use their own code, or c) provide an SDK that accepts only Kind=48 instances (and not Kind=16 instances). So section 5.9 is imprecise in its wording - it should say "to extract features from a Kind-48 [IRI07] reconstructed polar image and optionally a Kind=16 [IRI07] polar image, and match."</p>
7/28/08	You have specified the "bbox_*" parameters in function "convert_raster_to_cropped_rectilinear()" to be unsigned integers. However, the top left corner of the cropped centered iris image can be outside the image (thus, negative coordinates, when iris is close to or even slightly outside the top left corner of the original iris image). How to manage these cases?	<p>Yes, logically the top-left bounding box coordinates. The semantics of the bbox is now the rectangular region enclosing the iris.</p> <p>July 28 2008: The API function call in the IREX 08 test plan is changed to so that the bounding box has SIGNED integers.</p>
8/05/08	Can we merge class Y and Z implementations into one library?	Yes, because the API function names are distinct and can co-exist in one library.
8/06/08	Regarding the unsegmented polar format, if one looks at the input parameters in Table 10 of the July 28 IREX document, the sampling rates (num_samples_radially and num_samples_circumferentially) are included as input parameters of the IREX API. So NIST sets the sampling rate without giving IREX participants a chance to decide how best to choose these sampling rates. It might be appropriate to set different sampling rates for different images depending on the characteristics of the images at hand. Can you change the API?	<p>The API remains the same BUT the semantics change as follows:</p> <p>In Table 10, if either or both of the input values for num_samples_radially and num_samples_circumferentially are set to 0, the vendor implementation should choose appropriate values for these parameters internally.</p> <p>Some values here were discussed in [IRI07] and [BATH07].</p> <p>The API document has been updated with a "C++" comment to note this change in semantics.</p>
8/06/08	When will NIST release its promised forward and reverse polar transform code?	NIST intends to release code to support the forward and backward polar interpolation routines during the week of August 11-15, 2008.
8/06/08	Will NIST release binary examples of the proposed ISO/IEC 19794-6 record structure? What "kind" they be.	NIST intends to release one or more examples of each Kind of the Table 6 record in the period Aug 11-22. Note that these are not intended to be representative of the various cameras i.e. the purpose is not for development of image processing or segmentation algorithms. Instead the purpose is for checking that code can read and parse Table 6 records.

		NIST will expose its version-controlled software development tree for reading, writing and validation of the Table 6 records. This will use the PERFORCE server currently used for other NIST biometric projects - see http://biometrics.nist.gov/nigos
8/18/08	Is it correct that the raster images captured with LG 2200, LG 3000, LG 4000, Securimetics PIER are all of VGA type, i.e., 640 x 480 resolution with 8 bit gray scale value?	Yes. These will be the inputs to the API function calls of Tables 8, 9 and 10. But remember that the images records sent to the template generator in Table 11 will be of varying size due to cropping.
8/18/08	Do you anticipate using larger images in IREX 08?	No. But note that future IREX projects (beyond IREX 08) may include large images.
8/18/08	Where does polar encoding start?	The ISO/IEC 19794-6 (2005) standard indicates that the encoding of polar irises begins at the "6 o'clock" position, see clause 6.3.2.7 "Polar conversion". IREX 08 requires implementers to use the ISO/IEC 19794-6:2005 semantics. So the answer is "6'oclock" i.e. 270 degrees. Failing to implement this method is likely to cause problems for interoperability. Note the position of the eye lashes below.
8/28/08	The polar example image in the Table 3 of the API v7 document (and the M1070490 reference it cites) seems to have been encoded with <ul style="list-style-type: none"> – a clockwise manner, and – from outer boundary to inner boundary. Both of which are contrary to ISO/IEC 19794-6:2005. Please clarify.	The API document has been updated so that the polar example in Table 3 has been replaced with the correct bottom image below. The 2005 standard requires: 6 o'clock, counter-clockwise, and inner-to-outer. 
8/28/08	On the API for us to generate an IIR record in which you provide us a pre-allocated buffer, could you also provide us a length of the allocated buffer - just for safety's sake? I use a library which really likes knowing how big the buffer is I can provide it knowing width and height but [prefer not to].	API functions of Tables 8, 9 and 10 have been modified to accept a parameter saying how many bytes NIST allocated for the output record.
8/28/08	In Table 6 of the IREX08 spec, field 21 for "Eye" it states that NIST expects L/R will be known and supplied to the SDK. However the L/R information for the Eye is not passed to the API functions.	API functions of Tables 8, 9 and 10 have been modified to pass /U/L/R information to all three Table 6 "constructor" functions. ISO/IEC 19794-6:2005 semantics are used:

		<p>EYE_UNDEF = 0 (0x00)</p> <p>EYE_RIGHT = 1 (0x01)</p> <p>EYE_LEFT = 2 (0x02)</p>
8/28/08	In addition, Table 6 does not specify any value (such as 0, per ISO 19794-6) to indicate that Eye is unknown. Can you provide some guidance on how to proceed?	Modified table 6 allowed values, to follow the 2005 ISO standard. The structure of the Table 6 record has not changed.
8/28/08	Should the SDK verification function return 1.0 (maximum distance) for the pair of different position eyes (left vs. right eyes comparison)	<p>SDK should return a large distance for L-R comparisons.</p> <p>NIST is not requiring "1.0" to be a large distance. Implementations are free to report on [0,2.718], for example. Our only restriction is that comparison scores have smaller-is-more-genuine semantics.</p>
8/28/08	Regarding the sampling rate in table 10: A new note was added on Aug 6 to allow implementations to decide the best sampling rate. This is indicated when either of the input parameters is set to be 0. However, if both parameters are assigned positive values, the implementation has to use the given sampling rates even if this would affect performance. Is this correct?	Yes, that's the intention. The goal is to find out the best set of parameters, i.e. those that give small size and high accuracy.
8/28/08	Can NIST release some sample images of all three cameras mentioned in 7.2.2?	<p>No, because the images are privacy protected. If any participant has such images and would like to expose these to other participants, please let me know! I'll publicize a link.</p> <p>LG 2200 images were made available under the ICE program, and a number of public databases exist which may include these images.</p>
8/28/08	"Table 5 – Format for Freeman chain code" has on line 3 an incorrect length value: It says "N" but should be some integer number of bytes. What exactly length in bytes is it?	<p>2 bytes. The API document has been clarified text in two places:</p> <ol style="list-style-type: none"> 1. The Freeman chain code block of Table 5 2. Lines 51-54 of the IREX record structure in Table 6.
8/28/08	Is there any update on the availability of additional IREX validation data beyond the single Kind=1 record announced on 8-21-08?	We are working on this.
8/28/08	Also, in the timeline (page 4) in the July 28 version of the IREX test plan, in the row for Aug 18, 2008, it states that the window for submission of SDK's opens and then says "Please send draft SDKs for validation." Is NIST willing to test draft versions of SDKs? We are concerned right now about our limited ability to validate our SDK because of the lack of test data. A NIST validation would give us an extra level of assurance.	NIST will validate the basic function of the SDK by comparing records you produce in your lab with records we produce, here using common input images. Details will be provided in the coming days.

8/28/08	<p>The following type defines are used in the API specification for IREX08:</p> <pre>INT32; BYTE; UINT16; INT16;</pre> <p>Will these be defined in a NIST-supplied .h file or should they be defined in the vendor's .h file?</p>	<p>The ISO 9899 standard, known as C99, established <stdint.h> which is described here. http://en.wikipedia.org/wiki/Stdint.h</p> <p>http://www.ibm.com/developerworks/power/library/pa-ctypes3/index.html states that "not all systems provide all the new C99 types".</p> <p>We recognize that not all compilers will support them and that some people will have to typedef the equivalent types themselves in terms of what their own compiler supports. Something like the following might be appropriate in your ".h" files:</p> <pre>#include <stdint.h> typedef int32_t INT32; typedef uint32_t UINT32; typedef uint8_t BYTE; typedef uint16_t UINT16; typedef int16_t INT16;</pre>
8/28/08	<p>Could you let us know the JPEG2000 compression profile to be used in IREX08?</p>	<p>NIST will survey over appropriate compression parameters. A JPEG 2000 profile is more of an outcome of IREX 08 than an input.</p>
8/28/08	<p>According to section 5.3, NIST will call compression and decompression. Our understanding is that NIST will parse the cropped rectilinear records, ROI-masked records, and unseg polar records generated by vendor SDKs, apply compression and decompression to the image data (unsegmented polar rectangle for kind=16), then rewrite the corresponding records with the decompressed data and pass them to template creation functions. Is this correct?</p> <p>Furthermore, the kind=48 records to be passed to template creation are generated out of kind=16 records to which compression and decompression have be applied for. No compression and decompression would be applied to reconstructed polar image. Is this correct?</p>	<p>Yes</p> <p>Correct, compression will be applied by NIST in the polar domain.</p>
8/28/08	<p>(Table 13 on page 26: " NIST will assign the identifier that will uniquely identify the supplier and the SDK version number.").</p> <p>The window for submission of Participation Application and SDK was supposed to open on august 18.</p> <p>As there wasn't any notification by email and as NIST hasn't given us the SDK IDs for our submissions, we assume it has been delayed. If we are not wrong, can you give us an idea of the new timeline?</p>	<p>SDK IDs will be distributed shortly, via an automated mail.</p> <p>The submission window closes September 12. Given the number of changes in the Aug 28 additions to the FAQ, we will extend the submission window if we receive a request to do so.</p>

9/7/08	In the IREX API, can you modify it to allow asymmetric enrollment and verification templates?	Yes, two options in the modified API section 7.2.7.
9/7/08	We're intending to send more than one SDK. Do they need different IDs?	Yes, you should add sequentially to the IDs. Thus if the primary SDK has value 10090, then the subsequent SDK will have 10091.
9/7/08	We've accessed the "irexv" validation code and it's objecting to the value on line 18 of the Table 16 record. You have a required value for "1" everywhere, but for non-polar images that's surely incorrect?	Yes, my mistake. The API has been fixed so that line 18 of the Table 6 record can now be allow "0" or "1". The value should be set to "1" for "Kind = 16" images only, and 0 otherwise. (The field seems redundant in a future standard) The "irexv" program will be updated accordingly.
9/7/08	You're passing in labels for "left", "right" and "undef" for which eye it is. Which combinations should be attempt matching?	You should attempt matching for all possible combinations, including L-R. As stated in the API NIST expects to have solid ground-truth, and would not ordinarily compare L-R. If a L-R error is suspected in the ground-truth we will inspect the parent images.
9/10/08 Updated 9/11/08	In the "Description" row of the Tables 8, 9, and 10 (i.e. API functions) it says that when the software fails or elects not to produce an output image that "The result will nevertheless be a conformant instance with zero irides and zero images." But this is inconsistent with the required value 1 on Line 22 of Table 6.	Yes, again my mistake. The API document has been updated to: <ol style="list-style-type: none"> 1. Allow "0" on Line 6 of Table 6. 2. Indicate that "0" here means the record ends on line 20 (i.e. lines 21 onwards are not present). Width and height should then be set 0. Our "irexv" validation program will be updated accordingly. Note that a failure to process an image will be regarded as a "failure-to-enroll". Any core-dumps / segmentation violations / fatal exceptions will be reported to you, but at the end of the day these will be counted as FTEs. The proportion of images resulting in FTE will be reported.
9/11/08	Is it OK for you link to link our code using g++ as an alternative to gcc?	Yes. Please remind us on submission.
9/15/08	Where is the irexv code?	This answer repeats, for the record, email sent to you 9/2/08. NIST has released IREX software for reading and validating IREX Table 6 records. The code is under version control and is subject to daily change. There are two parts: A. BiomDI a library for handling standardized biometric data interchange records including an IREX library. B. BiomAPP a repository of applications. Of relevance here is the "irexv" validator. (B) depends on (A). The software is "C", specifically C99. To get the code there are two alternatives: <ol style="list-style-type: none"> 1. Zips are updated nightly - See the BiomDI and BiomAPP entries linked from here: http://biometrics.nist.gov/nigos 2. Subscribe to the Perforce server. Perforce is a software management suite, clients for which are free. They run under windows and linux. You'll see the code as it is updated, and you will receive email updates for each change.

		Follow the instructions here: http://biometrics.nist.gov/nigos Look at the "Main" branches.
9/15/08	The fields 33-38 are only mandatory for kind 16. From our point of view it could be also helpful for kind 48. As only NIST can generate this kind of images, is it possible to fill in these fields for kind 48 also?	The fields are mandatory for both Kind=16 and Kind=48, according to Table 6, last column, which says that the fields 33-38 "shall be populated if Kind=16 or Kind=48". So: <ul style="list-style-type: none"> - if someone makes a Kind = 16 record the fields must be there. - If NIST makes the Kind = 48 record then NIST will copy the fields from the Kind = 16 input record.
9/15/08	What information can we store in our templates?	The template content is entirely unregulated.
9/15/08	When you validate our SDK, what will you check? And by when do we have to submit our SDK for validation? (We are not talking about the final submission.)	See the validate.c program in the validation package. This was released to IREX 08 participants Sep 15, 2008.
9/15/08	Are you going to assign Quality Vendor ID? So far, you have not given one. We are thinking of using a string that we devised. As long as you do not delete it from the header file, we have no problem.	NIST will assign each participant a unique quality algorithm vendor ID line 26. This will be emailed individually. These are taken from the IBIA registrations when present here http://www.ibia.org/cbeff/organizations.php otherwise they're assigned by NIST. The contents of line 25 (the quality value) and line 27 (the algorithm version) are regulated only by the standard. Particularly Q = {[0,100], 254, 255}. If you do not compute a quality value, assign 254 here. We added 254 to Table 6. This necessitates a change to our irexv record validator. NOTE that, without a calibration program, quality values do not seem reliably interoperable.
9/15/08	About cropping, we believe that the cropping rectangle size may vary (i.e., not has to be 320 x 320). But we believe NIST should not allow rescaling. By rescaling, we mean a re-sampling; say, by 2:1 re-sampling (with or without interpolation) it is possible to turn 320 x 320 rectangle to 160 x 160 rectangle. We are proceeding under the assumption that this kind of re-sampling is not done.	If anybody sees value in rescaling contact us soon! Otherwise for IREX it is prohibited to up- or down-sample an image.
9/15/08	Is there any size restriction to SDK for electronic submission? We will have some auxiliary files whose size will be quite big.	No restriction at this time.

1
2