Concept of Operations (CONOPS) for
Evaluation of Latent Fingerprint Technologies (ELFT)

(Rev. D, 21 June 2007)

1. Overview

The *National Institute of Standards and Technology (NIST)* is conducting a series of tests for evaluating the state of the art in automated latent fingerprint matching. The intent of the testing is to quantify the core algorithmic capability of contemporary matchers. The testing will be conducted using software-only implementations, and utilizing NIST hardware.

The umbrella project for the series of tests has been named *Evaluation of Latent Fingerprint Technologies (ELFT)*. The scope and structure of these tests are based partly upon lessons learned from the April 2006 *NIST Latent Fingerprint Testing Workshop*, supplemented by technical interchanges with workshop participants and vendors. The initial round of tests was initiated in April 2007, and is on-going.

The principal objective of this report is to provide a "snapshot" of the thinking, analysis, and planning that went into ELFT. The report is not intended to provide any test results or conclusions. These will be presented in subsequent reports.

While the immediate goal of ELFT is to assess automated technology, the long-term goals go far beyond simply quantifying performance. It is fully expected that understanding the performance envelope and limitations of contemporary matchers will lead to improvements in technology. These in turn, will lead to enhanced performance for searches of ten-prints and plain impressions against unsolved latent databases/watchlists. Equally important, technology improvements will provide law enforcement the capability to search their unsolved latent fingerprints against ten-print files with greatly reduced effort. Figure 1 provides a schematic of the overall ELFT goals.
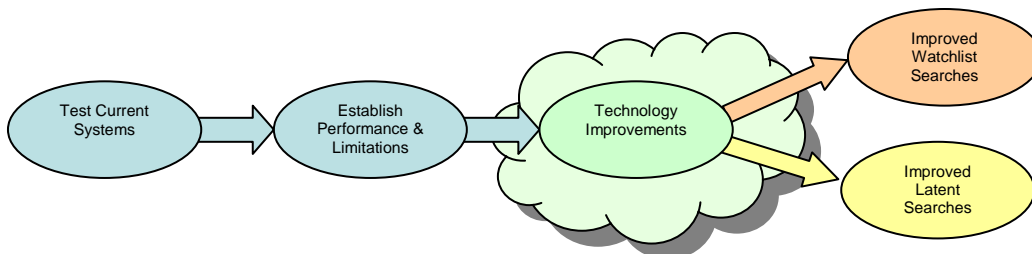

Figure 1 -- Overall Project Goals

EFTS is structured as a multi-year project, and the full impact of this work may not be felt for several years. The first part of this project, ELFT07 (07 is the year), consists of two tests, run in a "lights-out" environment, followed by a workshop. The two tests have been termed *Phase I and II*. Phase I is a *proof of concept test*, whose main purpose is to demonstrate integrity of the software in a lights-out environment. During Phase I the software will demonstrate: a) automated feature extraction from latent images; b) the ability to match these features against enrolled 10-print backgrounds; and c) generation of candidate lists. Phase II will then employ a larger database to quantify the achievable performance ("hit rate") for automated searches. A schematic diagram for ELFT07 is provided below.
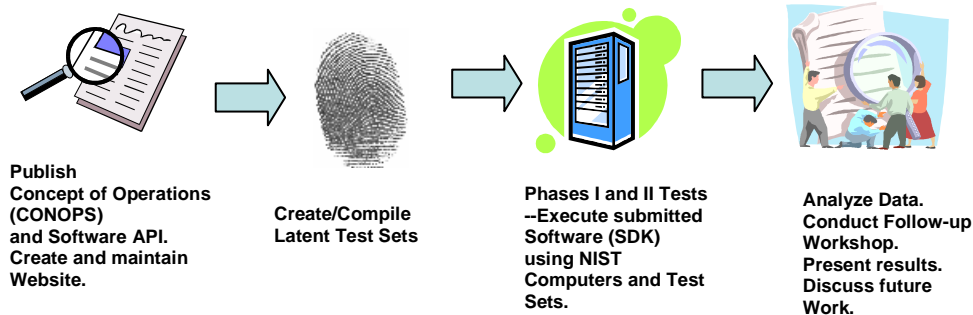

Publish Concept of Operations (CONOPS) and Software API. Create and maintain Website.

Create/Compile Latent Test Sets

Phases I and II Tests --Execute submitted Software (SDK) using NIST Computers and Test Sets.

Analyze Data. Conduct Follow-up Workshop. Present results. Discuss future Work.

Figure 2 – Schematic of ELFT07

In subsequent years (2008+) we plan to expand the above tests in several ways. First, we plan to augment the ten-print databases with a mix of rolled and plain impressions ("flats"). These will enhance NIST's and the latent community's understanding of the challenges of matching latents against flats. Continuing this line of investigation, we will then transition to searches of *plain impressions against databases of latent images* (sometimes referred to as *reverse searches*). Initially these tests will be restricted to single-finger searches, and subsequently will be enlarged to multi-finger latent searches. Figure 3 provides a diagram of the augmented latent testing plan.
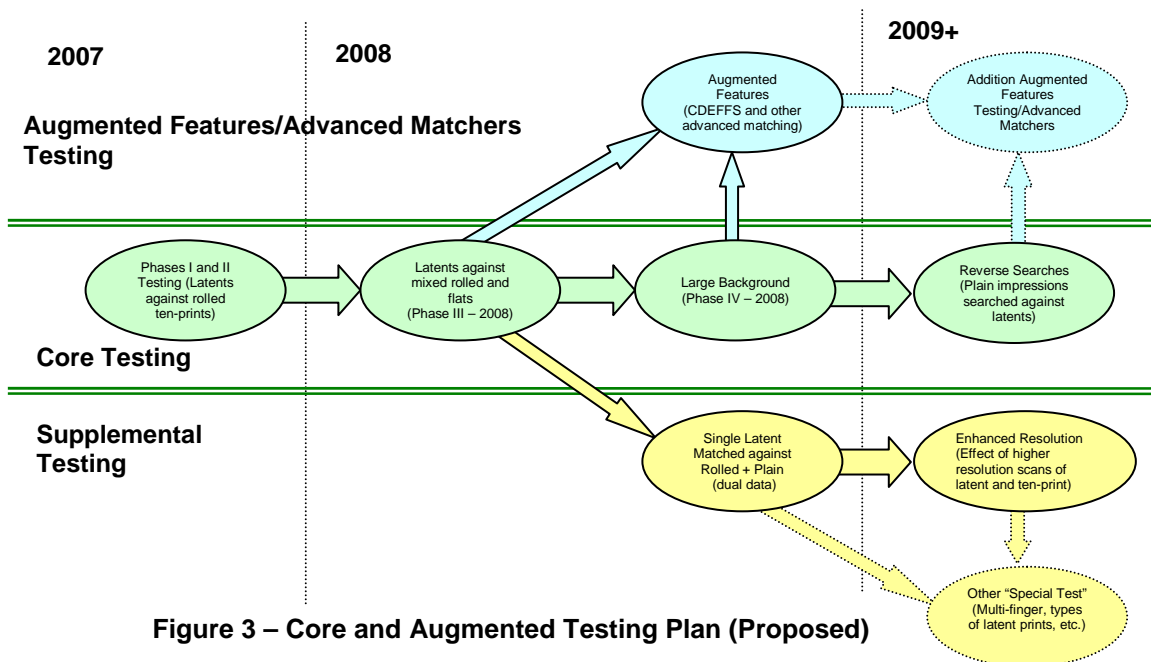
Figure 3 – Core and Augmented Testing Plan (Proposed)

As shown in the figure, envisioned tests include: a) latents and mates scanned at enhanced resolution (1000 and 2000 ppi); b) latents lifted/developed/processed in diverse manner (i.e., how was the image actually produced from the latent); c) latents matched against latents; and d) searches employing new or non-traditional features (e.g., level-3 features).

NIST is also looking into the development of *latent image quality measures (LIQM).* The principal function of the *latent image quality measure* is to provide a good indication of whether a latent is amenable to automated ("lights-out") matching. Only latent prints of higher quality measure would be submitted for lights-out matching. The development of a suitable quality measure might require addition testing, not encompassed in the above diagram.

We have outlined above the full scope of this project with a "broad brush."
In the remainder of this document we focus on this year's portion of the project, ELFT07. So as not to break up the flow, more detailed discussion of select topics have been moved to the end, Section 14.

2. Who Should Participate in the Tests?

Developers of latent fingerprint matcher software systems are strongly encouraged to participate in ELFT07. In addition, companies, research organizations, or universities that have developed mature prototypes, or stable research latent fingerprint matchers, are invited to participate. The latent fingerprint matching software submitted need not be a "production" system, nor be commercially available. However, the system must, at a minimum, be a stable implementation capable of being "wrapped" (formatted) in the *Application Programming Interface (API)* specification that *National Institute of Standards and Technology (NIST)* has published[1] for this evaluation. Anonymous participation will not be permitted. The Application

---

[1] The API is available at http://fingerprint.nist.gov/latent/elft07/elft07_api.pdf

form for participation is found on the NIST website. The Application form goes into further details regarding application and qualification.

3. Precedence of Documents

It is intended that this *Concept of Operations (CONOPS or ConOps)* be the single most comprehensive document covering *Latent Testing/ELFT* concepts. It will be periodically revised to reflect updates in NIST's planning, and to incorporate vendor comments and suggestions. However, this CONOPS is not guaranteed to be the most accurate for highly technical data. In the event of conflict with the *Application Programming Interface (API)* or the *Application Form* the latter two will take precedence. These two documents may be supplemented by additional reference documents in the future.

4. Test Objectives – What will be Tested during ELFT07?

As previously indicated, the primary purpose of the Phase I and II testing is to quantify the core algorithmic capability of contemporary matchers, in order to understand their strengths and limitations. In the initial tests the emphasis is on matching latents again ten-prints. (In subsequent tests this will be expanded to include other types of matching, for example latents against plain-impressions, and latents against latents.) The testing will be conducted using a software-only implementation, in a lights-out environment, and utilizing NIST hardware and datasets. During Phase I the software will demonstrate: a) automated feature extraction from latent images; b) the ability to match these features against enrolled 10-print backgrounds; and c) generation of candidate lists.

While Phase I is primarily intended to be a *proof of concept test*, it will nevertheless provide a certain amount of performance statistics. Only aggregate statistics (combining all successful participants) will be published. Individual results will be disclosed only to the owner of the respective SDK.

NIST will compute (and report as an aggregate statistic) the number of "hits" in each position (first, second, third, up to 50). A specimen of such a chart is shown below, followed by a specimen cumulative chart.
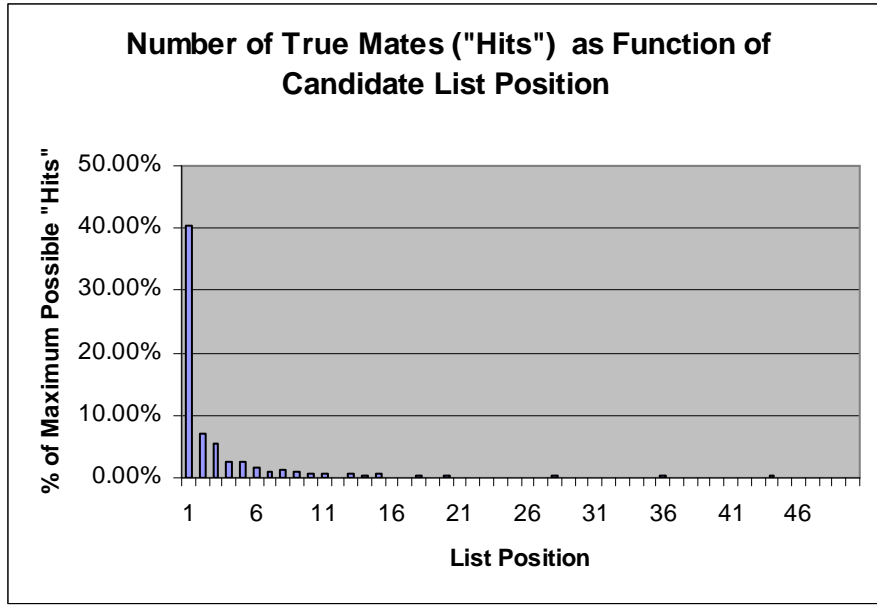
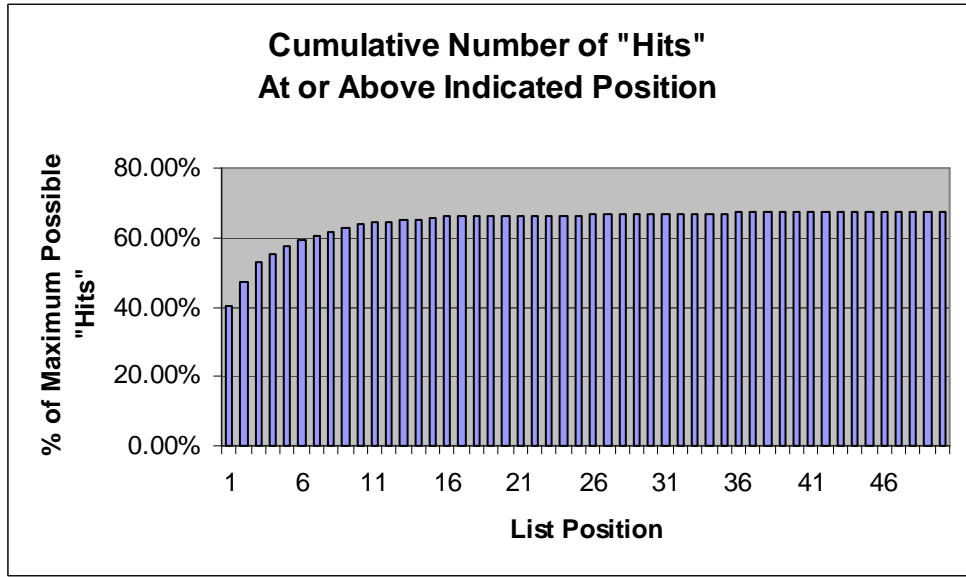**Figure 4 – Specimen Chart of "Hits" vs. Position**

Figure 5 – Specimen Cumulative Distribution

NIST will also report two overall performance metrics:

1) Metric #1 simply count the fraction of cases in which the correct mate appears in top position on the candidate list. This metric ignores all candidates below top position. This is basically the same as the first value in the above graph, in this case is 0.4. The highest possible value is 1.0 and the lowest is 0.

2) Metric #2 gives partial credit for mates appearing in lower than top position. It assign 1.0 to a mate in top position, 0.5 to one in second position; 0.3333 to third position, etc. The final score is the sum of these scores divided by the number of searches with mates. (For the validation data all searches have mates.) This score will always be at least as high as metric #1, and will generally be somewhat higher. The highest possible score is 1.0 and the lowest is 0.

NIST will also use DET performance metrics as a primary indicator of one-to-many identification search accuracy. This involves plotting *False Acceptance Rate (FAR)* and *True Acceptance Rate (TAR)* for all values of the threshold. (Equivalently, one can use false rejection and false acceptance rates.) An example of this procedure, employing hypothetical data, is shown in the following two figures.
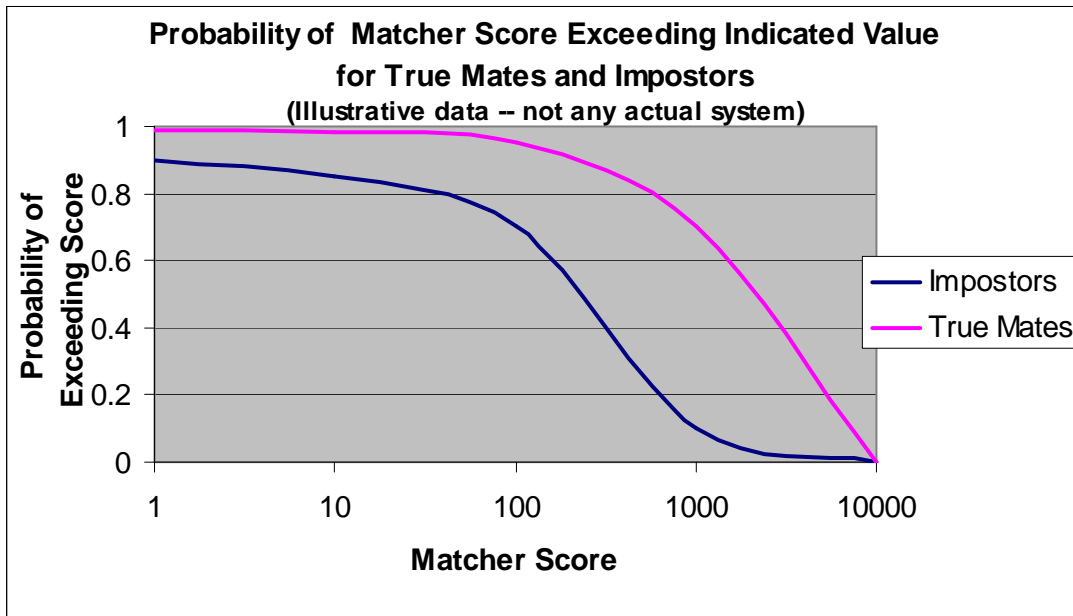

Figure 6 – Complement of Cumulative Distribution for True mates and Impostors

Figure 6 provides the complement of the cumulative distribution (1 – Cum_value) for two cases: a) when the search is matched against the true mate, and b) when the search is matched against an imposter (FAR). (The fact that both curves go to zero reflects the fact that very high scores become increasingly rare, for true mates as well as imposters.) By eliminating the score between the two curves of Figure 6 one obtains the DET curve, as shown in the next figure.
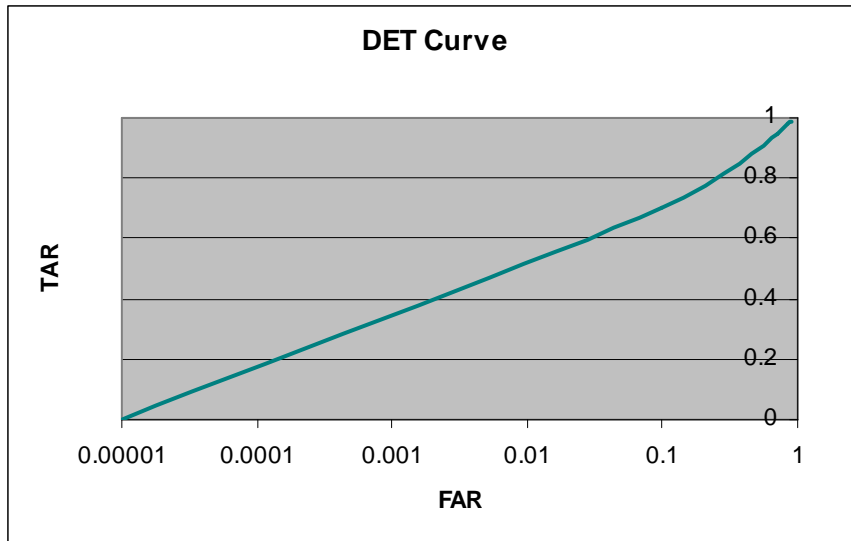
Figure 7 – Specimen DET Curve

For Phase II we will employ a larger database so as to improve the statistics. The results will be published as aggregate statistics, as well as individual participant statistics (for those participants electing to continue to Phase II).

NIST will also report enrollment and search timing information. Speed of execution, for both enrollment and latent search, is of secondary importance. However, in order to conduct these tests in a reasonable amount of time NIST must impose some limitations. These are covered in Section 11. In reporting timing measurements NIST will specify the exact hardware that the software was hosted on. NIST will in addition caveat timing measurements by noting that operational latent searching algorithms are likely to be implemented in more sophisticated hardware.

NIST recognizes that latent searches pose many special problems. One of these is that "strong hits" may have widely different matcher scores (for example, depending upon the number of minutiae). This may place challenges on the DET approach. Largely for this reason we have recommended the inclusion of a normalized score that attempts to compensate for the matching score variations. See Sections 13 and 14 for further discussion.

A comprehensive discussion of performance metrics and their definitions is found Patrick Grother, Ross Micheals, and P. Jonathon Phillips, ***Face Recognition Vendor Test 2002 Performance Metrics***, 31 March 2003. See also
Ł http://ois.nist.gov/nistpubs/technipubs/recent/search.cfm?dbibid=13808

5. Publication of Participation and Results

NIST understands that this project is entering a relatively unexplored field, and many challenges lie ahead. For this reason we have structured ELFT07 to include two phases. We consider Phase I to be a *Proof of Concept Test*. This means that the primary objective of Phase I will be to demonstrate that the submitted SDK executes on the Phase I data to completion, in a lights-out environment, and produces a

"meaningful" output. A "meaningful" output is basically an output in the correct format.

The detailed results of Phase I will be discussed with the Participant on a "one-on-one basis," but will not be published or submitted to other government agencies. The number, but not the names, of participants who attempted and completed Phase I will be disclosed. However, in the (likely) event there are a significant number of participants in Phase I, NIST is considering publishing the aggregate test results (under the premise that this may be a fair assessment of the state of the art). By "aggregate test results" we mean that results are "lumped," and that no specific candidate list or participant-specific scores will be mentioned. Participants will have the option to withdraw anonymously following participation in Phase I. (This means that their withdrawal will not become a public announcement.)

Participants who elect to continue to Phase II may resubmit their SDKs. These need not be identical to those of Phase I. Following completion of Phase II testing the Government will combine all results into a Final Report. The *Evaluation of Latent Fingerprint Technologies Test, Phase II Final Report* will contain, at a minimum, descriptive information concerning ELFT07, descriptions of each experiment, and aggregate test results. Should individual participant's scores be published, NIST will exercise care that any implied rankings are well supported by the underlying statistics. (I.e., two scores will be considered the same if the difference is not statistically meaningful.)

Participants will have an opportunity to review and comment on the Final Report. Participants' comments will be either incorporated into the main body of the report (if it is decided NIST reported in error) or published as an addendum. Comments will be attributed to the participant. After the release of the Phase II Final Report, Participants may decide to use the results for their own purposes. Such results shall be accompanied by the following phrase: "Results shown from the Evaluation of Latent Fingerprint Technologies Test (ELFT07) do not constitute endorsement of any particular system by the U. S. Government." Such results shall also be accompanied by the Internet address (URL) of the ELFT07 Final Report on the ELFT07 website.

For Phase III and beyond NIST intends to publish statements of the performance of all implementations submitted for testing. These will include measurements of identification error rates and throughput. These results will be attributed to participants. Accordingly, NIST will require an appropriately signed application form from all participants and NIST will not evaluate any implementation unless the participant consents to the disclosure of its performance. The NIST tests use sequestered images. These will not be provided to participants.

6. Protection of Participant's Software

NIST recognizes the proprietary nature of the participant's software and will take all reasonable steps to protect this. The software submitted will be in an executable library (SDK) format, and no algorithmic details need be supplied. NIST will use the Participants SDK software only for the agreed-upon testing, and in the event errors are subsequently found, to rerun prior tests and resolve those errors. NIST agrees not to use the Participants software for purposes other than indicated above, without express permission by the Participant.

7. Why "Lights-out"?

The term "lights-out," as used in this document, will indicate that no human assistance will be required in conducting the latent searches. In particular, all feature extraction steps, both for the enrolled images and for the latent images, must be performed entirely by the SDK under test.

There are good reasons why NIST selected the "lights-out" mode of testing:

1. It decouples the skill of the human expert from the intrinsic merits of the software.

2. It protects the privacy of the test data by keeping the data in house, and not requiring examination by non-government personnel. This mode of testing allows the use of *Sensitive but Unclassified* test data.

3. It encourages a forward-looking view of how latent searches might be done in the near future. It is anticipated this broader outlook will lead to technical innovations.

Algorithmically speaking, "lights out" consists of two separate concepts. The first is automated feature extraction (and of course matching). The second is *candidate list reduction*. NIST envisages that in the near future automated search capabilities will assist latent experts by reducing the size of candidate lists that they need to examine by eliminating the more obvious "nuisance" non-matches. We refer to this part of the automated matching as *candidate list reduction*. To achieve effective *candidate list reduction* may require additional computer processing, including the development of new algorithms.

Although this is a "lights-out" test NIST will use some human assistance in the data preparation phase. Any such assistance will be provided indirectly by NIST, and might include a) cropping and/or re-orienting of selected latent images, and b) specifying a region-of-interest in the from of a mask. The mask will be a byte image conformal with the size of the latent image. Initially only two values will be used for each pixel, 0 and 255. A zero value will indicate "do not use this pixel," while 255 will indicate a "good" pixel. In the future these two values may be augmented by other values to indicate finer gradations of quality. NIST will also involve latent experts for examining potential consolidations and for resolving contested or unclear results.

8. Test Data

NIST will select the test datasets from its internal sources. The Test Datasets are protected under the Privacy Act (5 U.S.C. 552a), and will be treated as **Sensitive but Unclassified** and/or **Law Enforcement Sensitive**. ELFT07 Participants will have no access to ELFT07 Test Data, neither before, during, or after the test, with the exception of the small Validation Dataset.

8.1 ELFT07 Datasets

The *Validation Dataset* is a very small data set, comprised of ten search latents and 100 background ten-prints. The purpose of this dataset is to demonstrate stable software, compliant with the API. Upon receiving the applicant's SDK and Validation Dataset results, NIST will rerun the applicant's software using the Validation Dataset.

For the applicant to be officially accepted (and designated a *Participant*) NIST must be able to reproduce the submitted results.

In the event of disagreement in the two outputs, or other difficulties, the Participant will be notified. Participants will be notified with a detailed description of the problem(s) and given reasonable opportunity to resubmit.

Both Phase I and Phase II Test Datasets will consist of latent images for searches and ten-prints for the background ("gallery"). For both tests, but especially for Phase I, NIST will attempt to make the test set as "benign" as reasonable. This means NIST will attempt to avoid including known problem cases such as: 1) very small latent area; 2) extremely busy or otherwise difficult background; 3) highly blurred image; 4) multiple fingerprint impressions; or 5) upside down or mirror images.

Latent images will be supplied uncompressed, and will have been scanned at either 500 ppi or 1000 ppi. The participant should be prepared to handle either resolution. Additional image characteristics may be found in the API.

Background ("gallery") data will consist of rolled ten-print impressions, scanned at 500 ppi, and presented in a decompressed form.

The nominal sizes of the Phase I and Phase II Datasets are as follows:

ELFT07 Dataset Characteristics		
Dataset	Number of Latent Searches	Number of Ten-prints
Validation	10	100
Phase I	100	1000
Phase II	1000	10,000

8.2 Size of Images

For ELFT07 the following guidelines apply:

- ν All mates (background) will be rolled impressions. There will be no "flats."
- ν In all cases rolled impressions will not exceed 1000 x 1000
- ν In all cases latent images will not exceed 2000 x 2000, though often will be significantly smaller.
- ν Minimum dimensions for latent images under 300 are possible, but never smaller than 150.

Over the entire series of planned tests the size of test images may vary considerably. For example, images scanned at 2000 ppi will contribute some very large sizes, potentially as large as 4000 x 4000 pixels. However, these very large sizes will only appear in the "downstream" tests.

9. Testing Platform

NIST will host the participant's software (SDK) on a high-end PC (workstation/server type). Although these PCs include of a mix of models, a "typical" PC will have the equivalent of a Pentium 4, 2.8 GHz processor, or higher; 2 GB of memory; and at least

50 GB of disk memory. The participant software must be able to reside and execute on this single PC. NIST, at its discretion, must be able to copy the software to several PCs to expedite or scale-up the testing. These computers are configured with either a Windows 2000 or Linux operating system.

10. Format of Participant Software

The software undergoing testing will be hosted on NIST-supplied computers. The executable modules will be built up from two sources: 1) participant-supplied software provided in the form of a Software Development Kit (SDK), and 2) NIST-supplied software. The core of the executable module is of course derived from the SDK. The part supplied by NIST is mainly concerned with the image retrieval and manipulation.
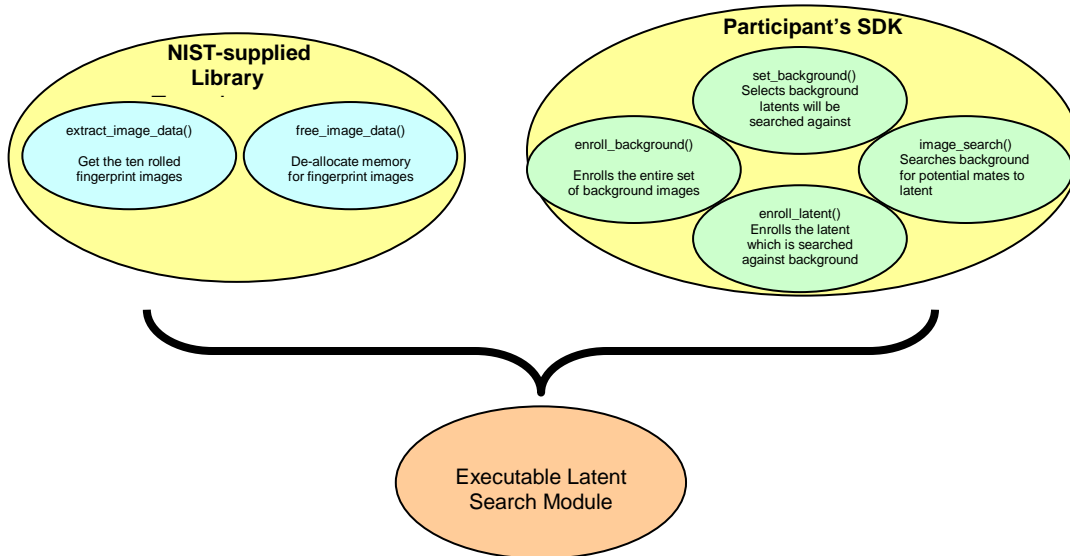

Figure 8 – Executable Software is a Combination of Participant-supplied (SDK) and NIST-supplied Software

The participant may supply their SKD in a Windows 2000 or Linux implementations. The Application Programming Interface (API) specification provides further details regarding the format in which participants must submit their software. See Ł http://fingerprint.nist.gov/latent/elft07.html

Execution occurs in two passes:

1. Enrollment of Background Pass: The first pass performs feature extraction of the background fingerprints. This process converts images into proprietary "feature" representations. The output of this enrollment is at the discretion of the participant, except that all extracted data must be written into a single directory. The background will consist entirely of rolled impressions of ten fingerprints during Phases I and II, and will be augmented by plain impressions ("flats") during some of the subsequent phases. The following diagram summarizes the background feature extraction pass.
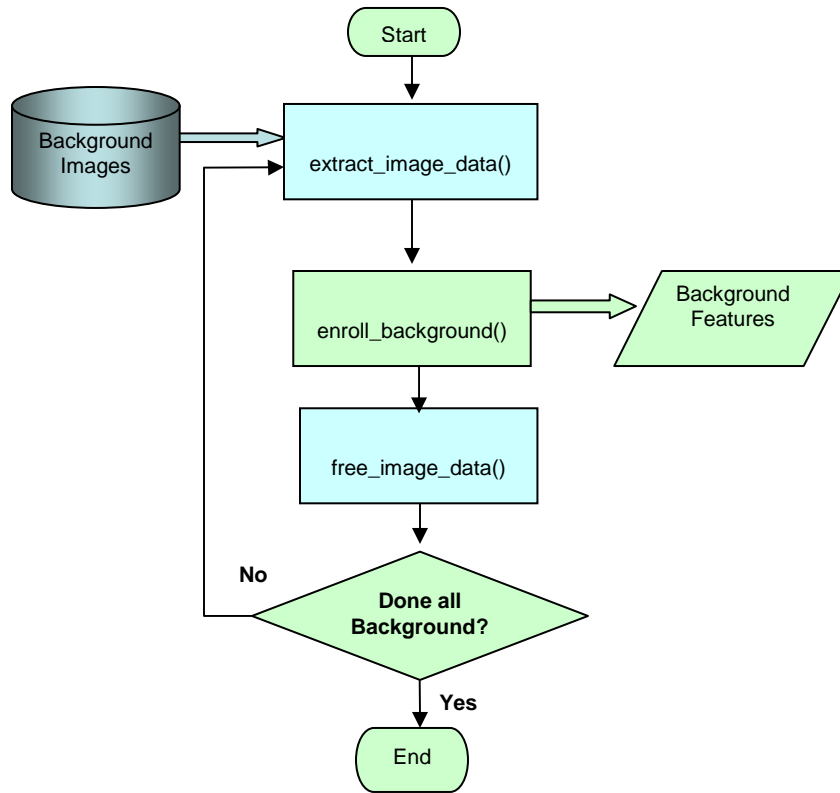
Figure 9 – Logic and Data Flow for Executable Enroll_Background
Module

2. Matching Pass: During this pass each latent search image is matched against the enrolled data and a candidate list is generated. The details of the searching process are unregulated by the API. In particular, participants may invoke multi-stage algorithms within the top level search function call by NIST.

Prior to the actual search, but part of pass two, is the enrollment of the latent image. The details of the latent enrollment process depend whether a latent *Region of Interest (ROI) Mask* is specified or not, as shown by the following figure.
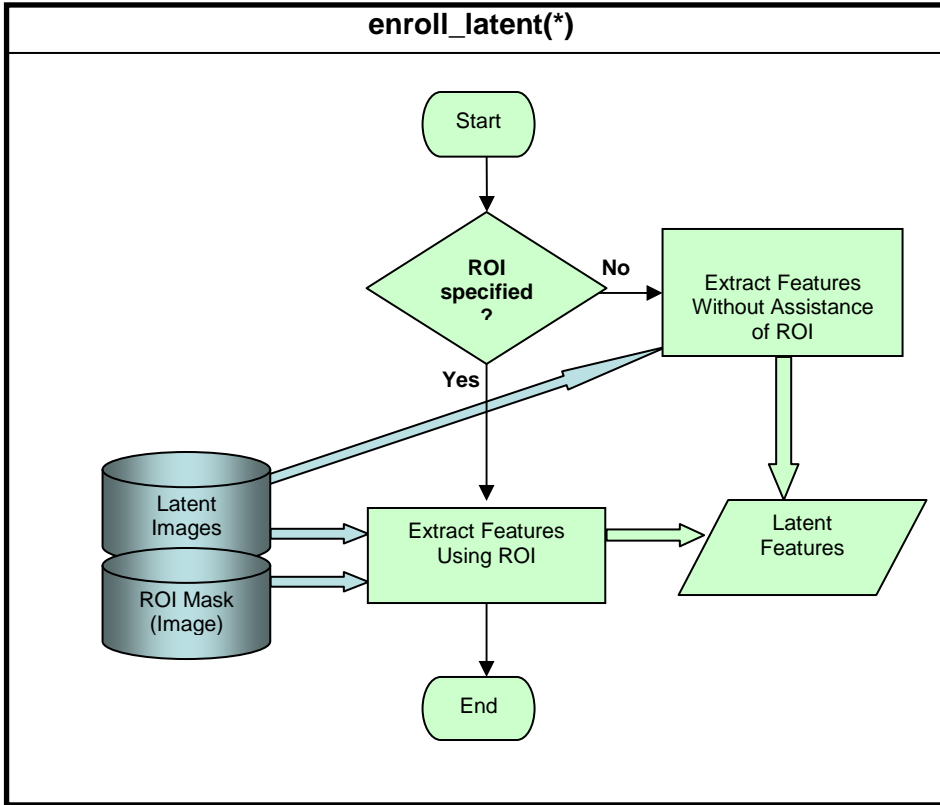
Figure10 – Logic and Data Flow for enroll_latent() Function

The following diagram summarizes the logic and data flow during the matching pass.
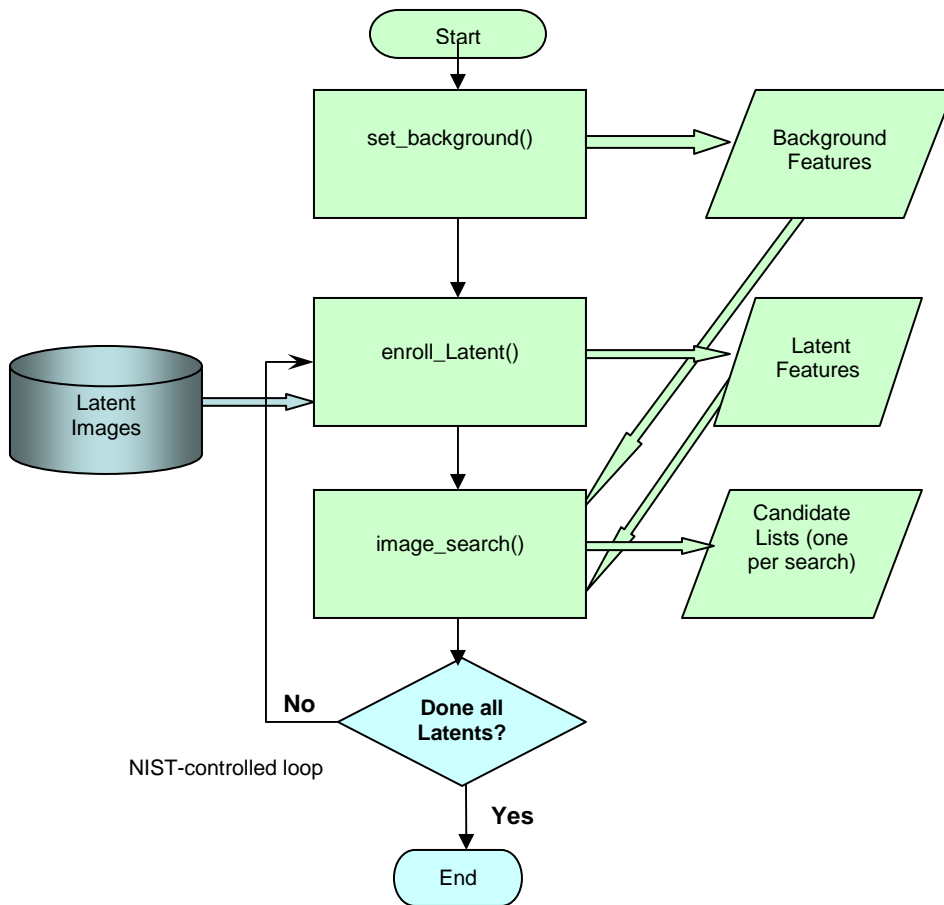
Figure 11 – Logic and Data Flow for Executable Search Module

Participants will be supplied a *Validation Dataset* to test their software before submitting it to NIST. Participants must submit candidate lists from the Validation Dataset to NIST. Once NIST has received participant implementations it will rerun the test and check its candidate lists against those submitted. If a difference is detected, or if other problems are encountered, the participant will be required to respond to NIST requests for updated implementations.

NIST may provide a latent image with an accompanying region of interest, in the form of a "mask." This mask will be a byte image conformal with the size of the latent image. Initially only two values will be assigned to each mask pixel, 0 and 255. A zero value will indicate "do not use this pixel," while 255 will indicate a "good" pixel. In the future these two values may be augmented by additional values to indicate finer gradations of quality. The participant's software should default to "no mask option" when no mask is supplied by NIST. The participant may use this information as they see fit. NIST plans to run the test with and without masks, but employing the same search and background. It is expected that searches with a mask will generally produce superior results. Quantifying the amount of improvement is a test objective.

NIST will also involve latent experts for examining potential consolidations and for resolving contested or unclear results.

11. Timing Requirements

For ELFT07 NIST will allow the simultaneous submission of two SDKs, a primary and secondary. It is envisioned that the first (primary) SDK would be a "slow but accurate" version of the software, while the second would be a "fast and but less accurate." (However, these are suggestions, and need not be followed. The participant may submit any two SDKs, subject to the timing restrictions below. The participant need not supply any details regarding the differences between the two SDKs which might reveal proprietary information.)

Execution time is a concern. While we plan to allow "generous" execution time, we request that participants remain within the following limits:

Phases I and & II Execution Time Limits	
Primary SDK	
Background Enrollment	Total Time (seconds) ≤ 150 * (number of ten-prints in background)
Latent Enroll	350 seconds per image
Search	Time per Search (seconds) ≤ 0.2* (number of ten-prints in background)
Secondary SDK	
Background Enrollment	Total Time (seconds) ≤ 80 * (number of ten-prints in background)
Latent Enroll	250 seconds per image
Search	Time per Search (seconds) ≤ 0.1* (number of ten-prints in background)

(Comments and questions should be submitted to the evaluation team at latent@nist.gov.)

12. Test Schedule
NIST has created a website to serve as the central repository for all information regarding NIST's latent fingerprint activity. The URL is Ł http://fingerprint.nist.gov/latent. It links to the ELFT-07 testing Ł homepage which contains:

1. Application form for ELFT07
2. API specifications for participants' implementations
3. This document (CONOPS)
4. Summary of NIST Latent Fingerprint Testing Workshop, and
5. FAQ

The schedule is tabulated below. The later entries are subject to certain contingencies. However, NIST intends to adhere to the schedule as closely as possible.

ELFT07 Master Schedule		
Action	**Actor**	**Date**
Announcement of Latent Evaluation	NIST	2006-11-17
Publication of draft API and Latent Evaluation Concept Paper, and solicitation of comments	NIST	2006-11-17
Deadline for submission of comments on API and CONOPS to NIST. These should be submitted to latent@nist.gov	Participants	2006-12-22
Publication of API	NIST	2007-01-26
Publication of Latent Evaluation Application Form	NIST	2007-01-26
Deadline for submission to NIST of Latent Evaluation Application Form	Participants	2007-02-23
Publication of Validation Data	NIST	2007-03-02
Submission of SDK and Validation Dataset results	Participants	2007-04-13
Completion of SDK conformance testing and integration	NIST	2007-06-08
Commencement of Phase I Test (NIST)	NIST	2007-06-11
Completion of Phase I Test	NIST	2007-07-13
Communication of Phase I Results to Participants		2007-07-20
Announcement of Phase II Test Schedule	NIST	2007-07-27
Submission of Phase II SDK (Withdrawal after this time will be reported in NIST publications.)	Participants	T
Completion of Phase II Test	NIST	T+6 months
Publication of formal evaluation report	NIST	T+6 months
Workshop to summarize Phase II and to plan subsequent phases	NIST	T+6 months

13. Format of Candidate List

The output candidate list should have a fixed length of fifty (50) candidates. We have selected this size because it is short enough to be convenient, yet long enough to give an indication of the number of "hits just out of reach." (We currently don't envision cases in which the background is less than 50 fingers. Should this situation arise, the candidate list could be suitably "padded.")

The candidate list consists of two parts, a required and an optional part. The required part consists of: 1) the index of the mating ten-print subject; 2) the matching finger number; 3) the absolute matching score; and 4) an estimate of the probability of a match (0 to 100, see also Section 14). The optional part consists of: 5) the number of good minutiae identified in the latent; 6) the number of latent minutiae which were

successfully matched; 7) the quality estimate of the latent (0 to 100, 100 is best); and 8) the quality estimate of the mate (0 to 100, 100 is best). The API provides further guidelines regarding the meaning of quality scores. The candidate list is ordered based upon the absolute score, highest score in first position.

Sample Candidate List							
Required Part				Optional Part			
	Mate Index	Finger No.	Abs. Score	Prob. Of True Match	No. Latent Minutiae	Matched Minutiae	Quality Latent	Quality Mate
1	731	2	3513	93	18	12	70	80
2	103	1	605	5	18	5	70	60
3	207	7	513	4	18	5	70	90
●					18		70	
●					18		70	
●					18		70	
50	107	9	422	3	18	4	70	70

The parameter *Probability of True Match* in the above table represents an initial step toward *candidate list reduction*. It is intended to be an estimate of the probability that the candidate is a true match. Its values range from 0 to 100. For the Phase I test candidate list reduction is not required, and the participant could always enter 100 into this filed, if so desired. For Phase II some type of candidate list reduction is preferred.

14. Supplemental Notes

14.1 Supplemental Notes to Section 1.0 -- Concepts of Operation for "Improved Watchlist Searches"

A major goal of this project is to improve searches of watchlist/lookout-lists. It is becoming increasingly common to capture live fingerprints of arriving passengers at ports of entry and similar venues. The capture fingerprints are then compared to fingerprints of the person on file (one-to-one comparison, or validation match); they may also be searched against selected watchlists one-to-many search). The diagram below shows capture of the two index fingerprints at a port of entry. (This will be supplanted by the capture of all ten fingerprints in the near future.)
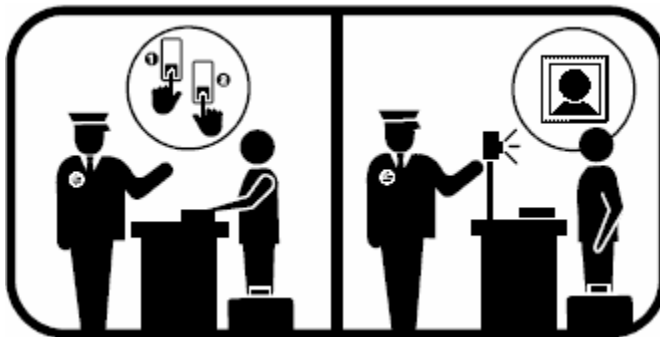

Figure 12 –Capture of Livescan Fingerprints at a Port of Entry

The newly acquired fingerprints may then be matched against any or all of three types of finger prints, as shown below. Verification, or one-to-one searches, are generally performed by matching with plain impression taken at a previous time. Watchlists or Lookout-lists contain the fingerprints of prior offenders, or persons of interest. They may be comprised of any of the types of fingerprints, though latent fingerprints are generally kept in separate files, exclusively dedicated to latents.
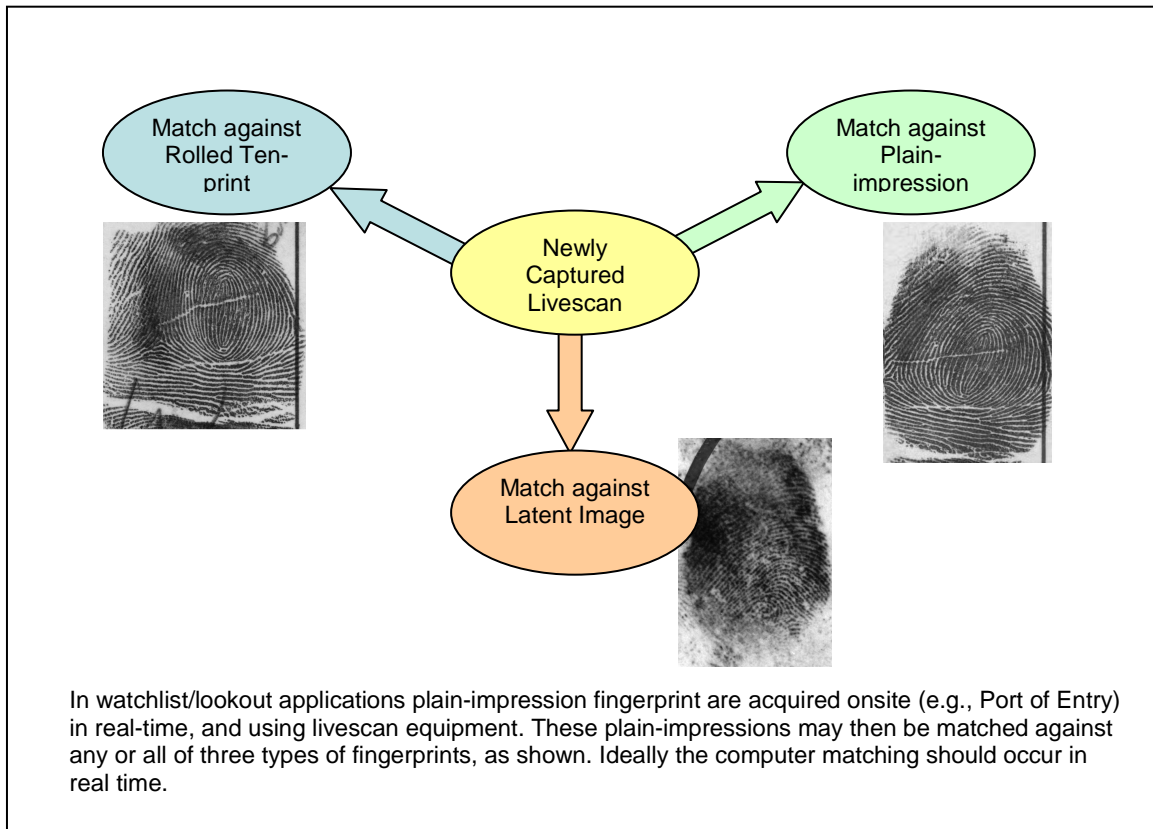

In watchlist/lookout applications plain-impression fingerprint are acquired onsite (e.g., Port of Entry) in real-time, and using livescan equipment. These plain-impressions may then be matched against any or all of three types of fingerprints, as shown. Ideally the computer matching should occur in real time.

Figure 13 – Plain-impression may be matched against any of three types of fingerprints

The simplest and clearest example of the applicability ELFT is to watchlists comprised of latents. However, there are several other ways in which ELFT may contribute to the point-of-entry scenario:

> Low-quality livescan images provide many of the same challenges as do latents. Improvements in latent matching should therefore transfer to real-time livescan matching. (Livescan images they are subject to "retake," but the number of retakes is necessarily very limited, because of the need to expedite the processing.)

To provide searches of watchlists in near-real-time, substantial algorithmic improvements are required. The multi-stage matching approach used by some latent matchers may offer a solution.

For increased search accuracy, additional features (e.g., level 3) might be required. A goal of ELFT to examine the performance increases provided by selected new features.

14.2 Supplemental Notes to Section 1.0 -- Concepts of Operation for "Improved Criminal Latent Searches"

A second major goal of ELFT is to provide "an automated latent search capability" to latent examiners. By this we mean that latent examiners should have the capability of *screening* their latent images with a minimum of effort. We use the term *screening* to emphasize that such searches are not fully equivalent to traditional searches. Figures 12a and 12b provide somewhat idealized schematics of how such automated screening might work.
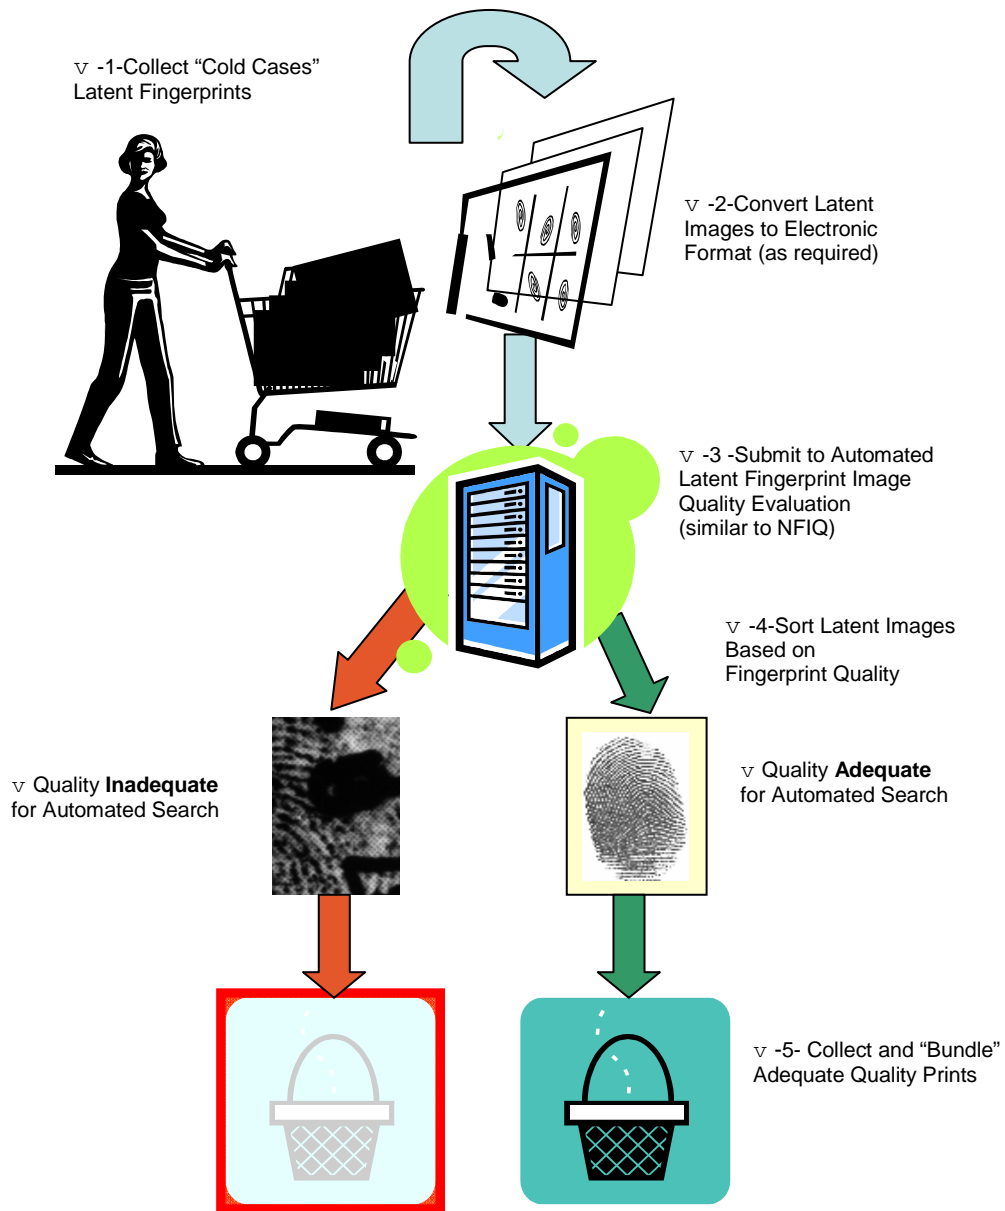
v -1-Collect "Cold Cases"
Latent Fingerprints

v -2-Convert Latent
Images to Electronic
Format (as required)

v -3 -Submit to Automated
Latent Fingerprint Image
Quality Evaluation
(similar to NFIQ)

v -4-Sort Latent Images
Based on
Fingerprint Quality

v Quality **Inadequate**
for Automated Search

v Quality **Adequate**
for Automated Search

v -5- Collect and "Bundle"
Adequate Quality Prints

Figure 12a – Automated Sort ("Triage") of Latent Fingerprints
based upon Suitability for Automated Searching

v - 6- Retrieve Collected "Bundled" Latent Prints of Adequate quality

v Quality **Inadequate** for automated Search --do not use!

v -7- Submit Prints for Automated Search

v -8-Perform Automated Search

No Candidate Found

1 Candidate Found:
Subject = M731
Finger = 2
Score = 3513
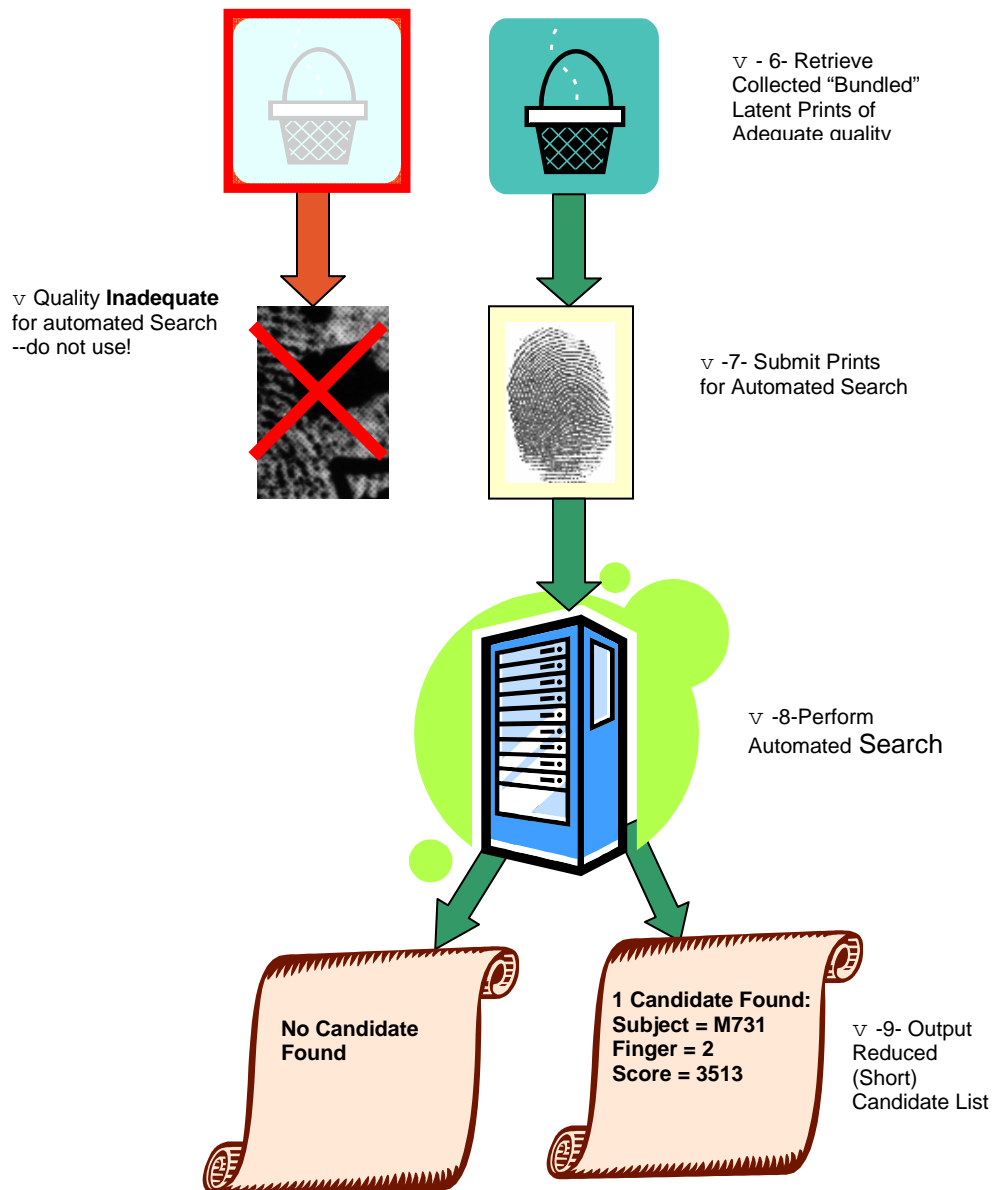
v -9- Output Reduced (Short) Candidate List

Figure 12b --Automated Search of Higher Quality Prints
with Candidate List Reduction

Following such a screening search the latent examiner may single out specific latents images for more traditional (and labor-intensive) searches.

The examiner can assist the search process by defining a *region of interest (ROI)* on the latent image. The examiner may also single out certain images for a more traditional search by improving the features extracted, starting either with the machine extracted features, or starting with a "clean slate."

14.3 Features for Use in Matching

Generally speaking, the selection of the features for use in the matching process is left to the participant. Matchers need not primarily be a "minutiae matcher." Matcher

architectures in which "advanced matchers" are selectively invoked depending upon initial results are allowed. For example the matcher might initially use certain core features in comparing the search (probe) with a background (gallery) subject. The result of this comparison might produce one of three possible outcomes: 1) the two fingerprints are too different, and no further effort should be expended on this candidate; 2) the two are so similar that this is definitely a mate; or 3) the two have points of similarity, but the match is not conclusive. The third case might then trigger a call to an "advanced matcher" for further resolution.

If using "advanced matchers", it is up to the participant to decide if the additional features (if any) required are to be extracted and stored on disk memory during the enrollment phase. Since it may not be possible to keep all gallery images in memory, it might be necessary for the software to retrieve the data from disk during searches. This extra fetch time will be included in execution time measurements.

Figure 3 shows that approximately a year-and-a-half downstream NIST intends to test the effect of using augmented *feature sets*. These will be largely based upon the CDEFFS feature sets, but are not necessarily limited to these. For a description of the proposed CDEFFS features please go to the following website
Ŀ http://fingerprint.nist.gov/standard/cdeffs/index.html

These augmented feature tests test will be run in a dual mode: a) first without employing any new feature, then b) employing designated new features. To allow this mode of operation a somewhat different format might be required for the SDK.

14.4 Supplemental Notes to Sections 7 and 13 -- Candidate List Reduction

NIST envisages that in the near future automated search capabilities will assist latent experts by reducing the size of candidate lists they need to examine through elimination of the more obvious "nuisance" non-matches (impostors). For example, assuming that one hundred latents are submitted for searches, and that each search produces a candidate list of twenty candidates, an examiner needs to look at 2000 candidates. Since a typical identification rate for latent searches might be around 4%, this means that 2000 candidates need to be examined to find the four true identifications. While it is true that skilled examiners can quickly dismiss "nuisance candidates," nevertheless it does take up valuable examiner time. An even larger concern is that too many nuisance candidates might result in the true mates being overlooked due to operator fatigue. It is therefore desirable to minimize these nuisance candidates. We refer to this part of the automated matching process as *candidate list reduction*. To achieve effective *candidate list reduction* may require additional computer processing, including the development of new algorithms.

Since candidate list reduction poses many challenges, we plan to implement it in stages. The initial stage is to introduce a new parameter called *Probability of True Match*. This is intended to give a numerical estimate that the candidate is a true mate of the latent. This parameter should be supplied as a number between zero and 100. The number 100 will be interpreted as an extremely high confidence "hit." The intent is to use this parameter as a key to candidate list reduction.

Certainly the raw matcher score by itself provides a strong clue regarding the merit of a given candidate. However, by itself it is insufficient. For one thing, there is no agreed-

upon standard for the range of matcher scores: Does a value of 5000 indicate a high score? A very high score? Secondly, whether a given score belongs to a true mate depends upon the size of the background. The larger the background the more likely it is that large impostor scores will be created. The *Probability of True Match* therefore needs to take matcher score and background size into account. Additional information might also be factored in, such as: a) the score gap to the next candidate; b) the quality of the latent; and c) the quality of the mate.

There does not appear to be any simple way of computing *Probability of True Match,* and participants are encouraged to develop their own procedures. Purely as an illustrative example, we offer the following procedure.

Assume a background size of N. Assume additionally that some candidate has achieved a score of S. To fix our ideas assume S = 5000. Assume further that data such as shown in Figure 6 is available. Then from this figure we obtain TAR = 0.3 and FAR = .00001. This may be interpreted as: the *a priori* probability of obtaining a score exceeding 5000 is 0.3 when matching against the true mate. Conversely, the *a priori* probability of exceeding 5000 when matching against an imposter is .00001. Note that this is for a single imposter. The probability that one or more imposters exceed 5000 in a background of N, assuming independence in match scores, is $1 - (1 - FAR)^N$. If this is taken as the *a priori* probability of an imposter exceeding 5000, then we can renormalize in the Bayesian sense to obtain the probability that the candidate is a true mate. The result is shown in the following figure
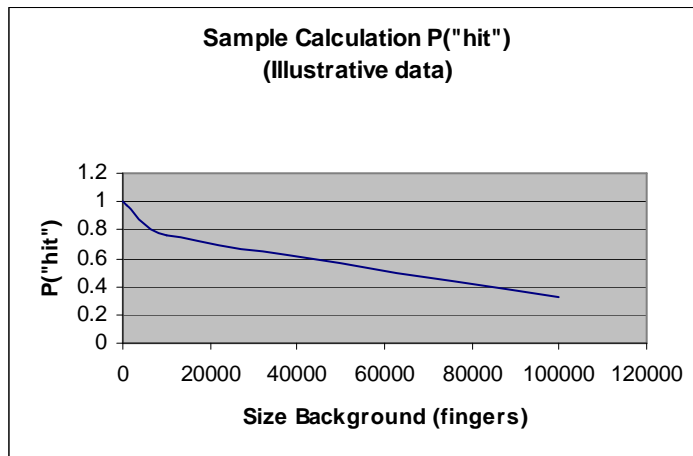

Figure 13 – Illustrative Calculation of P("hit") (See text)

The fact that the probability diminished rather rapidly as a function of background size is a result of the independence assumption. It is possible that in a real system impostor scores are not independent. One of the secondary objectives of ELFT will be to obtain more insight into this problem.