

PILL IMAGE RECOGNITION PILOT (PIR PILOT) REQUEST FOR INFORMATION (RFI) INSTRUCTIONS FOR RESPONDING

On February 2, 2015, the National Library of Medicine (NLM) issued a [Pill Image Recognition \(PIR\) Request For Information \(RFI\)](#). This RFI is a pilot for a Challenge whose goal is to develop smart phone apps that individuals can use to take pictures of prescription pills and then search for and retrieve pill images and associated data of likely matches in an NLM database. Examples of pills are capsules and tablets intended for oral use.

This document provides instructions to those interested in responding to this RFI and participating in the PIR Pilot. PIR Pilot is open to private industry, academia, national laboratories, and government agencies. NLM anticipates that respondents will include professionals and students, individually or in teams, in computer vision and computer graphics working on content-based image retrieval. Respondents to the PIR Pilot may also participate in the Challenge.

The PIR Pilot and the forthcoming PIR Challenge are part of NLM's [Computational Photography Project for Pill Identification \(C3PI\)](#).

WHY SHOULD YOU PARTICIPATE IN THE PIR PILOT?

From NLM's point of view, the primary reason for participating in the PIR Pilot is to help ensure that the forthcoming PIR Challenge instructions are accurate and clear, that the PIR Challenge judging will be impartial, and that PIR Challenge responses will include high quality software.

From the point of view of PIR Pilot participants interested in the PIR Challenge, they are likely to feel better prepared because they will:

- Have been able to exercise their algorithms and software on the PIR Pilot data that includes a key that tells which consumer images match which reference images – this is all described below
- Have been able to exercise their algorithms and software on both the PIR Pilot data and the forthcoming Challenge data –exercising over more data can only be helpful
- Have become familiar with the PIR Pilot submission process, anticipating that the PIR Challenge submission process will be similar
- Have been told if their PIR Pilot submission was properly assembled
- Through answering the questionnaire about the PIR Pilot process, have provided input for the PIR Challenge process
- Find that the summary statistics and a discussion of the responses to the PIR Pilot

that will be posted at the publicly-accessible web site pir.nlm.nih.gov are more informative, such as in comparing their results with the summary statistics

DATA PROVIDED TO PARTIES INTERESTED IN PARTICIPATING IN THE PIR PILOT

- Directory [DC](#) containing $M = 3,000$ files of JPEG consumer images taken with cell phone and consumer grade digital cameras. This directory can be considered an ordered set C of files sorted in lexicographical order by file name, where $c(i)$ is the i -th image file.
- Directory [DR](#) containing $N = 1,000$ files of high quality JPEG reference images and is the image of either the front or the back of a pill. This directory can be considered an ordered set R of files sorted in lexicographical order by file name, where $r(j)$ is the j -th image file.
- A comma-separated-value [file](#) that identifies which consumer image corresponds to which reference image, for all M images.
- A text file containing a [questionnaire](#) about the PIR Pilot process.

Download all four as a single compressed [file](#).

All data provided under the PIR Pilot is governed under the [Creative Commons Attribution 4.0 International License](#).

EXAMPLES OF PIR PILOT DATA

Examples of consumer images for four different pills – taken under different lighting conditions, using a variety of cell phone and consumer grade digital cameras:



Examples of reference images for four different pills – JPEG files of either the front or the

back of a single pill, centered on a neutral background:



SOFTWARE TO BE DEVELOPED FOR THE PIR PILOT


PIR Pilot participants are asked to develop algorithms and software for matching images of consumer pills to reference images – and to submit their executable software as described below. The software does not need to identify pills by name. The software can use any programming language(s).

The software should be a command-line batch-mode program or a script whose input and output are as follows:

- The first argument is the name of a directory containing one or more consumer images in JPEG format.
 - The images in this directory may be unknown, may be training or test images, and may be called query images.
 - Example: Directory DC described above
- The second argument is the name of a directory containing one or more reference images in JPEG format.
 - Example: Directory DR described above
- The third argument is the name of an output file containing an ASCII comma-separated-value (CSV) similarity matrix S in which the term $s(i,j)$ is a number in the unit interval $[0,1]$ that indicates how well the consumer image $c(i)$ matches reference image $r(j)$. A larger value should indicate a better match.
 - The matrix shall include row headers and column headers that correspond to the input file names.

SCHEMATIC OF THE PIR PILOT PROCESS

The following illustrates the relationship among input, software, and output:



HOW TO SUBMIT RESPONSES TO THE PIR PILOT

Participants are asked to submit a USB drive that contains:

- A virtual machine (VM) file or files containing the executable batch-mode program or script described above.
- A CSV file containing matrix S that is the output from running the executable software on input sets DC and DR .
- An ASCII text file containing the participant's answers to the PIR Pilot questionnaire.

The VM should open to a console or desktop with an open command window whose working directory contains the command-line-executable program.

The name of the executable should begin with the characters "pir", and no other files in the directory should begin with those characters.

The USB drive should be sent (UPS, Fedex, or USPS overnight preferred) to:

Pill Image Recognition Pilot
Computational Photography Project for Pill Identification (C3PI)
National Library of Medicine
Building 38A, Room B1-N30
8600 Rockville Pike
Bethesda, MD 20894

ABOUT THE VM

The VM can have any of the following 64-bit operating systems:

- CentOS 6.x
- CentOS 7.0
- Fedora 20
- Mac OS X 10.10.x

- Mac OS X 10.8.x
- Mac OS X 10.9.x
- Red Hat Enterprise Linux 6.x
- Red Hat Enterprise Linux 7.0
- Ubuntu 14.x
- Windows 7 SP1
- Windows 8.x

ABOUT NLM'S EVALUATION SYSTEM

NLM's evaluation system will be run with the following configuration:

- Apple Mac Pro
- 3.7GHz Intel Xeon E5 quad-core CPU, 10MB of L3 cache
- 64GB (4 x 16GB) of 1866MHz DDR3 ECC RAM
- 512GB PCIe-based flash storage
- Mac OS X 10.10.x
- VMware Fusion 7.0

Submissions must be in a format that is compatible with VMware Fusion 7.0, which can import virtual machines created with the following:

- Microsoft Virtual PC for Mac Version 7
- Open Virtualization Format (OVF and OVA) 1.0
- Parallels Desktop for Mac 2.5, 3.0, 4.0, 5.0, and 6.0
- Parallels Server 3.0

VMware and other third-party products have tools to convert between virtual machine file formats.

WHAT NLM WILL DO WITH THE PIR PILOT SUBMISSIONS

NLM will exercise the executable software on several input data sets. The first will be the PIR Pilot data sets DC and DR; NLM will check that the output matrix is the same as the submitted matrix S. NLM will post information about the PIR Pilot responses to pir.nlm.nih.gov no earlier than May 11, 2015.

These exercises will be conducted on a system that is disconnected from all Internet services. All submissions must therefore be self-contained.

NLM will also exercise the executable on at least one pair of input data sets each of which contains a single image, and NLM will exercise the executable on input data not in either DC or DR.

NLM will use these results to test, evaluate, and as needed refine the components of the PIR Challenge. The software and output submitted in response to the PIR Pilot will *not* themselves be judged.

EXAMPLES

Let directory DC contain the following two files of consumer images:

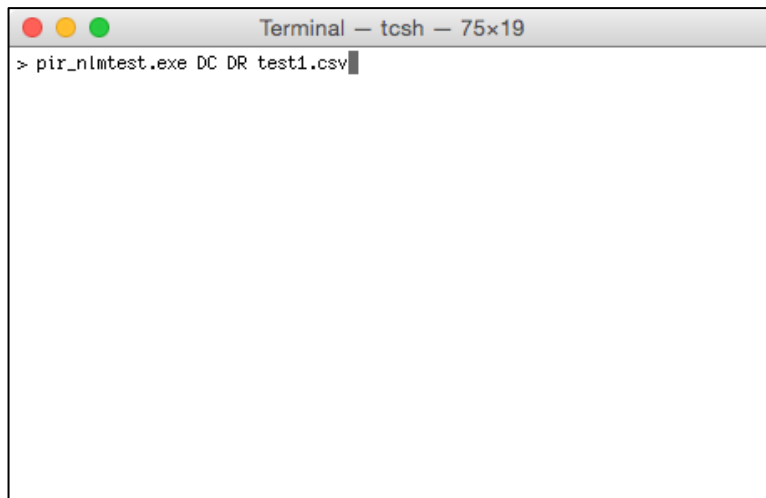
- 10.jpg
- 2.jpg

Note: In lexicographical order, file 10.jpg comes before 2.jpg.

Let directory DR contain the following two files of reference images:

- 00093-7155-98_NLMIMAGE10_4A21A50D.jpg
- 68180-0481-01_NLMIMAGE10_FA217D1B.jpg

Let the executable be in a file named “pir_nlmtest.exe”. One invokes the program with the output file “test1.csv” with the following command:

A screenshot of a terminal window titled "Terminal - tcsh - 75x19". The terminal shows a command prompt ">" followed by the command "pir_nlmtest.exe DC DR test1.csv" with a cursor at the end of the line.

```
Terminal - tcsh - 75x19
> pir_nlmtest.exe DC DR test1.csv
```

After execution, “test1.csv” contains:

```
,00093-7155-98_NLMIMAGE10_4A21A50D.jpg,68180-0481-01_NLMIMAGE10_FA217D1B.jpg
10.jpg,0.7,0.954334
2.jpg,0.457,0.6100923
```

This output is interpreted as the following similarity matrix:

	00093-7155-98_NLMIMAGE10_4A21A50D.jpg	68180-0481-01_NLMIMAGE10_FA217D1B.jpg
10.jpg	0.7	0.954334
2.jpg	0.457	0.6100923

Example of an Apple MacOSX command:

```
pir_besttest /Volumes/pir/con1 /Volumes/pir/ref1 /Volumes/pir/out/test1.csv
```

Example of a Windows command:

```
C:\> pirmytest.exe Z:\test1\consumer Z:\test1\reference Z:\output\test1.csv
```

QUESTIONS AND COMMENTS

Questions and comments about these instructions are welcome at pir@nlm.nih.gov.