

# Mahotas: Open source software for scriptable computer vision

Luis Pedro Coelho

Lane Center for Computational Biology, Carnegie Mellon University, Pittsburgh, USA  
Instituto de Medicina Molecular, Lisboa, Portugal

August 30, 2012

## Abstract

Mahotas is a computer vision library for Python. It contains traditional image processing functionality such as filtering and morphological operations as well as more modern computer vision functions for feature computation, including interest point detection and local descriptors.

The interface is in Python, a dynamic programming language, which is very appropriate for fast development, but the algorithms are implemented in C++ and are tuned for speed.

Mahotas is available under a liberal open source license (MIT License) and is available from <http://github.com/luispedro/mahotas> or the Python Package Index (<http://pypi.python.org/pypi/mahotas>).

**Keywords:** computer vision, image processing.

## 1 Introduction

Mahotas is a computer vision library for the Python Programming Language (versions 2.5 and up, including version 3). It operates on numpy arrays (Walt et al., 2011). Therefore, it uses all the infrastructure built by that project for storing information and performing basic manipulations and computations. In particular, unlike libraries written in the C Language or in Java (Marcel and Rodriguez, 2010), Mahotas does not need to define a data structure, but uses the numpy array structure. Many basic manipulation functionality that would otherwise be part of a computer vision library are handled by numpy, for example computing averages and other simple statistics, handling multi-channel images, converting between types (integer and floating point images are supported by mahotas, whenever it is meaningful). For the user, this has the additional advantage that they do not need to learn yet another set of functions.

Additionally, by integrating into the Python numeric ecosystem, users can use other packages in a seamless way. In particular, mahotas does not implement any machine learning functionality, but rather advises the user to use another, specialised package, such as scikits-learn or milk.

Python is a natural “glue” language: it is easy to use state-of-the-art libraries written in multiple languages (Oliphant, 2007). Mahotas itself is a mix of high-level Python and low-level C+. This achieves a good balance between speed and ease of implementation.

Version 1.0 of mahotas has been released recently and this is now a mature, well-tested package (the first version was made available over 4 years ago).

<sup>1</sup> Mahotas runs and is used on different versions of Unix (including Linux, SunOS, and FreeBSD), Mac OS X, and Windows.

## 2 Implementation and Architecture

### 2.1 Interface

The interface is a procedural interface, with no global state. All functions work independently of each other (there is code sharing at the implementation level, but this is hidden from the user).

The main functionality is grouped into the following categories:

**SURF** Speeded-up Robust Features (Bay et al., 2008). This includes both key-point detection and descriptor computation.

**features** Global feature descriptors. In particular, Haralick texture features, Zernike moments, local binary patterns, and threshold adjacency statistics (both the original (Hamilton et al., 2007) and the parameter-free versions (Coelho, Ahmed, et al., 2010)).

**wavelet** Haar and Daubechies wavelets. Forward and inverse transforms.

**morphological functions** Erosion and dilation, as well as some more complex operations built on these. There are both binary and grayscale implementations of these operators.

**watershed** seeded watershed and distance map transforms (Felzenszwalb and Huttenlocher, 2004).

**filtering** Gaussian filtering and general convolutions.

**polygon operations** convex hull, polygon drawing.

---

<sup>1</sup>Note for reviewers: the version currently available implements all functionality described in this manuscript. I will release it as version 1.0 when this manuscript is accepted to coincide with publication. Naturally, any bugs that are found and reported in the meanwhile will be addressed.

In general, functions operate on all data types. This is performed without any extra memory copies. Mahotas is heavily optimised for both speed and memory usage (it can be used with very large arrays).

There are a few interface conventions which apply to many functions. When meaningful, a structuring element is used to define neighbourhoods or adjacency relationships (morphological functions, in particular, use this convention). Generally, the default is to use a  $3 \times 3$  cross as the default if no structuring filter is given (the exception to this rule is the median filter, where the default is a  $3 \times 3$  square).

When a new image is to be returned, functions take an argument named `out` where the output will be stored. This argument is often much more restricted in type. In particular, it must often be a contiguous array.<sup>2</sup> Since this is a performance feature (its purpose is to avoid extra memory allocation), it is natural that the interface is less flexible (accessing a contiguous array is much more efficient than a non-contiguous one).

## 2.2 Example of Use

This is a simple example of mahotas usage. Code for this and other examples is present in the mahotas source distribution under the `demos/` directory. In this example, we load an image and find SURF interest point and descriptors on it.

We start by importing the necessary packages, including numpy and mahotas. We also use *milk*, to demonstrate how the mahotas output can integrate with a machine learning package.

```
import numpy as np
import mahotas
from mahotas.features import surf
import milk
```

The first step is to load the image and convert to 8 bit numbers. In this case, the conversion is done using standard numpy methods, namely `astype`:

```
f = mahotas.imread('luispedro.jpg', as_grey=True)
f = f.astype(np.uint8)
```

We can now compute SURF interest points and descriptors.

```
spoints = surf.surf(f, 4, 6, 2)
```

Using numpy operations and milk, we can select only the descriptors and cluster them into a group of five:

```
descrs = spoints[:, 6:]
values, _ = milk.kmeans(descrs, 5)
```

Finally, we can show the points in colour.

---

<sup>2</sup>Numpy supports non-contiguous arrays, which are most often slices into other, larger, contiguous arrays (e.g., given a  $128 \times 128$  contiguous array, one can build a  $64 \times 128$  non-contiguous array by taking every other row).




Figure 1: Example of Usage. On the left, the original image is shown, while on the right SURF detections are represented as rectangles of different colours.

```
colors = np.array(  
    [ 255,  25,   1],  
    [203,  77,  37],  
    [151, 129,  56],  
    [ 99, 181,  52],  
    [ 47, 233,   5])  
f2 = surf.show_surf(f, spoints[:64], values, colors)
```

The `show_surf` only builds the image as a multi-channel (one for each colour) image. Using `matplotlib` (Hunter, 2007), we finally display the image as Figure 1.

```
from matplotlib import pyplot as plt  
plt.subplot(1,2,1)  
plt.imshow(f)  
plt.subplot(1,2,2)  
plt.imshow(f2)
```

The easy interaction with `matplotlib` is another way in which we benefit from the `numpy`-based ecosystem.

### 2.3 Implementation

Mahotas is mostly written in C++, but this is completely hidden from the user as there are hand-written Python wrappers for all functions (automatically generated wrappers inevitably lead to worse error messages and are less flexible).

The main reason that *mahotas* is implemented in C++ (and not in C, which is the language of the Python interpreter) is to use templates. Almost C++ functionality is split across 2 functions:

1. A `py_function` which uses the Python C API to get arguments and check them.
2. A `template function<dtype>` which works for the type `dtype` performing the actual operation.

So, for example, this is how *erode* is implemented. `py_erode` consists mostly of boiler-plate code:

```
PyObject* py_erode(PyObject* self, PyObject* args) {
    PyArrayObject* array;
    PyArrayObject* Bc;
    PyArrayObject* output;
    if (!PyArg_ParseTuple(args, "OOO", &array, &Bc, &output) ||
        !numpy::are_arrays(array, Bc, output) || !numpy::same_shape(array, output) ||
        !PyArray_EquivTypenums(PyArray_TYPE(array), PyArray_TYPE(Bc)) ||
        !PyArray_EquivTypenums(PyArray_TYPE(array), PyArray_TYPE(output)) ||
        PyArray_NDIM(array) != PyArray_NDIM(Bc)
    ) {
        PyErr_SetString(PyExc_RuntimeError, TypeErrorMsg);
        return NULL;
    }
    holdref r_o(output);

#define HANDLE(type) \
    erode<type>(numpy::aligned_array<type>(output), \
               numpy::aligned_array<type>(array), \
               numpy::aligned_array<type>(Bc));
    SAFE_SWITCH_ON_INTEGER_TYPES_OF(array, true);
#undef HANDLE
    ...
}
```

This functions retrieves the arguments, performs some sanity checks, performs a bit of initialisation, and finally, switches in the input type with the help of the `SAFE_SWITCH_ON_INTEGER_TYPES()` macro, which call the right specialisation of the template that does the actual work. In this example *erode* implements (binary) erosion:

```
template<typename T>
void erode(numpy::aligned_array<T> res,
           numpy::aligned_array<T> array,
           numpy::aligned_array<T> Bc) {
    gil_release nogil;
    const int N = res.size();
    typename numpy::aligned_array<T>::iterator iter = array.begin();
    filter_iterator<T> filter(array.raw_array(), Bc.raw_array(), EXTEND_NEAREST, is_bool);
    const int N2 = filter.size();
}
```

```

T* rpos = res.data();

for (int i = 0;
     i != N;
     ++i, ++rpos, filter.iterate_both(iter)) {
    T value = std::numeric_limits<T>::max();
    for (int j = 0; j != N2; ++j) {
        T arr_val = T();
        filter.retrieve(iter, j, arr_val);
        value = std::min<T>(value, erode_sub(arr_val, filter[j]));
    }
    *rpos = value;
}
}

```

The template machinery makes the functions that use it very simple and easy to read. The only downside is that there is some expansion of code size when the compiler instantiates the function for the several integer and floating point types. Given the small size of these functions, this is not a big issue.

In the snippet above, you can see some other C++ machinery:

**gil\_release** This is a “resource-acquisition is object initialisation” (RAII)<sup>3</sup> object that release the Python global interpreter lock (GIL)<sup>4</sup> in its constructor and gets it back in its destructor. Normally, the template function will release the GIL after the Python-specific code is done. This allows several mahotas functions to run concurrently.

**array** This is a thin wrapper around `PyArrayObject` that knows its data type and has iterators which resemble the C++ standard library. This makes the code type-safer. This is also a RAII object in terms of managing Python reference counts.

**filter\_iterator** This is taken from `scipy.ndimage` and it is useful to iterate over an image and use a centered filter around each pixel (it keeps track of all of the boundary conditions).

The inner loop is as direct an implementation of erosion as one would wish for: for each pixel in the image, look at its neighbours, subtract the filter value, and compute the minimum of this operation.

## 2.4 Efficiency

Table 1 shows timings for different operations. These were normalized to multiples of the time it takes to go over the image and find its maximum pixel

<sup>3</sup>RAII is a design pattern in C++, or other languages with scope linked deterministic object destruction, such as D, where a resource is represented by an object, whose constructor acquires it and whose destructor releases it. This guarantees that the object is correctly released even if the scope is left through an exception (Stroustrup, 1994).

<sup>4</sup>In the CPython interpreter, the most commonly used implementation of Python, there is a global lock for many Python related functionality, which limits parallelism.

Operation	mahotas	pymorphy
erode	9.90	13.06
dilate	10.48	8.17
open	20.26	21.05
cwatershed	181.47	50211.56
center mass	6.39	NA
daubechies	16.92	NA
haralick	276.90	NA

Table 1: Efficiency Results for mahotas and pymorph. Shown are values as multiples of the time that `numpy.max(f)` takes.

value (this was done using `numpy.max(image)`). The measurements shown were obtained on an Intel 64 bit system, running Ubuntu Linux. However, due to the normalization, measurements obtained on another system (Intel 32 bits running Mac OS) were qualitatively similar (data not shown).

The comparison is against Pymorph (Dougherty and Lotufo, 2003), which is a pure Python implementation of some of the same functions. We can see that simple morphological operations (erosion, dilation, opening) take a similar amount of time in mahotas and pymorph (pymorph intelligently uses numpy operations for these). However, some more complex operations such as watershed are much faster in mahotas.

In keeping with the philosophy of blending in with the ecosystem, Mahotas uses the standard Python build machinery and distribution channels. Building and installing from source code is done using

```
python setup.py install
```

Alternatively, Python based package managers (such as `easy_install` or `pip`) can be used (mahotas works well with these systems).

## 2.5 Quality Control

Mahotas includes a complete automated suite of unit tests. These all functionality and include several regression tests. There are no known bugs in version 1.0. In fact, no releases have ever been performed with known bugs. Naturally, bugs were, occasionally, discovered in released versions, but corrected before the next release.

The development is completely open-source and development versions are available. Many users have submitted bug reports and fixes.

## 3 Availability

### Operating system

Mahotas runs and is used on different versions of Unix (including Linux, SunOS, and FreeBSD), Mac OS X, and Windows.<sup>5</sup>

### Programming language

Mahotas works in Python (minimal version is 2.5 and works in all more recent versions, including in the Python 3 series).

### Additional system requirements

None.

### Dependencies

It requires numpy to be present and installed.

### List of contributors

Luis Pedro Coelho (Carnegie Mellon University and Instituto de Medicina Molecular), Zachary Pincus (Stanford University), Peter J. Verveer (European Molecular Biology Laboratory), Davis King (Northrop Grumman ES), Robert Webb (Carnegie Mellon University), Matthew Goodman (University of Texas at Austin), K.-Michael Aye (University of Bern), Rita Simões (University of Twente), Joe Kington (University of Wisconsin), Christoph Gohlke (University of California, Irvine), and Sandro Knauss (University of Bremen).

### 3.0.1 Software location

#### Code repository

*Name:* Github

*Identifier:* <https://github.com/luispedro/mahotas>

*Licence:* MIT

*Date published:* Since 2009

## 4 Reuse Potential

Originally, this code was developed in the context of cellular image analysis. However, none of the functionality is specific to this context and many computer vision pipelines can make use of it.

This package (and earlier versions of it) have been used by myself (Coelho, Peng, et al., 2010; Coelho, Shariff, et al., 2009) and close collaborators in several publications (Cho et al., 2012). Other groups have used in published work, both in cell image analysis (Mashburn et al., 2012) and in other areas (Machlek and Oleviov, 2013).

---

<sup>5</sup>Christoph Gohlke has been instrumental in providing Windows packages as well as several fixes for that platform.



## 5 Discussion

Python is an excellent language for scientific programming because of the inherent properties of the language and because of the infrastructure that has been built around the numpy project. Mahotas works in this environment to provide the user with image analysis and computer vision functionality.

Mahotas does not include machine learning related functionality, such as  $k$ -means clustering or any classification. This is the result of an explicit design decision. Specialised machine learning packages for Python already exist (Demar et al., 2004; Pedregosa et al., 2011; Schaul et al., 2010; Sonnenburg et al., 2010). A good classification (or other) system can benefit both computer vision users and others. As these projects all use Numpy arrays as their data types, it is easy to use functionality from the different project seamlessly (no copying of data is necessary).

Python is an ideal language for fast development of both applications and scientific software. For this platform, Mahotas is a fast library of computer vision and image processing functions. It is implemented in C++, as the standard Python interpreter is too slow for a direct Python implementation. However, all of the Python interface code is hand-written, as opposed to using automatic interface generators like Swig Beazley, 2003. This is more work, but the end result is of much higher quality, especially when it comes to giving useful error messages (e.g., when a type mismatch occurs, an automatic system will often be forced to resort to a generic message as it does not have any knowledge of what the arguments mean besides their automatically inferred types).

Mahotas has been available in the Python Package Index since April 2010 and has been downloaded over 20,000 times. This does not include any downloads from other sources. Mahotas includes a full test suite. There are no known bugs.

## Acknowledgements

Mahotas includes code ported and incorporated from other projects. In particular, the SURF implementation is a port from the code from *dlib*,<sup>6</sup> a very good C++ library by Davis King. I also gleaned some insight into the implementation of these features from Christopher Evan’s OpenSURF library and its documentation (Evans, 2009).<sup>7</sup> The code which interfaces with the FreeImage library, was written by Zachary Pincus and some of the support code was written by Peter J. Verveer for the `scipy.ndimage` project. All of these contributions were integrated while respecting the software licenses under which the original code had been released. Robert Webb, a summer student at Carnegie Mellon University, worked with me on the initial local binary patterns implementation. Finally, I thank the several users who have reported bugs, submitted

---

<sup>6</sup>Dlib’s webpage is at <http://dlib.net>.

<sup>7</sup>OpenSURF is available at <http://www.chrisevansdev.com/computer-vision-opensurf.html>, where several documents describe details of the implementation.

small fixes, and participated on the project mailing list.

**Funding:** I was supported in my work by the Fundação para a Ciência e Tecnologia (grants SFRH/BD/37535/2007 and PTDC/SAU-GMG/115652/2008) and by a grant from the Siebel Scholars Foundation.

## References

- Bay, H., A. Ess, T. Tuytelaars, and L. van Gool (June 2008). “Speeded-up Robust Features (SURF)”. In: *Computer Vision and Image Understanding (CVIU)* 110.3, pp. 346–359 (cit. on p. 2).
- Beazley, D.M. (2003). “Automated scientific software scripting with SWIG”. In: *Future Generation Computer Systems* 19.5. `ce:title;Tools for Program Development and Analysis. Best papers from two Technical Sessions, at ICCS2001, San Francisco, CA, USA, and ICCS2002, Amsterdam, The Netherlands;ce:title;` pp. 599–609. ISSN: 0167-739X. DOI: 10.1016/S0167-739X(02)00171-1. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X02001711> (cit. on p. 9).
- Cho, Baek Hwan, Ivan Cao-Berg, Jennifer Ann Bakal, and Robert F Murphy (July 2012). “OMERO.searcher: content-based image search for microscope images”. In: *Nature Methods*, pp. 633–634. URL: <http://dx.doi.org/10.1038/nmeth.2086> (cit. on p. 8).
- Coelho, Luis Pedro, Amr Ahmed, Andrew Arnold, Joshua Kangas, Abdul-Saboor Sheikh, Eric P. Xing, William W. Cohen, and Robert F. Murphy (Jan. 2010). “Structured Literature Image Finder: Extracting Information from Text and Images in Biomedical Literature.” In: *Lecture notes in computer science* 6004, pp. 23–32. ISSN: 0302-9743. DOI: 10.1007/978-3-642-13131-8\_4. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2901994&tool=pmcentrez&rendertype=abstract> (cit. on p. 2).
- Coelho, Luis Pedro, Tao Peng, and Robert F. Murphy (2010). “Quantifying the distribution of probes between subcellular locations using unsupervised pattern unmixing”. In: *Bioinformatics* 26.12, pp. i7–i12. DOI: 10.1093/bioinformatics/btq220. eprint: <http://bioinformatics.oxfordjournals.org/content/26/12/i7.full.pdf+html>. URL: <http://bioinformatics.oxfordjournals.org/content/26/12/i7.abstract> (cit. on p. 8).
- Coelho, Luis Pedro, Aabid Shariff, and Robert F. Murphy (2009). “Nuclear segmentation in microscope cell images: A hand-segmented dataset and comparison of algorithms”. In: *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, pp. 518–521. ISBN: 978-1-4244-3931-7. DOI: 10.1109/ISBI.2009.5193098. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5193098> (cit. on p. 8).
- Demar, Janez, Bla Zupan, Gregor Leban, and Tomaz Curk (2004). “Orange: From Experimental Machine Learning to Interactive Data Mining”. In: *Knowledge Discovery in Databases: PKDD 2004*. Ed. by Jean-Francois Boulicaut, Flori-

- ana Esposito, Fosca Giannotti, and Dino Pedreschi. Vol. 3202. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, pp. 537–539. ISBN: 978-3-540-23108-0. DOI: [10.1007/978-3-540-30116-5\\_58](https://doi.org/10.1007/978-3-540-30116-5_58) (cit. on p. 9).
- Dougherty, Edward R. and Roberto A. Lotufo (2003). *Hands-on Morphological Image Processing*. Bellingham, WA: SPIE Press. ISBN: 0-8194-4720-X (cit. on p. 7).
- Evans, Christopher (2009). “Notes on the OpenSURF Library SURF : Speeded Up Robust Features”. In: 1 (cit. on p. 9).
- Felzenszwalb, Pedro and Daniel Huttenlocher (Sept. 2004). *Distance Transforms of Sampled Functions*. Tech. rep. Cornell University. URL: <http://hdl.handle.net/1813/5663> (cit. on p. 2).
- Hamilton, Nicholas A., Radosav S. Pantelic, Kelly Hanson, and Rohan D. Teasdale (2007). “Fast automated cell phenotype image classification.” In: *BMC bioinformatics* 8, p. 110. ISSN: 1471-2105. DOI: [10.1186/1471-2105-8-110](https://doi.org/10.1186/1471-2105-8-110). URL: <http://www.ncbi.nlm.nih.gov/pubmed/17394669> (cit. on p. 2).
- Hunter, John D. (2007). “Matplotlib: A 2D Graphics Environment”. In: *Computing in Science and Engineering* 9, pp. 90–95. ISSN: 1521-9615. DOI: <http://doi.ieeecomputersociety.org/10.1109/MCSE.2007.55> (cit. on p. 4).
- Machlek, Tom and Kamila Oleviiov (2013). “Decentralized Multi-Agent Algorithm for Translational 2D Image Alignment”. In: *Multimedia and Internet Systems: Theory and Practice*. Ed. by Aleksander Zgrzywa, Kazimierz Choro, and Andrzej Siemiski. Vol. 183. Advances in Intelligent Systems and Computing. Springer Berlin Heidelberg, pp. 15–24. ISBN: 978-3-642-32335-5 (cit. on p. 8).
- Marcel, Sbastien and Yann Rodriguez (2010). “Torchvision the machine-vision package of torch”. In: *Proceedings of the international conference on Multimedia*. MM '10. Firenze, Italy: ACM, pp. 1485–1488. ISBN: 978-1-60558-933-6. DOI: [10.1145/1873951.1874254](https://doi.org/10.1145/1873951.1874254). URL: <http://doi.acm.org/10.1145/1873951.1874254> (cit. on p. 1).
- Mashburn, David N., Holley E. Lynch, Xiaoyan Ma, and M. Shane Hutson (2012). “Enabling user-guided segmentation and tracking of surface-labeled cells in time-lapse image sets of living tissues”. In: *Cytometry Part A* 81A.5, pp. 409–418. ISSN: 1552-4930. DOI: [10.1002/cyto.a.22034](https://doi.org/10.1002/cyto.a.22034). URL: <http://dx.doi.org/10.1002/cyto.a.22034> (cit. on p. 8).
- Oliphant, Travis E. (2007). “Python for Scientific Computing”. In: *Computing in Science and Engineering* 9, pp. 10–20. ISSN: 1521-9615. DOI: <http://doi.ieeecomputersociety.org/10.1109/MCSE.2007.58> (cit. on p. 2).
- Pedregosa, Fabian, Gal Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and douard Duchesnay (Nov. 2011). “Scikit-learn: Machine Learning in Python”. In: *J. Mach. Learn. Res.* 999888, pp. 2825–2830. ISSN: 1532-4435. URL: <http://dl.acm.org/citation.cfm?id=2078183.2078195> (cit. on p. 9).

- Schaul, Tom, Justin Bayer, Daan Wierstra, Yi Sun, Martin Felder, Frank Sehnke, Thomas Rckstie, and Jrgen Schmidhuber (Mar. 2010). "PyBrain". In: *J. Mach. Learn. Res.* 11, pp. 743–746. ISSN: 1532-4435. URL: <http://dl.acm.org/citation.cfm?id=1756006.1756030> (cit. on p. 9).
- Sonnenburg, Sren, Gunnar Rtsch, Sebastian Henschel, Christian Widmer, Jonas Behr, Alexander Zien, Fabio de Bona, Alexander Binder, Christian Gehl, and Vojtch Franc (Aug. 2010). "The SHOGUN Machine Learning Toolbox". In: *J. Mach. Learn. Res.* 11, pp. 1799–1802. ISSN: 1532-4435. URL: <http://dl.acm.org/citation.cfm?id=1756006.1859911> (cit. on p. 9).
- Stroustrup, Bjarne (Jan. 1994). *The design and evolution of C++*. Addison-Wesley Professional. ISBN: 0-201-54330-3. URL: <http://portal.acm.org/citation.cfm?id=193198> (cit. on p. 6).
- Walt, S. van der, S.C. Colbert, and G. Varoquaux (Mar. 2011). "The NumPy Array: A Structure for Efficient Numerical Computation". In: *Computing in Science Engineering* 13.2, pp. 22–30. ISSN: 1521-9615. DOI: 10.1109/MCSE.2011.37 (cit. on p. 1).