

Building a Community around GNU Octave

John W. Eaton

NIPS Workshop on Machine Learning Open Source Software 2008
Whistler, BC
12 December 2008

- Brief discussion of Octave and its history
- Expanding the community of Octave developers
- Challenges for the future

Acknowledgment

First, I should thank some people who have made Octave possible.

- Jim Rawlings
- All the volunteers who have worked on Octave over the years
- Richard Stallman and the GNU project

Acknowledgment

First, I should thank some people who have made Octave possible.

- Jim Rawlings
- All the volunteers who have worked on Octave over the years
- Richard Stallman and the GNU project

Acknowledgment

First, I should thank some people who have made Octave possible.

- Jim Rawlings
- All the volunteers who have worked on Octave over the years
- Richard Stallman and the GNU project

What is Octave?

Octave is a free (GNU GPL) interactive system for numerical computations with a language that is mostly compatible with `MATLAB`¹

Features:

- N-d arrays and linear algebra
- Nonlinear equations
- Differential equations
- Image processing
- Signal processing
- Statistics
- Polynomials
- Sparse matrices
- Special functions
- Control theory
- Audio
- Graphics
- Finance
- Much more!

¹`MATLAB` is a registered trademark of The MathWorks, Inc.

Why Octave?

- **No license manager.** Octave is free software!
- **No black boxes.** The complete source code is available. You are encouraged to look at the internals.
- **Octave is portable** to POSIX systems with a standard C++ compiler. Octave can also be ported to other unusual systems (**you** have the source!).
- **Large support community.** Octave has an active support community with more than 1000 messages per month posted to the mailing lists.

- We did **NOT** start out to write a “MATLAB clone” (given some degree of compatibility, users wanted more).
- Octave was originally conceived as companion software for a textbook on chemical reactor design.
- We chose to name it after Octave Levenspiel, a pioneer in the field of chemical reaction engineering.
- In addition to writing some software, we hoped to create a community of Octave developers and users.

Development Timeline

- ??/1989 First discussions about textbook and software
- 2/1992 Development begins
- 1/1993 First announcement on web (version 0.60)
- 2/1994 First real release (ready for wider distribution?)
- 12/1996 Second major version (2.0) port to Windows (Cygwin)
- 3/1998 2.1 development branch
- 11/2004 2.9 branch in preparation for 3.0 release
- 12/2007 Version 3.0, major upgrade
- ??/2008 Version 3.2, improved graphics and MATLAB-style objects

Sources

- Releases available from `ftp.octave.org`
- Development version available via a Mercurial archive

Binary packages

- Included in Debian, part of Fedora Extras for Red Hat, etc.
- Mac OS X
- Windows

See `www.octave.org` for links to binary distributions.

Octave Forge (`octave.sf.net`)

- Collection of contributed functions.
- Originally created and managed by Paul Kienzle (NIST).
- Hosted at SourceForge.
- Place on the web for collaborative development.
- Currently includes approximately 70 domain-specific packages.

How Many Users?

Because Octave may be freely redistributed, it's impossible to know how many people are using Octave. My guess is in the hundreds of thousands.

How Many Users?

Because Octave may be freely redistributed, it's impossible to know how many people are using Octave. My guess is in the hundreds of thousands.

Some available data:

Octave web and ftp sites:

- More than **29,000** downloads of the Octave 3.0.0 sources from `ftp.octave.org` in **4 months**.
- Compare with 15,000 downloads of 2.1.71 over 5 months in 2005.

Octave Forge (contributed code):

- More than **65,000** downloads of the binary package of Octave 3.0.0 for Windows in **4 months**.
- Compare with 70,000 downloads of 2.1.50 binary over 27 months.
- About **280,000** page views per month (total for site; up from about 80,000 in 2005).

How Many Users?

Because Octave may be freely redistributed, it's impossible to know how many people are using Octave. My guess is in the hundreds of thousands.

Some available data:

Octave web and ftp sites:

- More than **29,000** downloads of the Octave 3.0.0 sources from `ftp.octave.org` in **4 months**.
- Compare with 15,000 downloads of 2.1.71 over 5 months in 2005.

Octave Forge (contributed code):

- More than **65,000** downloads of the binary package of Octave 3.0.0 for Windows in **4 months**.
- Compare with 70,000 downloads of 2.1.50 binary over 27 months.
- About **280,000** page views per month (total for site; up from about 80,000 in 2005).

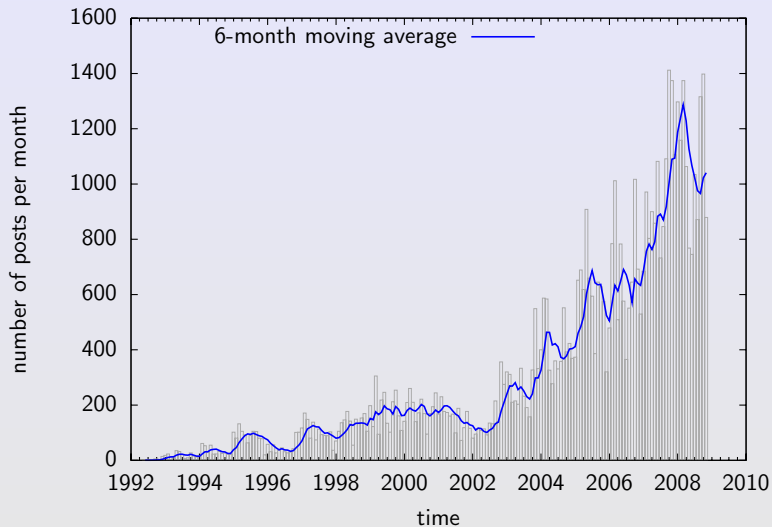
Postings to the primary Octave mailing lists for the last 12 months:

	subscribers	postings
<code>help@octave.org</code>	775	5300
<code>bug@octave.org</code>	185	3650
<code>maintainers@octave.org</code>	190	4450

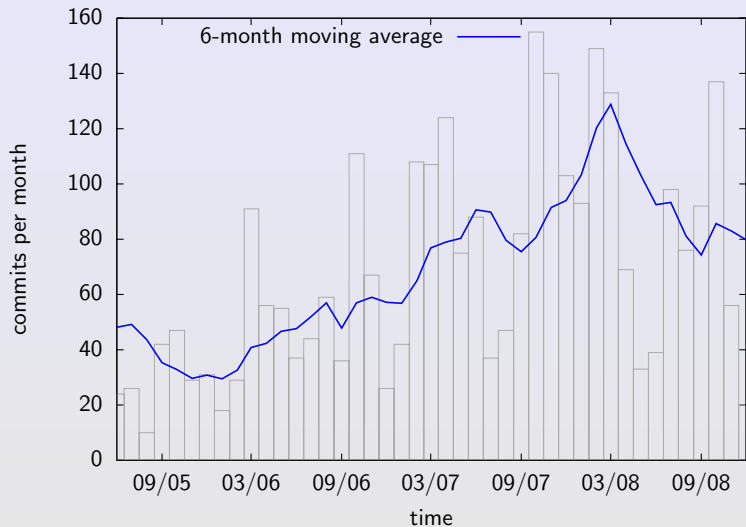
- Participation is increasing
- Overall traffic has approximately doubled in the last two years

Mailing List Activity

Posts to the primary Octave mailing lists (all time)



Commits to the Octave source archive per month (last three years)



Who is using Octave and what are they doing with it?

Primary uses include research, teaching, and industrial activities:

- Analyzing EEG data
- Atomic and plasma physics
- Bioinformatics
- Computer vision
- Control system design
- Data analysis
- Filter design
- Finite elements
- Fish stock assessment
- High school math
- Music
- Parallel processing
- Robotics
- Signal processing
- Teaching programming
- Telecommunications
- Ultrasound imaging

And we were originally just hoping that it could be used to solve a few chemical reactor design problems!

Version 3.0 Includes Many Significant Improvements

Major new features in 3.0 include:

- Reasonable start on MATLAB-compatible graphics
- Sparse matrix data type
- MATLAB-compatible integer data types
- Optimization functions (glpk, qp, sqp)
- New package system
- Native Windows port
- Updated user guide

Octave's interface to `gnuplot` (www.gnuplot.info) was a quick way to get graphics in Octave, but for many users it was a weak point (especially those who are familiar with the current graphics in `MATLAB`).

In 3.0, Octave provides more `MATLAB`-compatible graphics capabilities.

- Octave manages the data for the graphics system
- External packages will manage the actual plotting
- No direct interface to `gnuplot`

The upcoming 3.2 release will have more improvements to the graphics system.

Søren Hauberg and David Bateman have implemented a package system for Octave that allows the equivalent of “toolboxes” for Octave.

- Packages maintained and released independently of Octave
- Separate, smaller packages
- Easier maintenance
- More timely releases

For the Next Major Release

- Object-oriented features (MATLAB classes) (jwe)
- New OpenGL-based graphics renderer (Michael Goffioul and Shai Ayal)
- Single-precision data type (David Bateman)
- Improved array indexing (Jaroslav Hajek)
- 64-bit integer arithmetic (Jaroslav Hajek)
- Optimization of expressions like $a' * b$ (Jaroslav Hajek)
- Diagonal and permutation matrix optimization (Jaroslav Hajek)
- Additional improvements in MATLAB compatibility

Much Work Remains

- More graphics improvements
- IDE
- GUI tools
- Profiler
- Compiler (Octave to C++? Just-In-Time?)
- More and improved domain-specific packages
- Additional changes for MATLAB compatibility
- ...

Challenges for Free Software Projects

- Educating users
- Building a developer community
- Funding

As the user community grows, we must educate new users so they will see themselves as members of the Octave community rather than as consumers of Octave.

- We have a diverse user community
- Many users come to Octave expecting a “MATLAB clone”
- Many users expect to get Octave for free
- Some seem to also expect fast, expert support (also for free!)

The Octave Developer Community

- We have a diverse developer community
- Currently growing but still rather small (we'd love to have more participation)
- Most contributors make a few small changes and move on
- We must work to include more people in the community

Ways You Can Help

- Bug tracking
- Release management
- Web site work
- Mailing list maintenance
- Support users via the `help@octave.org` mailing list
- Core language interpreter development
- Create domain-specific packages: signal processing, image processing, control systems design, graphics, GUI development, etc.
- ...

- We have hundreds of thousands of users but very little direct funding. Is that what should be happening?
- Free software is not a charity project.
- Having a large user community expecting to get something for free from a small developer community is not sustainable.

- Independently wealthy developer
- Research project/government grant
- A few large donations from a small number of benefactors
- Many small donations from users
- Selling support contracts

- Octave is a widely used system for numerical computations
- Many exciting changes are currently underway
- Some of the development processes need work (bug tracking, etc.)
- Ensuring that development continues is a continual challenge

Recent Active Contributors

- David Bateman
- Michael Goffioul
- Jaroslav Hajek
- Søren Hauberg
- Ben Abbott
- Francesco Portotì
- Thorsten Meyer
- Thomas Weber
- Rafael Laboissiere
- Brian Gough
- Shai Ayal
- Your Name Here!