Rock solid, brand new, everyday, for free, not a joke: NeuroDebian

Michael Hanke

Debian Developer
Otto-von-Guericke University, Magdeburg

Max-Planck-Institute for Human Cognitive & Brain Sciences, Leipzig

Jan 31th 2012

The research software experience is not as great as it could be

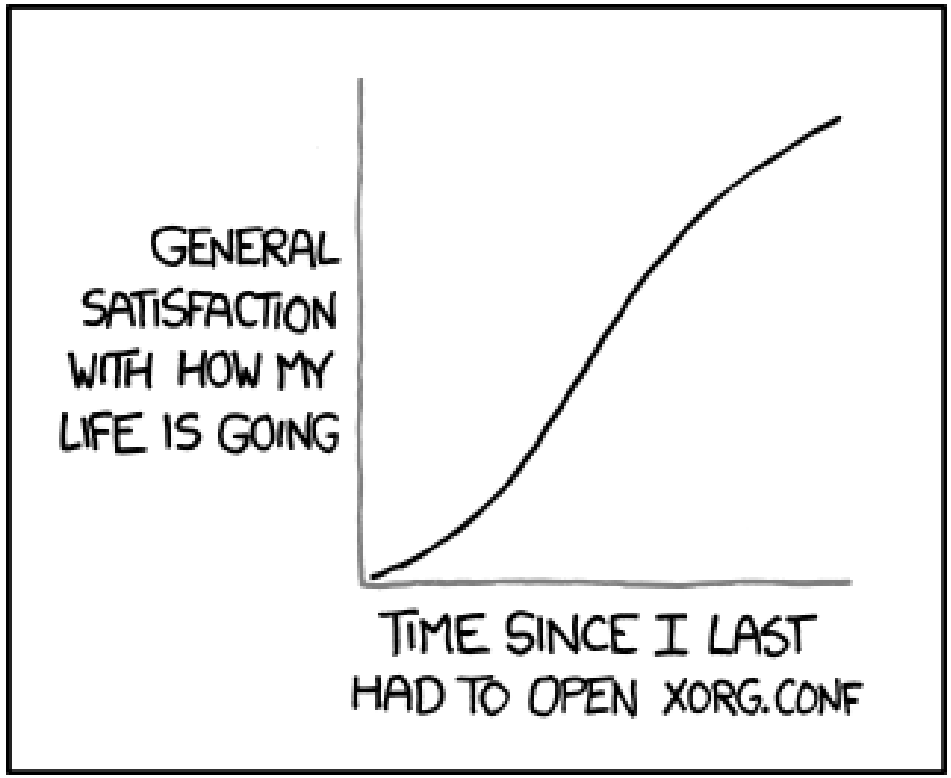

http://xkcd.com/963

Most research software is not *rock solid*

- Too few users, on too many platforms
- Bug reporting is heterogeneous, time-consuming, and painful
- Lack of professional programming training/experience
- Insufficient or inappropriate testing and quality assurance
- Death by Ph.D. phenomenon
- Opaque development procedures
  - No public version control system
  - No public bug tracker

Most research software is not *rock solid*

- Too few users, on too many platforms
- Bug reporting is heterogeneous, time-consuming, and painful
- Lack of professional programming training/experience
- Insufficient or inappropriate testing and quality assurance
- Death by Ph.D. phenomenon
- Opaque development procedures
  - No public version control system
  - No public bug tracker

Broken by design?

- Impossibility to obtain funding for software development and maintenance (alone)
- Development of software tools often not considered scientific progress

We are craving for *brand new* software, but are afraid of it

We want. . .

- latest research software to get access to bleeding edge technology and stay connected with the field
- latest tools for faster and "more interesting" publications

We are craving for *brand new* software, but are afraid of it

We want. . .

- latest research software to get access to bleeding edge technology and stay connected with the field
- latest tools for faster and "more interesting" publications

We **don't** want. . .

- to "loose results" with a new version for mysterious reasons
- to jeopardize system stability with buggy and unstable research software

We can't be bothered to make upgrades an *everyday* habit

- **It simply takes too much time!**

  The average neuroscientist on Windows spends about 14 h/month on non-research maintenance tasks (Hanke & Halchenko, Front. Neuroinf., 2011)
- Upgrading requires finding webpages, getting accounts, reading documentaion, downloading huge archives, running various installers, scripts [da capo al fine]

We can't be bothered to make upgrades an *everyday* habit

- **It simply takes too much time!**
  The average neuroscientist on Windows spends about 14 h/month on non-research maintenance tasks (Hanke & Halchenko, Front. Neuroinf., 2011)
- Upgrading requires finding webpages, getting accounts, reading documentaion, downloading huge archives, running various installers, scripts [da capo al fine]

But

Quick dissemination of new features and bug fixes is essential for efficiency

Vision

Why don't we all use the same platform...

- that works on all devices, operating systems, ...
- that is guaranteed to be available for as long as we want, wherever we want
- that we can freely share with anyone
- that makes manual maintenance trivial, or superfluous
- so all software is available in a single environment
- so we can share our experience with colleagues
- so we can share data processing workflows easily
- so developers can focus their scarce resources

Aww, come on!

Seriously, how do we get there?

Role model debian

- Origin of most active software distributions
- Vast archive of maintained software (≈30000 binary packages) – proven procedures
- Self-governed, "do-ocracy", no need to earn money, going strong for 20 years

Seriously, how do we get there?

Role model debian

- Origin of most active software distributions

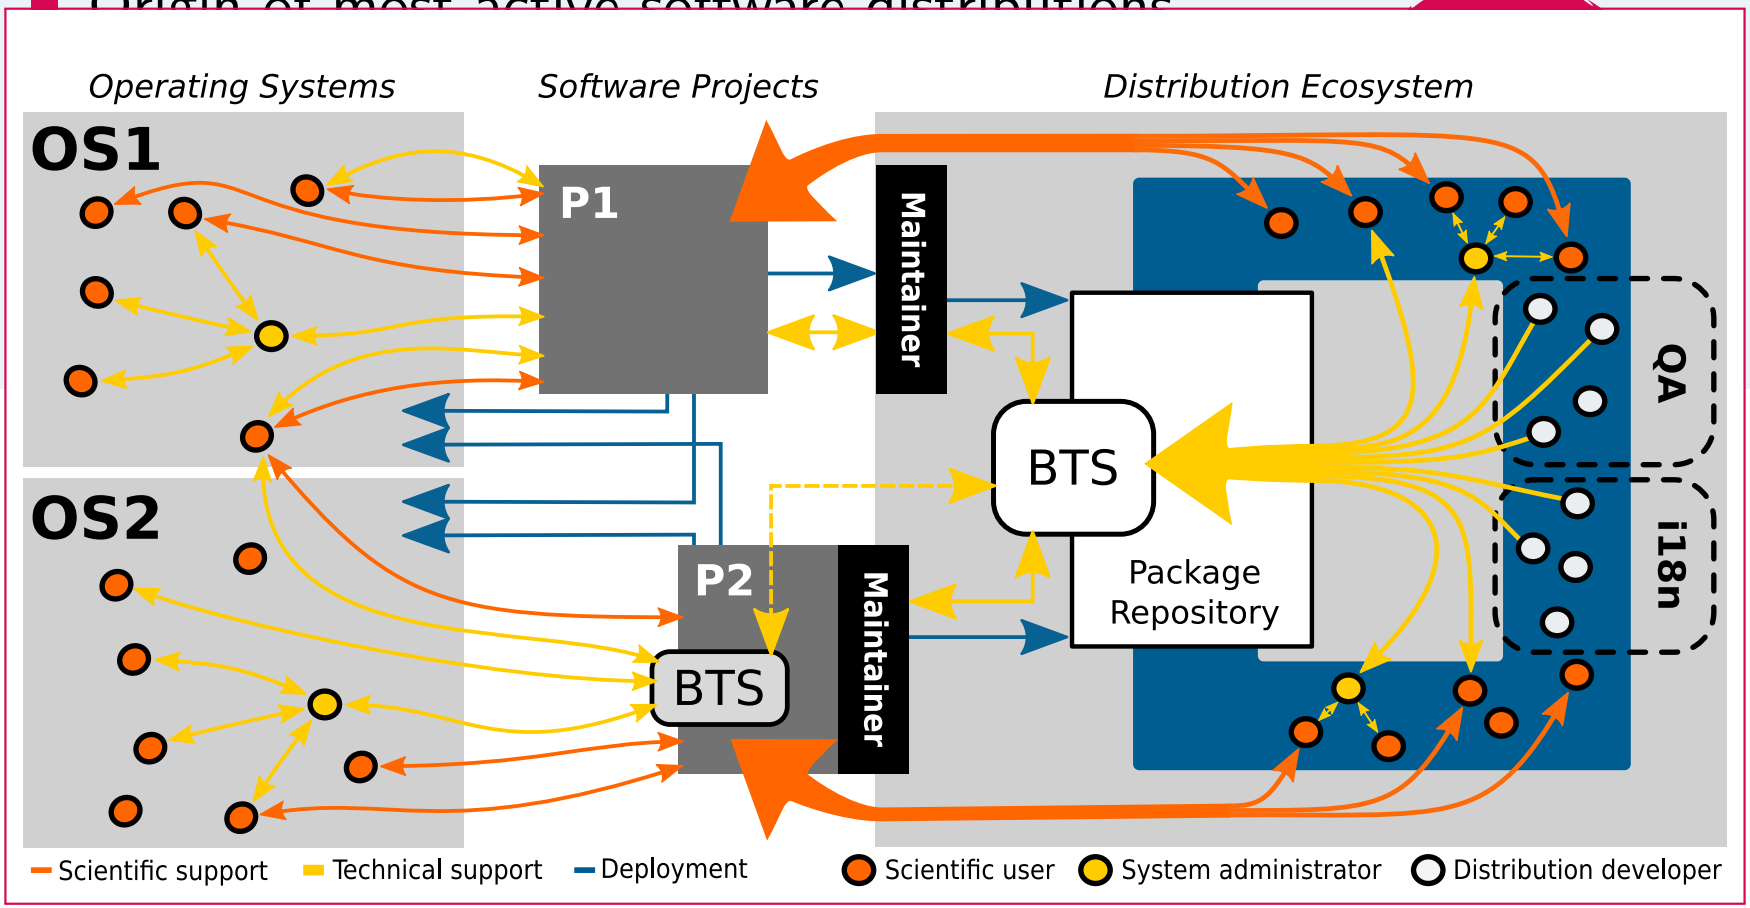

Seriously, how do we get there?

Role model debian

- Origin of most active software distributions
- Vast archive of maintained software (≈30000 binary packages) – proven procedures
- Self-governed, "do-ocracy", no need to earn money, going strong for 20 years

We could ...

- Adopt technology and procedures
- Participate in the Debian project and integrate all research software
- Benefit from the work of thousands of *additional* developers
- Call it NeuroDebian, add fancy logo

NeuroDebian from a researcher's perspective [artist's rendition]

Install simple editor

```
apt-get install gedit
```

Install complex MRI analysis package

```
apt-get install fsl
```

Install psycho-physics toolbox

```
apt-get install psychopy
```

Keep the whole system up-to-date

```
apt-get upgrade
```

Back in 2002...

This idea was born here

After ten years and the contributions of many people:

afni amide ants biosig brian caret cfflib condor connectomeviewer coop-computing dcmtp debruijn dicomnifti dipy freenect fsl fslview gdf gifti haxby2001 itksnap jist joblib kbibtex klustakwik libsvm lipsia matlab-support mdp medcon minc mipav mitools mlpy mni-autoreg-model mni-colin mni-icbm152-nlin mpi4py mriconvert mricron mrtrix nibabel nifti nipy nitime odin openelectrophy openmeeg openopt openpyxl opensesame psignifit psychofun psychopy psychtoolbox pydicom pyepl pymvpa pynn pyoptical pyxid r-cran-glmnet rorden-mri-tutorial scikits-learn scikits-statsmodels shogun sigviewer slicer spm statsmodels stimfit svgtune sympy via volpack voxbo xmedcon

NeuroDebian focus

Researchers/users

- Configure Debian/Ubuntu box or download virtual machine
  – visit `http://neuro.debian.net` for instructions
- `apt-get install mricron` (thousands of packages)
- `apt-get upgrade` (full system update)
- Get support at `neurodebian-users@lists.alioth.debian.org`

NeuroDebian: The virtual machine

- 32/64bit, multi-core
- Most convenient solution for Mac OS X, Windows
- Base image with setup wizzard, fully functional within minutes
- Great for teaching, workshops, development, analysis

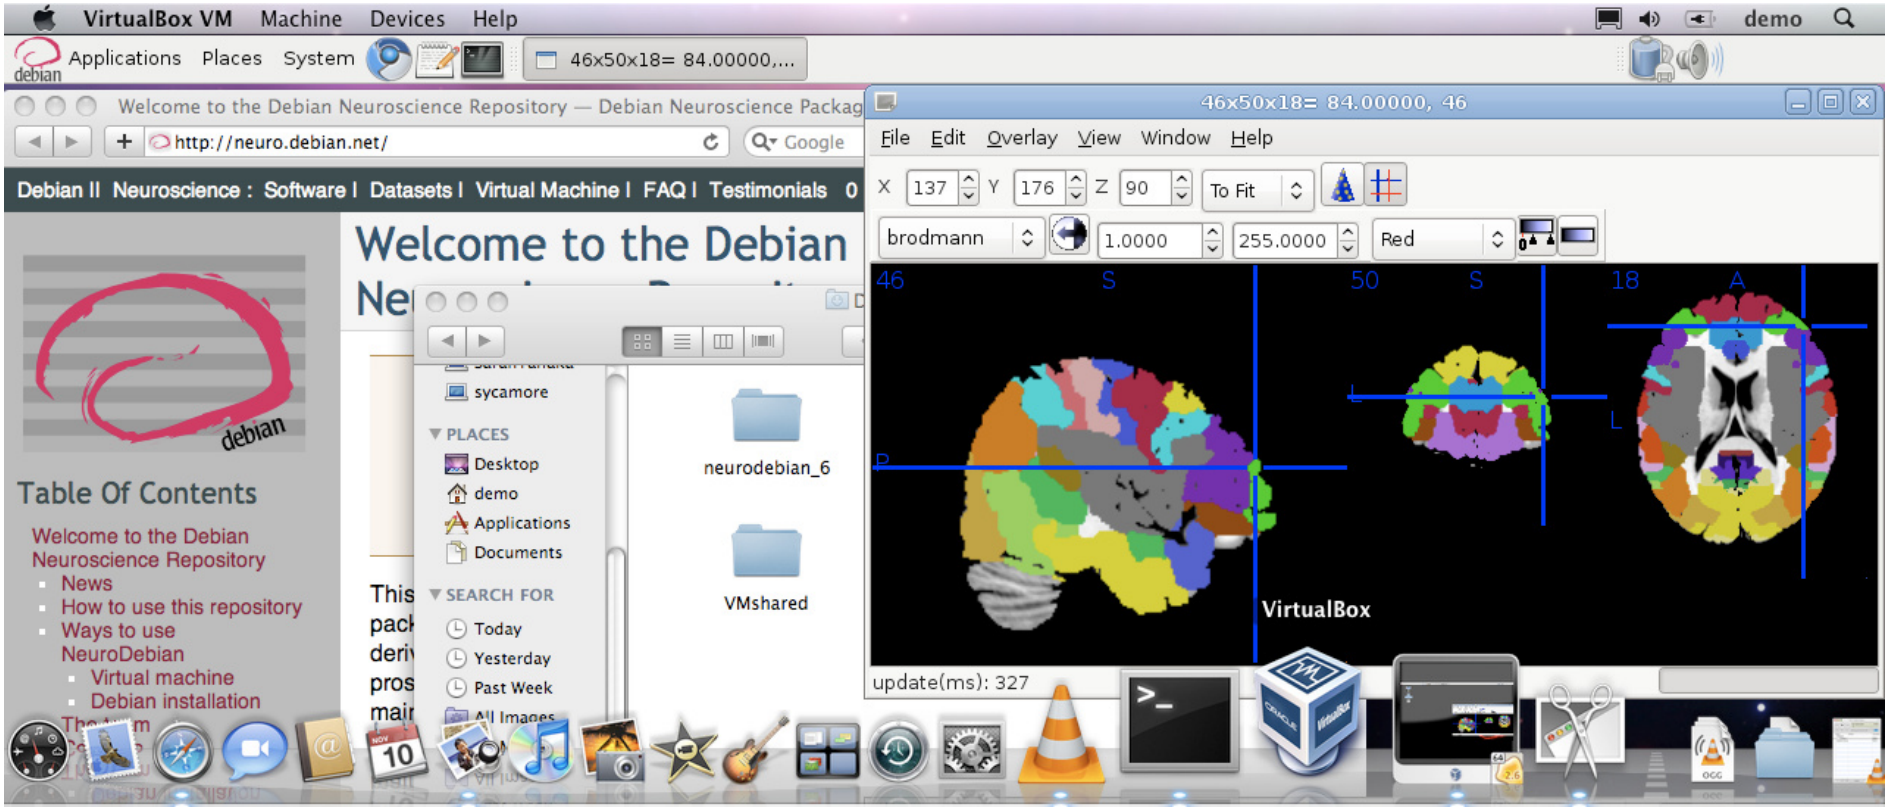

`http://neuro.debian.net/vm.html`

Walkthrough: NeuroDebian on Mac OS X

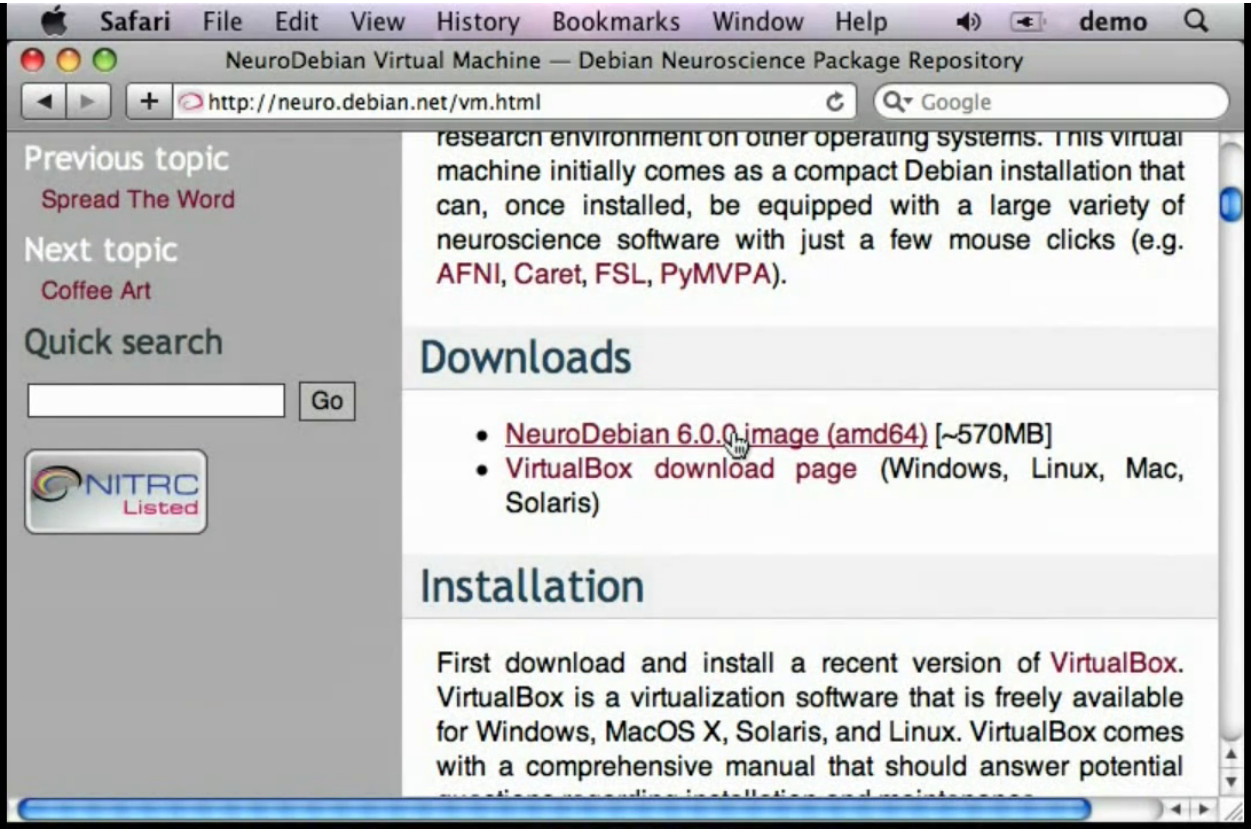

Walkthrough: NeuroDebian on Mac OS X

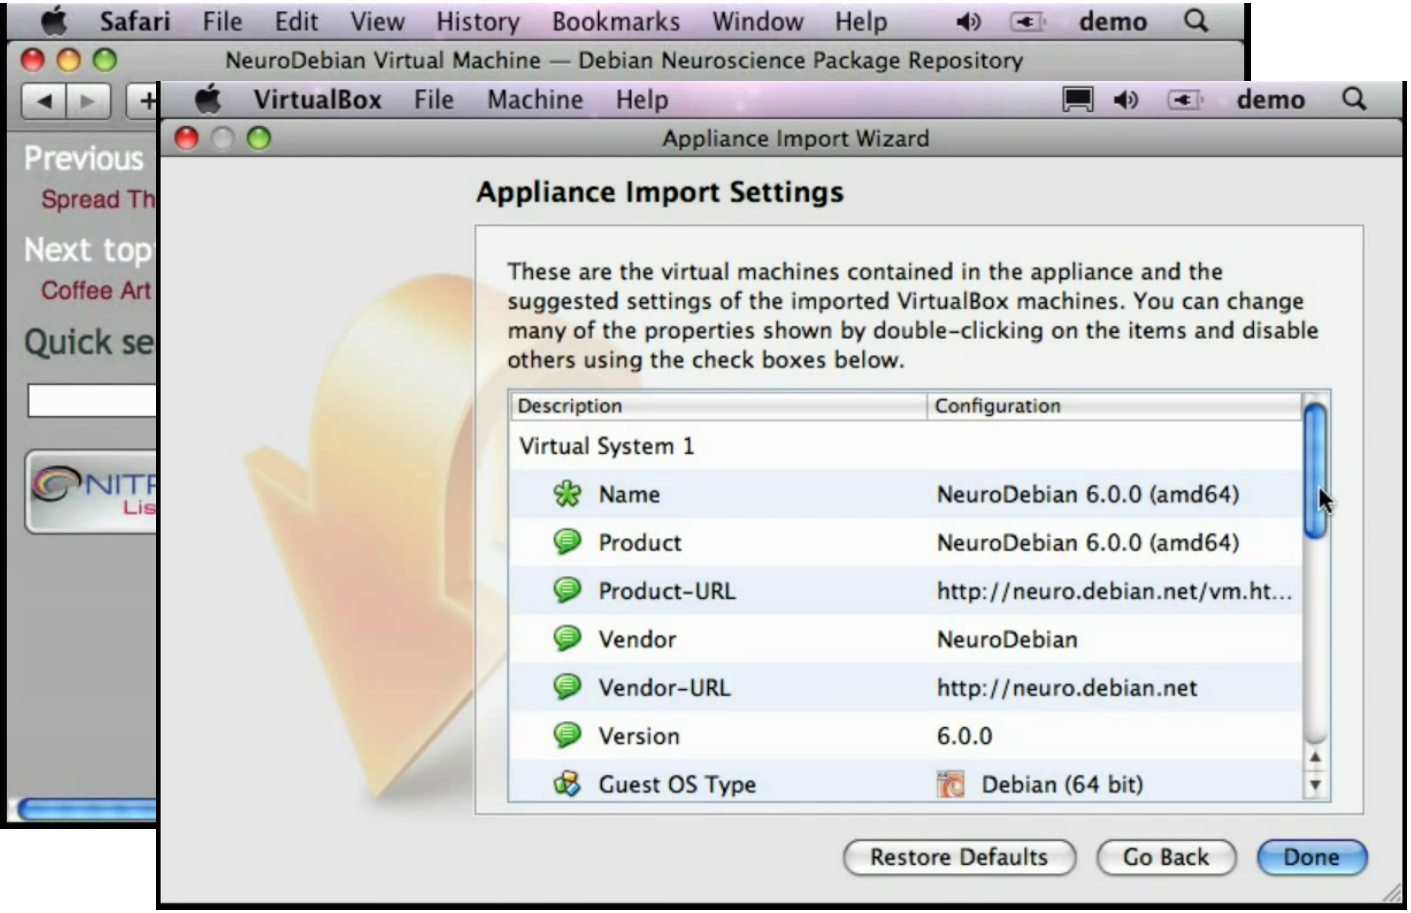

Walkthrough: NeuroDebian on Mac OS X

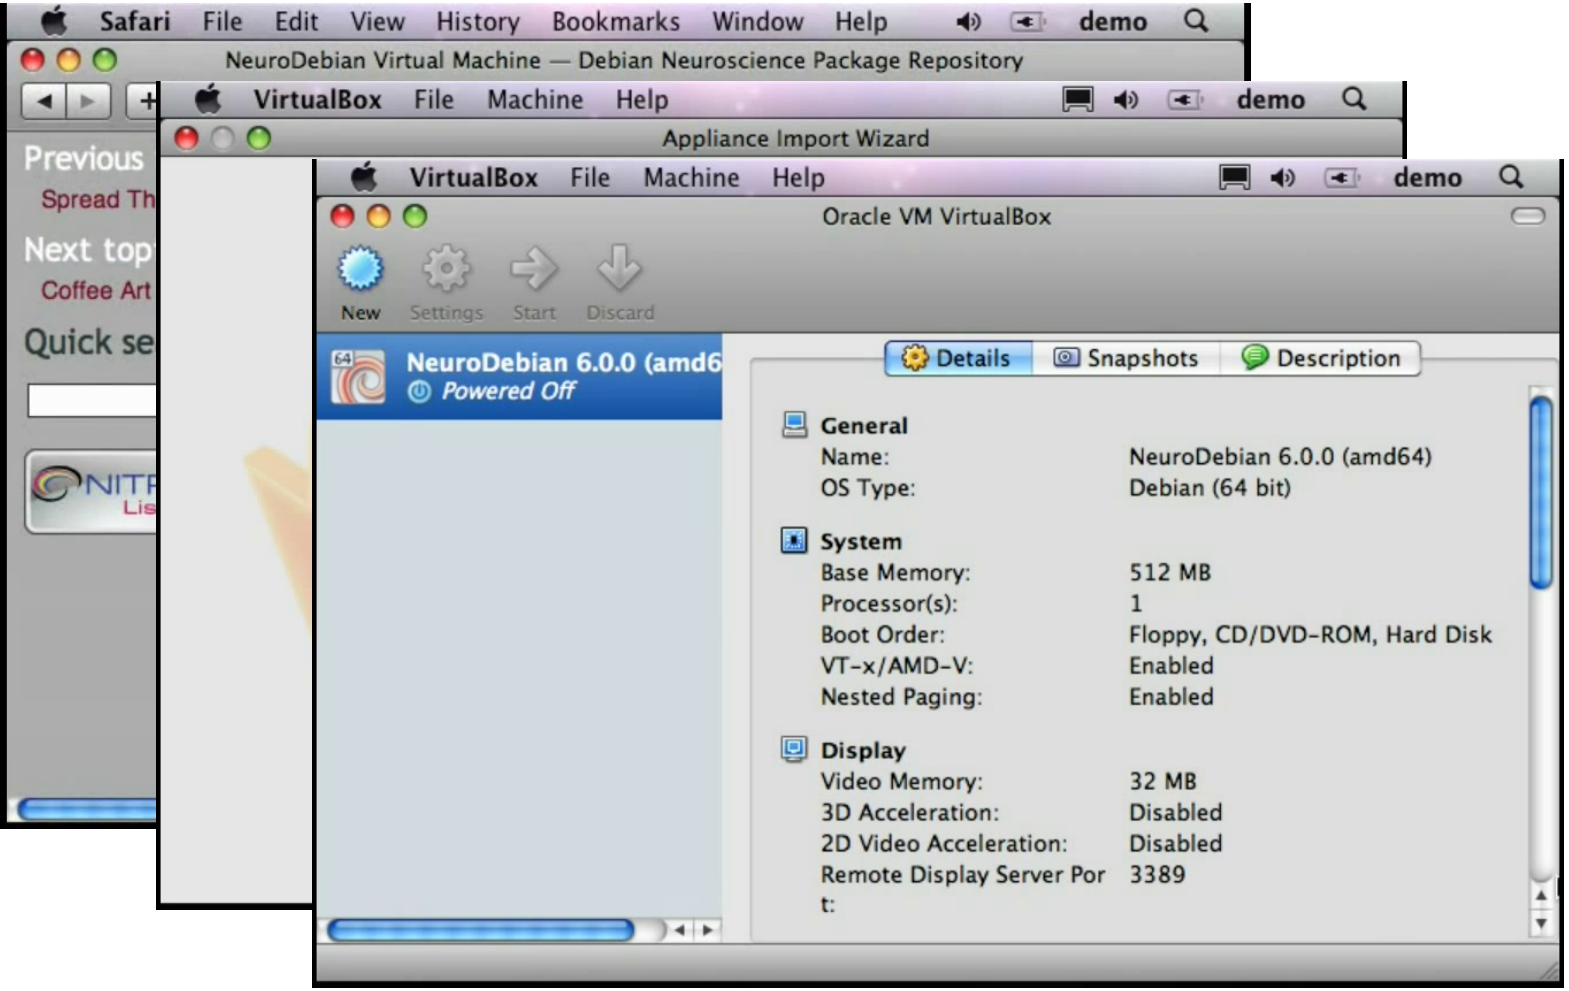

Walkthrough: NeuroDebian on Mac OS X

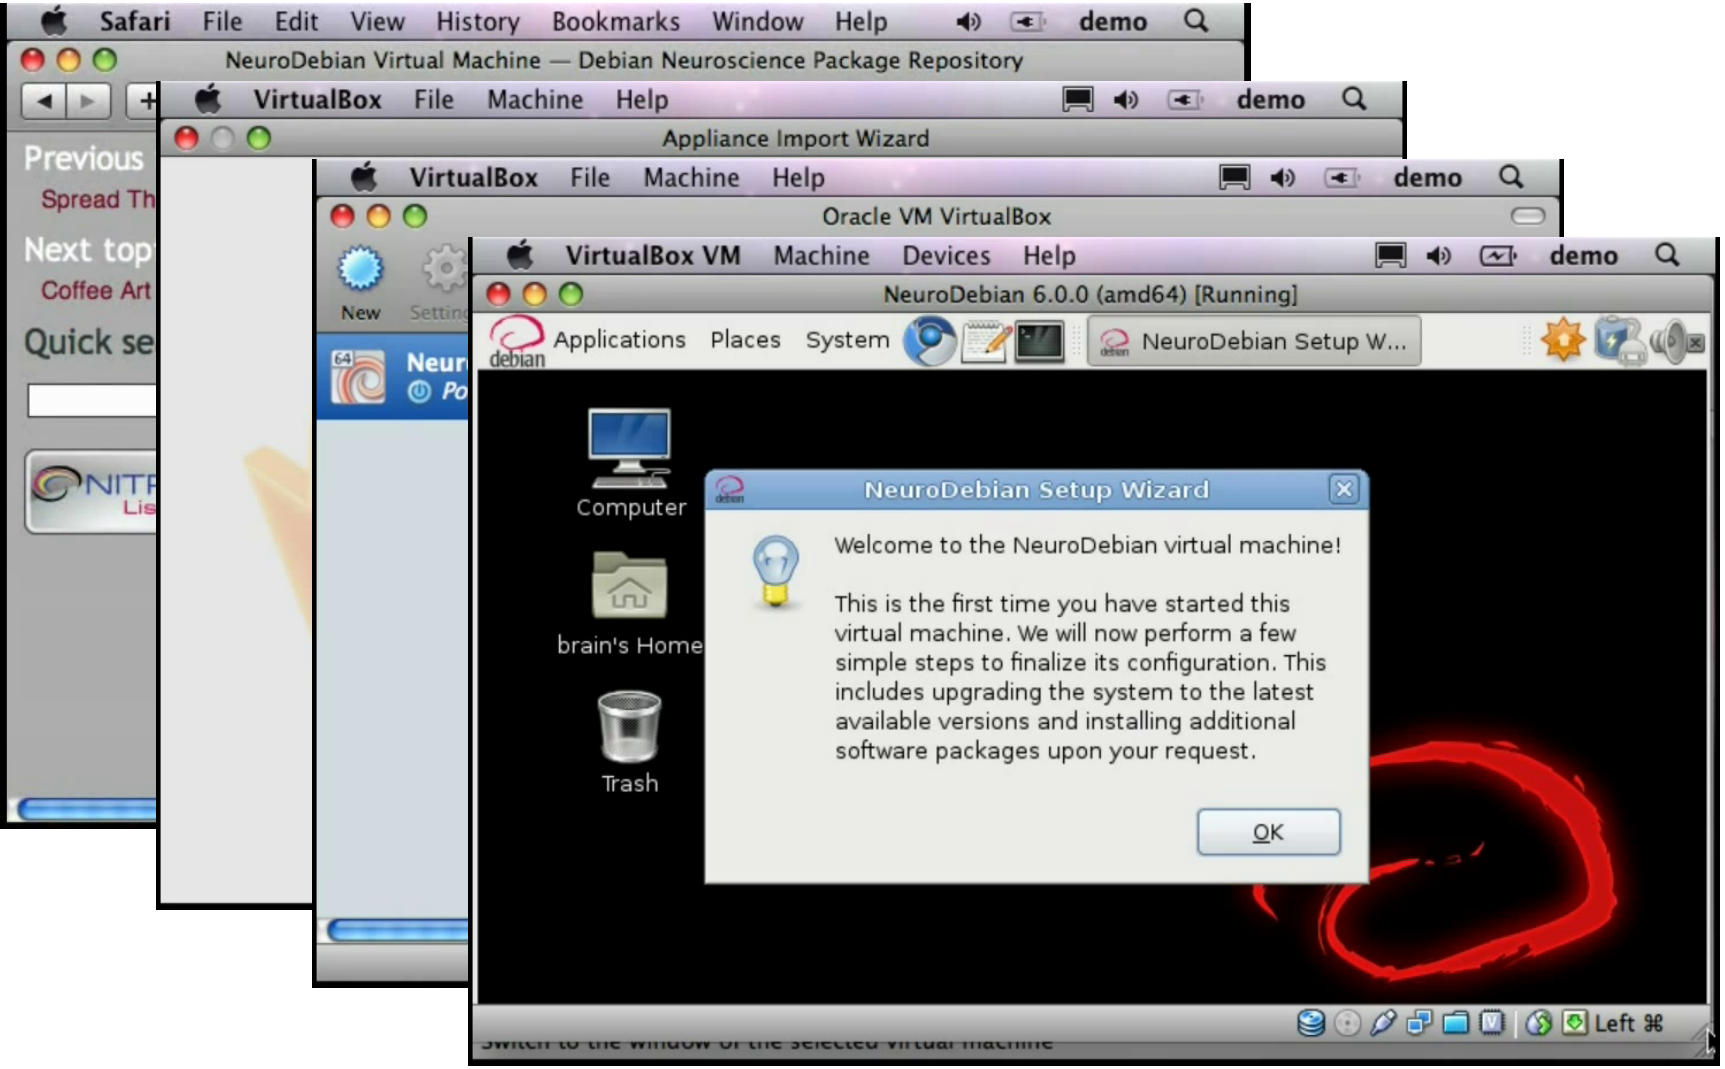

Walkthrough: NeuroDebian on Mac OS X

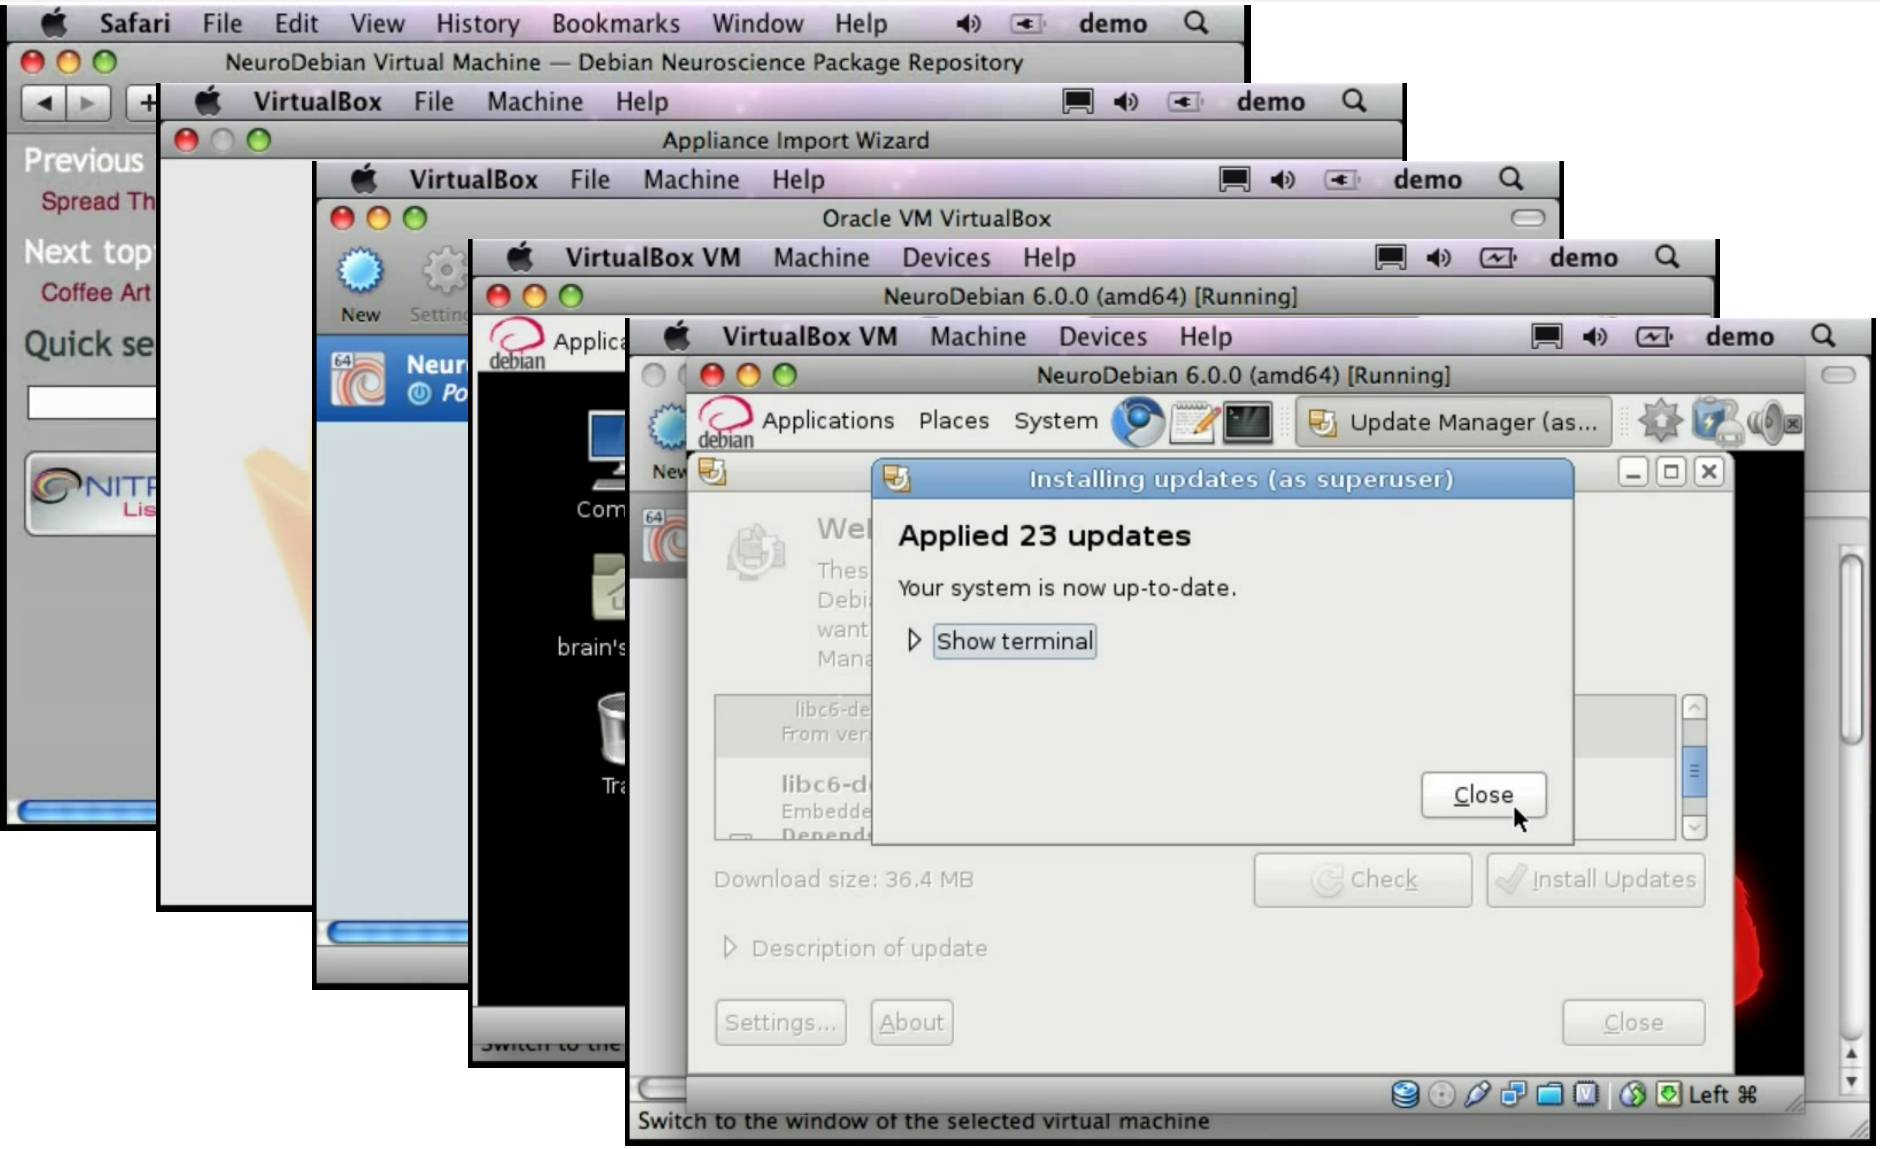

Walkthrough: NeuroDebian on Mac OS X

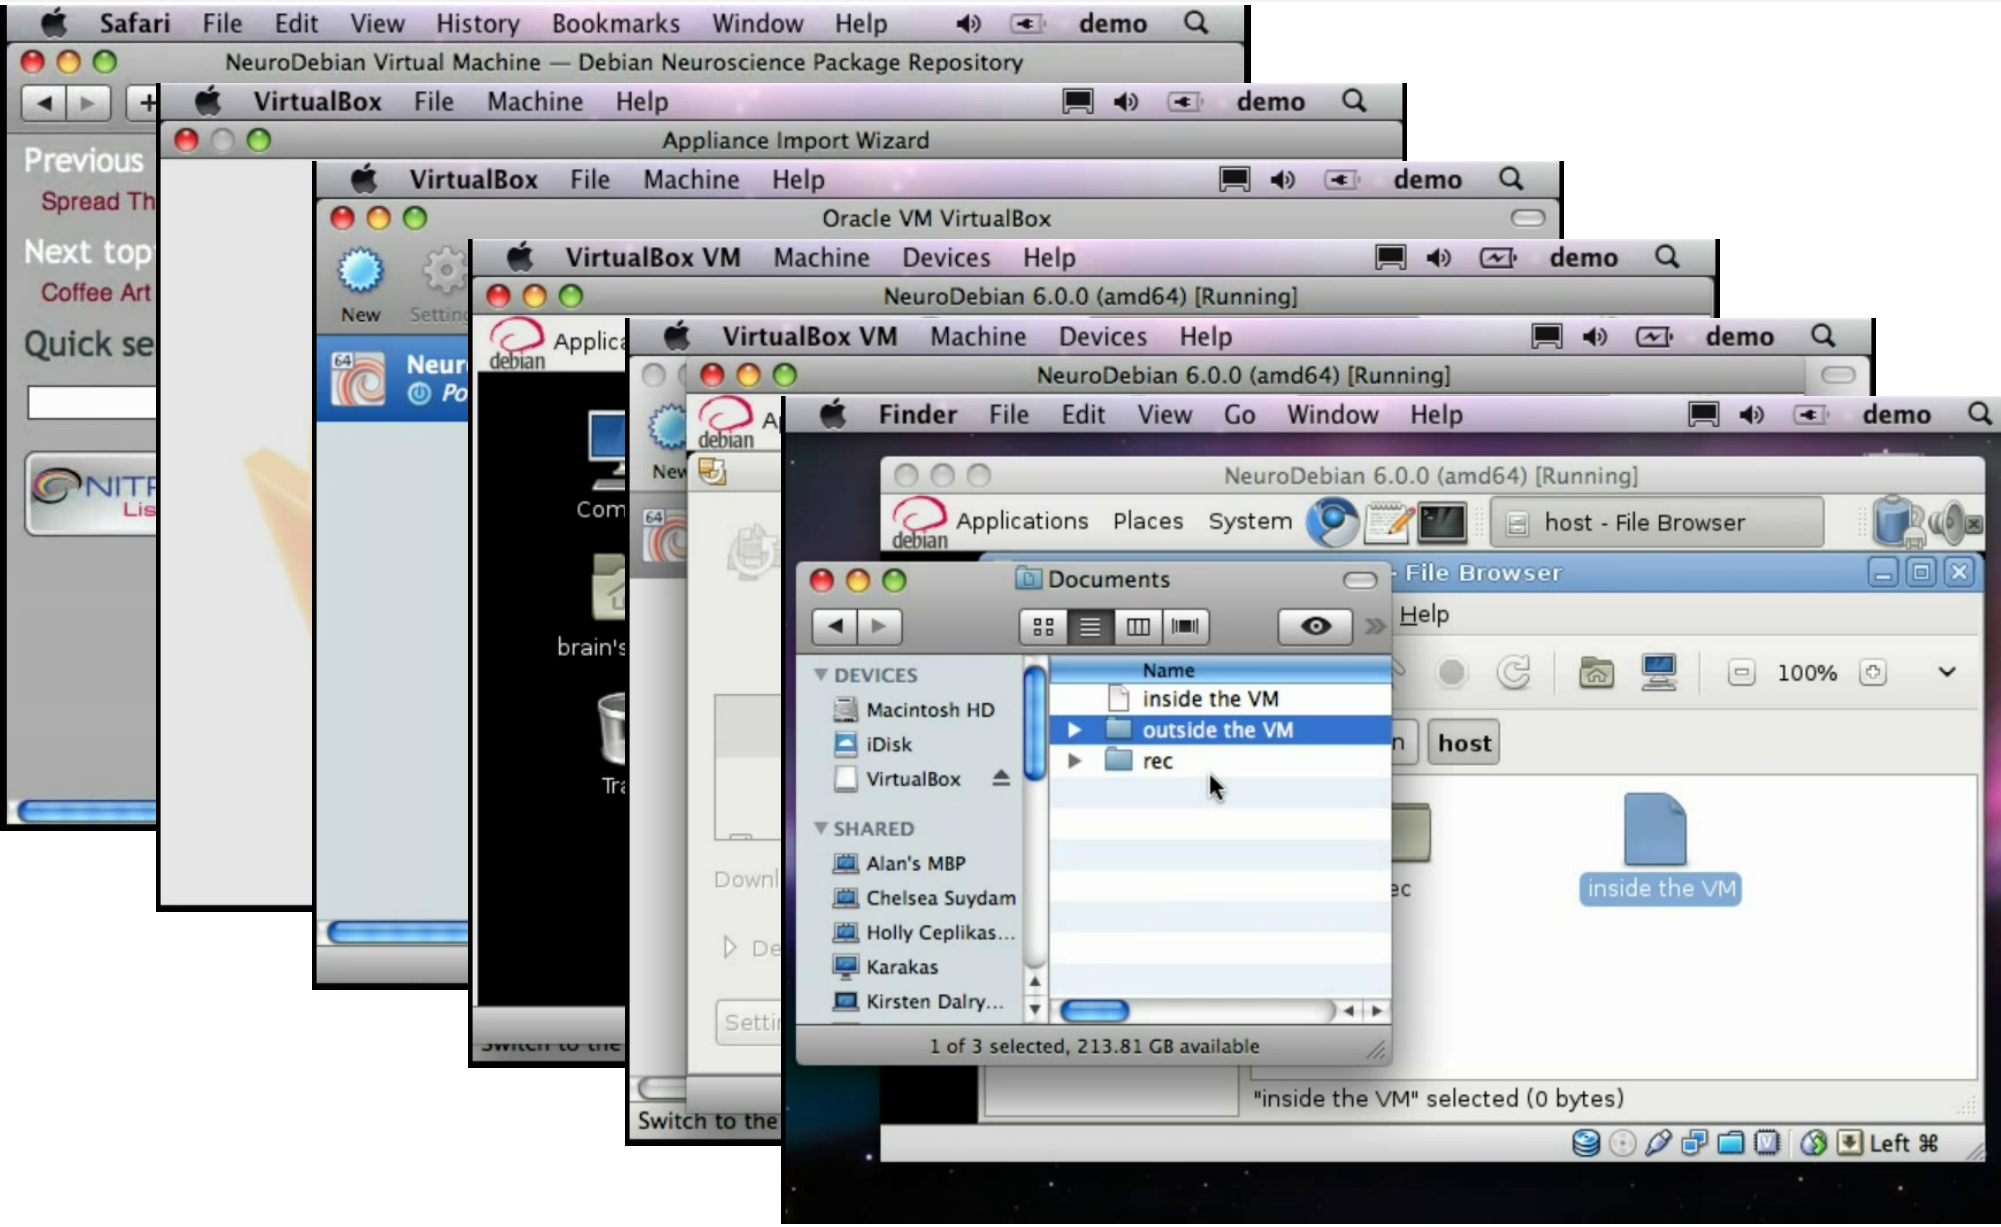

Walkthrough: NeuroDebian on Mac OS X

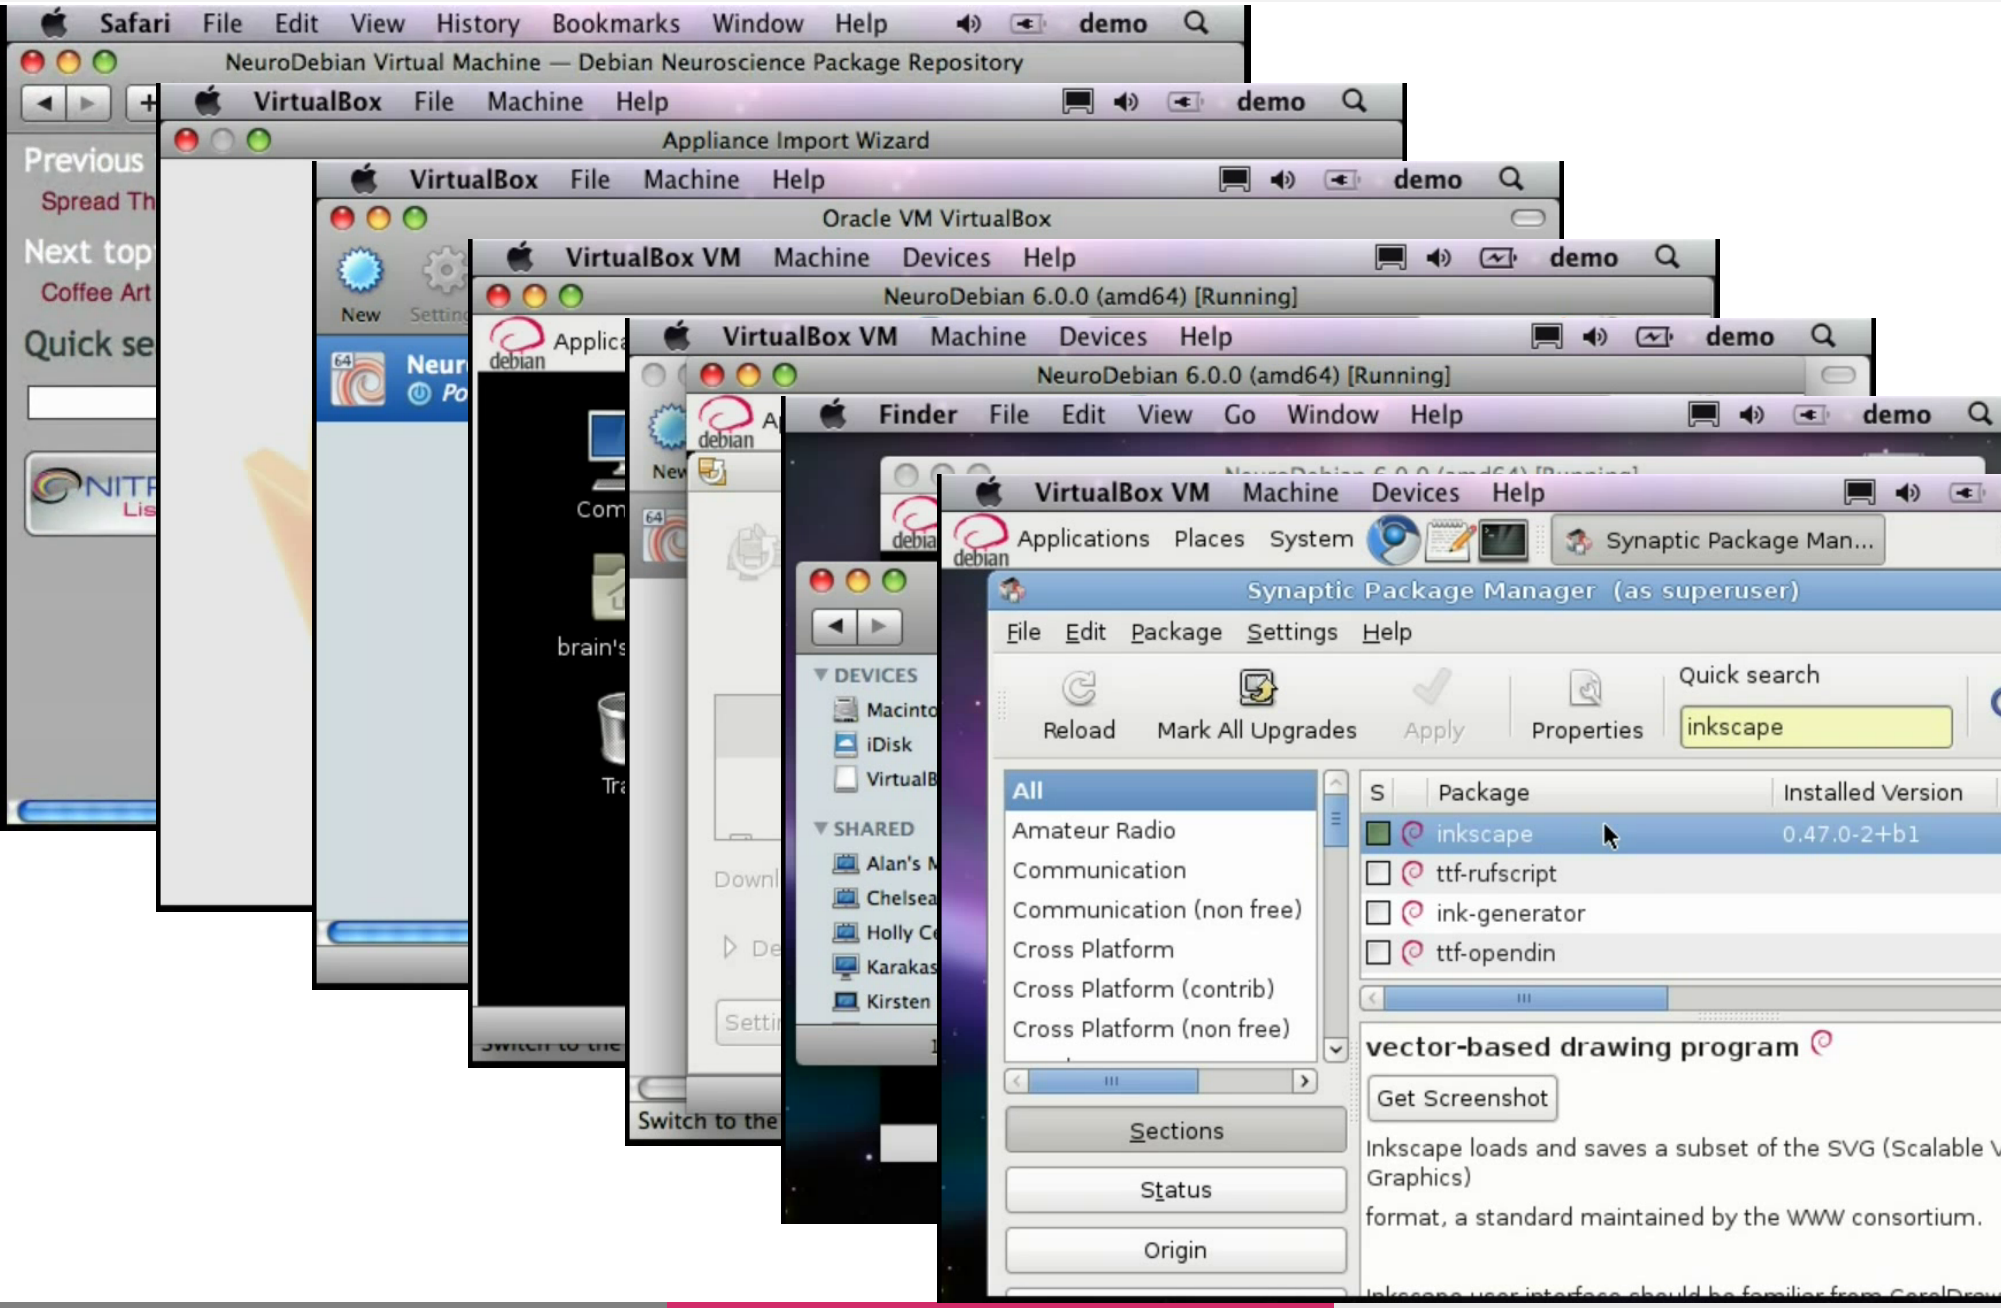

NeuroDebian focus

Researchers/users

- Configure Debian/Ubuntu box or download virtual machine
- `apt-get install mricron` (thousands of packages)
- `apt-get upgrade` (full system update)
- Get support at `neurodebian-users@lists.alioth.debian.org`

NeuroDebian focus

Researchers/users

- Configure Debian/Ubuntu box or download virtual machine
- `apt-get install mricron` (thousands of packages)
- `apt-get upgrade` (full system update)
- Get support at `neurodebian-users@lists.alioth.debian.org`

Developers

- Low-latency dissemination channel to thousands of users
- Platform-specific expertise available

NeuroDebian focus

Researchers/users

- Configure Debian/Ubuntu box or download virtual machine
- `apt-get install mricron` (thousands of packages)
- `apt-get upgrade` (full system update)
- Get support at `neurodebian-users@lists.alioth.debian.org`

Developers

- Low-latency dissemination channel to thousands of users
- Platform-specific expertise available

Maintainers

- Uniform environment for cross-project improvements and QA

This is not fundable

This is not fundable, so don't require any money!

Crowd-sourcing the software infrastructure

- Developers integrate software into Debian (with mentoring)
- Most software is in Debian, hence most scientist use Debian
- Users report bugs using built-in tools, they get tracked publicly
- Neuroscience software gets exposed and integrated in the open-source community
- Interested developers outside the projects can run QA efforts
- Software quality and longevity increases
- Provide the tools – not the service

And it seems to work. . .

Journal Info
Home
About the Journal
Editorial Board
Archive
Special Topics
View Some Authors
Review Guidelines
Search
Article Type
All
Publication Date
From
To
Go
Author Info

OPINION ARTICLE

Full Text | Provisional PDF | A A | Share this

Neuroscience runs on GNU/Linux

Michael Hanke[1, 2, 3*] and Yaroslav O. Halchenko[1, 2]

[1] Center for Cognitive Neuroscience, Dartmouth College, USA
[2] Department of Psychological and Brain Sciences, Dartmouth College, USA
[3] Department of Experimental Psychology, Otto-von-Guericke-University, Germany

In response to a recent grant application for a software development project, we received some reviewer comments that questioned the prevalence of GNU/Linux systems as a computing platform in neuroscience. Moreover, a concern was raised that virtualization is not a feasible solution to overcome limitations of any particular platform or to provide a convenient multiplatform working environment. We were surprised by these comments, because they are in contrast to what we experience daily while working with software developers worldwide to integrate neuroscience software into the NeuroDebian project.

Article Info
Full Text
Provisional PDF
Export Citation
The Authors in
Frontiers
Google
Google Scholar
PubMed
Related Article
in Frontiers
Google Scholar
PubMed

And it seems to work. . .

And it seems to work. . .

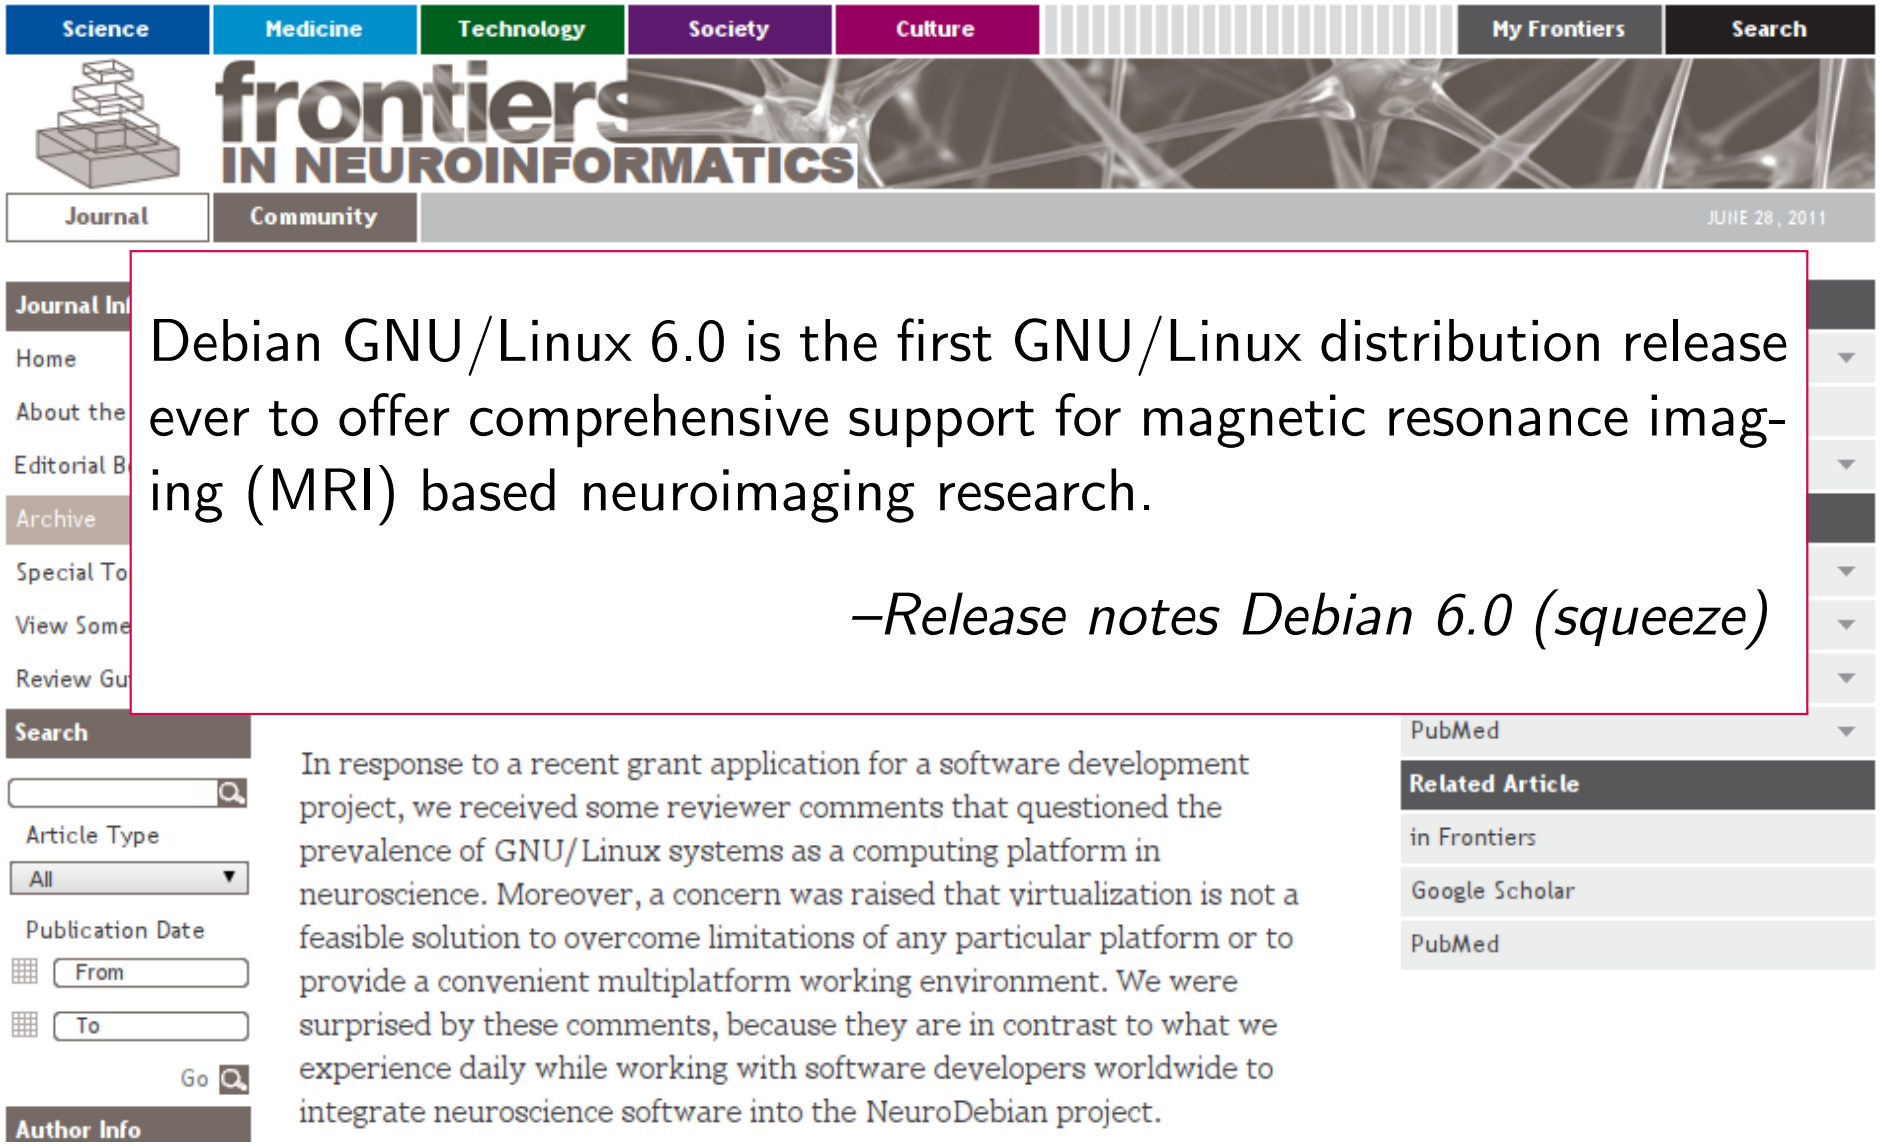

Debian GNU/Linux 6.0 is the first GNU/Linux distribution release ever to offer comprehensive support for magnetic resonance imaging (MRI) based neuroimaging research.

–Release notes Debian 6.0 (squeeze)

NeuroDebian popularity statistics

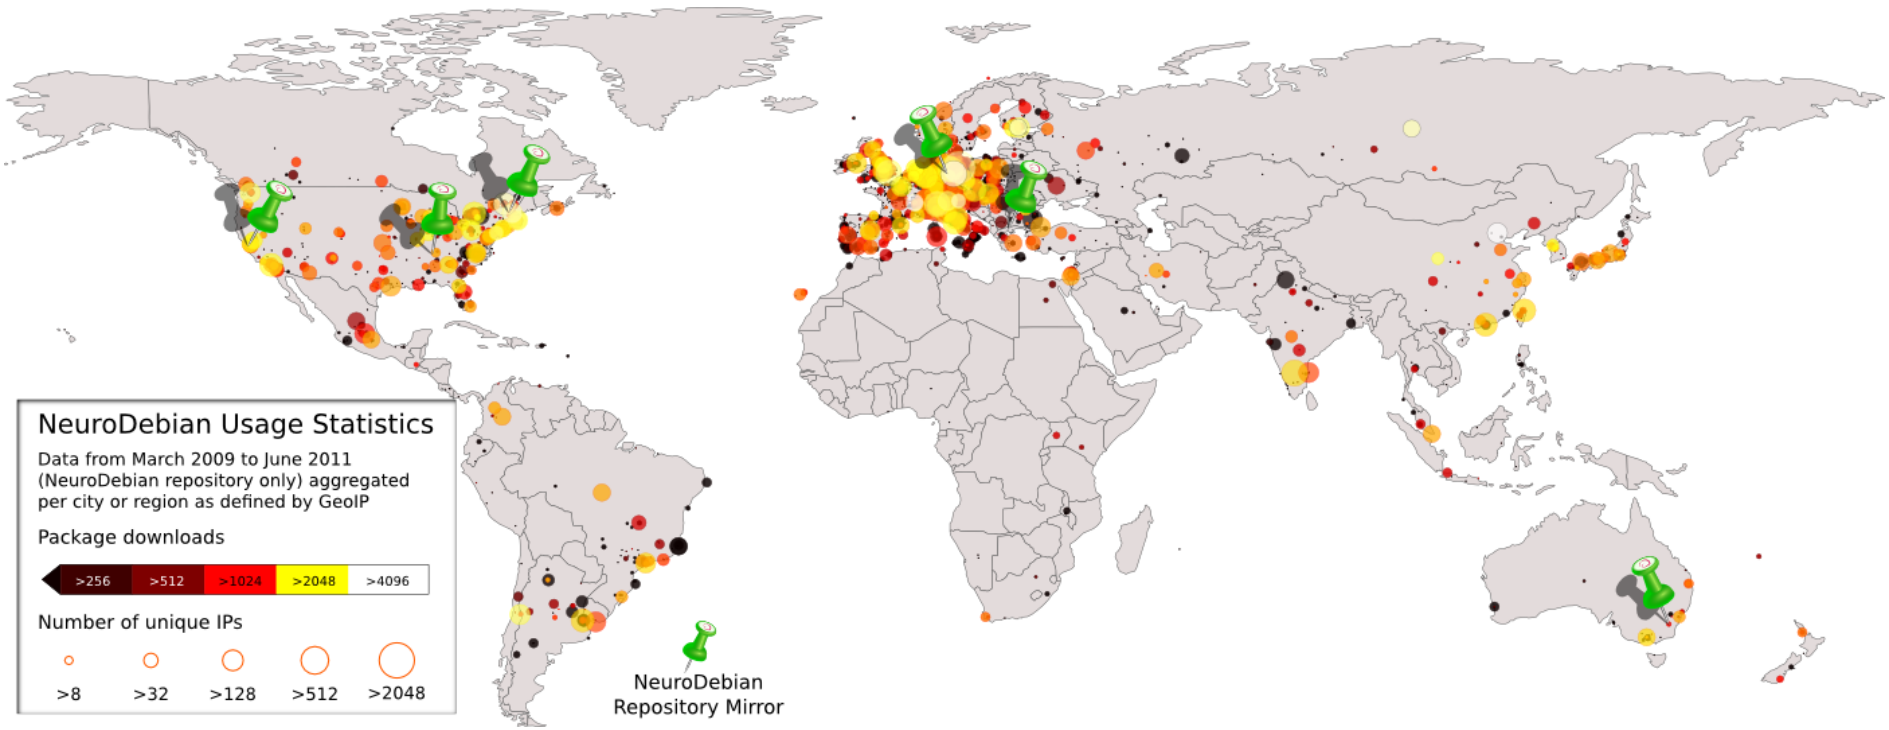

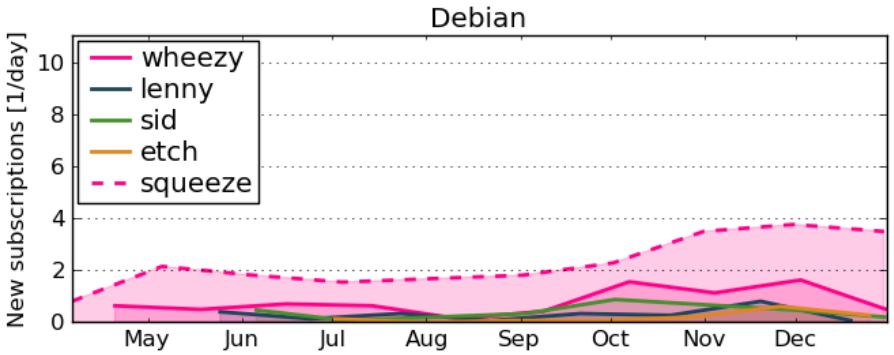

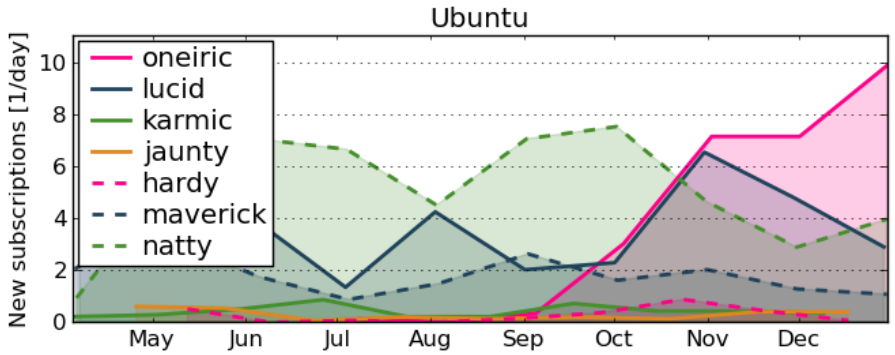

What is NeuroDebian?

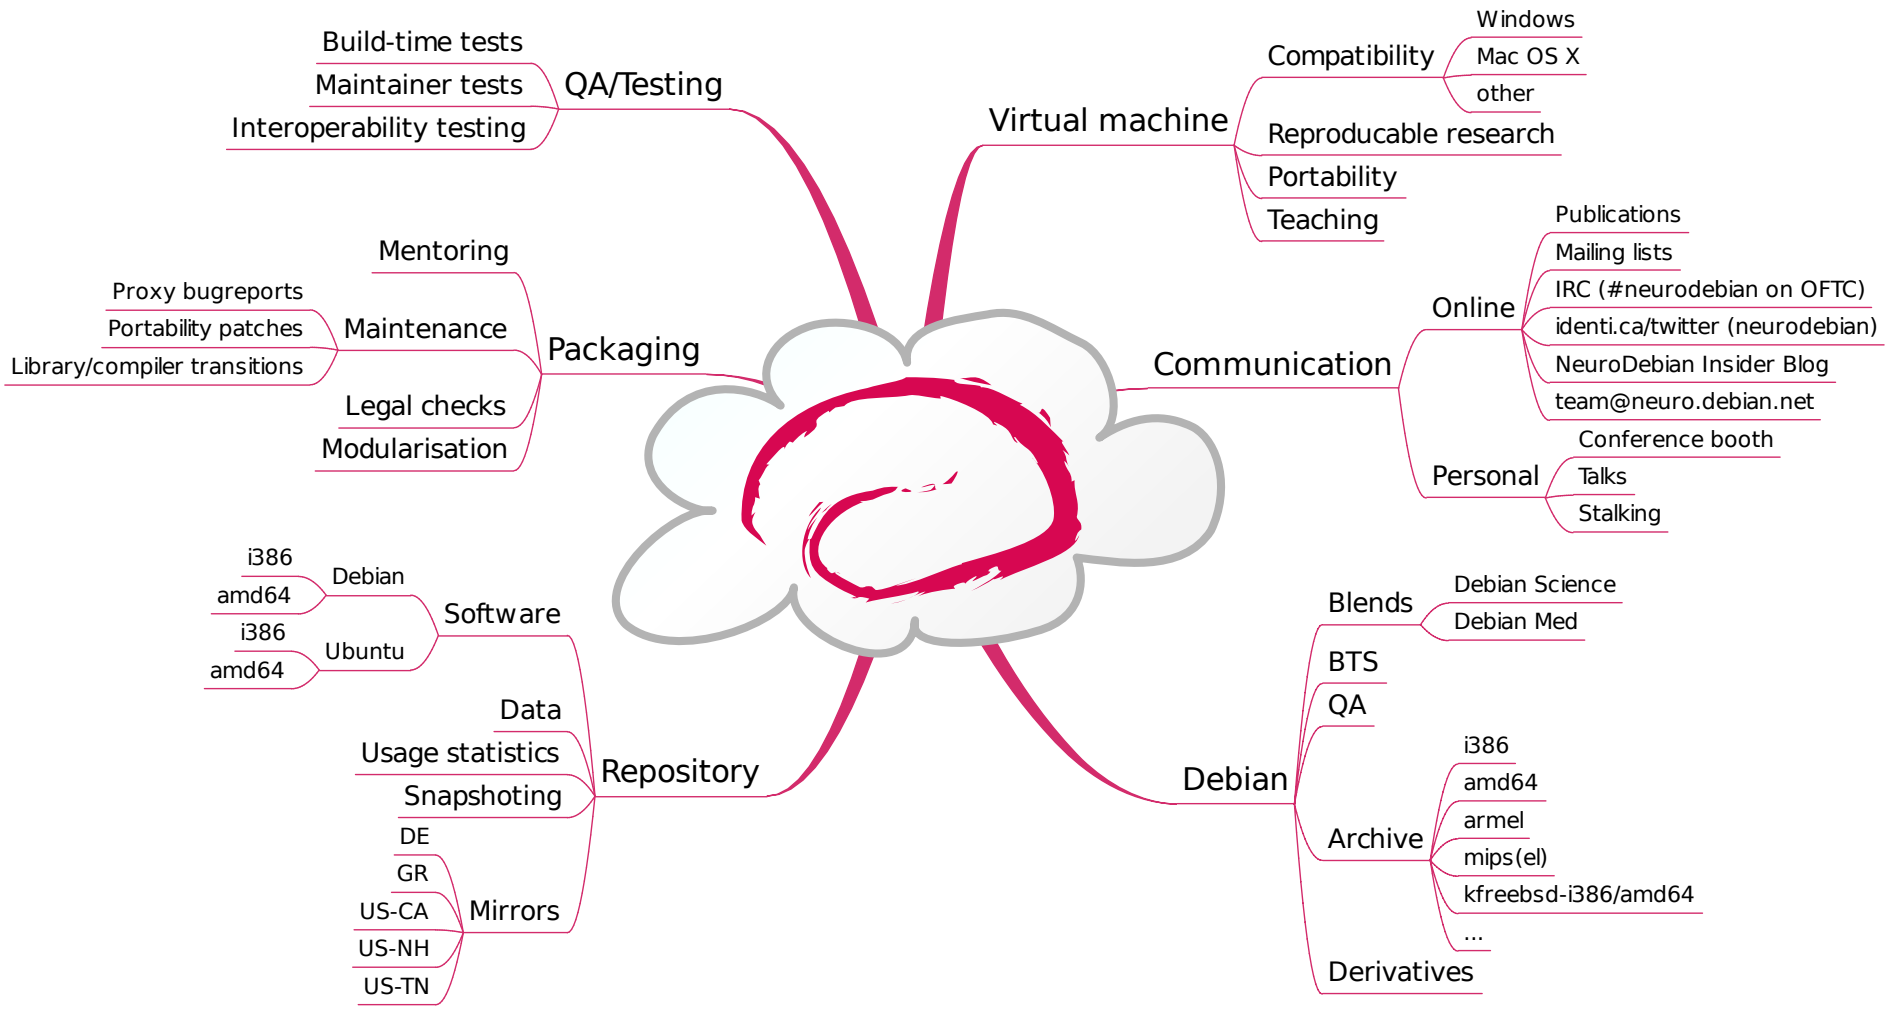

For whom is NeuroDebian?

You want to ...

- have **readily usable** software at your fingertips
- have the **latest developments** of research software
- use a **rock-solid** operating system
- **try something new**, without investing much time
- offer **students** a fully functional "take-away" research environment
- **efficiently collaborate** with other researchers
- **escape limitations** of an institutional computing environment
- **waste less time** maintaining computers
- have **your own software** easily available for other's to use
- **develop neuroscience software** without worrying about dependencies
- **help** make NeuroDebian more robust, and/or **built on top** of it

---

`http://neuro.debian.net`

Acknowledgements

Yarik Halchenko

NeuroDebian contributors

Mirror/bandwidth donors

Neuroscience free software developers

Debian Community

Jim Haxby

Stefan Pollmann

This was not a joke!

Michael Hanke

mih@debian.org

http://mih.voxindeserto.de

about the slides:
available at http://neuro.debian.net/#publications

Get involved!

- Find and evaluate software
- Report bugs, send patches
- Support: Mailing list, IRC (`http://neuro.debian.net/#contacts`)
- Post on the NeuroDebian blog
- Help (co-)maintain a package
- Package your own software
- Send us tests
- Spread the word

WE NEED HELP!

http://openhatch.org/+projects/NeuroDebian

How does software benefit from Debian?

- Extended reach
  - one stable release, two rolling "release" flavors
  - ≈130 derivative distributions (distrowatch.org)

How does software benefit from Debian?

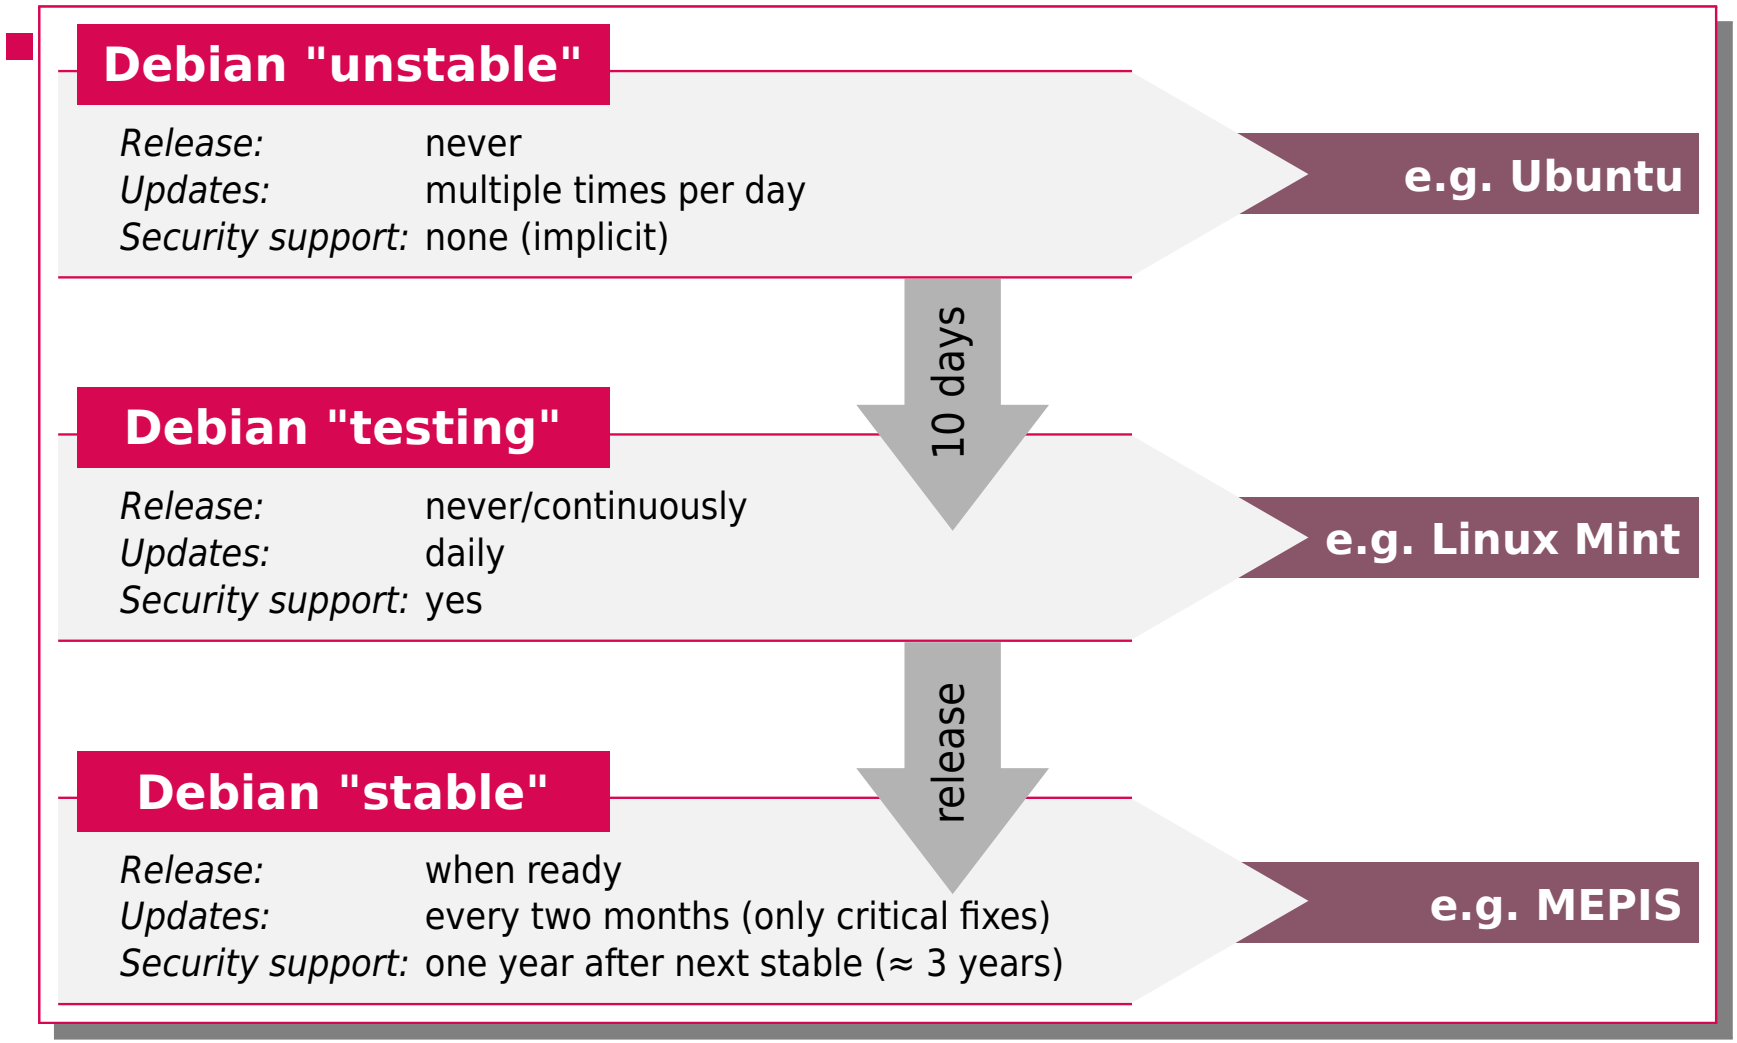

How does software benefit from Debian?

- Extended reach
  - one stable release, two rolling "release" flavors
  - ≈130 derivative distributions (distrowatch.org)
- Mutual awareness
  - Explicitly documented dependencies
  - Synchronized transitions
- Less maintenance work through modularity
  - 3rd-party software in dedicated packages maintained by someone else
- Continuous integration testing
  - 13 hardware architectures
  - Three kernels
  - Continuous automated testing for
    - Build success
    - Clean installation/de-installation, Availability of dependencies
    - Policy compliance
    - Package conflicts

But I only care about Ubuntu!

No, you don't!

- Most software we care about comes (almost) 1:1 from Debian (SciPy, VTK, ITK, ...)
- No LTS for neuroscience (NumPy only since 10.04)

Go Debian!

- Developers: Get it right in Debian, have it work in Debian/Ubuntu/Mint/aptosid/Mepis/... (at no additional cost)
- Users: Stable release with 3-4 years support for all software
- Scientists: Want your research tool to be found and used? Include it in the largest software archive in the world.

OS market share

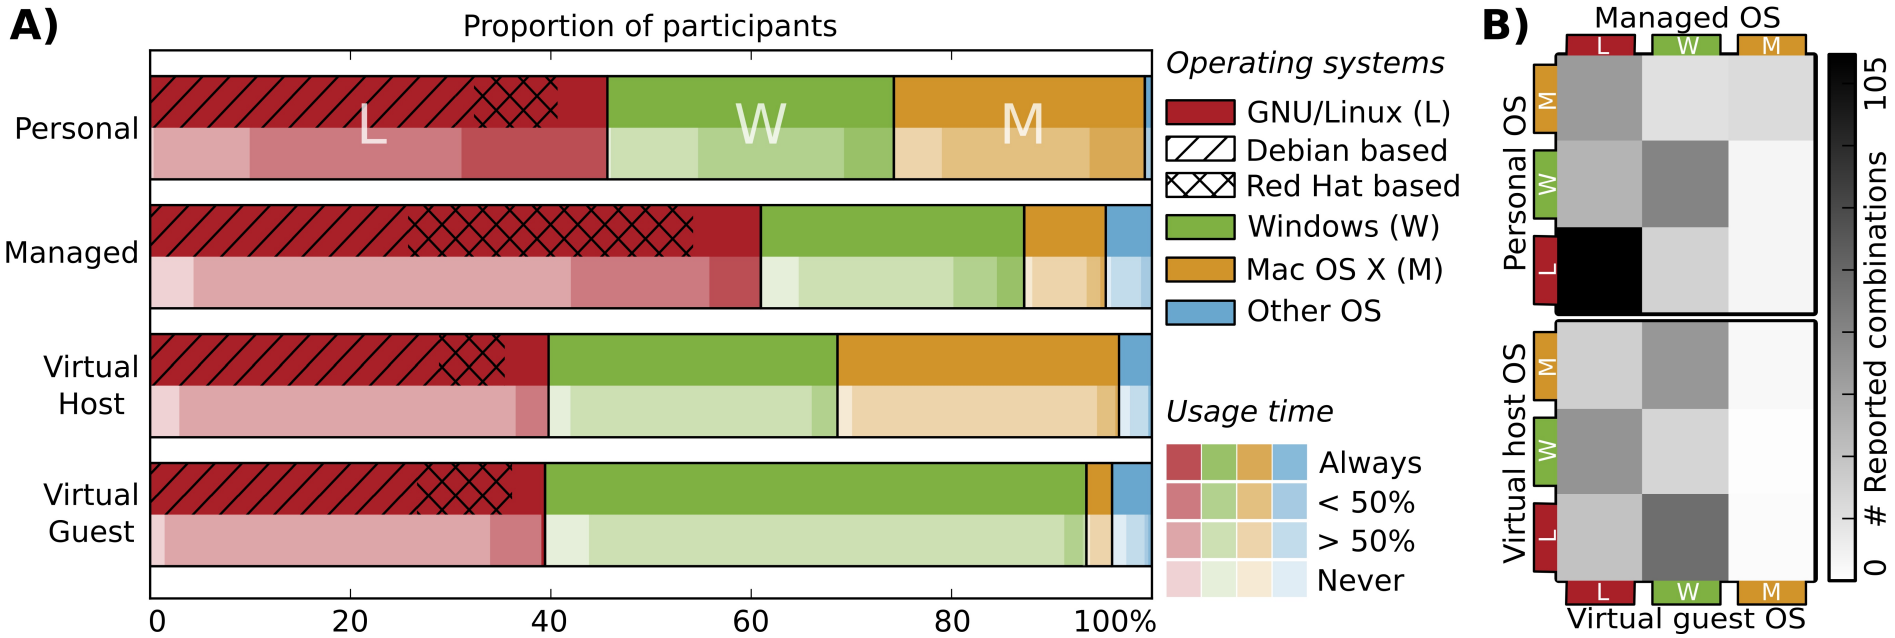

Hanke & Halchenko, 2011, Front. Neuroinf.