Lupe User's Guide

Oliver Thilmann, PhD
Dept. of Medical Radiation Physics
Lund University
S-221 85 Lund
Email: oliver.thilmann@web.de

1. Introductory remarks

Lupe (german for 'magnifying glass') is a software for the calculation of cerebral blood-volume and blood-flow parametric maps from dynamic-susceptibility-contrast (DSC) acquired images. Input to the program is a set of (typically 500 to 1200) DICOM images from your scanner - either from CD or transfered to your hard disk via a DICOM server. The calculated parameter maps can (a.o.) be saved as DICOM3.0 images and are thus suited for storage in e.g. a hospital's PACS system.

Currently the following calculation methods are implemented:

- Standard SVD as described in Ostergaard et al., MRM 36(1996), 715-725.

- Reformulated SVD, described in Smith et al., MRM 51(2004), 631-634.

- Blood-brain-barrier leakage correction, as published in Haselhorst et al. JMRI 11(2000), 495-505. Please note the comments below.

For a start, you might also want to have a look at the program's web page at

http://www.jubileum.lu.se/mrphysics/software/lupe/index.htm

The 'modules' section contains a more detailed description of the modules available in Version 1.2.

2. Legal issues

This software is published under the terms of the GNU public license, a copy of which you should have received with this distribution. This means the program comes *with no warranty* whatsoever. Even though the calculation techniques were implemented to the best of our knowledge and even though *Lupe* is used in daily clinical practice in Lund, we distribute the program exclusively for research purposes. *Do not base diagnostic decisions on the results of this program.* If you do so, it is your sole responsibility.

3. Feedback

That said, I'd be happy to hear from you. Please do not hesitate to send mail in case of questions or problems - I will provide as much support as time permits. Suggestions for further development,

bug reports, things that are wrong, missing or should be explained more thoroughly in the user's guide – any feedback is highly welcome. Furthermore a couple of experimental modules are available upon request. Please also consider including the author on the first published paper or abstract which uses *Lupe* for a significant portion of the data analysis. In any case I'd be happy to hear about it.

4. Installation instructions

Please follow the installation instructions in README.txt on the CD.

5. Running the program

5.1. Running a .sav file

If you do not have an IDL license, this is the only possibility (and the easiest one): simply start Lupe.sav by double-clicking.

5.2. Running from the development environment

After starting the development environment, type *mrlib_init* in the command line. This should compile the library. Then chose *File→Open Project...* and load Lupe.prj from the Lupe directory. Compile **twice** either via the *Compile project files* button or *Project→Compile→All Files*. Finally run the program via *Project→Run* (or the *Run Project* Button)

6. First start

Upon first start of the program, a couple of settings need to be done once. Please have a DSC perfusion image set in DICOM format from your scanner available (on CD or disk). If you should have a directory which contains many different series, you might want to resort the files first, so the perfusion images are located in a separate directory. If in doubt how to do this, please read the section about sorting of DICOM files in chapter 9.

- Start the program
- Select File→Preferences
- Select the preferred language. A language change will only become effective after the program is started anew. As this introduction refers to the English menu titles, it might be an idea to stick to English – at least for a start.
- Set the default working dir to the directory where your input perfusion series usually are located
- If used in a clinical environment, the option 'allow only one study in memory' should be selected
- Switch to tab *DICOM*
- Click 'Get from DICOM file' and select one DICOM file from your perfusion sequence (actually any DICOM file from your scanner should do). This will fill the 'Implementation class UID' field.
- Save preferences
- Click 'ok' to close the preferences window
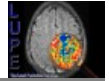
- Load a perfusion image set via File→Load perfusion series. The program expects all images of one series in the same directory. In the file-selector box, select **one** file from the perfusion

- input series. The program will load the sequence and calculate baseline and concentration images.

- If loading should fail, please send a bug report, including the error message and, if possible, one file of your input series.

- Select Global AIF methods → Select AIF

- Using the vertical slider, select a middle slice

- The two vertical lines in the two uppermost drawing windows denote the region used for baseline calculation. Moving the mouse cursor over the perfusion image will display the signal and concentration curves of the voxel under the cross-hair. Make sure that the region for baseline calculation does not extend into the signal drop. If necessary correct the baseline boundaries in File → Preferences, Tab Default (and save). Then recalculate the baseline and concentration images via Global AIF methods → Calculate concentration images (this will make the AIF module close itself). Reopen the AIF selection module.

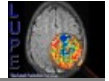- Done. If you changed the language, restart the program.

7. Image sets

The central data structure via which the individual parts of the program communicate is the '**image set**'. Examples for image sets in the program are the full input perfusion series, the set of calculated blood flow images, the set of baseline images, or even the global AIF (a somewhat peculiar image set, consisting of just one image of [1 x number of time points] pixels).

Each image set is identified via two 'names' - a *class identifier* and a *name identifier* or shorter *class* and *name*. *Class identifiers* are used for grouping image sets – e.g. all input data sets have *class identifier* 'raw', all calculated parametric maps are of *class* 'result', color coded image sets have *class* 'color'.

Within one class, the image sets are identified via their name. The following table contains the specifiers of all image sets in the program - where *name* is printed in *italics*, it is only a default name suggested by the program. You are free to chose these to your liking, e.g. when performing the same calculation with different calculation parameters.

Class	Name	Description
raw	signal	the DSC input data set from your scanner consisting of NxM images - N slices and M timepoints
raw	signal_xyt	the input data set slightly rearranged - internally this allows for more efficient calculation expressions. Don't bother. 'raw/signal' and 'raw/signal_xyt' hold identical data.

Class	Name	Description
raw	baseline	a set of N images with the baseline image for each slice.
raw	ctr_xyt	as signal_xyt, but contains the calculated concentration images
aif	global	the global AIF, just one image of size 1 x number of time points.
result	*rCBV sSVD*	regional relative cerebral blood volume parametric map, calculated via standard SVD
result	*rCBF sSVD*	regional relative cerebral blood flow parametric map, calculated via standard SVD
result	*MTT sSVD*	mean time of transit - the ratio of rCBV/rCBF calculated from standard SVD
result	*rCBV rSVD* *rCBF rSVD* *MTT rSVD*	as above, but uses Smith's reformulated SVD for calculation of the parametric maps
result	*rCBV BBBlc* *rCBF BBBlc* *MTT BBBlc*	maps are calculated using a correction algorithm for contrast agent that leaked into interstitial space - BBBlc stands for blood-brain-barrier leakage correction.
result	*T1* *Permeability*	the BBBlc algorithm also calculates a T1 and a Permeability map
color	*rCBV sSVD color*	a color coded version of *rCBV sSVD*
color	and so on...	the color coded versions of the result image sets have the same name with the word *color* appended.

8. Typical workflow

A typical workflow could look like the following. Please refer to the description of the single modules for more detailed information.

1. File → Load perfusion series

    → a file selector box comes up

2. Select one file from your perfusion series

    → the program loads the sequence, reorders the images and then calculates baseline and concentration image sets

3. Global AIF methods → Select AIF

    → The image browser opens

    → Moving the mouse over the image will display the time and concentration curves in the two topmost displays on the right. The two middle displays show the AIF curves from

points currently defined in this slice. The lower two displays show the currently defined AIF, averaged from all defined AIF points.

→ Define the AIF via the ROI menu above the display. An arbitrary number of AIF pixels in any slice may be defines.

→ Loading and saving of the AIF is possible via the load/save icons of the ROI UI.

→ Alternatively the program may suggest AIF pixels in the currently displayed slice via the button 'Suggest AIF'. Delete unwanted AIF pixels via the ROI UI. It might happen that the program is not able to find any AIF pixels - in this case chose the AIF by hand.

4. Global AIF methods → Standard SVD

→ Calculated result image sets will be stored with class name 'result' - the set name can be chosen by the user

→ 'calculate' performs the calculation

5. Tools → Color encoding

→ Chose the source data set for encoding

→ Set the destination name (class name for color encoded images is 'color') or accept the default

→ For logarithmic color encoding set a ROI that defines the mean value, all other color encodings are linear

→ 'Store' will store the currently defined ROIs for re-use in the next set, 'recall' will recall the last stored ROI set.

6. File → Save results

→ Chose the file format

→ When saving in DICOM format, set series numbers for all sets or set the first series number in the top-most set and click 'Set consecutive series numbers'

→ Set the file basename - filenames are fileBasename_sxx.yyy where xx denotes the slice number and yyy the file format.

→ For DICOM files one may also set a series description stored in the file.

→ Click the 'save' symbol to save and chose the destination directory.

→ 'Resize' images will optionally enlarge the saved images.

9. The single modules (in order of appearance)

9.1. File → Load perfusion series

This will load an image set from a perfusion sequence, that is a series consisting of several image slices, each acquired n times. Input format is DICOM3.0. All images of a perfusion series must be located in the same directory. A file-selector box is presented to the user who selects **one** Dicom file that belongs to the series. The module will then read in all files from the same directory and keep

the ones that belong to the same series, i.e. have the same series number. Number of slices and number of timepoints are automatically determined. Measurement data needed for the calculation (TE, TR) are also taken directly from the input files.

If set in the preferences, baseline and concentration images are automatically calculated after loading has completed.

9.2. File → Close patient

Closing a patient will delete all image sets of class 'raw' and 'aif' in memory.

9.3. File → Load result image set

Loads a result image set (that is an image set which consists of several slices i.e. files) in DICOM format. All DICOM files belonging to the image set have to be located in the same directory. The module presents a file-selector box to the user, who selects **one** file from the series. As when loading a perfusion series, the module will then load all files from the same directory and keep those that have the same series number as the selected file.

As explained above, image sets are identified via a *class* and a *name* identifier within the program. The loaded image set is attributed the class identifier 'result'. The *name* identifier has to be chosen by the user though. After loading, the module thus presents a small dialog box, where the user may enter the *name* identifier for this image set. Via the buttons 'Filename' and 'Description', the module may make suggestions based on the filename and the Dicom series description, respectively.

It is not possible to load color-encoded images.

9.4. File → Save results

Result image sets can be saved in DICOM3.0 and in several bitmap formats.

Saving in DICOM3.0 format:

·   Check 'anonymize patient data', if all patient related information shall be deleted from the written files.

·   Images in DICOM files have a value range from '0' to '4095'. However the calculated values are floating point values and possibly (esp. for MTT) quite small. To use the full dynamic range, the image values can be scaled to full Dicom range. All images within one set will be scaled by the same value, i.e. value ratios within one set are kept. Since the calculation routines provide only relative values anyway, this option is switched on by default.

·   If desired, the image matrix size can be enlarged by a constant factor before saving, using different interpolation methods. 'Nearest neighbour' will perform no interpolation, 'linear' and 'cubic' will apply a corresponding interpolation method ('cubic' may sometimes yield strange results – try it yourself).
    This option only effects the images written to disk. The images in memory remain unchanged.

·   Select a new series number for each set that shall be written to disk. To not confuse your PACS system (in case the result images shall be stored there) series numbers must not collide with other series numbers in the same study!
    To make it simple, set the first 'free' series number in the upper right text box and then press the 'Set consecutive series numbers' button.

- Select a file basename for each series or just accept the program suggestion. The filenames will be the basename plus '_sXX.dcm' where XX is replaced by the slice number.

- Chose a series description for each set, or be happy with the program suggestion.

- Save each image set separately by clicking the 'save' button to the right. Select the target directory in the dialog window that pops up and press OK. The images will be stored in that directory.

Saving in bitmap format:

- Chose one of the available bitmap formats – for JPEG files, the quality factor can be set.

- The image matrix size may be enlarged as described above.

- Chose a file basename as described above.

- All other options are irrelevant for bitmap formats. As above, use the 'save' button to the right to save.

9.5. File → Preferences

General:

Language: chose one. The modules in menu 'Dicom tools' have not been translated yet, so they'll always be in English.

The default working dir is the working directory used after program start. Set it to the place where your input files usually are located.

The logfile does not work, yet.

'Child windows are floating' means that extra windows will always be in front of the main window. Checking or unchecking has immediate effects on all newly opened windows. Try it.

'Allow only one study in memory' will ensure exactly that, if checked. This option should be switched on in clinical use to avoid confusion.

Defaults:

The baseline lower and upper boundaries should be set as described in 'things to do at first start' above. Default TR and TE are only needed if your DICOM files should not contain this information – leave the box 'Warn about usage of default values' checked. If you should get a corresponding warning, set these two values. Otherwise they're irrelevant.

'Warn before replacing existing results' will make the program bring up a dialog message, when image sets already exist. If unchecked, the sets will be replaced without further notice. It's save to leave this switched on.

'Automatically calculate concentration images after loading' will do exactly that if checked. Chances are, you want that box checked.

Dicom:

Every program that writes DICOM files has a worldwide unique identifier, the 'Implementation class UID' which is issued by a central organization. As Lupe does not have such an identifier, it uses the one of your machine. See above in 'things to do at first start'. It is a bad idea to leave that value at '0.0....' as the UID is used in the creation of further UIDs throughout the program.

SVD:

Calculation parameters. These values should be ok - only change them if you know what you are doing.

9.6. File → Quit

Guess. Once. Oh, and there is a question dialog box, too.

9.7. Calculations → Calculate baseline images

Calculates the baseline images over the region as specified in the preferences. Baseline images will only be recalculated if the boundaries have been changed – otherwise a message pops up, stating that recalculation is futile.

9.8. Calculations → Calculate concentration images

Calculates the concentration images, which will automatically recalculate the baseline image set, if needed (i.e. if boundaries have changed).

9.9. Calculations → Select AIF

Allows the user to define a global arterial input function from any number of points in any number of slices. If you should not be familiar with the ROI user interface, play around – it is hopefully relatively straight forward. Some comments nevertheless ;-)

· there's only one level of undo for **creation** of ROIs.

· selecting the 'cross' allows the user to delete ROIs. Clicking the 'cross' again while it is already selected will delete all currently displayed ROIs (after a confirmation dialog box).

· the two green blobs in the ROI user interface govern the display style: outlined/filled and transparent/solid. Admittedly the bitmaps could be nicer.

· Choosing the question mark and then a ROI will display some basic ROI information.

· Loading and saving the defined ROIs is possible via the load/save buttons in the upper left. This will actually load/save the full AIF, i.e. all ROIs from all slices, not just the ones currently on display.

The middle mouse button in the right 15% of the screen will change the zoom factor, in the remaining part of the screen it allows to pan the image. Double clicking will reset the zooming factor or center the image.

The right mouse button changes the gamma value (right 15% of screen) or center/width settings (remaining part). Double click resets either gamma or C/W value.

When the display has the focus (right click into it), the cursor keys may be used to change slice and time. 'c' will toggle some image information on and off. 's' will switch between the display of 'signal' and 'concentration' images.

While moving the cursor over the image, the two uppermost windows to the right will display the signal and concentration curves of the voxel under the cursor. The two middle windows display the corresponding curves for all ROIs defined in the currently displayed slice. The lower windows display the currently defined AIF – averaged over all ROIs (optionally weighted by the number of

points in each ROI) from all slices.

The two vertical lines in the signal window denote the upper and lower boundaries for the baseline calculation as they were used in the last calculation of baseline images. They do not reflect the settings from preferences. The small circle shows the current time point (try using the cursor keys).

Note that the middle windows also display a ROI **while** it is defined. That allows the user to compare the curve of the ROI he is about to set to the curves of already existing ROIs.

Suggest AIF will make the module suggest up to 10 AIF points, based on baseline, signal drop, as well as height, area and smoothness of the concentration curve. Usually one will remove all but a few of the suggested points.

Clicking 'Close module' will make your monitor implode.

9.10. Calculations → Standard SVD

Performs a perfusion calculation using simple singular value decomposition. Result image sets will be stored in class 'result' with the names as specified at the upper right. Otherwise, just hit the 'Calculate' button. Calculation will typically take on the order of a few seconds – depending on the number of slices and time points, the type and number of processors in your machine, the number of other programs you run in parallel, the amount of memory in the very same machine, the battery status, if you should run this on a laptop, and, most of all, how much you are in a hurry. Murphy's law applies.

### 9.11. Calculations → Reformulated SVD

Performs a perfusion calculation based on the reformulated singular value decomposition. Smith suggests a time shift around 10s – for a TR of 1.5s this implies a shift of ~ 7 time bins. Your mileage may vary. Depending on the time shift, this calculation allows you to go and get a coffee while waiting.

9.12. Calculations → BBB leakage correction

The blood brain barrier leakage correction algorithm tries to correct for T1 effects caused by contrast agent which leaked into interstitial space. Usually the correction will affect only few pixels in the image – please see Haselhorst's paper. Along with the flow and volume maps, additional approximate T1 and Permeability images are calculated. Otherwise 'Calculate' will do.

Please note that the blood-brain barrier leakage correction involves assumptions and approximations that are dependent on the pulse sequence, field strength and contrast agent. It is up to the user to adapt the algorithm to the locally used protocol:

  · The correction algorithm includes an approximate calculation of T1 in each pixel, and this approximation becomes less accurate at higher field-strength (e.g. 3T), due to the generally longer T1 relaxation times. Please also make sure that the very first acquired images in the time series are not disregarded by the system. Images required before reaching of steady state are required for the T1 calculation.

  · The correction algorithm also contains approximate values of the relaxivities (r1 and r2) for the employed contrast agent. These values depend on the specific contrast agent used, as well as on the employed field strength.
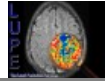
- The algorithm is designed for gradient-echo pulse sequences.

While rSVD allows you to take a coffee brake while calculating, BBBlc allows you to take a trip to Bolivia to pick the coffee. There's no 'cancel' button. You have been warned.

9.13. Tools → Image browser

Allows you to browse all image sets in memory, however only 2-dimensional images can be displayed. Changes in image sets are immediately reflected (try displaying a color image and then change the color encoding with the image browser open).

There's the possibility for some ROI analysis and the deletion of image sets. To save a set of ROIs, press the Store Button – the stored ROIs can then be applied to a different image set via the Recall Button. The cursor keys may be used to change slice and (if available) shot number, 'c' toggles the display of image information next to the cursor. As usual, the display needs keyboard focus for the keys to work – right-click into the display for this.

Mouse buttons work as in the AIF selection.

Dicom information may be displayed as well. When displayed in green, the image set is a Dicom set. If displayed in red, the shown information is taken from the image set from which the currently displayed set was derived.

9.14. Tools → Color encoding

A module for color encoding of result image sets. The color encoded images are stored in class 'color' with the name as specified in the field 'target name'. 41 out of the 42 color-schemes are linear and can be directly applied.

The *logarithmic* colorscheme requires a reference value defined via a ROI. In Lund we normalize the colorscheme to the cerebellum, i.e. draw two extended ROIs in the left and right hemisphere of the cerebellum. If the color encoded images should look too blueish or whitish, change the log stretch factor. The arrows will change in steps of 0.5. If smaller steps should be needed, it is possible to enter a stretch factor directly.

Don't use the 'Continuous update' option, unless you have machine faster than those currently available on the market.

'Storing' a ROI allows you to retrieve the same ROI for color encoding of the next set without laborious loading/saving of the ROIs.

9.15. Tools → Image arithmetics

This module will only work if you have an IDL license.

Image arithmetics allows for arbitrary arithmetic operations on single images or full image sets. Chose image sets A and B from the drop-down menu. Chose a result set name **and confirm by pressing return** (your machine will go 'bing' here). Chose a formula from the drop-down list or type in one and again confirm by pressing return. Optionally the resulting image(s) may be scaled in size as specified. Click 'calculate' to perform the calculation. NaN or Infinite values in the result image are automatically set to '0.'.

Variables that can be used in the formula are 'a' and 'b' for the image sets and 'maxa' and 'maxb'. When applying the formula to single images, 'maxa' and 'maxb' are the maximum of the respective

images. When applying the formula to a full image set, 'maxa' and 'maxb' represent the overall maximum of the respective image sets.

NaN (not a number) values in the result image are automatically set to '0'. That means that expressions like 'A/B' will work, even if B contains pixels of zero value.

This module can be used for a couple of things:

· Comparison of different calculation methods (obviously)

· Resizing of image sets: Enter 'A' as formula and chose the desired enlargement factor and interpolation method.

· Normalization of images a la 'A/maxa'

· Cut off of large values, via e.g. 'A < 100'. All images values greater than 100 will be set to 100.

9.16. Dicom tools → File browser

A simple DICOM file browser. When the file list has the focus (click into the list), the cursor keys may be used to browse through the files.

9.17. Dicom tools → Resort DICOM files

If a directory should contain DICOM files from different series, this module can be used to sort the files, based on their series number. Chose a source and a target directory – in case of the target directory you also have to provide a filename, which is however ignored (blame IDL for that one). Then click 'Sort'. The function will also add a '.dcm' extension to the filenames, if none is present.

9.18. Dicom tools → Anonymize directory

This module allows you to anonymize DICOM files. All files from the source directory are copied to the target directory while patient related data is removed.

9.19. Help → Modules

Displays the current versions of Lupe, the underlying Lund IDL library and all modules. Only useful in case of bug reports/version questions.

9.20. Help → About Lupe

Copyright and license.

- F I N -

(Closing remark: This document was prepared with OpenOffice Writer, which – after fighting Microsoft Word for too long about how to number paragraphs – I highly recommend. OpenOffice is available for free from http://www.openoffice.org and is a.o. able to read/write files in Word format and to directly generate PDF documents.)