

User Documentation for AppAffineRegistration

Jan Owoc

May 2006

There are three elements that are common to most projects requiring registration.

1. Pick points for a first guess
2. Run the registration
3. Transform the data

Picking Points

The picking of the points is currently best done with vis. The process of picking points is beyond the scope of this user documentation.

Registration

Following the point picking, the transformation needs to be calculated. The registration program will only take Euler angles for rotation, so converting is needed if the point-picking program does not return these by default. There is a converter program in the registration directory that will read in the 3x3 rotation/scale component of the matrix and using SVD return the three Euler angles and the scale factor. This transformation, along with other parameters, needs to be input into a configuration file. Below is a sample configuration file.

1. verbose	0
2. mode	0
3. generate_logfile	2
4. log_file_name	log.txt
5. fixed_image_filename	fixed.aim
6. moving_image_filename	moving.aim
7. output_image_filename	out.aim
8. crop_output	1
9. initial_guess_trafo_axis_x	0.026907
10. initial_guess_trafo_axis_y	0.053858
11. initial_guess_trafo_axis_z	0.214978
12. initial_guess_trafo_trans_x	2.54638
13. initial_guess_trafo_trans_y	-0.331382
14. initial_guess_trafo_trans_z	1.00001
15. initial_guess_trafo_scale	1.00017
16. optimizer_start_learningrate	1e-4
17. optimizer_max_learningrate	1e-2

18.optimizer_scales_translation	1
19.optimizer_scales_rotation	1
20.optimizer_scales_scale	1
21.metric_type	1
22.number_of_histogram_bins	50
23.number_of_spatial_samples	0.20
24.multi_resolution_levels	3
25.max_number_of_iterations	250
26.minimum_change	1e-6

General Settings

1. This setting adjusts how much information is sent to the screen during the normal operation of the program. If set to “1”, details about the metric value and transform are output for every iteration of the registration, as well as some other debugging information. Note that outputting information to the console tends to slow a program down, in many cases quite significantly. If set to “0”, only every tenth iteration has its parameters output (to facilitate the monitoring the progress) and then the final transformation.
2. The mode setting allows the user to run only registration (“1”), only the transformation (“2”), or both (“0”).
3. This setting allows the user to either generate a text log file of the registration (“1”), or not (“0”). The log file contains information about the registration, both the initial and final conditions. If set to “2”, it also outputs details on each iteration.
4. As the name implies, this is the name of the log file. If the file exists, the output will be appended to the file, so no data is lost.

I/O Settings

5. This is the path to the fixed image. It can be either an absolute path, or a path relative to the working directory of the program.
6. This is the path to the moving image. It can be either an absolute path, or a path relative to the working directory of the program.
7. This is the path to the output image. It can be either an absolute path, or a path relative to the working directory of the program. It is much faster to transform the image while it is already in memory. The output image will be cropped to the size of the fixed image.
8. In most of the examples of registration with ITK, the transformed moving image is written out such that it is put in the same format as the fixed image. This option would allow a user to not crop the image to the dimensions of the fixed image (potentially losing information) but rather fill the non-existing space with emptiness. There was insufficient

time to properly implement this, so changing the setting does nothing.

Transform Guess Settings

9,10,11. These are the Euler angle rotations about the x, y, and z axis respectively. Note that this is not a value in radians, but rather the sine of the angle. Setting this above 1 (or below -1) will not produce any warnings, but the registration will likely not converge.

12,13,14. These are translations in the x, y, and z directions respectively. These are in the same units as the spacing of the image, and therefore usually in mm.

15. The isotropic scaling factor.

Optimizer Settings

16. The starting learning rate for the registration. A typical value is 0.0001.

17. The maximum learning rate allowed, to prevent step sizes from getting too large. A typical value is 0.01.

18. The way to think of the scale for translation, is as a ratio to rotation. If left at “1”, this tells the optimizer that a translation of 0.1 mm is roughly equivalent in terms of step size to a rotation of 0.1 radians (about 6 degrees). This is the case for images on the order of a few millimeters (1-5 mm) across. For every order of magnitude the image is larger, this value needs to be decreased by an order of magnitude. For every order of magnitude the image is smaller, this value needs to be increased by an order of magnitude. For example, if imaging a human skull (around 200 mm across), this value should be around 0.01.

19. This is the scale for rotation. It is recommended to keep it at “1”, and adjust the other two scales. The only exception to this rule is if you know (for some reason) that there will be no rotation, in which case this value can be increased significantly, to about 10^{16} .

20. This is the scale for scale factor, and is to be thought of similarly to the scale for translation. If left at “1”, the optimizer assumes a change of 0.1 in scale factor is equivalent (in terms of step size) to a rotation of 0.1 radians. For most images, a value of “1” will be perfect. If a rigid registration is desired, the optimizer needs to be told that any changes to scale factor are an unacceptably large step. This is accomplished by setting the factor to 10^{16} .

Metric Settings

21. This setting allows the user to select a metric. The metrics currently implemented are Mattes Mutual Information Metric (“1”), Normalized Correlation Metric (“2”), and Mutual Information by Viola/Wells (“4”).

22. This is the number of histogram bins to be used with Mattes Mutual Information Metric. Studies have shown that “50” is the optimal number. Unless there is a very good reason to change it, the value should remain at “50”. This field is only read if the metric

is set to Mattes.

23. This is the number of spatial samples used by Mattes Mutual Information Metric, as a percentage of the image size. Increasing this number will increase the accuracy of registration, whereas decreasing this number will increase the speed. Studies have shown that a value of 20% (or “0.20”) gives the best trade off between speed and accuracy. It is also the absolute number of spatial samples used by Viola Wells Mutual Information Metric. A value of 50 is good. This field is only read if the metric is set to Mattes or Viola Wells Mutual Information.

Other Settings

24. This is the number of multi-resolution levels to use. Multi-resolution registration saves computer time by performing a coarse registration at a lower resolution, followed by fine-tuning at full resolution. Using more levels is better, but there is a maximum. The scale-down for the levels is in factors of 2 in each dimensions, or by a factor of 8 each time. Depending on the object being registered, there is a minimum resolution at which the down sampled object still resembles the original. As a general guideline, images over 1 GB in size can use 6 levels, images over 100 MB can use 5 levels, images over 10 MB can use 4 levels, images over 1 MB can use 3 levels. If the image is less than 1 MB, performing multi-resolution registration with more than 2 levels will yield unforeseen results.

25. This is a hard-coded maximum number of iterations per level, regardless of whether convergence has been reached or not. A good number is 500, as the hopefully the converge criteria will be met by this point.

26. This is the minimum change considered significant. Once the image has been registered as well as it can, the transform parameters will undergo “Brownian motion” about a central point. This number needs to be representative of the difference considered as “Brownian motion” as opposed to a valid step. A typical value would be 10^{-6} .

The program is run with the only parameter being “-config configuration.txt”. Everything else that the program needs will be given through the configuration file.

Transformation

Once the desired transformation is found, almost always one wishes to transform the moving image to the coordinates of the fixed image. As the image and transform are already in memory, it is convenient to write out the image. The registration program will automatically transform the moving image to the coordinates of the fixed image and write out an image file. The output image file will be cropped to the dimensions of the fixed image (until another option is implemented).