

The definitive version is available at <http://www.springerlink.com/>.

F.P. Vidal, J. Louchet, J.-M. Rocchisani, and É. Lutton: New genetic operators in the Fly algorithm: application to medical PET image reconstruction. In *Applications of Evolutionary Computation*. Volume 6024 of Lecture Notes in Computer Science, pages 292-301, Istanbul, Turkey, April 2010. Springer, Heidelberg. **Nominated for best paper award.**

DOI: 10.1007/978-3-642-12239-2_30

```
@inproceedings{Vidal2010EvoIASP,
  author = {F. P. Vidal and J. Louchet and {J.-M.} Rocchisani and \'E. Lutton},
  title = {New genetic operators in the {Fly} algorithm: application to medical
  {PET} image reconstruction},
  booktitle = {Applications of Evolutionary Computation},
  year = 2010,
  series = {Lecture Notes in Computer Science},
  volume = 6024,
  pages = {292-301},
  month = apr,
  address = {Istanbul, Turkey},
  annotation = {Apr-7--9, 2010},
  note = {Nominated for best paper award},
  abstract = {This paper presents an evolutionary approach for image
  reconstruction in positron emission tomography (PET). Our reconstruction
  method is based on a cooperative coevolution strategy (also called
  Parisian evolution): the ‘‘fly algorithm’’. Each fly is a 3D point that
  mimics a positron emitter. The flies’ position is progressively optimised
  using evolutionary computing to closely match the data measured by
  the imaging system. The performance of each fly is assessed using a
  ‘‘marginal evaluation’’ based on the positive or negative contribution of
  this fly to the performance of the population. Using this property, we
  propose a ‘‘thresholded-selection’’ method to replace the classical
  tournament method. A mitosis operator is also proposed. It is triggered to
  automatically increase the population size when the number of flies with
  negative fitness becomes too low.},
  doi = {10.1007/978-3-642-12239-2_30},
  publisher = {Springer, Heidelberg}
}
```

New genetic operators in the Fly algorithm: application to medical PET image reconstruction

F. P. VIDAL^{1,2,3}, J. LOUCHET³, J.-M. ROCCHISANI^{1,4}, and É. LUTTON¹

¹ INRIA Saclay - Île-de-France/APIS, 4 rue J. Monod 91893 Orsay Cedex, France

² member of Fondation Digiteo (<http://www.digiteo.fr>)

³ Now with University of California San Diego,

Rebecca and John Moores Cancer Center, Radiation Oncology, CA, USA

⁴ Artenia, 24 rue Gay-Lussac, 92320 Châtillon, France

⁵ Paris XIII University, UFR SMBH &
Avicenne hospital, 74 rue Marcel Cachin, 930013 Bobigny, France

Abstract

This paper presents an evolutionary approach for image reconstruction in positron emission tomography (PET). Our reconstruction method is based on a cooperative coevolution strategy (also called Parisian evolution): the “fly algorithm”. Each fly is a 3D point that mimics a positron emitter. The flies’ position is progressively optimised using evolutionary computing to closely match the data measured by the imaging system. The performance of each fly is assessed using a “marginal evaluation” based on the positive or negative contribution of this fly to the performance of the population. Using this property, we propose a “thresholded-selection” method to replace the classical tournament method. A mitosis operator is also proposed. It is triggered to automatically increase the population size when the number of flies with negative fitness becomes too low.

1 Introduction

Image reconstruction in tomography is an inverse problem that is ill-posed: a solution does not necessarily exist (e.g. in extreme cases of excessive noise), and the solution may not be unique. This problem can be solved as an optimisation problem, and on such cases, evolutionary algorithms have been proven efficient in general, and in particular in medical imaging [2, 4, 13]. We focus here on positron emission tomography (PET) reconstruction in nuclear medicine.

Nuclear medicine appeared in the 1950’s [1]. Its principle is to diagnose or treat a disease by administering to patients a radioactive substance (also called tracer) that is absorbed by tissue in proportion to some physiological process. When a pathology occurs, the metabolism most of the times increases: there are more molecules in the pathology area, i.e. the radioactivity also increases.

It is possible to reconstruct slices through the human body using methods similar to those used in conventional X-ray computed tomography [7]. In nuclear medicine, this method makes use of a gamma emitter as radio-tracer. It is called Single-Photon Emission Computed Tomography (SPECT). The reconstruction allows to recover the 3D distribution of the tracer through the body.

The other main tomographic technique in nuclear medicine is PET. Here a positron emitter is used as radionuclide for labelling, rather than a single gamma emitter. Positrons are emitted with high energy (1 MeV). After interactions, a positron combines with an electron to form a positronium. Then the electron and positron pair is converted into radiations. It is the annihilation reaction, which generally produces two photons of 511 keV emitted in opposite directions. Taking

advantages of this property, this radiation is detected in coincidence, i.e. using the difference in arrival times of the detected photons of each pair, and considering that each annihilation produces two photons emitted in exactly opposite directions. The line between the detectors that have been activated for a given pair of photons is called “line of response” (LOR). Prior to the reconstruction, the LOR data is often rebinned into a sinogram [5, 8]. This intermediate data representation corresponds to projection data that can be used by conventional tomographic reconstruction codes. A broad overview of reconstruction methods using projection data in nuclear medicine can be found in [8, 14].

The PET reconstruction methods are often divided into two classes: i) analytical methods, and ii) iterative statistical methods. Analytical methods are based on a continuous modelling and the reconstruction process consists of the inversion of measurement equations. The most frequently used is the filtered back-projection algorithm (FBP) [7]. Statistical methods are based on iterative correction algorithms. These include the most widely used techniques in SPECT and PET, such as the maximum-likelihood expectation-maximization method (ML-EM) [10] and its derivative, the ordered subset expectation-maximization algorithm (OS-EM) [6].

In a previous paper, we showed that a cooperative coevolution strategy (also called Parisian evolution) called “fly algorithm” [9] could be used in SPECT reconstruction [3]. Here, each fly corresponds to a 3D point that is emitting photons. The evolutionary algorithm is used to optimise the position of flies.

However, PET has taken over SPECT in routine clinical practice. Effort has therefore been made to propose an efficient evolutionary scheme that takes into account PET data acquisition principles [11, 12], but these were still restricted to low resolution PET scanners in 2D-mode. This paper describes our current research activities aimed at providing an effective method in both 2D or fully-3D mode, and it describes recent developments, such as i) the introduction of “thresholded-selection” replacing the traditional “tournament selection” and ii) taking advantage of the thresholded-selection to increase the population size (when the number of flies, whose fitness is negative, is too low), i.e. improve the statistics of the final image. The following section gives an overview of the methodology. The results and performance of our method using numerical phantoms are presented in Section 3. The paper ends with a conclusion that discusses the work that has been carried out and it provides directions for future work.

2 Material and Methods

2.1 Main Principles

Each individual, or fly, corresponds to a 3D point that mimics a radioactive emitter, i.e. a stochastic simulation of annihilation events is performed to compute the fly’s illumination pattern. For each annihilation event, a photon is emitted in a random direction. A second photon is then emitted in opposite direction. If both photons are detected by the scanner, the fly’s illumination pattern is updated. The scanner properties (e.g. detector blocks and crystals positions) are modelled, and each fly is producing an adjustable number of annihilation events. Each fly keeps a record of its simulated LORs. Therefore the result of these simulations consist of a list, per fly, of pairs of detector identification numbers that correspond to LORs. These lists are aggregated to form the population total illumination pattern.

Initially, the flies’ position is randomly generated in the volume within the scanner. Using genetic operations to optimise the position of radioactive emitters, the population of flies evolves so that the population total pattern matches measured data. The final population of flies corresponds to the tracer density in the patient, i.e. the reconstructed data. Note that cross-over operations are not used in this application. In our context, the result of such an operation may lead to meaningless results, e.g. in the case of cross-over between two flies of two distinct objects (the new fly will be wrongly located in between). Only mutation and immigration (i.e. a fly is created at a random position) are used.

2.2 Fitness Metrics

The fitness metrics corresponds to a distance measurement between the simulated data and the actual data given by the imaging system. City block distance provides a good compromise between accuracy and speed. Note that smaller the population’s cost is, closer the simulated data is to the actual data.

In [3], we showed that, when we were addressing the SPECT problem, if we defined the fitness of a fly as the consistency of the image pattern it generates, with the actual images, it gave an important bias to the algorithm with a tendency of the smaller objects to disappear. This is why we then introduced marginal evaluation ($F_m(i)$) to assess a given fly (i). It is based on the leave-one-out cross-validation method. We use a similar approach in PET:

$$F_m(i) = \text{dist}(pop, input) - \text{dist}(pop - \{i\}, input) \quad (1)$$

with $F_m(i)$ the marginal fitness of Fly i , $\text{dist}(A, B)$ the city block distance between two tables A and B , pop is the set of LORs simulated by the whole population, $input$ is the set of LORs extracted from the input data, and $pop - \{i\}$ is the set of LORs simulated by the whole population without Fly i . The fitness of a given fly will only be positive when the global cost is lower (better) in presence rather than in the absence of this fly. We therefore used a fixed threshold to operate selection.

2.3 Thresholded Selection

At each iteration of the evolution loop, a fly has to be killed, and a fly may be used during the mutation. We saw in the previous section that the fly’s fitness is its own contribution (positive or negative) with respect to the whole population. We take advantage of this principle as follows: i) any “bad” fly (its fitness is negative) is a candidate for death, and ii) any “good” fly (its fitness is positive) is a candidate for mutation. When a fly is killed, its LORs are removed from the total set of simulated LORs. When a new fly is created, its LORs are added. This process needs to be fast to be able to decrease the number of bad flies and increase the number of good flies as much as possible.

2.4 Mitosis

To obtain accurate, high resolution images it is necessary to use large populations of flies. However, as processing time is roughly proportional to the number of flies to be processed, we choose a simple scheme that begins with a small population, then multiply the population along the algorithm execution using a mitosis process: each fly is duplicated. Newly created flies will have their own illumination pattern.

It is triggered whenever the number of flies with a negative fitness gets too low. In practice, at each step in the steady state process, one fly is chosen randomly and its fitness tested: a genetic operator will only be applied if the fly is bad. We launch the population mitosis every time 50 consecutive flies are found with a positive fitness.

3 Results

We have developed numerical phantom models to assess the reconstruction algorithm. To date, no scattering and no tissue attenuation have been considered. Whilst this is not physically correct, it allows us to test and validate our approach in the simplest cases. First, we present quantitative results in 2D-mode, both in low and high resolutions. Then, qualitative results in 3D-mode are presented using a complex object. For each test case, the initial population is 5,000 flies. When the current number of flies of the population is above a given threshold, e.g. 1^6 or 2^6 , the evolution loop is stopped whenever the number of flies with a negative fitness becomes too low, i.e. the stopping criteria is similar to the mitosis criteria. 70 LORs per fly have been simulated in the 2D




Figure 1: Simulated PET System: a single ring of 72 linear blocks that include 8 crystals; two boxes ($7 \times 7 \times 0.4 \text{ cm}^3$ and $10 \times 10 \times 0.4 \text{ cm}^3$) with the same radioactivity concentrations ($\sim 930.000 \text{ counts/ml}$).




Figure 2: Slices (512×512 pixels) through the cubes.

test cases. The probability of LORs to be detected is much higher in the 3D case. Thus only 10 LORs per fly have been simulated in the 3D case.

3.1 2D-Mode

3.1.1 Test 1: large objects with different sizes and similar radioactivity concentration.

The purpose of this test was to assess the ability of the algorithm to retrieve relatively large objects, whose sizes are different, but with the same radioactivity concentration. Fig. 1 shows the simulated set up. The phantom is made of two boxes ($7 \times 7 \times 0.4 \text{ cm}^3$ and $10 \times 10 \times 0.4 \text{ cm}^3$) with the same radioactivity concentrations ($\sim 930.000 \text{ counts/ml}$). The simulated PET system is made of a single ring of 72 linear blocks that include 8 crystals. To evaluate the results, a 512×512 pixel slice is produced (see Fig. 2). Note that the typical image size in PET is 128×128 pixels. The slices are post-filtered using a gaussian convolution kernel, then linearly rescaled between zero and one. Profiles in this reconstructed image are compared to corresponding profiles in the phantom data (see Fig. 3). For both boxes, the reconstructed data seems to be close to the input data. Full width at half maximum (FWHM) is also measured to quantify errors (see Table 1). These results show that our evolutionary scheme is able to accurately recover the width of our test objects.

3.1.2 Test 2: small objects with different sizes and radioactivity concentrations.

This test case has been designed to assess the ability of our algorithm to detect small objects, and their relative radioactivity concentrations. Fig. 4(a) shows nine cylinders having two different radii (1 cm and 2.5 cm) and five different radioactivity concentrations ($C_1 = 114,590 \text{ count/ml}$,

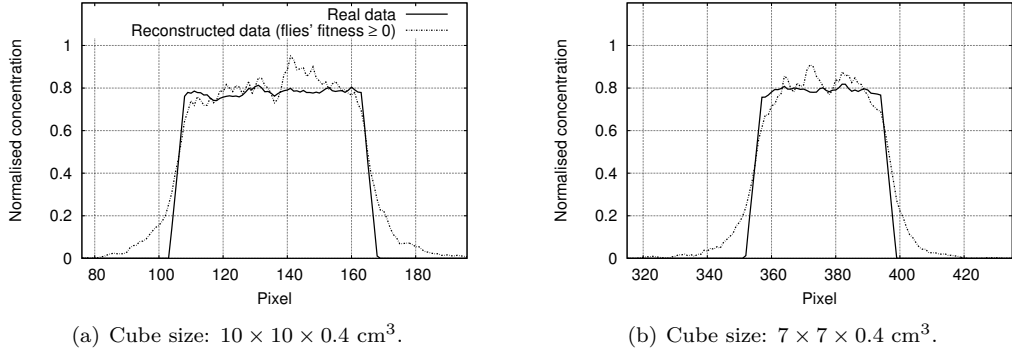


Figure 3: Profiles extracted from Fig. 2.


Table 1: FWHM estimated from Fig. 3

FWHM from phantom model (in mm)	FWHM from slice (in mm)	Relative difference (in %)
71	72	1
99	99	0

$C_2 = 2C_1$, $C_3 = 3C_1$, etc.). A low resolution PET system has first been considered. It is made of a single ring of 72 linear blocks that include only 1 crystal.

To evaluate the results, a 512×512 pixel slice is produced once again (see Fig. 4(b)). The reconstructed data appears to be visually close to the input data. In particular, the size and concentration of cylinders are visually well preserved.

To estimate the diameter of each cylinder, horizontal profiles have been extracted so that they crossed the cylinders in their respective centre (see Fig. 5). As the lower profiles are symmetrically similar to the upper profiles, they are not plotted here. FWHM is measured once again to quantify errors (see Table 2). Let Object i be the cylinder whose concentration is C_i in the phantom model and C'_i in the reconstructed slice. Whilst the profiles in the reconstructed slice seem to match respective profiles in the phantom model, error measurements in FWHM are relatively high for the



(a) Phantom model. (b) Data reconstructed from low res- (c) Data reconstructed from high res-
olution scanner. olution scanner.

Figure 4: Slices (512×512 pixels) through the cylinders.

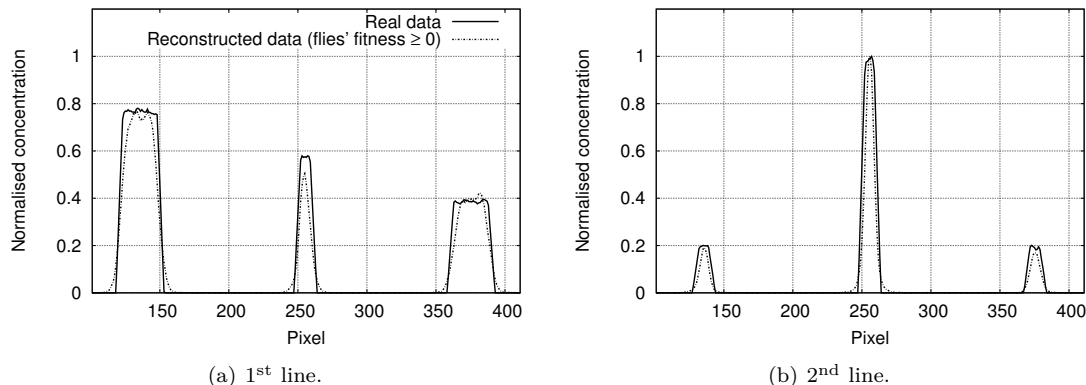


Figure 5: Profiles extracted from Fig. 4(b).

Table 2: FWHM estimated from Fig. 5 (using a low resolution PET system), and Fig. 6 (using a high resolution PET system).

Object	FWHM from phantom model (in mm)	FWHM in Fig. 5 (in mm)	Relative difference in Fig. 5 (in %)	FWHM in Fig. 6 (in mm)	Relative difference in Fig. 6 (in %)
1	19	13	31.6	18	5.7
2	49	43	12.2	48	2.3
3	19	14	26.3	18	6.8
4	49	44	10.2	47	3.2
5	19	12	36.8	17	8.6

smallest cylinders (up to 35%). To investigate the influence of the reconstructed slice resolution with respect to the low spatial resolution of the PET system, the test case presented in the next section makes use of similar objects and a PET scanner with higher spatial resolution.

To assess the validity of the radioactivity concentration within cylinders, the average value at the centre of each cylinder has been measured in Fig. 4(b) (see Table 3). We compare the respective ratio of the different concentrations to the lower value (C'_1), so that we can compare the reconstructed values with the theoretical values. In theory, we should get $C'_2 = 2C'_1$, $C'_3 = 3C'_1$, etc. Table 3 shows that the relative concentrations have been preserved in the reconstructed slice. However, the maximum relative error is about 16.50%.

3.1.3 Test 3: higher scanner resolution.

The previous test case shows that our algorithm is able to retrieve the respective size of objects and their respective concentration. However, relative errors can be as high as 35% for the FWHM and 16,5% for the concentration. In this test case, similar objects have been simulated. The size of crystals has been reduced so that their width matches the width of real crystals. A similar methodology is used to assess the results: i) Fig. 4(c) shows a 512×512 pixel slice that has been reconstructed, ii) profiles have been extracted (see Fig. 6), iii) FWHM estimated (see Table 2), and iv) the concentrations assessed (see Table 3). These results show that using a high resolution PET scanner reduced the maximum error in object size by a factor of 4. This is due to the improvement of the spatial resolution. On the other hand, errors in radioactivity concentration have not been reduced.

Table 3: Relative radioactivity concentration estimated from Fig. 4(b) (using a low resolution PET system), and Fig. 4(c) (using a high resolution PET system).

Object	Relative concentration in Fig. 4(b)	Relative error in Fig. 4(b) (in %)	Relative concentration in Fig. 4(c)	Relative error in Fig. 4(c) (in %)
1	C'_1	N/A	C''_1	N/A
2	$2.13 \times C'_1$	6.5	$2.17 \times C''_1$	8.4
3	$2.67 \times C'_1$	16.5	$3.19 \times C''_1$	9.5
4	$3.80 \times C'_1$	10.0	$4.40 \times C''_1$	19.9
5	$5.02 \times C'_1$	1.0	$5.35 \times C''_1$	17.5

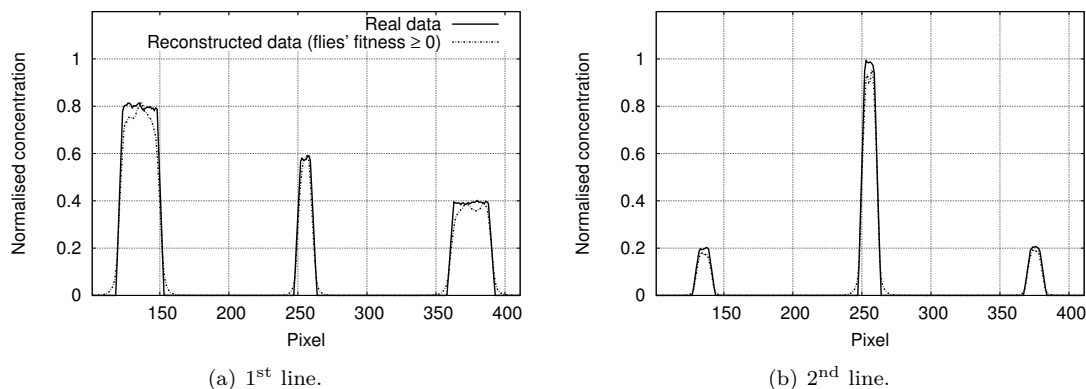


Figure 6: Profiles extracted from Fig. 4(c).

3.2 3D-Mode

The last case has been performed in fully-3D, i.e. the PET imaging system is made of a stack of detector rings. For a coincidence event, the two photons of a LOR can be detected onto different rings. Only visual results are presented here.

A complex 3D shape is used in this test. The simulation is performed using a polygon mesh (here we use the dragon model from *The Stanford 3D Scanning Repository*, <http://graphics.stanford.edu/data/3Dscanrep/>, last access 17 Jan 2010) that is uniformly filled with radio-tracers (see top row in Fig. 7). Then, LORs are recorded in fully-3D mode. Finally, we run our evolutionary reconstruction scheme. Note that the reconstruction algorithm is similar in both 2D and 3D modes. The only difference is the geometrical property of the simulated PET scanner. The bottom row in Fig. 7 presents the reconstructed dataset after volume rendering. One can visually distinguish the shape of the dragon from the population of flies.

4 Conclusion

It may occur that complex applications fuel fundamental technical developments. In the research presented here, we addressed a complex problem that had never been approached in the past using evolutionary computing, by transposing the Fly Algorithm technique originally developed in a stereovision context. We then faced several difficult issues which encouraged the development of new tools that can probably be used into other application fields in evolutionary computing. Using the ‘marginal fitness’ concept opened the way to using a simplified thresholded selection, which in turn allowed to introduce the mitosis operator that duplicates the population whenever




Figure 7: Top row: simulated object; bottom row: volume rendering of the reconstruction.

the proportion of individuals with a negative contribution to the global fitness becomes too low, thus periodically reviving the efficiency of the classical operators (mutation and immigration).

Preliminary results on tests objects show the validity of this approach in both 2D and fully-3D modes. In particular, the size of objects, and their relative concentrations can be retrieved in the 2D mode. In fully-3D, complex shapes can be reconstructed.

To date, only true coincidence events have been considered. Further work will therefore include the use of more realistic input data (including random events and scattering), which will finally lead to implement the correction of scattering within our algorithm. A comparison study against ML-EM and/or OS-EM methods will also need to be conducted.

References

- [1] R. D. Badawi. Nuclear medicine. *Phys. Educ.*, 36(6):452–459, 2001.
- [2] P. A. N. Bosman and T. Alderliesten. Evolutionary algorithms for medical simulations: a case study in minimally-invasive vascular interventions. In *Workshops on Genetic and Evolutionary Computation 2005*, pages 125–132, 2005.
- [3] A. Bousquet, J. Louchet, and J.-M. Rocchisani. Fully three-dimensional tomographic evolutionary reconstruction in nuclear medicine. In *Artificial Evolution 2007*, volume 4926 of *LNCS*, pages 231–242, 2007.
- [4] S. Cagnoni, A. B. Dobrzeniecki, R. Poli, and J. C. Yanch. Genetic algorithm-based interactive segmentation of 3D medical images. *Image Vision Comput.*, 17(12):881–895, 1999.
- [5] F. H. Fahey. Data acquisition in PET imaging. *J. Nucl. Med. Technol.*, 30(2):39–49, 2002.

- [6] H. M. Hudson and R. S. Larkin. Accelerated image reconstruction using ordered subsets of projection data. *IEEE Trans. Med. Imaging*, 13(4):601–609, 1994.
- [7] A. C. Kak and M. Slaney. *Principles of computerized tomographic imaging*. Society of Industrial and Applied Mathematics, 2001.
- [8] R. M. Lewitt and S. Matej. Overview of methods for image reconstruction from projections in emission computed tomography. In *Proc. of IEEE*, volume 91, pages 1588–1611, 2003.
- [9] J. Louchet. Stereo analysis using individual evolution strategy. In *International Conference on Pattern Recognition 2000*, page 1908, 2000.
- [10] L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Med. Imaging*, 1(2):113–122, 1982.
- [11] F. P. Vidal, D. Lazaro-Ponthus, S. Legoupil, J. Louchet, E. Lutton, and J. Rocchisani. Artificial evolution for 3D PET reconstruction. In *Artificial Evolution 2009*, LNCS. Springer, 2009. To appear.
- [12] F. P. Vidal, J. Louchet, E. Lutton, and J. Rocchisani. PET reconstruction using a cooperative coevolution strategy in LOR space. In *IEEE Medical Imaging Conference 2009*, 2009. To appear.
- [13] K. Völck, J. F. Miller, and S. L. Smith. Multiple network CGP for the classification of mammograms. In *EvoWorkshops 2009*, volume 5484 of *LNCS*, pages 405–413. Springer, 2009.
- [14] H. Zaidi, editor. *Quantitative Analysis in Nuclear Medicine Imaging*. Springer US, 2006.