

A genetic algorithm approach to enhancing the performance of a PET detector array

MNH Cook, SH Connell

University of Johannesburg

E-mail: martin@minpet.co.za

Abstract. The MinPET project aims to locate diamonds within kimberlite by activating carbon within kimberlite, then using Positron Emission Tomography (PET) to image carbon density. Distinguishing small diamonds from the background depends crucially on the accurate reconstruction of detector hit positions. This is subject to two kinds of errors: local errors, where the position of a hit within a particular detector is not accurately constructed from the incoming photomultiplier tube signals, and global errors, where the internal parameters that describe the physical location and orientation of detector pixels do not accurately match reality. Because of the large number of detectors in a full MinPET unit, there are too many parameters involved to feasibly adjust them by hand. We have therefore developed a custom genetic algorithm that iteratively evolves detector parameters in order to optimise the image quality. Results are presented from before and after the optimisation is performed, indicating that image accuracy and resolution are improved. This algorithm could be employed periodically in an industrial setting to automatically correct for detector movements or calibration drift.

1. Introduction

The MinPET project [1] images locked diamonds within coarsely crushed kimberlite (± 10 cm rocks), using Positron Emission Tomography (PET). The kimberlite is irradiated with high energy gamma rays, producing the unstable ^{11}C isotope via a photonuclear reaction. This beta decays, and the positron annihilation leads to back-to-back colinear 511 keV photons. These are detected in coincidence by two planes of position sensitive detectors above and below the kimberlite. A 3D carbon density image is created by back-projecting the lines of response formed by coincident events from opposite detector arrays.

A laboratory scale detector system has been created, with two arrays of pixelated scintillation crystals paired to position sensitive photomultiplier tubes. Fast electronics are utilised with good timing resolution, mostly eliminating any random background arising from separate positron annihilations being detected within the same coincidence time window. There is still a background present due to non-diamond positron annihilation events however. This is due both to homogeneously distributed carbon throughout the kimberlite, and oxygen-15, which is inevitably also produced in the irradiation process. Image fidelity is further degraded by Compton scattering, both within the kimberlite and the detectors. This leads to falsely reconstructed lines of response. The energy resolution is not sufficient to completely eliminate this by only accepting 511 keV photons.

The ability of the technique to differentiate small diamonds from this background depends crucially on the resolution of the detectors. This implies firstly that the detectors must accurately

reconstruct the positions of detected hits, and secondly that the detectors' orientation and position are accurately described within software. We employ a genetic algorithm approach to optimising various detector parameters, to improve performance in both of these areas. We refer to the optimisation of position reconstruction within each detector as "local", as opposed to "global" optimisation, which involves the relative position of detectors with respect to each other. Inspiration for this approach was taken from the track-based alignment procedure used in the ATLAS Transition Radiation Tracker [2], where detector elements are aligned so as to improve the matching of the reconstructed to known physics behaviour.

2. Detector system

The detector system consists of 16 Hamamatsu R2486 photomultiplier tubes (PMTs), with 16 crossed wire anodes coupled for each direction over a voltage divider network, furnishing two x and two y signals. Each scintillation crystal is a 1 cm thick, 5 cm diameter disk-shaped BGO ($\text{Bi}_4\text{Ge}_3\text{O}_{12}$) crystal that is further divided into individually wrapped 5 mm segments (or "pixels"). The detectors are split into an upper and a lower plane to allow coincidence detection.

The crossed wire anodes from the PMTs furnish two voltages for each of x and y , denoted x_a, x_b, y_c and y_d . These are amplified, and are passed to a spectroscopy amplifier. This has a slower output that integrates the incoming signal size to produce a shaped output, and a fast output for timing information. The fast outputs from the upper and lower detector planes are used to determine coincidence, which is then used to gate the data acquisition. In this manner, only coincident events are recorded, allowing a high event rate.

3. Genetic algorithm

Both local and global optimisation as described above are difficult tasks. Firstly, they need to find the global optimum, ruling out simple minimisation strategies which can get stuck in local minima. Secondly, they involve a very large parameter space. As we shall see below, local optimisation is within a 11-dimensional space, and global optimisation involves 6 parameters per detector plus another 10 collective parameters. Finally, the optimisation landscapes are highly chaotic due to statistical noise. This is shown in figure 1, which is a two-dimensional slice of the global optimisation landscape, showing how the detector system performs as the x and y position of one of the detectors is varied. The optimisation technique cannot rely on derivatives in a landscape of this nature.




Figure 1. Detector performance as a function of x, y position of one detectors (lower is better).

Due to the badly-behaved and poorly quantified nature of the optimisation landscape, we have created a custom genetic algorithm. Genetic algorithms are an example of evolutionary algorithms, which are inspired by biological evolution. Evolutionary computation has found increasing use in physics whenever the solution space has made more traditional approaches difficult. Examples include fields as diverse as astronomy and astrophysics ([3], [4]), heat transfer optimisation [5], event selection in high energy physics [6], spectrum fitting in X-ray fluorescence

[7], nuclear reactor design [8] and many others. Reference [9] contains an example where, as in our case, a genetic algorithm was used for image calibration.

In the biological analogy, “individuals” are points in parameter space, and their coordinates are their “genes”. A scoring function is defined that evaluates a given individual in terms of how well the detector system performs when calibrated according to the parameters that define that individual. The optimisation proceeds iteratively in “generations”, where the individuals within each generation are randomly generated from the previous generation, favouring parents with good scores. Genetic algorithms are highly robust, they are able to find global minima, they do not need derivatives and they require no assumptions about the structure of the landscape.

The MinPET genetic algorithm is built on the GALib framework [10], and its basic flow is shown in figure 2. In each generation, pairs of individuals are selected as parents, with a probability determined by their score. Offspring are then created from these parents by one of cloning, crossover or mutation. Cloning directly copies a single parent, in order to preserve a good solution. Crossover is the analogue of sexual reproduction, where genes are swapped between individuals. This allows the algorithm to explore different combinations of existing genetic information, and can break out of local minima to find new minima. Mutation takes a single individual and randomly perturbs each of its parameters. This allows the exploration of a given local minimum. A crossover and a mutation are shown in figure 3. The probability of crossover, mutation or cloning is dependent on the distance in parameter space between the parents, tuned such that parents from different local minima are more likely to create offspring via small mutations, reducing the probability that their offspring are kicked out of their respective minima. After two new individuals are created, there is a random chance of a mutation, randomly varying the individual, in order to introduce new genetic information into the population.




Figure 2. Flowchart of MinPET genetic algorithm.




Figure 3. Example of a crossover and a mutation.

4. Local optimisation

The output signals from the crossed-wire anodes, x_a , x_b , y_c and y_d , are discussed in section 2 above. Once these are suitably linearly calibrated and amplified, the position of the hit is given by the ratio of the signals:

$$x = \left(\frac{x_a}{x_a + x_b} - 0.5 \right) \times d, \quad y = \left(\frac{y_c}{y_c + y_d} - 0.5 \right) \times d \quad (1)$$

where d is the diameter of the detector. This reconstruction, however, can be significantly distorted, as can be seen in the left side of figure 4. Reconstructed hits can be up to 6 mm from

their true positions.

In order to improve performance, we introduce a series of parameters to adjust the reconstructed position. These consist of an overall scaling factor, a rotation around the origin, and parameters proportional to x^2 , y^2 , x^3 , y^3 , xy^3 and yx^3 . These were selected in order to correct for defects noticed in the reconstructed pixel locations, including an expansions of points near the centre, incorrect rotations, and asymmetries between x and y directions observed in some detectors.

Our initial approach to scoring the results relied on employing a 2D peak-finding algorithm to locate each pixel peak, then calculating a chi-squared for the deviation of each peak from the actual pixel position. This was however too unstable when the image was severely distorted and when adjacent pixels were not clearly separated. It was also too computationally expensive. An alternate scoring function was therefore defined as

$$S_{\text{local}} = \sum_{i \in H} \frac{(\mathbf{h}_i - \mathbf{p}_i)^2}{|H|} \quad (2)$$

where H is the set of all detector hits, $|H|$ is the number of hits, \mathbf{h}_i is the measured position of hit i adjusted according to the current parameter values and \mathbf{p}_i is the actual position of the closest pixel centre to \mathbf{h}_i .




Figure 4. Example detector pixel reconstruction before and after optimisation. Grey lines demarcate pixel boundaries. Red crosses denote pixels that should be empty. Axes are in cm.

Figure 5 shows an example of how the scores of a population migrate with successive generations for a detector. As generations go by, the population moves closer together, and there is an overall trend toward better scores.

The optimisation procedure was applied to each detector in turn, with a population size of 500 and 250 generations. In order to judge the effectiveness of the procedure, a sodium-22 point source was moved along a straight path through the detector array at a series of different positions, and the FWHM of the point source in the back-projected PET image was measured. This test was repeated for the best known parameters before the optimisation, and the same parameters with their values after optimisation. The genetic algorithm was able to reduce the FWHM by an average of 10% in the direction of movement, 6% in the vertical direction and 15% in the perpendicular direction, a substantial improvement.

5. Global optimisation

Global optimisation proceeds by varying detector element positions and orientations so as to improve the sharpness of the reconstructed image of a point source. Our initial approach was to

fit a three-dimensional gaussian to the point source peak, and evaluate the standard deviations along each axis. This did not yield good results, firstly because of computational complexity, and secondly because the point source peaks often did not resemble a gaussian closely. We therefore introduce a new method to quantify peak sharpness, S , defined as:

$$S_{\text{global}} = \sum_{(i,j,k) \in W} V(i, j, k)^2 \quad (3)$$

where $V(i, j, k)$ is a three-dimensional histogram that stores the sum of the tubes passing through the voxel labelled by coordinates i, j and k . The summation extends over a window W around the point source coordinates. By taking the square of the voxel counts, the score is increased both by sharpening the peak corresponding to the point source, as well as ensuring that as many tubes as possible pass through the fitting window W . In a sense the score can be thought of as a χ^2 fit for a flat background, and a good, sharp peak is one that fits this flat background as badly as possible. Minimisation proceeds by evaluating the inverse, $1/S$.

For a single point source in one position, some nonsensical solutions will give rise to good scores, such as all detectors placed exactly on top of each other. We therefore insist on covering the whole volume between the detector by sending a series of point sources on different constant velocity trajectories on a conveyor belt through the detectors. The algorithm simultaneously optimises all data sets, with the global score defined as the average of the individual track scores defined in equation 3. In order to form a point source image, each line of response's position must be extrapolated back in time based on the velocity and its time stamp. This is highly dependent on accurately quantifying the velocity. This was measured as accurately as possible, then included as a parameter in the optimisation.

The total set of parameters to be optimised is thus as follows:

- The belt velocity
- 3D rotation of the detector system relative to the conveyor belt
- 3D position and 3D rotation of lower detector array relative to upper detector array
- 3D position and 3D rotation of each detector (16 in total)

This corresponds to a total of 106 parameters. Optimising in a 106 dimensional space is not feasible, so we rather proceed in an iterative hierarchical fashion. The algorithm optimises each set of parameters in the list above separately while holding the remaining parameters constant. The parameters for a specific detector are optimised one by one, while keeping the other detectors stationary. The whole procedure is then repeated in a second iteration.

Figure 6 shows the change in global score as this procedure is followed. The optimisation parameters varied according to stage (more effort was put into the collective coordinates), with between 100 and 200 generations, and a population between 35 and 50. These numbers are limited by computational cost, where a back-projected 3D PET image must be created for each of a series of data sets, in order to evaluate each individual in each generation.

When the overall effect on image fidelity was examined, there were mixed results. While the overall orientation was improved, the algorithm had a bias towards moving all the detector elements closer to one another, in effect zooming out the entire image. An overall zoom factor improves the score as defined in equation 3. The next step would therefore be to use a known configuration of point sources to fix the scale, combining the distance between known and reconstructed points with the existing score.

6. Conclusions and discussion

We developed a custom genetic algorithm that successfully optimises detector parameters. Optimisation leads to significant improvements in the reconstruction of detected positions from




Figure 5. Change of score with local optimisation. Each dot represents the score of one individual in a given generation.




Figure 6. Change of score with global optimisation. The x axis shows successive optimisation stages.

detector signals. The procedure also showed promise for the accurate calibration of the position and orientation of detector elements. Further work is needed to remove a bias toward reducing the overall length scale.

This optimisation approach could easily be automated, and regularly deployed in an industrial setting as part of standard operating procedure, in order to keep a deployed MinPET system optimally configured. This would combat performance degradation due to calibration drift and detector element movement due to impacts and vibration.

References

- [1] Ballestrero S, Bornman F, Cafferty L, Caveney R, Connell S, Cook M, Dalton M, Gopal H, Ives N, Lee C A, Mampe W, Phoku M, Roodt A, Sibande W, Sellschop J P F, Topkin J and Unwucholaa D A 2010 Mineral-PET: Kimberlite sorting by nuclear-medical technology *12th International Conference on Nuclear Reaction Mechanisms* ed Cerutti F and Ferrari A (Varenna, Italy) pp 589–602
- [2] Bocci A and Hulsbergen W 2009 *ATL-INDET-PUB-2007-009*
- [3] Rajpaul V 2012 ArXiv:1202.1643 [astro-ph.IM] (*Preprint* 1202.1643)
- [4] Gutiérrez J A G, Cotta C and Fernández-Leiva A J 2012 *CoRR* abs/**1202.2523**
- [5] Gosselin L, Tye-Gingras M and Mathieu-Potvin F 2009 *International Journal of Heat and Mass Transfer* **52** 2169 – 2188 ISSN 0017-9310 URL <http://www.sciencedirect.com/science/article/pii/S0017931008006534>
- [6] Teodorescu L and Sherwood D 2008 *Computer Physics Communications* **178** 409–419
- [7] Brunetti A 2013 *Computer Physics Communications* **184** 573 – 578 ISSN 0010-4655 URL <http://www.sciencedirect.com/science/article/pii/S0010465512003438>
- [8] Pereira C M and Lapa C M 2003 *Annals of Nuclear Energy* **30** 555 – 565 ISSN 0306-4549 URL <http://www.sciencedirect.com/science/article/pii/S0306454902001068>
- [9] Dickens T P 1998 Image-calibration transformation matrix solution using a genetic algorithm *Industrial Applications of Genetic Algorithms* ed Freeman L M and Karr C L (Boca Raton, FL, USA: CRC Press, Inc.) 1st ed ISBN 0849398010
- [10] Wall M 2007 GAlib, A C++ Library of Genetic Algorithm Components <http://lancet.mit.edu/ga/> accessed: 2013-10-30