Tomographic Reconstruction with Search Space Expansion

Mohammad Majid al-Rifaie*
m.alrifaie@gre.ac.uk
School of Computing & Mathematical Sciences
University of Greenwich
London, United Kingdom

Tim Blackwell
t.blackwell.gold.ac.uk
Department of Computing
Goldsmiths College, University of London
London, United Kingdom

ABSTRACT

A search space expansion process is proposed in the context of tomographic reconstruction (TR). The idea is to widen the effective search space in a series of increasing sizes with clamping on the search space boundary. The technique was tested on four simple phantoms and on the clinically important Shepp-Logan phantom. Dispersive flies optimisation (DFO), a lightweight particle swarm optimisation (PSO) variant, is shown to produce lower reproduction errors compared to standard TR toolbox algorithms. The expansion technique demonstrably decreases salt-and-pepper noise. DFO with 50 subspace searches was found to be superior to differential evolution, PSO and, more importantly, a number of conventional reconstruction techniques. To the best of our knowledge, this is the first work where search space expansion, in its literal form, is introduced, discussed and applied to this problem.

CCS CONCEPTS

• **Computing methodologies** → *Continuous space search*; • **Theory of computation** → *Bio-inspired optimization.*

KEYWORDS

Swarm optimisation, tomographic reconstruction, high dimensional optimisation, search space resizing

ACM Reference Format:
Mohammad Majid al-Rifaie and Tim Blackwell. 2023. Tomographic Reconstruction with Search Space Expansion.

1 INTRODUCTION

The reconstruction of internal structure from projected images cast by penetrating radiation is a key technique in medical imaging and has widespread application through science, mathematics and industry [14, 18–20, 25, 28, 31, 36, 41]. The *few-view* regime is particularly important in cases where the number of projections, and hence the quantity of absorbed radiation, should be minimal, as in, for example, the imaging of children.

The standard reconstruction technique, filtered backprojection (FBP), requires only a single iteration but is not suitable for few-view imaging [27]. Algebraic Reconstruction Techniques (ART) have recently come on stream by virtue of increased computation power. ART is an iterative algorithm based on Kacmarz's method [41] and is applicable to the few-view scenario but it can introduce image artefacts due to overfitting and is not proven in large patient populations [26].

Iterative statistical methods such as maximum-likelihood expectation maximisation (MLEM) and the more computationally efficient ordered subset expectation maximisation (OSEM), have been trialled and found to be superior to FBP in some cases e.g. [43, 47].

In an alternative to conventional reconstruction techniques, the task can be cast as an optimisation problem, and thereby opening up the possibility of population-based algorithms and other metahuristics: harmony search [38], tabu search [33], simulated annealing [34], memetics [22] and evolutionary algorithms [13]. Swarm algorithms have developed for binary reconstruction [35], geophysical reconstruction [44], electrical capacitance and impedance tomography [30, 48] and surface reconstruction from 3D data [24]. An ant algorithm has been developed for binary reconstruction [5].

The few-view tomographic reconstruction (TR) problem is highly underdetermined and the solution set is numerous. Only some members of the solution set are medically feasible; often noise and artefacts obscure important detail. Underdetermination implies that zero reconstruction error (the difference between measured and reconstructed projections) does not necessarily equate to zero reproduction error (how far the reconstructed object differs from the actual object) [6].

This paper builds on the findings of reference [6] in which a comparative study of reconstructions by swarm algorithms and conventional techniques on standard phantoms (TR test cases) in the few-view regime indicated that DFO-TR, an algorithm that formally interpolates between differential evolution and particle swarm optimisation, produces reconstructions with lower reproduction error than state-of-the art ART. One drawback, however, is that DFO-TR images are visibly peppered with noise.

Insight on the importance of clamping at the search space boundary, led to a consideration of a novel search space expansion (SSE) technique which is reported here. We find that the expansion mechanism can lead to a significant reduction in salt-and-pepper noise.

The paper continues with a review of search space resizing; a formal statement of the reconstruction problem and associated algorithms follows before a report on a series of experiments on artificial phantoms that were designed to provide insight on SSE, and test results on the medically pertinent Shepp-Logan phantom.

2 SEARCH SPACE RESIZING

Search space resizing is a technique primarily designed to combat premature convergence when the optimal solution(s) x_{opt} lies/lie within the search space X, or, rarely, to mitigate search bias when x_{opt} lies/lie on the boundary, ∂X.

Consider a search space $X = \bigcup_{d=1}^{n} X_d$ of dimension n, decision variables (DV's) $x_d \in X_d$ and an objective function $f(x_1, x_2 \ldots, x_n)$.

*Corresponding author

X might define a region of feasibility or an initialisation volume believed to contain the desired optimum; f may or may not be defined for $x \notin X$. DV's that attempt to leave X during the optimisation process might be evaluated if the objective function is so defined, or they might be returned to X according to a predefined scheme and then evaluated; indeed, algorithm performance can depend critically on bounds policy [29]. Resizing can be employed to deal with search bias arising from interaction between the boundary policy and the defined region of f; more commonly, x_{opt} lies within X and resizing can mitigate premature convergence by, for example, re-initialising the search population within a smaller volume that contains the most promising solution.

Search space resizing refers to either reduction or expansion of X, or to dynamic combinations of the two. Reduction might mean the progressive narrowing of each X_d, either on a DV by DV basis, or for all DV's at once, or it might mean the culling of unhelpful and unnecessary dimensions. Similarly, expansion might refer to the broadening of each X_d and/or the addition of extra DV's. Resizing can be triggered automatically, or at preset intervals.

In practice only several of the possibilities have been explored.

2.1 Reduction

The main desideratum for reduction is to prevent premature convergence by re-invigorating search within a smaller volume.

There has been a steady trickle of papers investigating search space reduction (SSR) over the past twenty years. The technique has been renamed, perhaps on rediscovery, as 'compression', 'squeezing' and 'pruning' but the essential idea is the same. SSR is most often explored in the context of a real-world problem and schemes are problem-specific, and often lever known properties of the solution space.

Che and Smith applied SSR to the flowshop problem, but the reduction was applied to the initial population [21]. A more sophisticated procedure to the same problem was suggested by Yong and Sannomiya [51]. In this work, a system of reduced search spaces by a series of included subsets was proposed; genetic algorithm (GA) performance over fixed search spaces was enhanced.

Search space reduction has been investigated in the context of water distribution optimisation [46]. The motivation for their technique was to exclude unpromising regions of the original search space where solutions are either infeasible or impractical.

Izzo and co-authors devised a search space pruning step in their solution to the problem of gravity assisted spacecraft trajectories within the solar system [32]. The technique was specific to particular properties of the objective function but global optimisation was demonstrably enhanced for DE, GA and PSO.

SSR has been considered for optimising quality of service in mobile ad hoc networks [12].

A search space 'squeezing' was trialled by Barisal et al [11] in the context of power generation. Resizing is applied at, or near, stagnation and is governed by the relative distance of the global best of a particle swarm to dynamic feasible search space limits.

Rajarathinam et al experimented with search space resizing in dynamic systems governed by differential equations. Their technique was able to place the elite population of a GA within feasible boundary regions of the search space [39].

Search space 'compression' has been proposed by Zang and collaborators [50]. Compression is triggered when the fraction of the population of solutions close to the best solution is reached: the search space is resized and the population is randomised within the reduced volume. The motivation behind the mechanism was to prevent premature convergence and the authors find enhancement in GA, PSO and DE.

More recently, Das and Pratihar [23] devise a search space reduction technique that kicks in at low population diversity providing the search has not stagnated; GA diversity is claimed to be balanced with selection pressure so that convergence is improved without sacrificing global search capacity. The idea is favourably demonstrated with real-coded GAs on a small benchmark.

There is a short history of SSR for the solution of water drainage problems [15]; Bayas-Jimenéz et al, in this work, consider pruning unnecessary decision variables, but the method is restricted to the particular problem under scrutiny.

2.2 Expansion

Search space expansion (SSE), according to our research, has only been occasionally investigated.

Tsutsui and Goldberg analysed boundary extension methods designed to deal with real-valued GA sampling bias near the corners of a search space [45]. The authors suggested two mechanisms. In boundary extension by mirroring (BEM), out-flying individuals are evaluated by placing them at the mirror image of the function for the purposes of evaluation. Boundary extension by selection is similar to BEM except that selection can be made without evaluation for a small number of 'helper' individuals that reside in borders beyond the search space. Tsutsui and Goldberg found that either extension method could mitigate search bias and improve performance in not only cases when the optimum is at a search space corner, but also for objective functions with optima at the centre of X. This result on the importance of boundary policy for optimisation algorithms was later confirmed by Helwig et al [29].

Expanding the search space to a torus, again to counteract GA crossover bias, has also been considered, although this technique is better classified as a boundary policy rather than a search space extension [42]. A more refined toroidal boundary method was proposed by Yoon et al [49]. In this work, ambiguities in function value at the boundary were resolved by extension.

3 TOMOGRAPHIC RECONSTRUCTION

Incident radiation is modelled by a projection matrix $A \in \mathbb{R}_{\geq 0}^{m \times n}$ where m is the total number of projections, and n is the number of pixels in the reconstructed image. The discrete reconstruction problem is:

$$\text{find } x \in \{0, 1, \ldots, k-1\}^n, k > 1$$
$$\text{such that } Ax = b$$

where $b \in \mathbb{R}^m$ is the vector of detector values. In the greyscale problem considered in this paper, x is a vector of reconstructed pixel values and $k = 256$ i.e. $x \in \{0, 1 \ldots, 255\}^n$. The original object, denoted x^*, also satisfying $Ax^* = b$, is either a phantom (artificial test case) or an imaged subject. The equation $Ax = b$ cannot be inverted if $m < n$ and there are multiple solutions.

Suppose y is a trial solution; y is forward projected:

$$Ay = c$$

and a *reconstruction* error evaluated

$$e_1(y) = ||b - c||_1. \quad (1)$$

An iterative scheme will produce a sequence of candidate solutions of decreasing e_1 error but, due to underdetermination, low reconstruction error does not imply faithfulness to the original object x^*. The proximity of y to x^* can be measured by a *reproduction* error:

$$e_2 = ||y - x^*||_1 \quad (2)$$

in cases where x^* is known. e_2 provides a test of the ability of an algorithm to find a faithful reconstruction.

Facilities such as the Astra toolbox supply the forward projection operator A as well as standard reconstruction algorithms such as ART, FPB and SIRT. In the absence of real-world data, and for algorithm development, a phantom x^* is designed and $b = Ax^*$ is computed by forward projection. The objective function for optimisation is defined as $e_1(y) = ||b - Ay||$ where Ay is computed, for trials y, by the toolbox.

Discrete TR can be translated into a real valued problem suitable for optimising algorithms such as PSO and DE by allowing y to take continuous rather than discrete values, $y \in [0, 255]^n$, and to allow e_1 and e_2 to accept real inputs. The original object, x^*, remains discrete. Final solutions y can be discretised for visualisation purposes if necessary.

4 DISPERSIVE FLIES OPTIMISATION AND THE TR LANDSCAPE

Dispersive flies optimisation (DFO) is a lightweight particle swarm optimisation (PSO) variant, distinguished by the abolition of particle memory: updates are computed from current, rather than historical, position [2]. The algorithm's exploration and exploitation behaviour is studied in [3]. DFO includes component-wise particle jumps which have been shown to be beneficial in bare bones PSO [17].

DFO particles collaborate in a ring social network. A DFO iteration at $t + 1$ starts by determining the best position g^t of the N particles in the swarm and the position of the best particle $n_i^t = \arg\min\left(f(x_{i-1}^t), f(x_{i+1}^t)\right)$ in each particle's neighbourhood (the neighbourhood does not include self) for objective function f and mod N index arithmetic. Arbitrary choices are made in the case of ties in g or n_i. The particle with the best position, g, is not updated; the position components (corresponding to the decision variables in Sec. 2) x_{id} of all other particles update according to

$$\begin{aligned} &\text{if } u \sim U(0,1) < \Delta \\ &\quad x_{id}^{t+1} \sim U(X_d) \\ &\text{else} \\ &\quad x_{id}^{t+1} = n_{id}^t + u'\phi(g_d^t - x_{id}^t) \end{aligned} \quad (3)$$

where $d = 1, 2 \ldots, n$, Δ is a preset jump probability, $U(X_d)$ is the uniform distribution along axis d of the search space X, $u, u' \sim U(0, 1)$ and $\phi \in [0, \sqrt{3}]$. The constraints on the convergence controller ϕ are derived from a convergence analysis for stochastic difference equations [17].

Rule 3 mixes aspects of local and global PSO topologies. Search is focused on a (usually) suboptimal position, n_i; however search spread is governed by separation from g. DFO has two arbitrary parameters, N and Δ, set to 100 and 0.001 respectively, and ϕ is set to 1 in the original paper [2].

The DFO mechanism depends, like differential evolution (DE), on instantaneous position and not memory, yet retains PSO's inter-particle communication network. In this sense, which is only formal and does not imply intermediate performance, DFO interpolates between PSO and DE. The algorithm has been applied to a wide range of problems in computer vision, aesthetics measurement and art, optimising food processes, electronics, data science and neuroevolution [1, 4, 7–10, 16, 37].

A comparative few-view study of PSO, DE and DFO on five standard test phantoms of sizes 32^2 and 64^2 for medical TR revealed that, at standard settings of these algorithms, as derived by benchmarking on low dimensional real-valued problems over several decades, DFO was the superior algorithm, in terms of both reconstruction and reproduction error, to DE and global PSO (GPSO) in all cases, and to local PSO (LPSO) in 80% of the trials [6]. The ability of any off-the-shelf algorithm that had not been adapted to high dimensional optimisation to reconstruct images was unexpected, but function profiling suggested that the landscape of Eq. 1 consists of a single wide funnel with a small-scale, multi-modal bottom. Successful optimisers are therefore expected to marry good downhill descent behaviour with a diversity mechanism to aid small-scale search at a valley bottom.

The parameters of the most promising population algorithm, DFO, were tuned, in an earlier study [6], on a single phantom, yielding optimal settings $N = 2, \Delta \approx 0.001$ and $\phi \approx \sqrt{3}$. Only one particle of the two particle DFO swarm updates. Suppose that particle 1 is at g and particle 2 is at x. Then x is updated: $x_d^{t+1} = g_d^t + u'\phi(g_d^t - x_d^t)$ with probability $1 - \Delta$; therefore, DFO reduces to a simple hill descender. Diversity, however, is maintained by the jump mechanism. Since the problem dimension is 1024 for a 32^2 phantom, one component of x restarts, on average, at each iteration, and four components for the 64^2 phantom. The mixture of fast descent coupled with occasional dimensional restart is unique to DFO and might account for its superiority for these TR problems.

A final feature which might have a bearing on the ability of the relative success of population algorithms in this context is the boundary policy. Since the search space is strictly $X = [0, 255]^n$, out-flying individuals cannot be evaluated. Boundary clamping was enforced in the study of Eq. 1. Therefore, any component d straying outside X_d was moved to the nearest ∂X_d. The net effect is that decision variables can saturate at boundary values 0 and 255 if one of the optima is near a boundary. Saturation is advantageous for binary reconstruction where all pixel values are either 0 or 1, and for non-binary reconstruction where image background pixel values are 0 because incident radiation that does not penetrate the subject is not attenuated (images are negatives so that salient features appear bright against a dark background).

The idea that saturation of decision variables at intermediate values between 0 and 255 might aid the discrete case where there are several distinct pixel values was the motivation for the search space expansion technique delineated in the next section.

5 TR SEARCH SPACE EXPANSION

The idea is to conduct the search in a series of progressively larger subspaces $\Xi_p \subseteq X$, $p = 1, 2 \dots P$ such that the final subspace is identical to X, the full search space. Trial solutions are clamped to $\partial\Xi_p$. Expansions from Ξ_p to Ξ_{p+1} kick-in at preset intervals during the optimisation.

This scheme differs from the expansions reported in Sec 2.2 because, rather than artificially extending the search space to raise sampling frequency at the boundary, the effective search space is expanded sequentially into previously unprobed territory of the complete search space.

A simple prescription for empirical trials is to define a series of boxes $\Xi_p = [0, \frac{p}{P} \times 255]^n$ with expansions at equal divisions of the total budget of function evaluations. For example, with four subspaces and a budget of 100000 function evaluations (FEs), search is conducted in the following boxes:

$$\begin{aligned}
\Xi_1 &= [0, \frac{255}{4}]^n, & 0 < FE \leq 25000 \\
\Xi_2 &= [0, \frac{255}{2}]^n, & 25000 < FE \leq 50000 \\
\Xi_3 &= [0, \frac{3 \times 255}{4}]^n, & 50000 < FE \leq 75000 \\
\Xi_4 &= [0, 255]^n, & 75000 < FE \leq 1000000.
\end{aligned}$$

6 PRELIMINARY EXPERIMENTS

As a proof of feasibility of TR SSE, a series of reconstructions with differing numbers of boxes, P, was conducted on four test phantoms in the few-view regime. The test cases, 32×32 phantoms with uniform square internal structure, are depicted in the rightmost columns of Figs. 1-4. The phantoms were imaged with just six projections so that the reconstruction remains challenging despite the seeming triviality of the test cases.

The DFO optimiser was run with the empirical optimal TR settings from ref. [6], $N = 2, \Delta = 0.001$ and $\phi =$ (nearest double precision value to) $\sqrt{3}$, for 100000 FEs on each phantom, and for box subspaces $\Xi_p = [0, \frac{p}{P} \times 255]^n$ ranging from $P = 1$ (i.e. the reference case of no expansions at all) to $P = 100$. Reconstructions of single sample runs are visualised in Figs. 1-4. Each experiment was run 30 times and the median end point reconstruction and reproduction errors for each test phantom and each P are tabulated in Table 1. The results of Wilcoxon rank sum statistical analysis at a confidence level of 95% are reported in Table 2.

Phantom W, with an internal square of pixel value (PV) 255 surrounded by a frame of 0 PV, is perfectly ($e_1 = e_2 = 0$) reconstructed for 1, 2 and 3 subspace boxes (the rows in Fig. 1 are not complete because runs are terminated at $e_1 = 0$). SSE does no better than default optimisation in the entire search space ($P = 1$) because clamping DV's at the edge of the search space is sufficient. The images show the effects of clamping to the edges of the subspace by a lightening of the central square during the corresponding FE interval.

Phantom G, with the white internal square replaced by grey (PV = 128) is not well reconstructed in the default, $P = 1$, case. The final image shows the kind of salt-and-peppering that was noticed

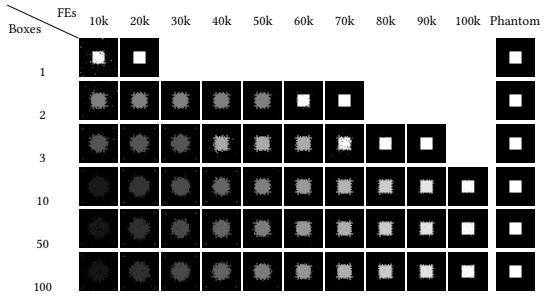

Figure 1: 'White' phantom (W), 32×32, 6 projections. Row images depict reconstructions every 10k FEs throughout the run for differing numbers of box subspaces.

in a previous study [6]. Near perfect reconstruction is achieved by a single expansion ($P = 2$); clamping to the first subspace boundary at PV = 0 and PV = $\frac{255}{2}$ is apparently sufficient.

Phantom WG, with its double square internal structure, is a more difficult problem at low projections. A larger number of boxes, $P = 10$, are optimal, and, although there is some noise, salt-and-peppering is reduced compared to the default control (no subspaces, 1 box).

Phantom GG is similar to G except the background is set to PV = 200. The problem is harder than W and G because clamping at PV = 0 will not aid reconstruction. 10 boxes again performs well in this case. The default algorithm is quite poor, showing extensive salt and pepper noise in the background.

For W, $P = 1$ is optimal, for G, $P = 2$ is optimal, and for the more complex, WG and GG, $P = 10$ is optimal. The larger optimal P for the more complex phantoms, WG and GG, possibly indicates that larger P are more suited for more complex phantoms or phantoms with more levels of pixel value.

The series of trials on test phantoms confirms the intuition that clamping on the edges of a series of expanding boxes can reduce the noise that plagued previous trials with a single search space. The hypothesis is that driving DV's to saturation, via clamping, is advantageous because internal structures are composed of regions of similar PV. If the box sizes were to exactly match the intensity of the salient internal regions, then we might expect good reconstructions, but of course algorithms do not have access to PV values of internal structure in advance, and the boxing scheme remains empirical and arbitrary, but there is a hint that more complex images objects require more boxes.

7 SHEPP-LOGAN PHANTOM

The Shepp-Logan phantom, introduced in 1974, is a schematic representation of a cranial slice [40]. It has become a standard phantom for TR algorithm testing.

DFO Shepp-Logan reconstructions in the few-view regime were studied for various expansions. Fig. 5 depicts representative reconstructions for various search space expansions. Boxing appears to reduce noise and produces visually more respectable reconstructions, and more expansions is preferred.

Table 3 lists e_1 and e_2 median errors for 30 runs. Although a single box produces the smallest median reconstruction error, the

Table 1: Median errors for various search space expansions $\Xi_p = [0, \frac{p}{P} \times 255]^n$ for four phantoms, W, G, WG and GG. The smallest errors are highlighted for each phantom. The values in parentheses show the number of function evaluations before reaching the optima.

	W		G		WG		GG	
Boxes	e_1	e_2	e_1	e_2	e_1	e_2	e_1	e_2
1	0	0 (27604)	1600	8246	5560	13532	3670	71740
2	0	0 (74624)	23	75	5068	9673	3959	63859
3	0	0 (89899)	970	3235	5784	9616	3820	49998
10	1163	395	203	445	7053	8809	2739	12752
50	3409	1534	390	758	8476	8910	2676	12999
100	3619	1645	389	743	8972	9408	2528	12855

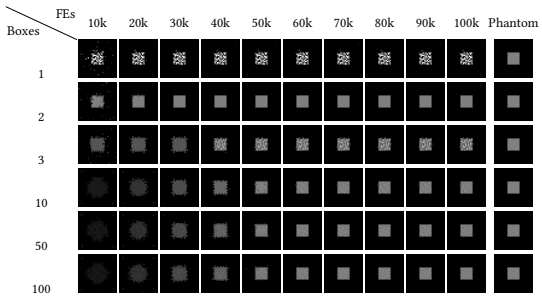

Figure 2: 'Grey' (G), 32×32, 6 projections

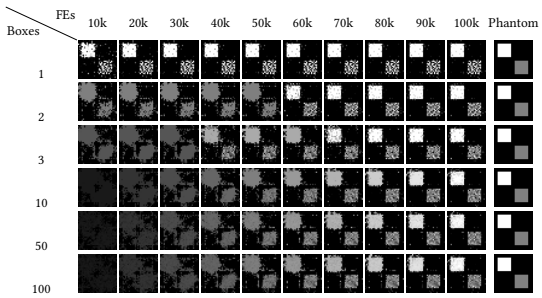

Figure 3: 'White and Grey' (WG), 32×32, 6 projections

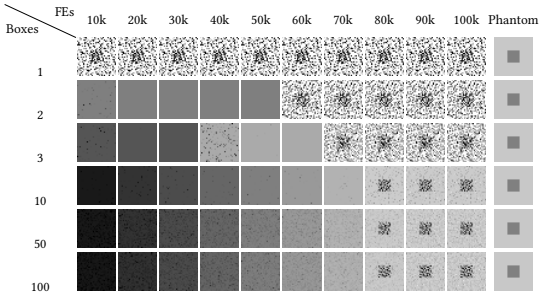

Figure 4: 'Grey on Grey' (GG), 32×32, 6 projections

result is very noisy (Fig. 5); 50 boxes, however, produces a more uniform image with better reproduction error i.e. the reconstruction is closer to the original.

Table 2: Statistically significant (Wilcoxon, 95% conf.) wins of one boxing scenario over the others for reproduction error e_2 for phantoms W, G, WG and GG. For example, entry 3 (G, WG, GG) for 10 boxes (row) and 3 boxes (column) means that the 10 box algorithm ($P = 10$) produced significantly lower e_2 values on phantoms G, WG and GG than $P = 3$. $P = 10$, followed by $P = 50$ are the top expanding box scenarios.

Boxes	1	2	3	10	50	100
1	NA	0	0	1 (W)	1 (W)	1 (W)
2	3 (G,WG,GG)	NA	1 (G)	2 (W,G)	2 (W,G)	2 (W,G)
3	3 (G,WG,GG)	1 (GG)	NA	1 (W)	1 (W)	1 (W)
10	3 (G,WG,GG)	2 (WG,GG)	3 (G,WG,GG)	NA	3 (G,WG,GG)	3 (G,W,WG)
50	3 (G,WG,GG)	2 (WG,GG)	3 (G,WG,GG)	0	NA	2 (W,WG)
100	3 (G,WG,GG)	2 (WG,GG)	2 (W,GG)	0	0	NA

Statistical significance was tested by Wilcoxon comparison at 95% confidence. Results for e_1 and e_2 are shown in Table 4. A single box (no expansion) is confirmed to produce the lowest reconstruction error; increasing the number of expansions is significantly favourable in reducing e_2 up to 50 boxes. There is no significant difference between 50 and 100 boxes.

Comparisons to classical reconstruction algorithms (ART, FBP and SIRT), DE, G/LPSO are illustrated in Fig. 6 and median reproduction errors are tabulated in Table 5. DFO refers to the algorithm with TR-tuned parameters, and DFO with 50 subspace searches is labelled 'DFO-50'.

DFO-50 is evidently the better population algorithm for the single instance of Fig 6 and is comparable visually to the classic algorithms. Also, DFO-50 has the lowest median reproduction error for all projections.

Wilcoxon significance results for 30 runs for four projections are given in Tables 6 and 7. Unsurprisingly, the best classical algorithm, SIRT, produces the best e_1 reconstruction and DFO, although worse than the toolbox methods, is the best population algorithm. Reproduction errors – which quantify how closely the reconstructions appear irrespective of e_1 error – tell a different story. DFO-50 outperforms SIRT even for larger views (an entry of 4 in Table 7 indicates significantly better performance over all projections). The beneficial nature of SSE is manifested in the 4 wins for DFO-50

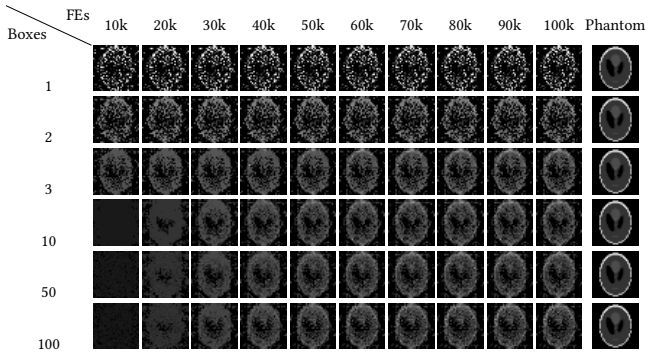

Figure 5: Shepp-Logan phantom reconstruction for varying numbers of expanding boxes over 100k function evaluations, with a 32×32 phantom and 6 projections.

Table 3: Median error values for 30 runs on each of five boxing scenarios for 32×32 Shepp-Logan phantom imaged with 6 projections. The smallest e_2 error is highlighted.

Boxes	e_1	e_2
1	10397	40392
2	11362	32779
3	12086	27131
10	12031	24684
50	11501	23852
100	11869	23884

when competing with DFO. DFO-50 produces lower reproduction errors for all projections.

8 CONCLUSIONS

This study proposes a particular search space expansion scheme for tomographic reconstruction: a progressively wider search in a series of larger clamped boxes. Experiments in the few-view regime on four test phantoms and the Shepp-Logan phantom show that artefact-free reconstructions are possible. In particular, DFO, a population optimiser similar to PSO but without memory and with component jumps, produces lower reproduction errors than the standard TR toolbox algorithms. Clamping to a series of expanding walls appears to eliminate the salt-and-pepper noise that had plagued previous reconstructions.

A tuned version of DFO with 50 boxes was found to be superior to single-box differential evolution and particle swarm optimisation with conventional parameter settings. The combination of boxing and clamping might be specific to DFO or it might generalise to other population algorithms. The implementation of an adaptive boxing on the individual rather than the population level, application of boxing and clamping to standard global optimisation benchmarks and other few-view reconstructions are the subject of future research.

REFERENCES

[1] Bishwa Babu Acharya, Sandeep Dhakal, Aayush Bhattarai, and Nawraj Bhattarai. 2021. PID speed control of DC motor using meta-heuristic algorithms. *International Journal of Power Electronics and Drive Systems* 12, 2 (2021), 822.
[2] Mohammad Majid al-Rifaie. 2014. Dispersive flies optimisation. In *2014 Federated Conference on Computer Science and Information Systems*. IEEE, 529–538.
[3] Mohammad Majid al-Rifaie. 2021. Exploration and exploitation zones in a minimalist swarm optimiser. *Entropy* 23, 8 (2021), 977.
[4] Mohammad Majid al-Rifaie and Ahmed Aber. 2016. Dispersive Flies Optimisation and Medical Imaging. In *Recent Advances in Computational Optimization*. Springer, 183–203.
[5] Mohammad Majid al-Rifaie and Tim Blackwell. 2016. *Applications of Evolutionary Computation: 19th European Conference, EvoApplications 2016, Porto, Portugal, March 30 – April 1, 2016, Proceedings, Part I.* Springer International Publishing, Cham, Chapter Binary Tomography Reconstruction by Particle Aggregation, 754–769. https://doi.org/10.1007/978-3-319-31204-0_48
[6] Mohammad Majid al Rifaie and Tim Blackwell. 2022. Swarm Led Tomographic Reconstruction. In *Proceedings of the Genetic and Evolutionary Computation Conference* (Boston, Massachusetts) *(GECCO '22)*. Association for Computing Machinery, New York, NY, USA, 1121–1129. https://doi.org/10.1145/3512290.3528737
[7] Mohammad Majid al-Rifaie and Marc Cavazza. 2022. Evolutionary Optimisation of Beer Organoleptic Properties: A Simulation Framework. *Foods* 11, 3 (Jan 2022), 351. https://doi.org/10.3390/foods11030351
[8] Mohammad Majid al-Rifaie, Anna Ursyn, Robert Zimmer, and Mohammad Ali Javaheri Javid. 2017. On Symmetry, Aesthetics and Quantifying Symmetrical Complexity. In *Computational Intelligence in Music, Sound, Art and Design: Evo-MUSART 2017*, João Correia, Vic Ciesielski, and Antonios Liapis (Eds.). Springer International Publishing, 17–32. https://doi.org/10.1007/978-3-319-55750-2_2
[9] Haya Abdullah Alhakbani and Mohammad Majid al-Rifaie. 2017. Optimising SVM to classify imbalanced data using dispersive flies optimisation. In *2017 Federated Conference on Computer Science and Information Systems (FedCSIS)*. IEEE, 399–402.
[10] Prashant Aparajeya, Frederic Fol Leymarie, and Mohammad Majid al-Rifaie. 2019. Swarm-Based Identification of Animation Key Points from 2D-medialness Maps. In *Computational Intelligence in Music, Sound, Art and Design*, Anikó Ekárt, Antonios Liapis, and María Luz Castro Pena (Eds.). Springer International Publishing, Cham, 69–83.
[11] AK Barisal. 2013. Dynamic search space squeezing strategy based intelligent algorithm solutions to economic dispatch with multiple fuels. *International Journal of Electrical Power & Energy Systems* 45, 1 (2013), 50–59.
[12] Leonard Barolli, Makoto Ikeda, Giuseppe De Marco, Arjan Durresi, Akio Koyama, and Jiro Iwashige. 2007. A search space reduction algorithm for improving the performance of a GA-based qos routing method in ad-hoc networks. *International Journal of Distributed Sensor Networks* 3, 1 (2007), 41–57.
[13] Kees Joost Batenburg and Walter A Kosters. 2009. Solving Nonograms by combining relaxations. *Pattern Recognition* 42, 8 (2009), 1672–1683.
[14] Kees Joost Batenburg and Willem Jan Palenstijn. 2004. On the reconstruction of crystals through discrete tomography. In *International Workshop on Combinatorial Image Analysis*. Springer, 23–37.
[15] Leonardo Bayas-Jiménez, F Javier Martínez-Solano, Pedro L Iglesias-Rey, and Daniel Mora-Meliá. 2021. Search space reduction for genetic algorithms applied to drainage network optimization problems. *Water* 13, 15 (2021), 2008.
[16] J. Mark Bishop and Mohammad Majid al-Rifaie. 2016. Autopoiesis in Creativity and Art. In *Proceedings of the 3rd International Symposium on Movement and Computing* (Thessaloniki, GA, Greece) *(MOCO '16)*. ACM, New York, NY, USA, Article 27, 6 pages. https://doi.org/10.1145/2948910.2948948
[17] Tim Blackwell. 2011. A study of collapse in bare bones particle swarm optimization. *IEEE Transactions on Evolutionary Computation* 16, 3 (2011), 354–372.
[18] MD Butala, RJ Hewett, RA Frazin, and F Kamalabadi. 2010. Dynamic three-dimensional tomography of the solar corona. *Solar Physics* 262, 2 (2010), 495–509.
[19] José-Maria Carazo, CO Sorzano, Eicke Rietzel, R Schröder, and Roberto Marabini. 1999. Discrete tomography in electron microscopy. In *Discrete Tomography*. Springer, 405–416.
[20] Bruno Carvalho, Gabor Herman, Samuel Matej, Claudia Salzberg, and Eilat Vardi. 1999. Binary tomography for triplane cardiography. In *Information Processing in Medical Imaging*. Springer, 29–41.
[21] Stephen Chen and Stephen F Smith. 1999. Improving genetic algorithms by search space reductions (with applications to flow shop scheduling). In *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation-Volume 1*. 135–140.
[22] Marco Cipolla, Giosuè Lo Bosco, Filippo Millonzi, and Cesare Valenti. 2014. An island strategy for memetic discrete tomography reconstruction. *Information Sciences* 257 (2014), 357–368.
[23] Amit Kumar Das and Dilip Kumar Pratihar. 2019. A new search space reduction technique for genetic algorithms. In *Contemporary Advances in Innovative and Applicable Information Technology*. Springer, 111–119.

Table 4: Statistical analysis of e_1 (top) and e_2 (bottom) for varying boxing scenarios for 32×32 Shepp Logan with 6 projections. A '1' indicates a Wilcoxon win for the algorithm over 30 runs in the left hand column versus the algorithm in the top row. '0' indicates no significant difference in the algorithms.

e_1

Boxes	1	2	3	10	50	100
1	NA	1	1	1	1	1
2	0	NA	1	1	0	1
3	0	0	NA	0	0	0
10	0	0	0	NA	0	0
50	0	0	1	0	NA	0
100	0	0	0	0	0	NA

e_2

Boxes	1	2	3	10	50	100
1	NA	0	0	0	0	0
2	1	NA	0	0	0	0
3	1	1	NA	0	0	0
10	1	1	1	NA	0	0
50	1	1	1	1	NA	0
100	1	1	1	1	0	NA

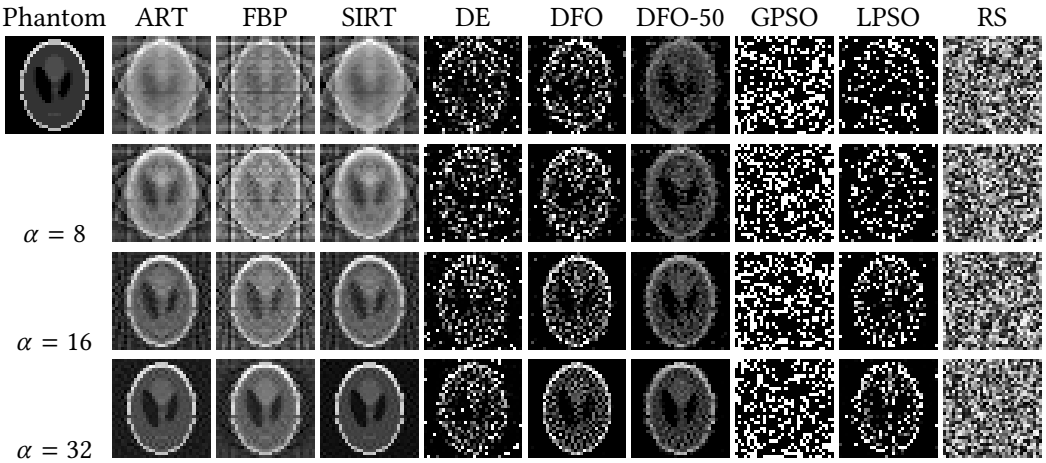

Figure 6: Comparison with other methods. Size: 32×32, projections: 6 (top), and 8, 16 and 32. See Table 5.

Table 5: Median reproduction (e_2) error over 30 runs on the 32×32 Shepp Logan with $\alpha = 6, 8, 16, 32$ projections for each algorithm. Lighter shading indicates the proximity of the reconstructions to the phantoms. The largest errors in the phantom is 255×32^2.

	ART	FBP	SIRT	DE	DFO	DFO-50	GPSO	LPSO	RS
$\alpha = 6$	81771	84742	80748	48691	41078	23852	93344	52980	103306
$\alpha = 8$	69001	93809	71061	46981	37521	21547	92871	51388	103155
$\alpha = 16$	45268	66107	46304	42407	27781	16894	93687	45457	103397
$\alpha = 32$	26132	45887	15953	38162	19702	11681	93373	42364	103167

[24] Akemi Gálvez and Andrés Iglesias. 2012. Particle swarm optimization for non-uniform rational B-spline surface reconstruction from clouds of 3D data points. *Information Sciences* 192 (2012), 174–192.

[25] Richard J Gardner. 1995. *Geometric tomography*. Vol. 1. Cambridge University Press Cambridge.

[26] Lucas L Geyer, U Joseph Schoepf, Felix G Meinel, John W Nance Jr, Gorka Bastarrika, Jonathon A Leipsic, Narinder S Paul, Marco Rengo, Andrea Laghi, and Carlo N De Cecco. 2015. State of the art: iterative CT reconstruction techniques. *Radiology* 276, 2 (2015), 339–357.

[27] A. Giussani and C. Hoeschen. 2013. *Imaging in nuclear medicine*. Springer.

[28] U Hampel, A Bieberle, D Hoppe, J Kronenberg, E Schleicher, T Sühnel, F Zimmermann, and C Zippe. 2007. High resolution gamma ray tomography scanner for flow measurement and non-destructive testing applications. *Review of scientific instruments* 78, 10 (2007), 103704.

[29] Sabine Helwig, Juergen Branke, and Sanaz Mostaghim. 2012. Experimental analysis of bound handling techniques in particle swarm optimization. *IEEE Transactions on Evolutionary computation* 17, 2 (2012), 259–271.

[30] Gang Hu, Min-you Chen, Wei He, and Jin-qian Zhai. 2011. Clustering-based particle swarm optimization for electrical impedance imaging. *Advances in Swarm Intelligence* (2011), 165–171.

[31] R.W. Irving and M.R. Jerrum. 1994. Three-dimensional data security problems. *SIAM J. Comput.* 23 (1994), 170–184.

[32] Dario Izzo, Victor M Becerra, Darren R Myatt, Slawomir J Nasuto, and J Mark Bishop. 2007. Search space pruning and global optimisation of multiple gravity assist spacecraft trajectories. *Journal of Global Optimization* 38, 2 (2007), 283–296.

[33] Fethi Jarray and Ghassen Tlig. 2010. A simulated annealing for reconstructing hv-convex binary matrices. *Electronic Notes in Discrete Mathematics* 36 (2010), 447–454.

[34] Fethi Jarray, G Tlig, and A Dakhli. 2010. Reconstructing hv-convex images by tabu research approach. In *International Conference on Metaheuristics and Nature Inspired Computing*. 3.

[35] Póth Miklós. 2014. Particle swarm optimization approach to discrete tomography reconstruction problems of binary matrices. In *Intelligent Systems and Informatics (SISY), 2014 IEEE 12th International Symposium on*. IEEE, 321–324.

Table 6: Algorithm comparison based on e_1 performance for 32×32 Shepp Logan phantom. The numbers indicate statistically significant (Wilcoxon, $p < 0.05$) wins over four experiments (6, 8, 16 and 32 projections) of 30 runs for the algorithm in the left hand column versus the algorithm in the top row.

e_1	ART	FBP	SIRT	DE	DFO	DFO-50	GPSO	LPSO	RS
ART	NA	4	0	4	4	4	4	4	4
FBP	0	NA	0	2	0	0	4	2	4
SIRT	4	4	NA	4	4	4	4	4	4
DE	0	2	0	NA	0	0	4	0	4
DFO	0	4	0	4	NA	3	4	4	4
DFO-50	0	4	0	4	0	NA	4	4	4
GPSO	0	0	0	0	0	0	NA	0	4
LPSO	0	2	0	0	0	0	4	NA	4
RS	0	0	0	0	0	0	0	0	NA

Table 7: Algorithm comparison based on e_2 performance for 32×32 Shepp Logan phantom. The numbers indicate statistically significant (Wilcoxon, 95% conf.) wins over four experiments (6, 8, 16 and 32 projections) of 30 runs for the algorithm in the left hand column versus the algorithm in the top row.

e_2	ART	FBP	SIRT	DE	DFO	DFO-50	GPSO	LPSO	RS
ART	NA	4	2	1	0	0	4	1	4
FBP	0	NA	0	0	0	0	3	0	4
SIRT	2	4	NA	1	1	0	4	1	4
DE	3	4	3	NA	0	0	4	4	4
DFO	4	4	3	4	NA	0	4	4	4
DFO-50	4	4	4	4	4	NA	4	4	4
GPSO	0	1	0	0	0	0	NA	0	4
LPSO	2	4	3	0	0	0	4	NA	4
RS	0	0	0	0	0	0	0	0	NA

[36] Guust Nolet et al. 2008. A breviary of seismic tomography. *Imaging the Interior* (2008).

[37] Hooman Oroojeni, Mohammad Majid al-Rifaie, and Mihalis A. Nicolaou. 2018. Deep Neuroevolution: Training Deep Neural Networks for False Alarm Detection in Intensive Care Units. In *European Association for Signal Processing (EUSIPCO) 2018*. IEEE, 1157–1161. https://doi.org/10.23919/EUSIPCO.2018.8552944

[38] Ahlem Ouaddah and Dalila Boughaci. 2014. Improving reconstructed images using hybridization between local search and harmony search meta-heuristics. In *Proceedings of the Companion Publication of the 2014 Annual Conference on Genetic and Evolutionary Computation*. ACM, 1475–1476.

[39] Kumaran Rajarathinam, J Barry Gomm, DingLi Yu, and Ahmed Saad Abdelhadi. 2017. An improved search space resizing method for model identification by standard genetic algorithm. *Systems Science & Control Engineering* 5, 1 (2017), 117–128.

[40] Lawrence A Shepp and Benjamin F Logan. 1974. The Fourier reconstruction of a head section. *IEEE Transactions on nuclear science* 21, 3 (1974), 21–43.

[41] Abe R. Shliferstein and YT Chien. 1977. Some properties of image-processing operations on projection sets obtained from digital pictures. *IEEE Trans. Comput.* 26, 10 (1977), 958–970.

[42] Hiroshi Someya and Masayuki Yamamura. 2005. A robust real-coded evolutionary algorithm with toroidal search space conversion. *Soft Computing* 9, 4 (2005), 254–269.

[43] Ana Carolina Trevisan, Michel David Raed, Vitor Tumas, Leonardo Alexandre-Santos, Felipe Arriva Pitella, Emerson Nobuyuki Itikawa, Jose Henrique Silvah, Mery Kato, Edson Zangiacomi Martinez, Jorge Alberto Achcar, et al. 2020. Comparison between OSEM and FBP reconstruction algorithms for the qualitative and quantitative interpretation of brain DAT-SPECT using an anthropomorphic striatal phantom: implications for the practice. *Research on Biomedical Engineering* 36, 1 (2020), 77–88.

[44] Jens Tronicke, Hendrik Paasche, and Urs Böniger. 2012. Crosshole traveltime tomography using particle swarm optimization: A near-surface field example. *Geophysics* 77, 1 (2012), R19–R32.

[45] Shigeyoshi Tsutsui and Devid E Goldberg. 2001. Search space boundary extension method in real-coded genetic algorithms. *Information Sciences* 133, 3-4 (2001), 229–247.

[46] Kalanithy Vairavamoorthy and Mohammed Ali. 2005. Pipe index vector: a method to improve genetic-algorithm-based pipe optimization. *Journal of Hydraulic Engineering* 131, 12 (2005), 1117–1125.

[47] Floris HP van Velden, Reina W Kloet, Bart NM van Berckel, Saskia PA Wolfensberger, Adriaan A Lammertsma, and Ronald Boellaard. 2008. Comparison of 3D-OP-OSEM and 3D-FBP reconstruction algorithms for High-Resolution Research Tomograph studies: effects of randoms estimation methods. *Physics in Medicine & Biology* 53, 12 (2008), 3217.

[48] P Wang, JS Lin, and M Wang. 2015. An image reconstruction algorithm for electrical capacitance tomography based on simulated annealing particle swarm optimization. *Journal of applied research and technology* 13, 2 (2015), 197–204.

[49] Yourim Yoon, Yong-Hyuk Kim, Alberto Moraglio, and Byung-Ro Moon. 2012. A theoretical and empirical study on unbiased boundary-extended crossover for real-valued representation. *Information Sciences* 183, 1 (2012), 48–65.

[50] Shui-ping Zhang, Bi Wang, and Xue-jiao Wang. 2017. A dynamic search space strategy for swarm intelligence. *International Journal of Computational Science and Engineering* 15, 1-2 (2017), 21–31.

[51] Yong Zhao and Nobuo Sannomiya. 2001. An improvement of genetic algorithms by search space reductions in solving large-scale flowshop problems. *IEEJ Transactions on Electronics, Information and Systems* 121, 6 (2001), 1010–1015.