



Evolutionary background-oriented schlieren tomography with self-adaptive parameter heuristics

ANDREAS UNTERBERGER^{1,*}  AND KHADIJEH MOHRI^{1,2,3} 

¹University of Duisburg-Essen, Institute for Combustion and Gas Dynamics - Tomography Group, 47048 Duisburg, Germany

²University of Duisburg-Essen, Institute for Combustion and Gas Dynamics - Fluid Dynamics, 47048 Duisburg, Germany

³University of Duisburg-Essen, , Center for Nanointegration Duisburg-Essen (CENIDE), 47048 Duisburg, Germany

*andreas.unterberger@uni-due.de

Abstract: For volumetric reconstruction of the refractive index field in a flow, background-oriented schlieren (BOS) imaging which measures the deflection of light rays due to refractive index variations is combined with an evolutionary tomographic algorithm for the first time, called evolutionary BOS tomography (EBOST). In this work application to reactive flows is presented. Direct non-linear ray-tracing of the reconstruction domain is used to evaluate the fitness of solution candidates during the evolutionary strategy that was implemented to run on a multi-GPU system. The use of a diversity measure and its consideration in a migration policy was tested against a simple scheme that distributes the best chromosome (solution candidate) in an island-based genetic algorithm. The extensive set of control parameters of the presented algorithm was harnessed by a self-adaptive strategy taking into account the fitness function and operator rates. Quantitative characterisation of the EBOST via numerical phantom studies, using flame simulations as ground truth data is presented. A direct comparison to a state-of-the-art BOST algorithm demonstrates similar accuracy for a turbulent swirl flame phantom reconstruction. A series of experimental applications of the EBOST on several unsteady and turbulent flames is also presented. In all cases, the instantaneous and time-averaged flame structure is revealed, proving the benefit of EBOST for volumetric flow diagnostics.

© 2022 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

Technological advances in computation, image processing and data acquisition have fostered the use of improved methods for three dimensional (3D) flow imaging diagnostics via tomographic techniques in different areas of research ranging from engineering to medical applications. Projection measurements are gathered from different directions around a target volume and are modelled in an inverse relation to the discretised target field. The inverse problem is typically ill-posed and several algorithms exist for solving the resulting set of equations. Obtaining the instantaneous 3D distributions of different quantities, such as velocity and density, in unsteady or turbulent flows from experiments is extremely insightful, and has been demonstrated by combining a tomography algorithm with different types of measurements such as particle image velocimetry (PIV) [1], schlieren [2] and X-ray attenuation [3,4]. In case of reacting flows, luminescence from added tracers or natural chemiluminescence that is due to the intermediary reactions have also been used to reconstruct the flame shape [5–14]. The 3D data from reacting flows allows us to advance our fundamental understanding of different processes and to improve burner systems, for example to achieve higher efficiency and reduce harmful emissions. Additionally, experimental 3D datasets can be used to validate and/or enhance complex combustion modelling like in large eddy simulation (LES). Algorithms that can provide successful reconstructions using a moderate

number of projection measurements are particularly interesting since the experimental costs will be reduced for the simultaneous measurement of transient flows. Furthermore, optical access is typically limited in practical technical applications. Examples of iterative methods that perform better than the conventional filtered back-projection algorithm when the number of projections is small [15] include the algebraic reconstruction technique (ART) [16] and its variants, and Bayesian approaches [17,18], where additional assumptions about the target field can be incorporated into the inversion process. Generally, the above mentioned algorithms struggle with solving non-linear tomographic problems, due to the formulation of the inverse problem as a typically very large and ill-posed system of linear equations. Algorithms have been presented to overcome this restriction, e.g. by Ma et al. [19], which relies on the assumption that the non-linear dependency of the measurement on the reconstructed quantity can be modelled as a factor. However, approaches featuring heuristics such as simulated annealing [20] or evolutionary concepts are increasingly interesting due the steep increase in computing power over the recent years and their advantages over standard iterative or gradient-based methods for solving non-linear tomographic problems. A genetic algorithm (GA) for tomography has been introduced for the first time by Kihm [21] where a 2D optical tomography problem was investigated with numerically generated fields. Other evolutionary strategies have been developed in different areas like electrical impedance tomography [22], multi-phase flow analysis [23] and medical applications [24,25].

An evolutionary reconstruction technique (ERT) was developed by Unterberger et al. [13] and used for volumetric reconstruction of the flame chemiluminescence for the first time. Like other flame chemiluminescence reconstruction methods, the measurement model works with the assumption that density is homogeneous inside the tomographic volume. However, an actual flame zone contains variations in the refractive index which can be visualised using background-oriented schlieren (BOS) measurements, and the 3D refractive index field can be estimated using BOS tomography (BOST) [26–34]. For the BOST methods to date, the so called paraxial approximation is made to allow the use of straight ray paths inside the reconstruction domain; this is further explained in Sec. 2. Due to direct non-linear ray-tracing capability of EBOST, this assumption is not made and thus EBOST can be applied to spatially constrained lab environments as well as to flows where the density gradient is significantly large compared to standard BOST applications.

Applications of GAs, especially in engineering, can involve a lot of parameters to control the system, which is often different from investigations that study its fundamental nature. This typically increases the complexity for the user and tuning the parameters is often only possible with experience and expert knowledge of the model. Deterministic parameter tuning, where mostly one parameter is changed at a time does not consider the complex interaction between parameters and the likelihood to achieve an overall optimal set of parameters is therefore very small. Parameter studies are also limited when the parameter space is too large and might only be performed on subsets. A way to solve this problem is the introduction of self-adaptive parameter control, which has become a common technique in evolutionary computation [35,36]. For example a two-chromosome individual can be defined that is used to evolve a solution candidate and a set of control parameters. Hinterding [37] applied such a two-chromosome approach for self-adaption of parameters in a Cutting Stock Problem. Multi-chromosome methods were used previously, e.g. by Juliff [38] on the Pallet Loading Problem and by Pierrot and Hinterding [39] on a simple mixed integer problem. Here, we demonstrate the application of a meta-learning system that holds a large set of control parameters to calculate the 3D refractive index field in gaseous flows by EBOST. A two-chromosome approach is used where the first chromosome represents the target field to be reconstructed and the second chromosome holds a staged system of self-adapted control parameters. We apply a combined form of control, considering rates and step size of the genetic operators and a weighting factor in the fitness function for regularisation.

The EBOST is based on the ERT [13] which can incorporate beam-steering in the ray-tracing scheme and allows the definition of a non-linear fitness function. Besides presenting the EBOST algorithm, where ERT is combined with the BOS measurement model for the first time, the ERT was conceptually improved by the introduction of new genetic operators. A hierarchy of differently scaled operators is used to capture the natural occurrence of different scales in the target field. Additionally, an efficient implementation for a GPU-cluster/multi-GPU system was achieved by exploiting the inherent parallel nature of GAs that allows the population to be split and lets them evolve separately on islands. Island-based GAs (IGAs) with specific migration policies can leverage the evolutionary process and are suitable to minimise communication time between the GPU devices which decreases the total run-time of the algorithm significantly. We have investigated an IGA, which aims to increase the diversity of the island populations and was proposed by Araujo et al. [40,41] under the name "MultiKulti Algorithm". In our application, the comparison to a simple IGA that distributes the best global chromosome to all islands showed a superior result in favour of the MultiKulti-strategy.

2. Background-oriented schlieren tomography

BOS is an established technique that is used for visualising the density field (derived from refractive index gradients) [29] of a target flow. The BOS setup consists of a camera that images an illuminated background pattern. Keeping the background illumination fixed, a reference image at time $t = 0$ is captured without the target flow in place and a second image at time $t = 1$ is taken with the target flow that contains refractive index variations. Comparing the image pairs reveals the deflections caused by the target flow in the background pattern due to the refractive index variations, as illustrated in Fig. 1. The deflections can be calculated by using a dotted background pattern and applying a cross-correlation algorithm as used in PIV, for which advanced commercial systems are available at a high cost. Other techniques feature different types of structured backgrounds, e.g. wavelet noise [28], in combination with an optical flow method, such as the methods proposed by Lucas and Kanade [42] or Horn and Schunck [43]. The 3D refractive index field can be derived by a suitable tomographic inversion technique using measurements from different perspectives. Typical algorithms used for reconstructions are based on filtered back-projection [27], ART [28] or based on a Bayesian approach [32,33]. These methods use the paraxial assumption [27,28] to approximate the actual curved ray path in the reconstruction volume by a straight ray. A single refraction event per ray is then assumed close to the centre of the domain. The assumption is appropriate if the deflection magnitude in the probe




Fig. 1. BOS principle and ray deflection due to refractive index variations in the target flow. For the reference (ref.) image the target flow is not present and points 1 and 2 are mapped linearly on the sensor at locations 1' and 2', respectively. For the deflection (def.) image the gradient in the density of the target flow distorts the ray path and a pattern at location 1 appears visually shifted by a distance δ to location 1'' on the sensor.

volume is very small compared to the voxel dimensions used for discretising the volume and the distance of the volume to the background. Hence, a practical recommendation to increase the measurement quality in these models is to increase the object-to-background distance [31]. Typically, distances range from 1.5 m to 2.50 m [27,28,31–33]. An exception to this, with a distance of 0.5 m is presented by Liu et al. [34], but with a comparatively small volume of interest of 27 mm cubed.

3. Reconstruction methodology

3.1. Generic ERT concept

Details of the ERT that we first developed to reconstruct the 3D chemiluminescence field of turbulent flames can be found in [13] and here only a summary of the key elements is presented: For the fitness evaluation, the ERT features volumetric ray-tracing, using a pin-hole camera model, of a 3D scalar field to be evolved by a GA.

- A chromosome is defined as the reconstruction domain (3D scalar field) that is discretised using cubic voxels.
- A gene is defined as the value with uniform distribution inside a voxel.
- A population of chromosomes is evolved by a GA where voxel-wise alterations (mutations) are used in combination with selection and arithmetic averaging of chromosome pairs (crossover).
- The fitness of a chromosome is measured by a suitable distance measure, e.g. L_1 or L_2 norm of the volume-rendered images of the chromosome and the measurement images.
- A stochastic mask is applied to incorporate *a priori* information and to increase the efficiency of the optimisation.

Masking the reconstruction domain is a known technique that can improve the reconstruction quality [30,33]. In contrast to the masking method presented by Nicolas et al. [30], the applied mask in the ERT does not rely on the definition of a threshold value based on the measurement data, which can typically be very noisy, making it difficult to distinguish real signal. The stochastic mask that was developed for the ERT can incorporate prior information about the target field based on the measurement images and a Metropolis sampling algorithm, before initialising the reconstruction process. The stochastic mask provides voxel coordinates where evolutionary operators are applied and is an efficient way to reduce the search space for the GA and hence the runtime. The bin count of the 3D histogram of the sampled voxel coordinates is proportional to the probability that a respective voxel will be selected for mutation. Figure 2 shows an example of a stochastic mask generated from a flame simulation phantom. For example, voxel A in Fig. 2 has a lower bin count compared to voxel B and hence voxel B will be chosen for mutation with a higher probability. The stochastic mask in our ERT is adapted on a low frequency during the evolution, eg. every 10% of the total number of generations. Adaption consists of recalculating the stochastic mask based on the best solution calculated so far. The sampling from an intermediate solution requires a certain degree of convergence of the algorithm that can be inferred from the slope of the stored fitness values of the best individual in each generation.

3.2. Evolutionary background-oriented schlieren tomography

For the EBOST, an optical flow (OF) calculation in combination with a wavelet-noise background pattern was found to be a suitable low-cost solution. The OF method that is used here is based on




Fig. 2. Example of a stochastic mask. Coordinate axes in voxels. Left: iso-surface at 0.5 of the SwB1 DNS [44,45] phantom, downsampled to $43 \times 50 \times 43$ voxels. Middle: slice of the 3D histogram of 9 million random voxel coordinates, sampled by a Metropolis algorithm from 20 rendered views of the phantom. The dashed line shows the isoline at 0.5 of the phantom. Right: normalised 3D histogram for values above 1.6×10^{-5} .

the work of Horn and Schunck [43] where an image brightness function $E(x, y, t)$ is assumed to fulfil a conservation equation in the form of

$$\frac{dE}{dt} = \frac{\partial E}{\partial x} \frac{dx}{dt} + \frac{\partial E}{\partial y} \frac{dy}{dt} + \frac{\partial E}{\partial t} = 0, \quad (1)$$

where x and y are the image coordinates and t is time. The terms dx/dt and dy/dt are the deflections, which are abbreviated with u and v , respectively, and form the deflection vector $\mathbf{d} = [u, v]$. The partial derivatives of E with respect to x , y and t are abbreviated with E_x , E_y and E_t , respectively. Horn and Schunck proposed the deflections in Eq. (1) can be determined by minimising the following functional:

$$F = \int \int \left(\frac{dE}{dt} \right)^2 + \alpha \left(\|\nabla u\|^2 + \|\nabla v\|^2 \right) dx dy. \quad (2)$$

The second term in Eq. (2) is an additional regularisation term as the plain optimisation problem is ill-posed and doesn't have a unique solution. Thus, the second term promotes a smooth solution with a regularising parameter α . A Discrete version of Eq. (2) using partial derivatives of E as suggested by Horn and Schunck [43], is the basis of the fitness function used in our GA. The complete functional that is minimised by the GA and implemented in the EBOST is given as

$$\mathcal{F} = \sum_{i=1}^{NC} \tilde{F}_i, \quad (3)$$

where \tilde{F}_i is the discrete functional for the measurement data of the i -th camera and NC is the number of cameras. For chromosome selection by stochastic universal sampling (SUS) [46], the absolute fitness value \mathcal{F} is rescaled to a relative fitness value according to

$$relFit(i) = 1 - \frac{\mathcal{F}(i) - \min(\mathcal{F})}{\max(\mathcal{F}) - \min(\mathcal{F})}. \quad (4)$$

The fitness value of a chromosome is determined by ray-tracing the refractive index field that is associated with the chromosome. The floating point value of each gene in the chromosome represents the local refractive index $n(x, y, z)$. The deflection vector \mathbf{d} is calculated by tracing a ray from the centre of a pixel in the sensor plane of a pinhole camera to the background. Between the sensor and entry to the reconstruction domain and from the exit of the domain to the

background plane, the ray is assumed to follow a linear path. Inside the reconstruction domain the propagation is described by the non-linear ray equation of geometric optics

$$\frac{d}{ds} \left(n \frac{d\mathbf{p}}{ds} \right) = \nabla n, \quad (5)$$

where the vector \mathbf{p} describes the position of a hypothetical particle travelling along a path s through the domain and s is an infinitesimal path length [32]. A full derivation of Eq. (5) can be found in [47,48], and it can be cast into a system of two first order differential equations

$$n \frac{d\mathbf{p}}{ds} = \mathbf{q}, \quad (6)$$

$$\frac{d\mathbf{q}}{ds} = \nabla n, \quad (7)$$

where \mathbf{q} gives the local direction of the particle's motion. These two equations are discretised by an Euler-forward scheme for the ray-tracing by

$$\mathbf{p}^{i+1} = \mathbf{p}^i + \frac{\Delta s}{n} \mathbf{q}^i, \quad (8)$$

$$\mathbf{q}^{i+1} = \mathbf{q}^i + \Delta s \nabla n, \quad (9)$$

where Δs is a chosen step size. The change in the propagating particle's direction is determined by the gradient of the refractive index, which is evaluated by linear interpolation. When the intersection of a deflected ray with the background has been found, it can be linearly back-projected onto the camera sensor to determine the deflection vector \mathbf{d} by the difference of the starting and end points of the traced rays. For example, in Fig. 1, the non-linear ray path starts at location 2' and ends at location 1 at the background. The linear projection of location 1 is at location 1' on the sensor, hence the pattern at location 1 appears to be shifted by the distance $\delta = 1'' - 1'$ on the sensor.

One of the target flow fields investigated in this paper is the highly turbulent reacting flow of the Cambridge-Sandia SwB1 flame presented by Sweeney et al. [49,50], that has been subject to experimental and computational studies, for example the DNS simulations by Proch et al. [44,45]. The turbulent flow contains different length scales, ranging from the geometry of the burner (large) down to very small Kolmogorov-dimensions [51]. To aid the capturing of different scales, a hierarchy of differently sized mutation operators was introduced. The scaled mutation operator (SMO) contributes to the random changes in the evolutionary process. The SMO sets a cubic region of a particular size to a new value that is drawn from a truncated normal distribution with standard deviation $\kappa \cdot (b - a)$, where κ is a free parameter and $[a, b]$ is a predefined interval of estimated refractive index values also used as cut-off values (expected minimum and maximum values). The mean of the normal distribution equals the average refractive index value of the nearest neighbours of the voxel that has been drawn from the stochastic mask. Similarly, the scaled annihilation operator (SAO) sets a cubic region to a defined ambient refractive index value. The annihilation operator was introduced in the original ERT formulation [13] to set back gene values to their ground state, where the ground state was zero intensity for the chemiluminescence applications, to help in sharpening the flame front structures (against the black background from the imaging) and promoting hollow regions. For a premixed flame, the density is typically lower on the burnt side of the flame front, separating the hot gas products from the premixed unburned gases. In the EBOST application, the ground state is the ambient refractive index value surrounding the open flame and hot gas plume. The SMO and SAO hierarchies are defined as a series of operators acting on cubic regions with different edge lengths $l_1 = 1, l_2 = 2, l_3 = 3, \dots, l_d = s$ voxels. Each of the operators is applied with a different

probability mp_i , with $i = 1, \dots, dm$ and ap_i with $i = 1, \dots, da$ for the mutation and annihilation, respectively. The choice of the depths dm and da , that are the numbers of different operators contained in the hierarchies, depends on the target field, but also on the total number of voxels in the domain (NX by NY by NZ corresponding to width, height and depth, respectively). The hierarchy depth also influences the convergence, which is further investigated in Sec. 4.1.

To simplify the application of the code for the user, and to achieve a consistent and self-adaptive system the operator rates are integrated into the definition of the individual and an additional chromosome par is defined, see Fig. 3. This chromosome does not only contain the rates mp_i and ap_i , but also the parameters κ and α that were specified earlier. To self-adapt these parameters, par is structured in 3 levels. A Gaussian operator is used to mutate a gene of a specific subset u of par with a probability mr_u . Different values of u specify the inclusion of different genes of level 0: for $u = 1$ the dm mutation operators, when $u = 2$ the da annihilation operators, and when $u = 3$ and $u = 4$ the κ and α parameters are included, respectively. A new value for a gene is sampled from a normal distribution $\mathcal{N}(\mu, \sigma)$, where μ is the current value of the gene (level 0 gene) and σ_{mru} is a standard deviation. The rates mp_i and ap_i , as well as the parameters κ and α are expected to change at different rates mr_u and with different step sizes σ_{mru} for adaption during the evolutionary process. To achieve this, the parameters mr_u and σ_{mru} are defined to constitute the level 1 genes of the meta-learning system. To reduce the set of free parameters further, a third stage (level 2) is added to par that contains a single meta-learning rate mlr as a gene. In the end, there is only one external parameter left for tuning, that is the standard deviation σ_{mlr} and it defines the Gaussian mutation for mlr and level 1 parameters. The influence of σ_{mlr} on the underlying levels is discussed in Sec. 4.2. Recombination of the chromosome par is performed by a two-point crossover. With an adaptive rate for each operator, the total number of genes in par depends on the user's choice. For the phantom study in Sec. 4.1, the maximum number of genes in par was 25 with $dm = 9$ and $da = 5$. In summary, an evolutionary strategy for a population of combined individuals was defined where each individual constitutes a chromosome box representing a discrete refractive index field and a chromosome par containing evolution parameters.




Fig. 3. Definition of an individual with two chromosomes box and par . The chromosome par is structured in three levels, called the meta-learning system. The free parameter σ_{mlr} and the initial values of each gene are defined by the user. The parameters in level 0 are influenced via a Gaussian mutation by level 1 and level 2, respectively. Only the level 0 genes can directly influence the chromosome box and its fitness value in an evolution step.

The stochastic mask that is used for the EBOST is generated by sampling from $|E_t|$, normalised by its maximum. This differs from the original stochastic mask sampling outlined in [13] because here, the measurement is different. The calculated deviations of the curved ray paths from the straight rays are very small and result in deflections of magnitude in the pixel to sub-pixel range, as will be discussed when presenting Fig. 16 in Sec. 6. Hence, beam steering is not considered in the Metropolis sampling and it is assumed to have minor importance for the EBOST algorithm, since it only results in a very small underweighting effect of regions in the domain that are visited less often by the genetic operators compared to others.

3.3. Implementation of EBOST on a multi-GPU system

Ray-tracing features a very high degree of data parallelism and hence the EBOST algorithm was implemented to run on a GPU server. NVIDIA's CUDA API [52] and C programming were used to achieve a runtime-optimised ERT scheme. The thread parallelism available on GPUs greatly speeds up the rendering, arithmetic average calculation for offspring chromosome formation and the fitness evaluation. For the implementation of fast reduction techniques and massive parallel rendering, GPU-specific memory types such as shared, constant and texture data types are available for cache-optimised memory access.

The MPI parallelisation allows us to run the code on a GPU cluster with multiple computing nodes for high-resolution calculations. To alleviate the communication bottleneck between the separate GPU memories an island-based genetic algorithm (IGA) was implemented and tested. The global population of size $P = np \cdot N$ was split on np islands, where each island holds a population of $c_i, i = 1, \dots, N$ individuals and is evolved by an MPI rank. Depending on the number of available GPU devices ng , the MPI ranks can share a device or perform calculations individually on a single GPU. A flowchart of the EBOST algorithm is presented in Fig. 4.

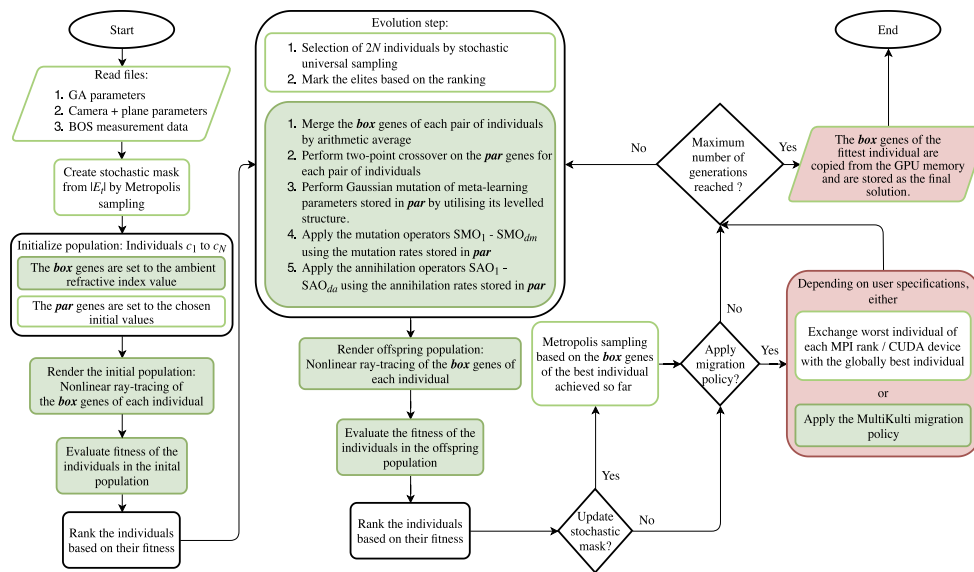


Fig. 4. Flowchart for the EBOST algorithm. The algorithm can be initialised on a host system with multiple MPI ranks (np) and GPU devices (ng). The flowchart shows the EBOST code from the perspective of a single island that is computed on an MPI rank. Minimal requirements are $np = 2$ and $ng = 1$, as multiple ranks can perform computations on a shared GPU device. Green boxes indicate computations that are performed mainly on the GPU. A box with a green edge indicates a necessary memory transfer between GPU and host memory. A red box indicates communication between MPI ranks and/or GPU devices.

The decision of how many islands to use and the population size is not trivial. It can affect the performance of the GA and the quality of the reconstruction and investigating these topics is out of the scope of this paper. Based on a decision on the number of islands and the population size P , the computing resources have to be chosen for optimal runtime performance and utilisation rates of the GPUs. Therefore, the runtime of EBOST depends on the available type and number of GPUs. For the presented reconstructions with real flame data in Sec. 6 and for the phantom reconstructions in Sec. 5, several tests were performed and the results of the phantom studies in Sec. 4 were considered, to ensure high quality reconstructions. The fitness value of the best solution discovered in each generation is tracked for all reconstructions. This information was used to define a suitable stopping criterion by a preset number of generations for the phantom and experimental reconstructions. We ensured that the increase in the tracked fitness values was marginal for at least the last 10% of the total number of generations. The starting generation for the second stage Metropolis sampling, see Sec. 3.1, was chosen based on the same tracked fitness values. For convenience and future applications adaptive stopping criterions can be investigated based on the variance or the slope of tracked fitness values, the same holds for the starting generation of the second stage Metropolis sampling.

3.4. Island-based GA (IGA)

Using distributed populations, where each island-population can explore different regions of the search space, can improve the performance of a GA. For our particular application, the IGA comes with the additional benefit of strongly reducing the run-time of the program. This is the case as long as the migration policy of choice is not too costly and can be computed mostly on one device with minimal requirements for information exchange. Different possibilities for implementing parallel GAs can be found in the literature [53,54]. The MultiKulti strategy proposed by Araujo et al. [41] was chosen as a foundation for further investigations. It offers a sound paradigm by considering diversity in its migration policy and is intuitive in its application. We will outline the MultiKulti approach, with its specific features for the EBOST here: All islands are arranged in a ring topology. In a synchronous version, each island sends its average fitness-value \bar{f} to its preceding node every ν_m generations. The average fitness-value \bar{f} serves as a threshold for the sending node. The sending node will transfer a maximum number of k random individuals (the complete chromosome c) with a fitness above the threshold. The individual with the highest diversity will be accepted by the receiving node to replace the node's individual with the lowest fitness. Diversity is measured by the average L_2 norm of the chromosome box between the sent individual and each individual of the receiving node. The choice of the migration frequency ν_m and the number of exchanged individuals k is predefined and the influence of ν_m on the IGA is discussed in Sec. 4.3.

4. Phantom study on turbulent flame DNS data

It is very important to benchmark a new algorithm with phantom data, that is exactly known ground truth where a direct and quantitative comparison is possible. In accordance with our usual phantom study strategy, we use flame simulation results as ground truth, as a realistic numerical representation of a flame. The phantom refractive index field was calculated from DNS results of the Cambridge-Sandia SwB1 flame [44,45]. The high-resolution DNS field was downsampled to a suitable domain size and the synthetic scene was generated comprising 20 cameras that are arranged equidistantly in a semi-circle, with a background plane facing each camera. The distance from each camera pinhole to the reconstruction domain centre, which is the defined origin of the world coordinate system, is 300 mm, and to the background plane is 950 mm. We evaluate the Pearson-correlation coefficient ρ [55] and the L_2 norm error distance ϵ

between the reconstructed data \mathcal{D}_r and the ground truth data \mathcal{D}_{gt} for the 3D fields via

$$\epsilon = \frac{L_2(\mathcal{D}_r - \mathcal{D}_{gt})}{L_2(\mathcal{D}_{gt})}. \quad (10)$$

Here, the L_2 norm yields for the ideal case of perfect data reconstruction $\epsilon = 0$. The correlation measure ρ indicates, whether a linear relationship between phantom and reconstructed data is present. A correlation value close to $\rho = 1$ means that the data is trend wise almost identical.

Synthetic measurement data was produced by ray-tracing camera sensors with specifications matching those in the experiment, see Sec. 6, using a step-size of $\Delta s = 0.1$ mm. From the rendered images the partial derivatives E_t , E_x and E_y from Eq. (1) were calculated and the respective data was downsampled by a factor of 0.5 and cropped to reduce the ray-tracing costs. This resulted in measurement data used by the EBOST algorithm to represent sensors containing 163 by 163 pixels of size 19.8 μm . A step-size of $\Delta s = 0.5$ mm was used for the ray-tracing in the reconstruction process. The scaled-down phantom field size is $NX = 56$ by $NY = 65$ by $NZ = 56$ voxels, with are solution of 1.08 mm. Each phantom test was repeated 10 times to obtain statistics. The settings of the IGA used for the tests are summarised in the following points:

- Four islands, each containing a population of $N = 12$ chromosomes
- For each sub-population one elite was used, which was transferred unchanged to the next generation
- Maximum number of generations is 200000
- Second-stage Metropolis sampling starts from 100000 generations, with resampling every 10000 generations

The population size typically affects the convergence of a GA, and the number used was chosen based on previous tests that showed an acceptable trade-off between quality and total run time for the above settings. The number of generations and starting generation for the second-stage Metropolis sampling were selected based on tracked fitness values, see Sec. 3.3.

The investigations in Sec. 4.1 and 4.2 used a simple IGA, which distributes the globally best fitting individual with a frequency of 100 generations to all islands to replace their worst chromosome. The initialisation values of the chromosome *par* used for all the runs presented are given in Table 1.

Table 1. Initialisation values of chromosome *par* for the phantom study

Level 0	Level 1	Level 2
$mp_i = 0.01$	$mr_1 = 0.02$	$mtr = 0.02$
	$\sigma_{mr1} = 0.02$	
$ap_i = 0.01$	$mr_2 = 0.02$	
	$\sigma_{mr2} = 0.02$	
$\kappa = 0.2$	$mr_3 = 0.02$	
	$\sigma_{mr3} = 0.02$	
$\alpha = 4.0e - 3$	$mr_4 = 0.02$	
	$\sigma_{mr4} = 0.02$	

4.1. Investigations of the operator hierarchy

The results of multiple runs that were performed using different operator hierarchies, for the mutation and annihilation, are discussed here. The operators of size 1, 2, 3, 5, 7, 8, 9, 11, 13 and 15 voxels cubed were used. The standard deviation of the meta-learning rate was set to $\sigma_{mlr} = 5.0e - 4$ (this choice is discussed further in Sec. 4.2). The cases, which are defined in Table 2, were run ten times to compare the average of the performance measures (error distance and correlation). In preliminary tests we observed, that the annihilation hierarchy does not need to be expanded in the same manner as the mutation hierarchy. The choice of the hierarchies influences the convergence rate. A hierarchy that is too small, e.g. using only single voxel operators leads to a solution that is not converged well enough within the 200000 generations limit. One long run per case with 10^6 generations was performed, which also lead to converged solutions for case 1, but was in a premature state compared to the other cases, see performance measures in Fig. 5. The hierarchy test results shown illustrate that even with a 5-fold longer runtime the performance of the algorithm with a small hierarchy, e.g. cases 1,2 and 3, is not as good as with an expanded one. Therefore, the additional degrees of freedom provided by the expanded hierarchies (increasing case number) are valuable for improving the reconstruction quality and runtime. As the data shows, this advantage tapered at the size of case 5. It can also be observed that the standard deviation shrinks with an expanded hierarchy, which means more stable results and better convergence properties in general.

Table 2. Case definitions for the mutation and annihilation hierarchies

Case	Mutation (voxels cubed)	Annihilation (voxels cubed)
1	1	1
2	1,2	1,2
3	1,2,3	1,2
4	1,2,3,5	1,2,3
5	1,2,3,5,7	1,2,3
6	1,2,3,5,7,9	1,2,3
7	1,2,3,5,7,9,11	1,2,3,5
8	1,2,3,5,7,9,11,13	1,2,3,5
9	1,2,3,5,7,9,11,13,15	1,2,3,5,7




Fig. 5. Performance of the operator hierarchy, the mean and standard deviation (std) for the correlation and error distances are shown. Increased case number reflects expanded hierarchies.

4.2. Investigations of the meta-learning system

The meta-learning system was investigated by varying the parameter σ_{mlr} , which was described in Sec. 3.2. Statistics were obtained by running each parameter case ten times. The hierarchy case 6 from the previous study was used for all the tests in this study. The study revealed that generally a value below 0.001 lead to results of similar accuracy, see Fig. 6. As described in Sec. 3.2, the relevant parameters for the reconstruction process are stored in level 0 (the levels were detailed in Fig. 3), where self-adjustments should lead to a converged solution within an acceptable runtime. The initial values for level 1 and 2 parameters were set such that this goal can be achieved without any additional tuning by the user. For low values of σ_{mlr} , the meta-learning rate mlr (the parameter in level 2) remains relatively close to its well chosen initial value, which in turn allows enough variation in the preceding levels 1 and 0 to support increased reconstruction accuracy. When σ_{mlr} is too large the variation in level 1 and 0 parameters is strong which can potentially lead to premature convergence and hence lower reconstruction accuracy and fitness. We could conclude that for σ_{mlr} values below 0.001, the system is able to achieve the desired results. Our tests showed that with a value of $\sigma_{mlr} = 5.0e - 4$, enough dynamics (variation) is introduced to move the entire system away from initialisation also for level 1 parameters. Chosen mutation operator rates from level 0 were averaged for all the hierarchy cases from Table 2 (which were run using $\sigma_{mlr} = 5.0e - 4$) on a single island for all generations and all 10 runs per case are presented in Fig. 7. It can be seen that in general, the SMO operator rates in level 0 that were obtained by averaging over the ensemble and then time, shown in Fig. 7 are lower for a more expanded hierarchy, i.e. higher number cases. Our studies showed that although a low σ_{mlr} value suppresses the variations in the meta-learning rate and level 1, the suppression influence reduces in the preceding level 0. Further investigations on the dependency of the meta-learning scheme on the initial values should be made in future, because the benefit of adaption in level 1 and 2 cannot be fully understood by the variation of σ_{mlr} alone. The accuracy of the final reconstruction solution, resulting from the choice of parameters from our studies is shown later for two different flame phantoms (unsteady Bunsen and turbulent SwB1) in Sec. 5.1, and also for experimental results in Sec. 6.




Fig. 6. Variation of the meta-learning system parameter σ_{mlr} . The mean and standard deviation (std) of the correlation and error distance are shown.

Finally, Fig. 8 shows the ensemble average $\bar{\alpha}$ of the regularisation parameter α for the first 2000 generations, calculated for all cases listed in Table 2. The $\bar{\alpha}$ rapidly drops from its initial value, tapering off to a value of about $\bar{\alpha}^* = 3.67e - 4$. The average correlation and error distance between the original and reconstructed phantoms, for the runs with the α parameter active and using case 6 of Table 2 are 0.9709 ± 0.0017 and $1.8982e - 5 \pm 0.0113e - 5$, respectively. Test runs with the regularisation parameter deactivated ($\alpha = 0$) and keeping everything else the same, showed a similar average correlation and error distance, 0.9717 ± 0.0013 and $1.8195e - 5 \pm 0.0659e - 5$ respectively. Therefore, it appears that at least for the cases considered, the regularisation parameter has no significant effect on the reconstruction quality. This is in-line with arguments




Fig. 7. Exemplary averaged rates from level 0 of the meta-learning system, for all the hierarchy cases listed in Table 2. The mean and standard deviation (std) of the SMO operator rates from level 0 are shown

made by Grauer and Steinberg [33] where a completely different reconstruction algorithm for BOST, based on the Bayesian formulation, is used.




Fig. 8. Averaged regularisation parameter $\bar{\alpha}$ for the first 2000 generations, for all the cases listed in Table 2

At the beginning of the evolutionary process the 3D field is initialised with the ambient refractive index and the mutation operators might induce different spurious deflection vectors before the onset towards convergence and a positive effect from the regularising factor could be imposed on the evolutionary process. After all, the complete data for all generations, that is partially shown in Fig. 8, shows a non-zero $\bar{\alpha}$ for the complete reconstruction process. In the end, the effect of the regularisation on the reconstruction quality is hard to analyse since α is influencing the fitness function directly - especially in experimental reconstructions where no ground truth data is available.

4.3. Investigations of the IGA

For the IGA implementations two sets of data were created and are compared, one for a simple migration scheme and one for the MultiKulti scheme. For both sets, the values of the migration frequency ν_m tested were 5, 10, 50, 100, 500, 1000 and 5000 generations. It can be seen in Fig. 9 that the MultiKulti scheme generally leads to better fitness and correlation values, hence better reconstructions. However, the error distances do not indicate a clear superiority for either schemes. In both schemes no systematic dependency on the tested values for the migration frequency is visible.

We also investigated the effect of a switched-off meta-learning scheme, which means that the genes listed in level 0 will stay at their initialisation value, listed in Table 1. The performance measures, displayed in Fig. 9, that were obtained as average of 10 runs using a migration frequency of 100 generations, show that the meta-learning scheme is beneficial for the reconstruction not

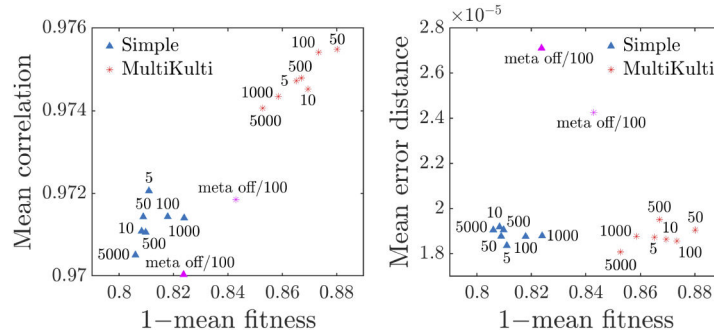


Fig. 9. Effect of various migration frequencies ν_m for the IGA schemes. The mean correlation (left) and error distance (right) are presented for a simple migration and the MultiKulti schemes. Additionally, one case for each migration policy with switched off meta-learning system is shown.

only in terms of the correlation, but also for the error distance. Figure 9 also indicates that the MultiKulti migration benefits more from the meta-learning than the simple scheme. For both schemes the complete chromosome c is transferred, where the fitness/diversity was measured with respect to the chromosome box . These findings hold at least for the runtime of 200000 generations and the hierarchies defined for case 6 in Table 2 that was applied for all runs in this section.

4.4. Effect of imaging noise on the reconstruction quality

Imaging with electronic cameras is typically prone to different types of noise, e.g. shot noise, which affect the measurement data, hence corrupting the information that is passed on to the reconstruction algorithm. This is an experimental feature which can only be modelled in a quantified phantom study. Different levels of noise, signal-to-noise ratio (SNR), were implemented on the synthetic reference and deflection images. The SNR is defined by

$$SNR = 10 \cdot \log_{10} \left(\frac{\sum_i^w \sum_j^h \hat{f}(i,j)^2}{\sum_i^w \sum_j^h (\hat{f}(i,j) - f(i,j))^2} \right), \tag{11}$$

where \hat{f} is the noisy image, f is the original noise-free image, w and h are the image width and height, respectively [56]. Different SNR values were achieved by rescaling the variance of a standard normal distributed random variable which was added to the noise-free images.

For each SNR test, 10 runs were conducted using the initialisation settings in Table 1 and the hierarchy case 6 given in Table 2. Figure 10 shows the correlation and error distance between




Fig. 10. Mean and standard deviation (std) of the correlation (left) and error distance (right) as a function of image signal-to-noise ratio (SNR). Statistics are obtained from 10 runs.

the original and reconstructed phantoms as a function of the image noise level. The estimated SNR from experiments is not expected to be higher than the values tested here. Furthermore, the correlations increase to about 0.97 and plateau for $SNR \geq 6.79$ dB indicating the robustness of EBOST against this type of error.

5. EBOST reconstructions of flame phantoms

5.1. SwB1 and Bunsen flames

The turbulent SwB1 [44,45] and unsteady Bunsen [32] flame phantoms were reconstructed using sensors that are equivalent in size to the CCD sensors in the experiment (659 by 494 pixels of size $9.9 \mu\text{m}$), and 20 different views of the synthetic scene.

For the SwB1 phantom, the chosen region of interest (ROI) from the rendered synthetic images that contains the flame was 329 by 329 pixels in width and height, and hence the images were cropped to this size. The reconstruction domain was discretised into $NX = 108$ by $NY = 125$ by $NZ = 108$ voxels (genes) in width, height and depth, respectively. The voxel resolution in the domain was 0.56 mm/voxel in all directions. The hierarchy case 9 from Table 2, which features 9 differently sized mutation operators and 5 annihilation operators was used. Exemplary slices from the original and reconstructed phantoms are presented in Fig. 11, which had a correlation of 0.98 and an error distance of $\epsilon = 2.53e - 5$ between them, that was achieved within 800000 generations.

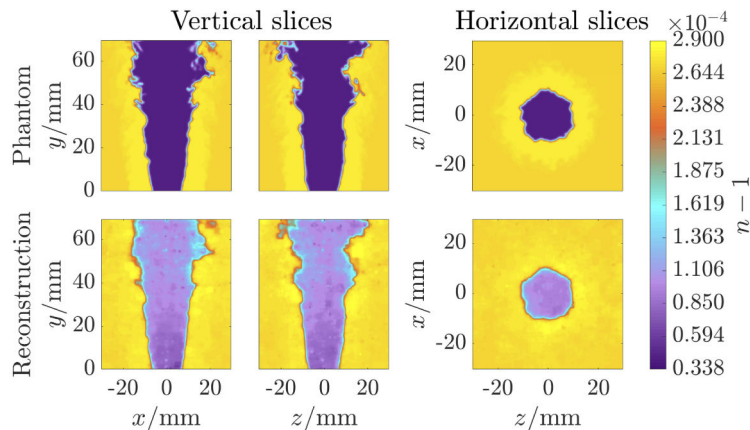


Fig. 11. Exemplary vertical slices at the burner centre from two perspectives, and horizontal slices at $y = 27 \text{ mm}$ height above the burner (2.13 bluff body diameters in the downstream direction), for the original and reconstructed turbulent SwB1 flame phantom's refractive index field n .

For the unsteady Bunsen flame phantom the ROI, and hence synthetic images were of size 260 by 230 pixels in width and height. The reconstruction domain was discretised into $NX = 96$ by $NY = 86$ by $NZ = 96$ voxels, each with a resolution of 0.5 mm/voxel , corresponding to the original resolution of the LES simulation. The same EBOST settings were used for this phantom reconstruction, except the largest 15 cubed mutation operator was omitted. The original and reconstructed phantom's refractive index fields are compared in Fig. 12. Within 800000 generations a correlation of 0.99 and error distance of $\epsilon = 2.76e - 5$ was achieved between them.

Both flame phantom reconstructions show an underestimation of the refractive index, which is symptomatic of other BOST efforts [32,33]. This could be due to numerical issues in the creation of the phantom deflection data, i.e. the reference and deflection images. Other contributors to the errors could be the discretisation of the field, the non-linear ray-tracing scheme and the finite

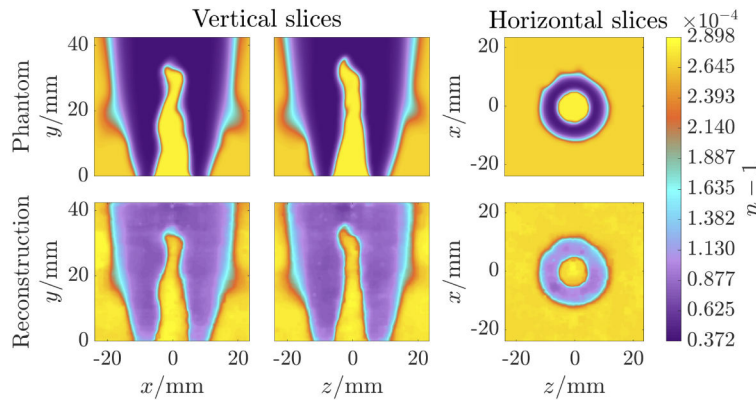


Fig. 12. Exemplary vertical slices at the burner centre from two perspectives, and horizontal slices at $y = 1.5$ mm height above the burner (0.17 nozzle diameters in the downstream direction), for the original and reconstructed unsteady Bunsen flame phantom's refractive index field n .

resolution of the background texture. Since the scale of the deflections is in the sub-pixel to pixel range, slight errors could lead to the observed deviation. Nevertheless, the reconstructions clearly recover the structural detail with good accuracy and exhibit high correlations with the ground truth.

5.2. Swirl flame phantom

Here we present a test case that uses an LES phantom of a central bluff-body swirl flame [11]. The 3D data of the refractive index field was downsampled and cropped to a field of $NX = 75$ by $NY = 75$ by $NZ = 75$ voxels. The same set of data was provided to Grauer and Steinberg [33] to evaluate a recent state-of-the-art Bayesian method for BOST and we devised the exact synthetic tomography scene; 9 cameras with lenses of focal length 10.0 mm and sensors with 500 by 500 pixels and backgrounds with a sine wave pattern at a distance of 1.5 m opposite to each camera. In Fig. 13 slices of the EBOST reconstruction are shown. A correlation of 0.97

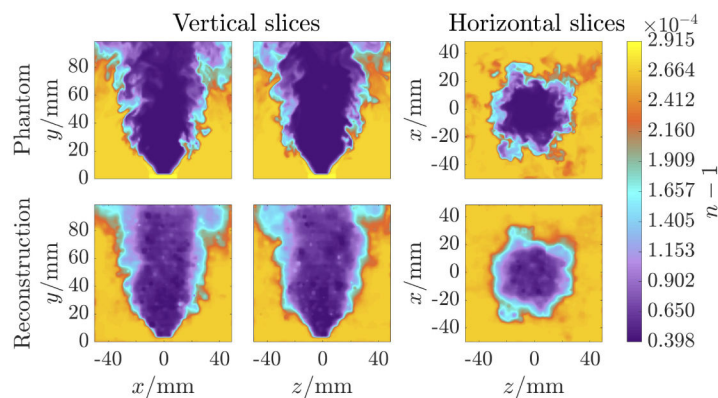


Fig. 13. Exemplary vertical slices at the burner centre from two perspectives, and horizontal slices at $y = 50.7$ mm height above the burner (1.69 bluff body diameters in the downstream direction), for the original and reconstructed turbulent swirl flame phantom's refractive index field n .

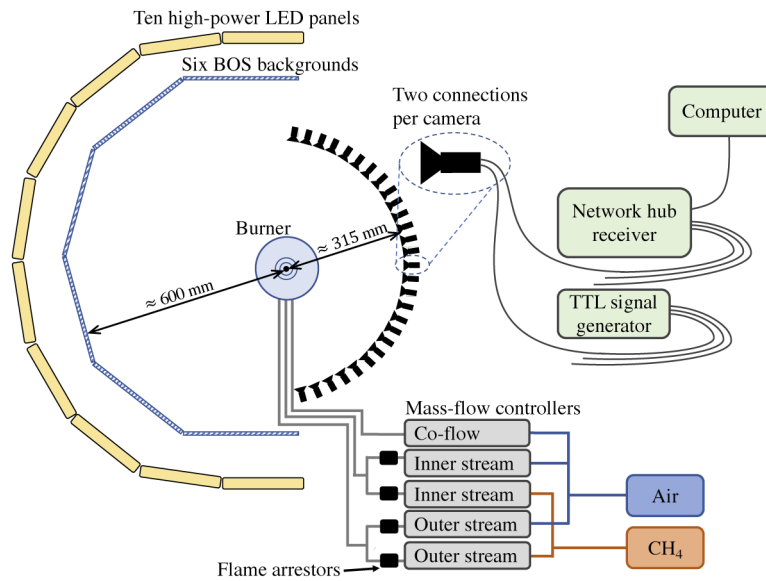


Fig. 14. Schematic of the experimental BOST setup: 6 background patterns, which are back-illuminated by 10 high power LED panels opposite 28 cameras. The gas supply system and mass flow controllers shown are exemplary for the Cambridge-Sandia stratified burner. The cameras are triggered simultaneously by a TTL generator and images are stored via a network-hub (Ethernet connection) on a computer.

between phantom data and reconstruction was achieved. For the EBOST, the error distance was $\epsilon = 2.38e - 5$; Grauer and Steinberg [33] reported an error distance of approx. $\epsilon = 2.33e - 5$.

It is interesting to note that the EBOST reconstructions do not exhibit the errors that are shown in [33] at the upper and lower boundaries of the domain. In addition, it is worth noting that the EBOST algorithm does not use a mask that requires thresholding, see Sec. 3.1. Setting all voxel values below a threshold to the ambient refractive index can increase the reconstruction quality, especially for phantom reconstructions because it eliminates noise in the respective areas completely. However, it can be challenging with experimental data to define a threshold signal.

6. Experimental application of EBOST

6.1. Experimental procedure

Our generic multi-camera setup comprising 28 CCD cameras distributed equiangularly within a total fan angle of 168° on one side of the burner, and 6 BOS backgrounds on the opposite side, as shown in the schematic in Fig. 14 and the image in Fig. 15, was used. The Bunsen and Cambridge-Sandia burners were exchanged for each experiment. Both burners were operated with a premixed CH₄/air mixture, using Bronkhorst mass flow controllers to regulate the flows. The Cambridge-Sandia SwB1 has two streams; the inner with 11.4 CH₄ / 144.0 air and the outer with 34.8 CH₄ / 441.7 air (units in SLPM). The co-flow for the SwB1 condition is half of the actual rate specified by Sweeney et al. [49,50] (382.8 SLPM), because of limited supply of pressurised air. The Bunsen is a regular single stream containing 1.5 SLPM CH₄ and 7.0 SLPM air.

The CCD cameras are from Basler, acA645-100 gm, containing 1/2" Sony ICX414 monochrome sensors with 659 by 494 pixels of size 9.9 μm . They were equipped with Kowa C-mount objectives of 12 mm focal length. The BOS background (wavelet noise) pattern was printed on six large white sheets of paper. The sheets were sandwiched between two flat




Fig. 15. Generic multi-camera setup for multi-angular BOS measurements using 28 CCD cameras and six background planes surrounding the Cambridge-Sandia burner operated with the turbulent SwB1 condition.

plexiglas plates (each one 2 mm thick) of size 310 mm by 450 mm, and were mounted portrait on the optical table. The backgrounds were back-illuminated using 10 high-power LEDs (Osram 200 W, 20000 lm). No problems were encountered with reflection issues from the plexiglas surfaces. The f-stop used on all lenses was $f/8$, providing a compromise between deviation from the ideal pinhole model and ability to measure enough light intensity for the low exposure times used. All cameras were triggered simultaneously by one signal generator. For each test, 200 instantaneous images per camera were captured. For reference images (without flame), the average of 20 images were used for each test. An exemplary image from one camera perspective that includes the flame is shown in Fig. 16 (bottom row, left). It should be noted that the chemiluminescence of the flame is not visible in the image, due to the light conditions that are necessary for BOS measurements. The magnitude of the calculated OF deflection vector field for this camera view and the absolute value of the derivative E_t from Eq. (1) are shown in Fig. 16 (bottom row, middle) and (bottom row, right), respectively. The calculated E_t , E_x and E_y from Eq. (1) are the inputs for the algorithm's optimisation target, as explained in Sec. 3.2. Furthermore, $|E_t|$ serves as the input for the stochastic mask sampling. The visual comparison of the experimental measurement data and ray-traced phantom data in Fig. 16 shows similar deflection magnitudes. Differences can result for example from the numerics, the camera model, the exact distance of the camera to the object and to the background-plane, the gas composition that was used to derive the phantom refractive index field and measurement noise.

For the scene registration, we used our in-house camera calibration algorithm, that is based on an evolutionary optimisation method and uses a special 3D target [57]. The extrinsic parameters of the cameras were calculated using one image of the target, which was placed on top of the burner. For the intrinsic camera parameters, no pixel skew and no shift of the optical centre were assumed. The distance of the camera pinhole to the image plane was assumed to be equal to the focal length given by the manufacturer. In addition to the camera parameters, the exact location and orientation of the BOS background planes must be determined with respect to the common world coordinate system that includes all cameras. An optimisation strategy that is similar to our camera calibration code was applied, but based on images of a 2D calibration pattern, as shown in Fig. 17. The pattern was printed on paper and sandwiched between the same plexiglas sheets, in a similar manner to the printed BOS background targets. The location of the plexiglas sheets was fixed at all times. The BOS plane calibration algorithm minimises the Euclidean distance between a reference image of the 2D target and a rendered view of the target. Image distortions due to the optical system in the camera were not considered in this calibration. The complete calibration procedure constitutes first finding the camera parameters for the scene, and consecutively the calibrated BOS plane parameters. Successful application of any experimental BOST requires an accurately calibrated measurement setup. Misaligned




Fig. 16. Experimental image of the flame and ray-traced phantom scene from one camera perspective (left) including region of interest (red box), magnitude of the OF calculated deflections in pixel units (middle), and absolute value of the image brightness function derivative $|E_t|$ (right). The camera exposure time used for capturing the experimental image was $t_{exp} = 300 \mu\text{s}$.

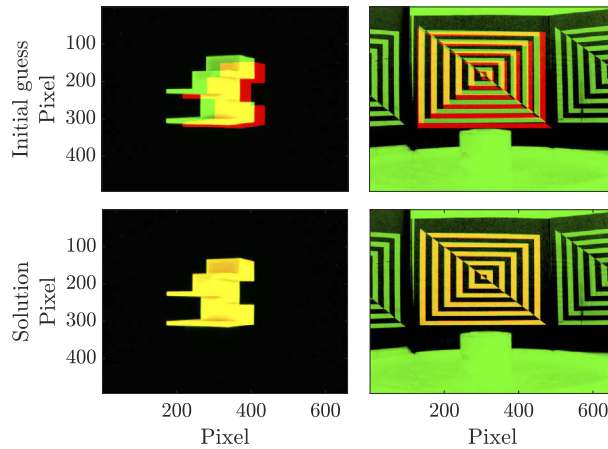


Fig. 17. Exemplary images of the 3D (left) and 2D (right) targets for the camera and a BOS plane, respectively, both shown in green. The rendered image based on the calibration result for each case is shown in red and superimposed onto the actual target image, for the initial guess and the final calibration solution.

cameras, which are cameras where the line-of sight vectors of the pinhole model do not match the location and orientation of the real camera will cause artefacts in the reconstruction. An additional source of error can be wrong information about the location and orientation of the background planes. Since the plane calibration relies on the cameras' extrinsic parameters, both errors are not independent.

6.2. Unsteady Bunsen & turbulent SwB1 reconstructions

The instantaneous and time-averaged refractive index fields of the Bunsen flame were reconstructed using a domain size of $NX = 115$ by $NY = 100$ by $NZ = 115$ voxels in width, height and depth, respectively. The cuboid voxels had a size of 0.5478 mm and 0.5 mm in the width/depth and height, respectively. In total, measurement from 27 of the cameras were used (one camera's view was blocked in the experiment). The exposure time was $t_{exp} = 350 \mu\text{s}$. Examples of the slices




Fig. 18. Slices from reconstruction of the instantaneous (top row) and time-averaged (bottom row) refractive index field n of the unsteady Bunsen flame. The vertical slices from two perspectives at the burner centre and horizontal slices at $y = 12.5$ mm height above the burner are shown. Images were captured with a camera exposure of $t_{\text{exp}} = 350 \mu\text{s}$.




Fig. 19. Slices from reconstructions of the instantaneous (top row) and time-averaged (bottom row) refractive index field n of the turbulent SwB1 flame. The vertical slices from two perspectives at the burner centre and horizontal slices at $y = 27$ mm height above the burner are shown. Images were captured with a camera exposure of $t_{\text{exp}} = 300 \mu\text{s}$.

from the instantaneous and time-averaged reconstructions are presented in Fig. 18, illustrating a sharp hot plume edge. The expected conical shape of the averaged Bunsen is also visible in the slices.

The reconstruction domain for the SwB1 flame case contained $NX = 108$ by $NY = 125$ by $NZ = 108$ cubic voxels of edge length 0.56 mm. All 28 camera measurements with an exposure time of $t_{\text{exp}} = 300 \mu\text{s}$ were used for the reconstructions. Exemplary slices from the instantaneous and time-averaged flame reconstructions are shown in Fig. 19. Comparing the results with the DNS phantom and its reconstruction from Fig. 11 qualitatively reveals that the EBOST faithfully estimates the flame structure from the experiments.

Iso-surfaces of the Bunsen and SwB1 flame reconstruction are visualised for the instantaneous cases in Fig. 20 revealing detailed structures of the refractive index fields.




Fig. 20. Iso-surfaces from the the instantaneous refractive index field n of the Bunsen (a) and SwBI (b) flame reconstructions.

7. Conclusions

We have successfully demonstrated the application of an evolutionary strategy to background-oriented schlieren tomography for the first time by using a set of adaptive parameters, called the meta-learning scheme, and an island based genetic algorithm. Detailed studies on parameters of importance, reconstructions of several phantom and experimental data were presented showing good quantitative agreement between phantoms and their reconstructions and also good qualitative agreement between experimental reconstructions and phantom data. In comparison to the most recent advanced BOST algorithm, that was tested on phantom data only, the EBOST performed convincingly, and quantitative error measures were at the same level. The fundamental difference to other advanced BOST reconstruction methods is first of all the evolutionary approach to tomography. The differences in the masking procedure as well as in the ray-tracing method were explained and the benefits of the EBOST approach for practical applications were highlighted.

The reconstructed 3D refractive index field is proportional to the density that can be derived by the Gladstone-Dale equation in a post-processing step. Density is typically lower on the burnt side of the premixed flame, the hot gas region, and hence the reconstructed fields will help to distinguish the burnt and unburnt zones. This provides very useful information for the analysis of reactive flows. The applicability also extends to non-reactive flows such as general compressible and multi-phase flows. From a modelling standpoint, the ERT method offers a very high degree of flexibility for applications where the model cannot be linearised and standard reconstruction methods like the ART are not applicable anymore.

The generic ERT method was adapted and improved by introduction of operator hierarchies, and the computational efficiency was improved by a GPU implementation and deployment of an IGA. A migration policy for the IGA was benchmarked against a simple migration scheme showing results in favour of the non-trivial migration policy. Additionally, the complexity for the user was reduced by the meta-learning system for adaptive parameter control. The meta-learning scheme showed superior results to cases that used the same fixed set of initial values for the parameters.

Funding. University of Duisburg-Essen by the Open Access Publication Fund.

Acknowledgments. We gratefully acknowledge the computing time provided by the Paderborn Center for Parallel Computing (PC²). The funding by the Ministerium für Innovation, Wissenschaft und Forschung des Landes Nordrhein-Westfalen is also appreciated. We thank Andreas Kempf for providing the flame simulation data, and Simone Hochgreb for lending the Cambridge-Sandia burner. Zachary Emuong and Cheau Tyan Foo are thanked for helping with the wavelet noise generation and lab experiments.

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

References

1. F. Scarano, "Tomographic PIV: principles and practice," *Meas. Sci. Technol.* **24**(1), 012001 (2012).
2. A. K. Agrawal, N. K. Butuk, S. R. Gollahalli, and D. Griffin, "Three-dimensional rainbow schlieren tomography of a temperature field in gas flows," *Appl. Opt.* **37**(3), 479–485 (1998).
3. E. Boigné, P. Muhunthan, D. Mohaddes, Q. Wang, S. Sobhani, W. Hinshaw, and M. Ihme, "X-ray computed tomography for flame-structure analysis of laminar premixed flames," *Combust. Flame* **200**, 142–154 (2019).
4. E. Boigné, N. R. Bennett, A. Wang, K. Mohri, and M. Ihme, "Simultaneous in-situ measurements of gas temperature and pyrolysis of biomass smoldering via x-ray computed tomography," *Proc. Combust. Inst.* (2020).
5. H. M. Hertz and G. W. Faris, "Emission tomography of flame radicals," *Opt. Lett.* **13**(5), 351–353 (1988).
6. J. Floyd and A. M. Kempf, "Computed tomography of chemiluminescence (ctc): High resolution and instantaneous 3d measurements of a matrix burner," *Proc. Combust. Inst.* **33**(1), 751–758 (2011).
7. J. Floyd, P. Geipel, and A. M. Kempf, "Computed tomography of chemiluminescence (ctc): 3d time resolved measurements of a turbulent opposed jet flame," *Combust. Flame* **158**(2), 376–391 (2011).
8. M. Hossain, G. Lu, and Y. Yan, "Optical fiber imaging based tomographic reconstruction of burner flames," *IEEE Trans. Instrum. Meas.* **61**(5), 1417–1425 (2012).
9. N. A. Worth and J. R. Dawson, "Tomographic reconstruction of OH* chemiluminescence in two interacting turbulent flames," *Meas. Sci. Technol.* **24**(2), 024013 (2012).
10. J. Wang, Y. Song, Z. hua Li, A. Kempf, and A. zhi He, "Multi-directional 3d flame chemiluminescence tomography based on lens imaging," *Opt. Lett.* **40**(7), 1231–1234 (2015).
11. K. Mohri, S. Görs, J. Schöler, A. Rittler, T. Dreier, C. Schulz, and A. Kempf, "Instantaneous 3d imaging of highly turbulent flames using computed tomography of chemiluminescence," *Appl. Opt.* **56**(26), 7385–7395 (2017).
12. A. Unterberger, M. Röder, A. Giese, A. Al-Halbouni, A. Kempf, and K. Mohri, "3D instantaneous reconstruction of turbulent industrial flames using computed tomography of chemiluminescence (CTC)," *J. Combust.* **2018**, 1–6 (2018).
13. A. Unterberger, A. Kempf, and K. Mohri, "3d evolutionary reconstruction of scalar fields in the gas-phase," *Energies* **12**(11), 2075 (2019).
14. C. T. Foo, A. Unterberger, J. Menser, and K. Mohri, "Tomographic imaging using multi-simultaneous measurements (times) for flame emission reconstructions," *Opt. Express* **29**(1), 244–255 (2021).
15. W. Cai and C. F. Kaminski, "Tomographic absorption spectroscopy for the study of gas dynamics and reactive flows," *Prog. Energy Combust. Sci.* **59**, 1–31 (2017).
16. R. Gordon, "A tutorial on art (algebraic reconstruction techniques)," *IEEE Trans. Nucl. Sci.* **21**(3), 78–93 (1974).
17. K. J. Daun, S. J. Grauer, and P. J. Hadwin, "Chemical species tomography of turbulent flows: Discrete ill-posed and rank deficient problems and the use of prior information," *Journal of Quantitative Spectroscopy and Radiative Transfer* **172**, 58–74 (2016).
18. S. J. Grauer, P. J. Hadwin, and K. J. Daun, "Bayesian approach to the design of chemical species tomography experiments," *Appl. Opt.* **55**(21), 5772 (2016).
19. Q. Lei, Y. Wu, W. Xu, and L. Ma, "Development and validation of a reconstruction algorithm for three-dimensional nonlinear tomography problems," *Opt. Express* **24**(14), 15912–15926 (2016).
20. W. Cai, D. J. Ewing, and L. Ma, "Application of simulated annealing for multispectral tomography," *Comput. Phys. Commun.* **179**(4), 250–255 (2008).
21. K. D. Kihm, K. Okamoto, D. Tsuru, and H. S. Ko, "Adoption of a genetic algorithm (GA) for tomographic reconstruction of line-of-sight optical images," *Exp. Fluids* **22**(2), 137–143 (1996).
22. R. Olmi, M. Bini, and S. Priori, "A genetic algorithm approach to image reconstruction in electrical impedance tomography," *IEEE Transactions on Evolutionary Computation* **4**(1), 83–88 (2000).
23. C. Wu, Y. Cheng, Y. Ding, F. Wei, and Y. Jin, "A novel X-ray computed tomography method for fast measurement of multiphase flow," *Chem. Eng. Sci.* **62**(16), 4325–4335 (2007).
24. A. Bousquet, J. Louchet, and J.-M. Rocchisani, "Fully three-dimensional tomographic evolutionary reconstruction in nuclear medicine," in *Artificial Evolution*, N. Monmarché, E.-G. Talbi, P. Collet, M. Schoenauer, and E. Lutton, eds. (Springer Berlin Heidelberg, Berlin, Heidelberg, 2008), pp. 231–242.
25. Z. Ali Abboud, J. Lavauzelle, Évelyne Lutton, J.-M. Rocchisani, J. Louchet, and F. P. Vidal, "Voxelisation in the 3-d fly algorithm for pet," *Swarm and Evolutionary Computation* **36**, 91–105 (2017).
26. M. Raffel, H. Richard, and G. Meier, "On the applicability of background oriented optical tomography for large scale aerodynamic investigations," *Exp. Fluids* **28**(5), 477–481 (2000).
27. E. Goldhahn and J. Seume, "The background oriented schlieren technique: sensitivity, accuracy, resolution and application to a three-dimensional density field," *Exp. Fluids* **43**(2-3), 241–249 (2007).
28. B. Atcheson, I. Ihrke, W. Heidrich, A. Tevs, D. Bradley, M. Magnor, and H.-P. Seidel, "Time-resolved 3d capture of non-stationary gas flows," *ACM Trans. Graph.* **27**(5), 1–9 (2008).
29. M. Raffel, "Background-oriented schlieren (bos) techniques," *Exp. Fluids* **56**(3), 60 (2015).

30. F. Nicolas, V. Todoroff, A. Plyer, G. Le Besnerais, D. Donjat, F. Micheli, F. Champagnat, P. Cornic, and Y. Le Sant, "A direct approach for instantaneous 3d density field reconstruction from background-oriented schlieren (bos) measurements," *Exp. Fluids* **57**(1), 13 (2015).
31. H. M. Lang, K. Oberleithner, C. O. Paschereit, and M. Sieber, "Measurement of the fluctuating temperature field in a heated swirling jet with bos tomography," *Exp. Fluids* **58**(7), 88 (2017).
32. S. Grauer, A. Unterberger, A. Rittler, K. J. Daun, A. Kempf, and K. Mohri, "Instantaneous 3D flame imaging by background-oriented schlieren tomography," *Combust. Flame* **196**, 284–299 (2018).
33. S. J. Grauer and A. M. Steinberg, "Fast and robust volumetric refractive index measurement by unified background-oriented schlieren tomography," *Exp. Fluids* **61**(3), 80 (2020).
34. H. Liu, J. Huang, L. Li, and W. Cai, "Volumetric imaging of flame refractive index, density, and temperature using background-oriented schlieren tomography," *Science China Technological Sciences* (2020).
35. R. Hinterding, Z. Michalewicz, and A. E. Eiben, "Adaptation in evolutionary computation: a survey," in *Proceedings of 1997 IEEE International Conference on Evolutionary Computation (ICEC '97)*, (1997), pp. 65–69.
36. A. E. Eiben, Z. Michalewicz, M. Schoenauer, and J. E. Smith, "Parameter control in evolutionary algorithms," *Studies in Computational Intelligence* **54**, 19–46 (2007).
37. R. Hinterding, "Self-adaptation using multi-chromosomes," in *Proceedings of 1997 IEEE International Conference on Evolutionary Computation (ICEC '97)*, (1997), pp. 87–91.
38. K. Juliff, "A multi-chromosome genetic algorithm for pallet loading," in *Proceedings of the 5th International Conference on Genetic Algorithms*, (Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1993), pp. 467–473.
39. H. J. Pierrot and R. Hinterding, "Using multi-chromosomes to solve a simple mixed integer problem," in *Advanced Topics in Artificial Intelligence*, A. Sattar, ed. (Springer Berlin Heidelberg, Berlin, Heidelberg, 1997), pp. 137–146.
40. L. Araujo, J. Merelo Guervós, C. Cotta, and F. Vega, "Multikulti algorithm: Migrating the most different genotypes in an island model," *CoRR* **abs/0806.2843**, (2008).
41. L. Araujo, J. J. Merelo, A. Mora, and C. Cotta, "Genotypic differences and migration policies in an island model," in *Proc. 11th Annu. Genet. Evol. Comput. Conf. GECCO-2009*, (2009), 1331–1338.
42. B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'81*, (Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1981), pp. 674–679.
43. B. K. Horn and B. G. Schunck, "Determining optical flow," *Techniques and Applications of Image Understanding* **0281**, 319–331 (1981).
44. F. Proch, P. Domingo, L. Vervisch, and A. M. Kempf, "Flame resolved simulation of a turbulent premixed bluff-body burner experiment. part i: Analysis of the reaction zone dynamics with tabulated chemistry," *Combust. Flame* **180**, 321–339 (2017).
45. F. Proch, P. Domingo, L. Vervisch, and A. M. Kempf, "Flame resolved simulation of a turbulent premixed bluff-body burner experiment. part ii: A-priori and a-posteriori investigation of sub-grid scale wrinkling closures in the context of artificially thickened flame modeling," *Combust. Flame* **180**, 340–350 (2017).
46. J. E. Baker, "Reducing bias and inefficiency in the selection algorithm," in *Proceedings of the Second International Conference on Genetic Algorithms on Genetic Algorithms and Their Application*, (L. Erlbaum Associates Inc., USA, 1987), pp. 14–21.
47. J. Stam and E. Languénoü, "Ray tracing in non-constant media," in *Rendering Techniques '96*, X. Pueyo and P. Schröder, eds. (Springer Vienna, Vienna, 1996), pp. 225–234.
48. I. Ihrke, G. Ziegler, A. Tevs, C. Theobalt, M. Magnor, and H.-P. Seidel, "Eikonal rendering: Efficient light transport in refractive objects," *ACM Trans. Graph.* **26**(3), 59 (2007).
49. M. S. Sweeney, S. Hochgreb, M. J. Dunn, and R. S. Barlow, "The structure of turbulent stratified and premixed methane/air flames i: Non-swirling flows," *Combust. Flame* **159**(9), 2896–2911 (2012).
50. M. S. Sweeney, S. Hochgreb, M. J. Dunn, and R. S. Barlow, "The structure of turbulent stratified and premixed methane/air flames ii: Swirling flows," *Combust. Flame* **159**(9), 2912–2929 (2012).
51. S. B. Pope, *Turbulent Flows* (Cambridge University Press, 2000).
52. J. Nickolls, I. Buck, M. Garland, and K. Skadron, "Scalable parallel programming with cuda," *Queue* **6**(2), 40–53 (2008).
53. E. Alba, G. Luque, and S. Nesmachnow, "Parallel metaheuristics: recent advances and new trends," *International Transactions in Operational Research* **20**(1), 1–48 (2013).
54. Y.-J. Gong, W.-N. Chen, Z.-H. Zhan, J. Zhang, Y. Li, Q. Zhang, and J.-J. Li, "Distributed evolutionary algorithms and their models: A survey of the state-of-the-art," *Applied Soft Computing* **34**, 286–300 (2015).
55. M. G. Bronstein, I. N. Semendjajew, K. A. , and M. H. , *Taschenbuch der Mathematik* (Harri Deutsch GmbH, 2006), 6th ed.
56. R. Gonzalez and R. Woods, *Digital Image Processing* (Prentice Hall, 2006), 3rd ed.
57. A. Unterberger, J. Menser, A. Kempf, and K. Mohri, "Evolutionary camera pose estimation of a multi-camera setup for computed tomography," in *2019 IEEE International Conference on Image Processing (ICIP)*, (2019), pp. 464–468.

DuEPublico

Duisburg-Essen Publications online

UNIVERSITÄT
DUISBURG
ESSEN

Offen im Denken

ub | universitäts
bibliothek

This text is made available via DuEPublico, the institutional repository of the University of Duisburg-Essen. This version may eventually differ from another version distributed by a commercial publisher.

DOI: 10.1364/OE.450036

URN: urn:nbn:de:hbz:465-20220831-144114-6

All rights reserved.