euronoise 2021
25th to 27th of October
Madeira, Portugal - online

Applications of Machine Learning Bounding-Breps for Optimised Acoustical Reflectors

John O'Keefe

O'Keefe Acoustics
john@okeefeacoustics.com

Abstract
Traditional Machine Learning (ML) algorithms shape geometries such as an acoustical reflector inside a rectilinear six-sided box, often referred to as a Bounding-Box. The space available for acoustical reflectors in a performing arts venue rarely takes the form of a rectilinear box. Designers must either limit the ML optimisation to a box that is smaller than desired or risk the result of a reflector that interferes with sight or lighting lines. A series of geometric algorithms has been developed that are being referred to as Bounding-Breps. These have been applied to two different buildings: a 175 seat multi-purpose "Black-box" theatre and a 670 seat proscenium arch theatre. Reflectors in the rooms have been optimised with a multi-objective Fitness Function that has been developed to distribute laterally reflected sound as uniformly as possible on a receiving surface. Pareto-Optimal sets of reflector design options suggest improvements in Lateral Energy Fractions of one or, depending on initial conditions, several Just Noticeable Differences.

Keywords: Room Acoustics; Multi-Objective Optimization (MOO); Genetic Algorithms(GA); Pareto Set; Epsilon-Constraint Method

1 Introduction

The techniques of Machine Learning (ML) are making their presence felt in the world of acoustics research. A very thorough summary is presented in [1], noting the application of ML in areas such as audiology, speech recognition, underwater acoustics and more. For architectural acoustics, though, there is an unfortunate paucity of publications. The one notable exception being Sato et al.'s work in 2002 [2]. Pioneering work by Kathirchelvan [3] encouraged the author to explore the field further. A series of what might be called "geometric algorithms" was developed in [4] that allows a computer to build and iteratively modify acoustic reflecting panels or solid reflecting objects. As opposed to the rectilinear Bounding-Boxes used in traditional geometric optimisation, these methods are based on Non-uniform Rational B-Splines (Nurbs) and Boundary Representations (Breps). The result is what might be called, not a Bounding-Box but, rather, a Bounding-Brep (or B-Brep). In that same study, the author posited that Genetic Algorithms (GA) might be more appropriate for architectural acoustics design, rather than the Simulated Annealing algorithm [5] used by Kathirchelvan. Most Simulated Annealing algorithms are better suited to single objective optimisation and, as such, render a single "optimal" design. Design for architectural acoustics – and architecture in general – is a multi-faceted challenge that is best met with algorithms that can deal with multi-objective optimisation. The Genetic Algorithm [6] is one such routine, the most popular version being the Non-dominated Sorting Genetic Algorithm (NSGA-II) developed by Deb, et al. [7]. The goal of this paper is to experiment with some practical applications of the concepts developed by the author in [4] using NSGA-II.

2 Objective Function

An Objective Function, used to guide the GA through the optimisation process, has been proposed in [4]. It shall be briefly described here. First, it is important to ensure a uniform distribution of reflection intersections on the receiver surface. Otherwise, the GA could end up concentrating the reflections in one corner or along one edge of the receiver surface. A grid is constructed on top of the receiver surface and the reflection intersections are attracted to its nodes or vertices. The spacing of the vertices (d_{MEAN}) is based on the optimum packing of circles on a surface.

$$d_{mean} = \sqrt{\frac{Area\ of\ Receiving\ Surface}{\#\ of\ Receiving\ Points}} \tag{1}$$

The "lateralness" of a given reflection is quantified using a method developed by Protheroe and Day [8], which will be referred to here as the Single Reflection Lateral Fraction (sLF). Ref. [8] has found a reasonable correlation between the sLF and the more familiar Lateral Fraction [9]. The assumption is made that, as the iteration process evolves, the distances between the received reflection points and the grid vertices will form a normal distribution. Likewise the sLFs at the points of intersection are assumed to be normally distributed with respect to the desired value of sLF_{GOAL}. The two dimensional Object Function is thus:

$$d_{MEAN}\ Fitness = \frac{1}{N_{grid}} \sum_{i=0}^{N_{grid}} 1 - e^{\frac{d_{NN}}{\tau_{dist}}} \tag{2}$$

where: N_{grid} is the number of grid points
d_{NN} is distance from the i^{th} grid point to its nearest neighbour (NN)
τ_{dist} is the convergence coefficient associated with distance

and:

$$sLF_{GOAL}\ Fitness = \frac{1}{N_{grid}} \sum_{i=0}^{N_{grid}} 1 - e^{\frac{\Delta sLF}{\tau_{sLF}}} \tag{3}$$

where: ΔsLF is the difference between the sLF at the NN and sLF_{goal}
τ_{sLF} is the convergence coefficient associated with associated with sLF

3 NSGA-II

3.1 Brief Description

The Non-dominated Sorted Genetic Algorithm (NSGA-II) [7] has proved very useful in solving multi-objective problems. It is based on the concept of non-dominated Pareto-Optimal Fronts and Crowding Distance. A Pareto Front is the set of optimum solutions (in our case, a set of reflector design iterations) where one of the two Objective Functions cannot be improved without reducing the other Objective Function, and vice versa. In Figure 1, we see that, after 50 generations, the NSGA-II routine has found three Pareto Fronts. When one has to decide which solutions to propagate the next generation, it is best to encourage diversity. So, for example, in Figure 1, the two solutions in the middle of the third front (at

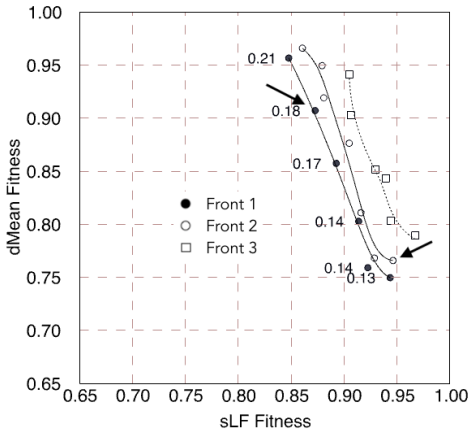

Figure 1 Pareto Fronts for the Nia ceiling reflector, optimised with the Uniform Cartesian Grid method. sLF values are shown for solutions on the 1st Front.

0.94, 0.85) are close to each other. One of them might be useful to the next generation but not both. Their properties are too much alike.

Thus, the way to interpret a representation of the data as shown in Figure 1 is; because we are trying to minimize the functions, the better solutions are the ones closer to the origin., i.e. the ones on the first front. Along that front any improvement in sLF will mean a compromise in the uniformity of the sLF on the receiver surface. That may be true of any point on the graph but this is the set where one will find the best compromises. Crowding Distance (i.e. two solutions close together) is not a factor in the decision making. It is a very useful component of the algorithm as it progresses towards better solutions but is not helpful in the eventual selection of a good reflector design.

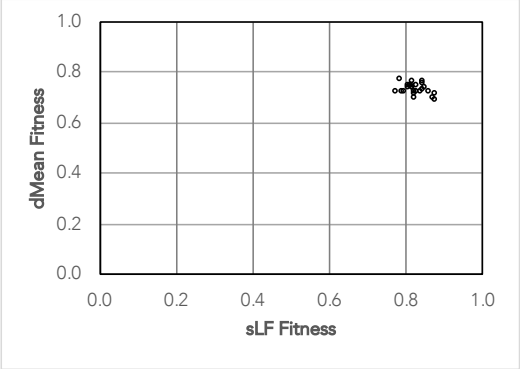

Figure 2. The estimated Pareto Front for the model shown in Figure 4, found by the Epsilon-Constraint method (ε = 0.04).

3.2 Pareto-Optimal Front Identification

In the current study, we are trying to minimise two fitness functions simultaneously towards the ultimate Pareto-Optimal Front. The problem is, we don't know where that front lies beforehand. One way to identify the Pareto Front is with the so-called Epsilon-Constraint method [10]. In this method, the Genetic Algorithm iteratively minimises one of the fitness functions whilst the second is constricted. Then the second fitness is minimised while the first is constricted. Figure 2 shows the results when the search of the fitness functions is spanned across their entire range, i.e. from 0.0 to 1.0. This was done with the ε increment set at 0.04, resulting in something that looks more like a bundle than a front. A further "fine-grain" study was performed around this region, this time with the ε increment set at 0.014. The results were similar. It should be stressed that what has been found here is only an estimate of an unknown Pareto-Optimal Front. These searches are computationally expensive, taking several hours for each fitness function. It was decided that the estimate shown in Figure 2 and its "finer grain" partner were probably accurate enough. This was confirmed in the subsequent experiments, to be described in the following sections.

3.3 Parameter Settings for Experiments

For the Epsilon-Constraint method Pareto Front search shown in Figure 2 and for all the other studies quoted in this paper, the Genetic Algorithm (GA) and Ray Tracer routines were configured as shown in Table 1.

Table 1 – Experimental Parameter Settings

Tool	Parameter	Value
Ray Tracer	# of Rays	19,700
Reflector Grid	U-V-W Counts	6-6-20
GA	# of Generations	50
GA	Population	25
GA	Mutation Rate	1.5%
GA	Elitism	0
Object Function	d_{MEAN}	Eqn.(2)
Object Function	sLF_{GOAL}	0.35
Object Function	τ_{sLF}	0.05
Object Function	τ_{dMean}	500

4 Project Examples

4.1 Nia Centre Project Description

The Nia Centre for the Arts is a multi-disciplinary arts institution serving the Afro-Canadian diaspora in Toronto, Canada. The newly renovated building will include, among other facilities, a 175 seat flexible theatre space. One of the acoustic design goals was to keep the room as lively as possible, without compromising Speech Intelligibility or musical Clarity. For good speech intelligibility, the early reflected sound must be stronger than the late reflected sound. As outlined in the ISO 3382-1 standard [9], the accepted threshold between early and late sound is 50 ms. The traditional approach in these so-called Black-Box theatres is to provide the appropriate early to late balance by absorbing the late energy with adjustable acoustic curtains. This can, and often does, create an acoustically "dead" room. An acoustically claustrophobic space that lacks room Presence, as described by Marshall [11]. The design of the Nia Centre's acoustics represents a reversal of the traditional approach to the early to late energy balance problem. Rather than attenuate the late energy, the goal has been to encourage the early energy as much as possible and, in so doing, reduce the need to absorb the later reflected sound. Indeed, when the building opens it will not have any acoustic curtains. Although, curtain tracks will be provided if and when an acoustically dead space might thought to be theatrically or otherwise appropriate.

An early version of the acoustic model for the Nia Centre's main auditorium is shown in Figure 3. It was found that, like most other rooms, only small areas of the walls provide the early reflected sound required for good Speech Intelligibility or musical Clarity. Figure 3 shows the room in its End-Stage format but similar studies show similar results in a Centre-Stage configuration. For a sound source at the front of the room, the yellow trapezoids indicate the very limited zones of the side and upstage walls that cast 1st order reflections to the receiving surface. With a limited budget and the multi-purpose realities that a room like this must satisfy, a complicated geometry of side wall reflectors was thought inappropriate. Rather, it was decided to slope the upper reaches of the walls down such that they form a tilted connection to the ceiling above. A version very close to the final design of those reflectors is seen in Figure 3.

This design was not arrived at using the Machine Learning methods that are the subject of this paper but, rather, by the author's manual manipulation of the computer model using Method of Images sources to cast reflections towards the receiving surface(s). For each of the eight reflector components that make up the tilted ceiling/wall, a virtual image source was created and then the reflector was manually tilted or rotated. The computer model and its manipulation were done using C# programs developed in the Rhino/Grasshopper Version 6 software environment, as were the subsequent GA optimisations. Having committed this design to construction, it was later decided to use it as a test-bed for the Machine Learning methods developed in [4].

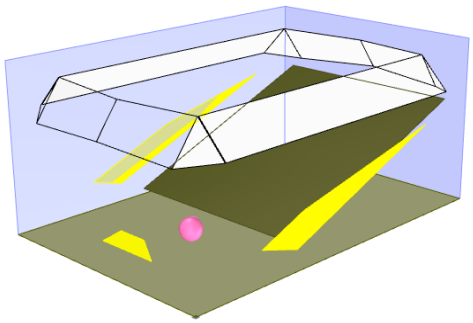

Figure 3 Nia Centre computer model showing the effective (yellow) reflection zones on the walls and the ceiling tilted down around its perimeter.

4.2 Optimisation – Cartesian Grid Methods

Two of the seven geometric algorithms developed in [4] were employed in this first experiment with the Nia Centre; the Uniform Cartesian Grid method and the Shaped Cartesian method. In both methods a panel-like reflector is created from a 3-dimensional grid of control points. The control points are perturbed randomly along the gridlines

created by either one of the two methods. In the NSGA-II genetic algorithm, this grid is used in concert with a Simulated Binary Crossover [12] to create each generation's progeny.

The grid-line structure used by the computer for the two experiments with the Nia auditorium is shown in Figure 4a. The black oval-like curves are the 2-dimensional Bounding Nurb curves (C_{BN}) that have been projected on to 20 planes located on and perpendicular to the Spinal curve that surrounds the top perimeter of the room. There are 6 pairs of U-V gridlines, shown with the white dashed lines. The U-V gridline sets are also projected on to the 20 perpendicular planes. The white dashed lines between the 20 gridline bundles are the W-lines. A further explanation of this geometric structure is given in [4].

The initiating reflector mesh for the Genetic Algorithm is seen in the white lines near the top of Figure 4b. It is a flat surface, a triangular mesh parallel to the X-Y plane of the room's ceiling. The vertices of this mesh are perturbed along the grid lines and a Nurb based Boundary Representation (or Brep) is formed. In this and the following figures it is seen as the semi-transparent undulating grey surface. The pink sphere at the front of the room is the sound source and the sloped brown rectangle is the receiver surface. The small white spheres on the receiving surface are the vertices of the receiver grid used in the Objective Function, described in Section 2. One of the goals of the Objective Function is to get the reflections to arrive as close as possible to these grid points. The other goal is to have those reflections arrive from a lateral direction. The yellow arrows on top of the receiver surface indicate three things: the base of the arrow is the location where the reflection has intersected the receiver surface, the direction of the arrow indicates where it came from, and the length of the arrow indicates the sLF value associated with it.

4.3 Uniform Cartesian Grid Method Results

After 50 generations of optimisation searching, the NSGA-II algorithm has selected and sorted our population of reflector geometries, as shown in the graph originally presented in Figure 1. There are 25 reflector designs in the population, each of which is on its way to optimisation. The ones on the first Pareto Front are closest to that. There is no single "best" reflector design but, rather, a collection of good designs that have been optimised. From these, we can select the most appropriate.

Figure 1 can be read as follows. The "better" designs are found on the first Pareto Front, indicated by the solid circles. To the left of the first front, there is a series of numbers, stating the sLF value for the adjacent solution.

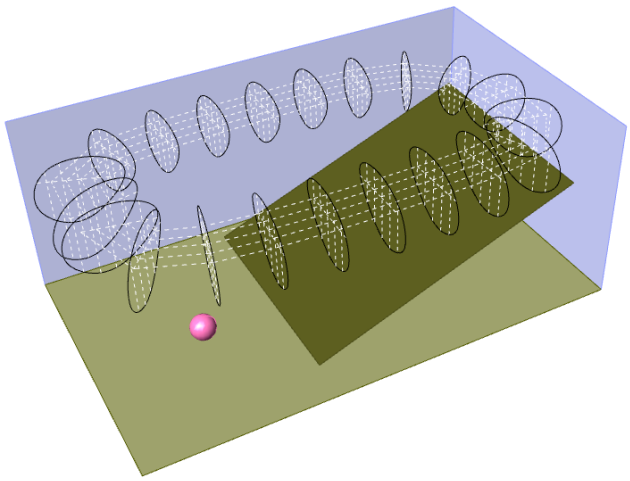

Figure 4a Gridline structure used by the computer for reflector optimization of the Nia Centre ceiling.

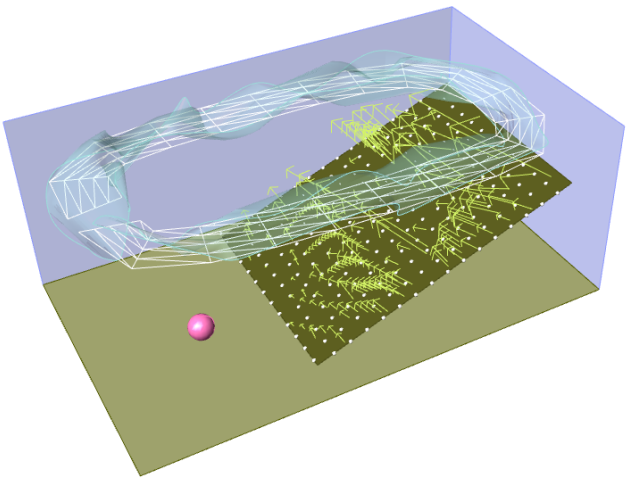

Figure 4b The initiating mesh, seen as the flat surface parallel to the ceiling and a Brep generated from a perturbation of this mesh, shown in transparent grey.

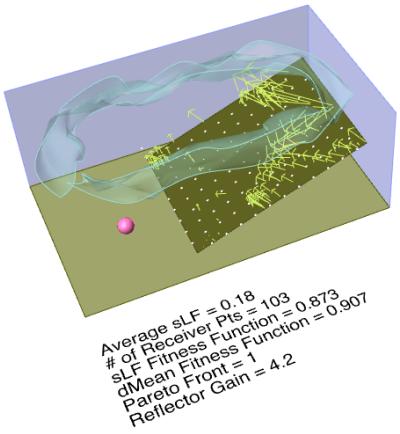

Figure 5a Representation of the reflection field for a 1st Pareto Front solution from Figure 1.

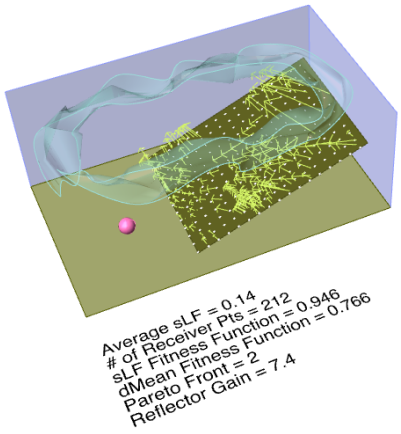

Figure 5b Same as Figure 5a for a 2nd Pareto Front solution from Figure 1. The sLF is slightly lower than in Figure 5a but the reflector casts more and better distributed reflections to the receiver surface.

We might point out here that the predicted sLF values should not be confused with the sLF Fitness. The sLF value gives an indication of what a listener might hear in the room, correlating as it does with the Lateral Fraction. It is, in a sense, the thing that the acoustical designer is most interested in. The sLF Fitness is merely a construct for the Genetic Algorithm (GA), it is a part of the Objective Function study – a tool used to guide the GA towards the optimum solutions.

The software developed for this study allows the user to scroll through a solution chart, such as the one shown in Figure 1, to review other aspects of a given solution. So, to demonstrate, we have chosen two points from the solutions in Figure 1. One on the first Pareto Front and the other on the second. The black arrows in Figure 1 point towards the two solutions. The solution on the first Pareto Front has an sLF value of 0.18. A representation of its reflection field is shown in Figure 5a. The solution on the second front has an sLF of 0.14. Its representation is shown in Figure 5b. We shall consider the differences between the two examples.

Insofar as Spatial Impression is concerned, the difference in sLF values is 0.04, slightly less than a Just Noticeable Difference (JND) for the Lateral Fraction [13] (which is a correlate of sLF). But a visual comparison of the two images clearly indicates that the reflector generating the higher sLF (Figure 5a) delivers fewer total reflections than the solution delivering the lower sLF (Figure 5b). In Figure 5a, the solution with the higher sLF, there is a noticeable "dead zone" in the middle of the receiving surface. One way to quantify the difference between these reflection count efficiencies is with a measure we have called Optimisation Gain. It is a simple logarithmic ratio of the Gain of the optimised reflector to the Gain of the original or initiating reflector.

$$Optimisation\ Gain = Gain_{optimised} - Gain_{original} \tag{4}$$

where: $Gain = 20log \sum_i^N \frac{Reflection\ Path\ Length_i}{Direct\ Path\ Length_i}$

N = # of successful reflections cast by the reflector

The difference between the Optimisation Gains for the two examples on the first and second Pareto Fronts (i.e. Figures 5a and 5b) is 3.2 dB. The reflector with the lower sLF is the one that will encourage a higher overall acoustic Strength (G) [9]. The JND for Strength (G) is generally taken to be 1.0 dB. While it should be clear that the 3.2 dB difference in Optimisation Gain is for individual panels and the JND of 1.0 dB is for an entire

room of reflecting surfaces, the difference is worth noting. A detailed study of the two reflectors inside a complete computer model of the room is warranted.

Perhaps the most cogent question is, however, did the optimisation exercise actually produce better acoustical results? It has, indeed. But that is due in good part to the initial conditions. The panel started out as a flat surface, parallel to the X-Y plane, generating a space averaged sLF of 0.00. Optimised sLF values on Figure 1's first Pareto Front range from 0.13 to 0.21. This represents an improvement of three to four JNDs, which, of course, would be a clearly audible improvement.

4.4 Shaped Cartesian Grid Method Results

The exercise described in Section 4.3, using the Uniform Cartesian Grid method, was repeated using the Shaped Cartesian method. The difference between the two methods is that, in the latter, the gridlines used to govern the movement of the B-Brep's control points are based on the shape of the initiating panel as drawn by the designer. In the former, the Uniform Cartesian method, the gridlines are uniformly distributed inside the search space and the initiating panel starts out flat and parallel to the X-Y plane. The Shaped Cartesian Grid benefits from the designer's initial input but may not always explore the full extents of the search space. The Uniform method explores the search space better but may not converge as efficiently.

In this experiment, using a Shaped Cartesian Grid, the initiating reflector was slightly curved and tilted at an angle of 22.5° down and in towards the centre of the room. Not unlike the original reflector design shown in Figure 3. After a optimisation run of 50 generations, and as shown in Figure 6, the population of 25 panel designs have lined up well and are distributed evenly along only two Pareto Fronts. Of the 25 potential solutions only one is not on one of these first two fronts and most are on the first front. This is the hallmark of the NSGA-II algorithm. It converges quickly, in part because it avoids the "crowding" of solutions along a given Pareto Front. Less crowding means more diversity in the gene pool as the algorithm moves from one generation to the next.

It is also clear, at least in these experiments, that the Shaped Cartesian Grid method is providing better results than the Uniform Grid. Both in terms of Fitness scores and, more importantly, the sLF values that have been achieved.

As mentioned above, one of the advantages of a Pareto set of optimised solutions, as opposed to a single optimum solution that a Simulated Annealing algorithm might provide, is that one can study and perhaps question the solutions proposed by the computer. This time, we'll consider the two extremes of the first Pareto Front shown in Figure 6. At one extreme we have a solution with an sLF of 0.20. An image of the reflection field for this solution is shown in Figure 7a. At the other end of the first Pareto Front the solution has an sLF of 0.27, shown in Figure 7b. Note the difference in the number and distribution of the yellow arrows indicating the incident reflections. The difference between the Optimisation Gains at the two extremes of the first Pareto Front (i.e. Figures 7a and 7b) is 3.7 dB. Again, this is higher than the JND for a normal acoustic Strength (G) measurement but, as noted, our Optimisation Gain parameter and an ISO 3382-1 Strength (G) measurement are not the same thing. The 3.7 dB difference is worth noting but further study is required.

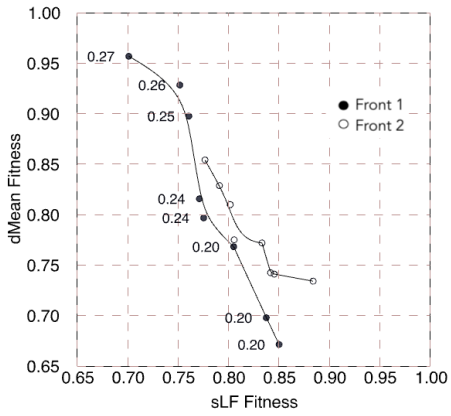

Figure 6. Pareto Fronts for the Nia Centre using the Shaped Cartesian method.

Finally, we return to the question of acoustically audible improvements. Has the optimisation of the reflector improved the acoustics? In the previous experiment it was easy to make an improvement

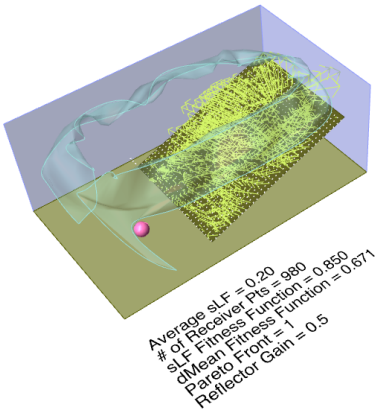

Figure 7a. Representation of the reflection field for the (0.70,0.96) point in Figure 6.

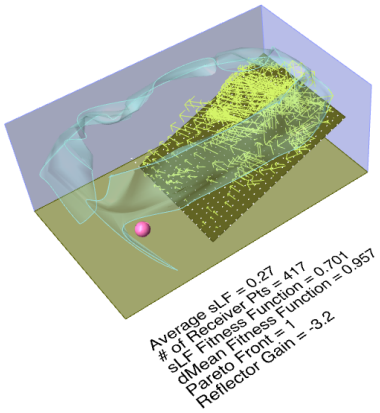

Figure 7b. Same as Figure 7a, for the (0.85,0.67) point in Figure 6.

because the initiating reflector wasn't generating any lateral reflections at all. In this experiment, the original reflector design was indeed generating lateral reflections and had an initial average sLF of 0.17. In Figure 6 we see that on the first Pareto Front, the NSGA-II optimisation has found solutions with sLFs in the range of 0.24 and as high as 0.27. This represents an improvement of slightly more than one Just Noticeable Difference (JND). Does this small improvement justify the fabrication of such a complicated geometry? Probably not. Further discussion will follow.

4.5 Polar Grid Project Description and Results

The Playhouse Theatre is a 670 seat room, part of the Queen Elizabeth Theatre complex in Vancouver, Canada. Originally opened in 1960, it now faces either renovation or demolition. As part of a proposal for a feasibility study, the author prepared a computer model of the room. This will serve as the second test-bed for the geometric algorithms developed in [4], this time using the Polar Grid method. Similar to many rooms built in its era, the Playhouse is very wide, fan shaped at the front, and lacks laterally reflected sound. The author conceived the idea of a central annulus installed above the stalls seating. The open centre of the annulus might be used for lighting positions while its perimeter ring could act as an acoustic reflector providing lateral reflections. Objective Function and Pareto Front studies similar to those described above were carried out. The initiating ring reflector was circular in cross-section. An image of one of the optimisation designs developed by the Genetic Algorithm is shown in Figure 8.

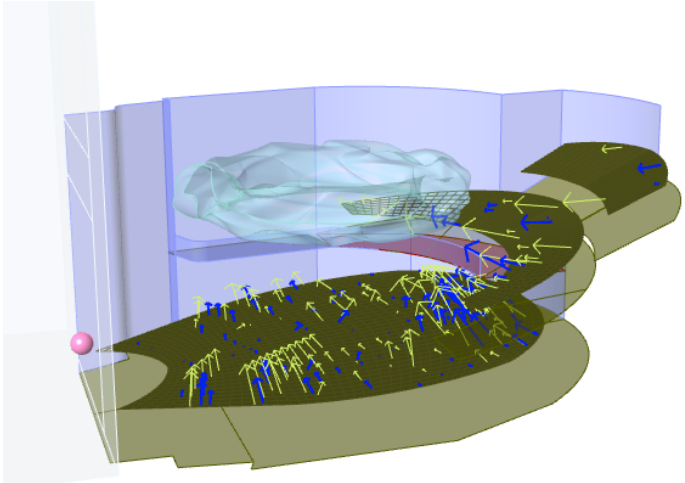

Figure 8 Playhouse Theatre, Vancouver, showing an optimised circular reflector and the reflection arrival vectors. Arrows indicate reflection point of reception, direction of arrival and sLF value. Original reflector in blue, this iteration in yellow.

Without the annulus in place, the sound headed in its direction would hit the flat ceiling, generating only frontal reflections and an sLF of 0.00. Adding the (initiating) circular annulus generated space averaged sLFs in the range of 0.10. The GA optimised reflector produced sLFs in the range of 0.11 to 0.15 on the first

Pareto Front. This suggests a Lateral Fraction improvement of 2 Just Noticeable Differences (JNDs) just for the installation of the cylindrical annulus and a further single JND when it has been optimised by the computer. We might note here, that is often difficult to improve on the scattering created by circular or cylindrical object such as the initiating annulus, yet the NSGA-II algorithm has managed to do so.

5 Discussion

The first thing to point in this discussion, of course, is that some of the reflector designs developed by the Machine Learning algorithms are not entirely practical insofar as economic constructability or theatrical practicalities are concerned. Nor were they intended to be. The intention of this study has been to apply the algorithms in theoretical rooms to discover how they might be refined for more practical application in the real world.

The complicated curvaceous geometry generated for the Nia reflectors led to sLFs in the range of 0.20 to 0.27. The average of the solutions on the first Pareto Front was sLF = 0.233. The Nia reflectors that will actually be built are simply 8 flat surfaces of gypsum board, tilted down towards the centre of the room. Their predicted sLF is 0.165. The difference between the two is only slightly beyond a single JND. This would hardly justify the more complicated construction. The designs for the Playhouse Theatre, however, are more practical and may be worth consideration. For example, at the Nia Centre, the reflectors are above the lighting grid and cannot be seen. For the Playhouse Theatre, the open centre of the reflector offers the opportunity for theatrical lighting positions. The reflector itself could be easily fabricated off-site from glass fibre reinforced gypsum (GRG) to be installed as an artistic feature of the room's ceiling.

Secondly, it's worth comparing the first two methods we have applied. The Uniform Cartesian Grid method is better at exploring the search space but, in a sense, relies more on the computer to find the best solution. The Shaped Cartesian method benefits from the designer's knowledge, initiating with what might already be a better solution. In this and other experiments the Shaped Cartesian method – the one that relies more on the designer – produces better acoustic results.

Finally, Figure 9 shows a closer view of one of the solutions generated by the Shaped Cartesian experiment described in Section 4.4. The yellow arrows on the brown receiving surface indicate the direction of arrival for each reflection. Note how the arrows are not all pointing in the same direction. This is a good result but it is not always the case. Some solutions, ones which have reasonably high average sLFs, receive most of their lateral energy from one side of the room, i.e. most of the yellow arrows point in the same direction. This might suggest a potential image shift or, at the very least, a slightly peculiar spatial image. The diversity of incidence directionality, as seen in Figure 9, is desirable and perhaps might be considered as an addition to the Objective Function. As might Optimisation Gain. The NSGA-II algorithm however is best suited to no more than 3 optimisation dimensions.

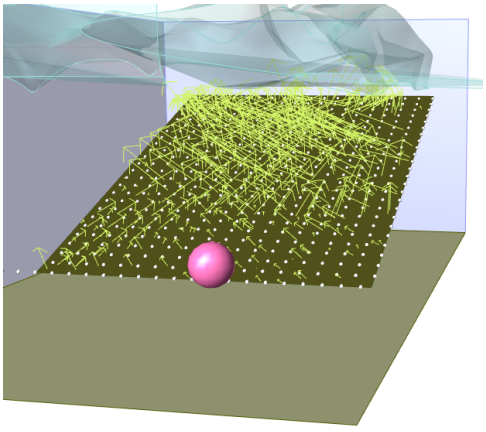

Figure 9. Nia Centre reflection field showing equally distributed directions of arrival.

6 Conclusion

Traditional optimisation of geometric objects has been done inside rectilinear Bounding-Boxes. This is a limitation because the space available for an acoustical reflector inside a theatre or concert hall rarely takes the form of a six-side rectilinear box. A series of geometric algorithms based on Nurb curves has

been developed in [4]. These produce Bounding-Breps (or B-Breps) inside which a computer can perturb and optimise the shape of a reflector. The current study has applied three of these geometric algorithms to real world challenges. The results suggest optimisations that would be audible, i.e. beyond Just Noticeable Differences. The constructability of some of the solutions is questionable. Some, though, may have merit. Judicious application is encouraged.

Acknowledgements

Once again, the author would like to acknowledge the numerous universities, institutions and private individuals currently making their lectures and tutorials freely available on-line. In particular, the NPTEL network and the Government of India's Global Initiative of Academic Networks (GAIN).

References

[1] Bianco, M.J., Gerstoft, P. et al. Machine learning in acoustics: Theory and applications, : *Journal of the Acoustical Society of America* Vol 146, 2019, pp. 3590–3628

[2] Sato, S. Otori, K. Takizawa, A. Sakai, H. Ando, Y. Kawamura, H. Applying Genetic Algorithms To The Optimum Design of a Concert Hall, *Journal of Sound and Vibration,* Vol 258(3), 2002, pp. 517-526.

[3] Kathirchelvan, T. Robust shape optimization of NURBS based acoustic reflectors using stochastic search techniques, Proc. *International Symposium on Room Acoustics*, 2013

[4] O'Keefe, J. Geometric Algorithms for Machine Based Optimisation of Acoustic Reflectors, *Proc I3DA,* Bologna, 2021.

[5] Kirkpatrick, S. Gelatt Jr., C.C. Vecchi, M.P. Optimization by Simulated Annealing. *Science.* Vol. 220 (4598), 1983, pp. 671-680.

[6] Holland, J, *Adaptation in natural and artificial systems*, Ann Arbor, MI: University of Michigan Press, 1975

[7] Deb, K. Pratap, A. Agarwal, S. Meyarivan, T. A fast and elitist multi-objective genetic algorithm: NSGA-II *IEEE Transactions on Evolutionary Computation.* 6 (2), 2002, p.182

[8] Protheroe, D. Day, C. Validation of lateral fraction results in room acoustic measurements, *InterNoise* Melbourne, 2014.

[9] ISO 3382-1, 2009. Acoustics - Measurement of room acoustic parameters - Part 1: Performance spaces, *International Organization for Standardization*, Geneva.

[10] Laumanns, M., Thiele, L., Zitzler, E. An Adaptive Scheme to Generate the Pareto Front Based on the Epsilon-Constraint Method, *Practical Approaches to Multi-Objective Optimization*, Dagstuhl Seminar Proceedings, 2005.

[11] Marshall, A.H., Acoustical Dimensions for Production of Presence in Music Spaces, *Proc. IOA,* Vol. 40. Pt. 3. , 2018, pp. 137-146.

[12] Deb, K. Agrawal, R. B. Simulated Binary Crossover for Continuous Search Space, *Complex Systems* Vol. 9, 1995, pp. 115-148.

[13] Cox, T.J. Davies, W.J. and Lam, Y.W. The sensitivity of listeners to early sound field changes in auditoria, *Acustica* Vol. 79, 1993, pp. 27-41.