Evolutionary Optimization of a Focused Ultrasound Propagation Predictor Neural Network

Jakub Chlebik
Faculty of Information Technology
Brno University of Technology
Centre of Excellence IT4Innovations
Brno, Czech Republic
ichlebik@fit.vut.cz

Jiri Jaros
Faculty of Information Technology
Brno University of Technology
Centre of Excellence IT4Innovations
Brno, Czech Republic
jarosjir@fit.vut.cz

ABSTRACT

The search for the optimal treatment plan of a focused ultrasound-based procedure is a complex multi-modal problem, trying to deliver the solution in clinically relevant time while not sacrificing the precision below a critical threshold. To test a solution, many computationally expensive simulations must be evaluated, often thousands of times. The recent renaissance of machine learning could provide an answer to this. Indeed, a state-of-the-art neural predictor of Acoustic Propagation through a human skull was published recently, speeding up the simulation significantly. The utilized architecture, however, could use some improvements in precision. To explore the design more deeply, we made an attempt to improve the solver by use of an evolutionary algorithm, challenging the importance of different building blocks. Utilizing Genetic Programming, we improved their solution significantly, resulting in a solver with approximately an order of magnitude better RMSE of the predictor, while still delivering solutions in a reasonable time frame. Furthermore, a second study was conducted to gauge the effects of the multi-resolution encoding on the precision of the network, providing interesting topics for further research on the effects of the memory blocks and convolution kernel sizes for PDE RCNN solvers.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; *Genetic programming*; • **Applied computing** → Consumer health.

KEYWORDS

Evolutionary Optimisation, Evolutionary Design, Ultrasound Propagation Predictor, Cartesian Genetic Programming

ACM Reference Format:
Jakub Chlebik and Jiri Jaros. 2023. Evolutionary Optimization of a Focused Ultrasound Propagation Predictor Neural Network. In *Genetic and Evolutionary Computation Conference Companion (GECCO '23 Companion), July 15–19, 2023, Lisbon, Portugal.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3583133.3590661

GECCO '23 Companion, July 15–19, 2023, Lisbon, Portugal

ACM ISBN 979-8-4007-0120-7/23/07.
https://doi.org/10.1145/3583133.3590661

1 INTRODUCTION

A focused ultrasound treatment, also known as a focused ultrasound surgery [4, 11, 12], is a technique of sending a focused beam of ultrasound into the tissue and causing a variety of reactions by energy exchange inside the focus. The effects vary significantly depending on the intensity of the waves - whether it is the high-intensity waves (HIFU) to treat a variety of solid malignant tumours in a well-defined volume, such as the pancreas, liver, prostate, or breast; or the neurostimulation techniques carried out by low-intensity waves (LIFU) to treat Parkinson's Disease, obsessive-compulsive disorder or a vagus nerve stimulation for epilepsy and depression [1, 2, 7].

Using ultrasound modelling and knowledge of the properties of the medium, it is possible to predict the ultrasound field inside the tissue after propagating through it, and thus account for subject-specific dose and targeting variations [5, 10]. However, existing models based on conventional numerical techniques typically take tens of minutes to several hours to complete due to the large size of the computational domain, in some cases generating models requiring tens of thousands of iterations to solve [8, 10]. This makes them too slow to be used for online calculations and corrections, i.e., while the subject is undergoing therapy.

The recent renaissance of machine learning technologies could provide a solution to this problem, as a recently published article [10] presented a Physics Informed Neural Net to predict Acoustic Propagation through the human skull. While the utilized UNet is reasonably small, multiple redundant parts are present within the design. Furthermore, the authors themselves suggested more experimentation was needed with the architecture to explore the effects of different parts of the design.

To use this net in an ultrasound treatment plan optimization loop, precision and delivery speed are of the highest importance. In this spirit, we attempt to optimize the architecture, while preserving or increasing the precision. With the emergence of modern Neural Architecture Search (NAS) methods, using evolutionary approaches managing to outperform hand-designed architectures [6, 9, 13], we employ genetic programming techniques to carry out our experiments.

2 HELMNET

The solution, we are attempting to improve, was presented by Stanziola et. al. in [10] and its aim is solving the 2D version of the Helmholtz equation. In this implementation, the boundary condition is satisfied by the use of a perfectly matched layer [3]. The

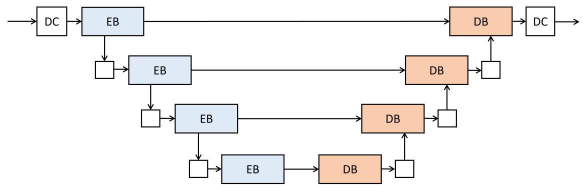

Figure 1: Scheme of Helmnet - a 4 level deep UNet proposed by the authors of [10]. Four dimension layers, each encoding at different spatial dimensions. Every layer consists of an Encoding and a Decoding Block, with a skip connection between them. The network is lightweight, with only 8 channels per convolution block at every scale and a total of $47k$ trainable parameters.

training is guided using a physics-based loss, formed by the residual of the Helmholtz equation - this allows the solution to avoid labelled training data and, in turn, alleviates the need for supervised learning with large preexisting datasets. One distinctive feature of Helmnet is the utilization of a replay buffer, which enables the model to be trained by unrolling for a large number of iterations.

3 SETUP

The entire net is encoded into a genotype as a repeated sequence of integers. Each sequence contains genes for the memory operation, encoder block operation, encoder block connections, decoder block operations and decoder block connections. This sequence is then repeated for every resolution layer of the UNet.

3.1 Crowds

In addition to optimizing the precision of the existing net, one of the goals of this work is also to explore how different resolution layers perform and interact. Each resolution layer is assigned a bit, expressing whether or not this resolution layer is enabled. For this purpose, a crowding idea is introduced to separate solutions based on the enabled and disabled resolution layers. Given 4 resolution layers, the 4 bit binary combination gives us the crowd number for each individual.

To explore both pure precision optimization and crowds exploration, two separate experiments were run:

- **Crowd Self-Adaptation Experiment** - the main purpose of this experiment was to optimize the precision through exploitation. Here, the evolution process itself was given a chance to change the crowd of any solution to a different one.
- **Heuristics Assisted Experiment** - this experiment's main purpose is to explore different crowds. As such. we start with an overall much bigger population and a uniform spread of individuals between possible crowds. With each generation, a step of cooperating Simulated Annealing is executed on the ratio of the population space occupied by each crowd, taking some members of a badly-performing crowd and assigning their spots inside the population to better-performing ones.

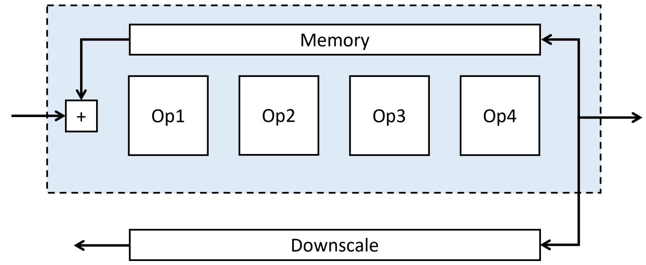

Figure 2: Base encoder block with memory (EB) phenotype.

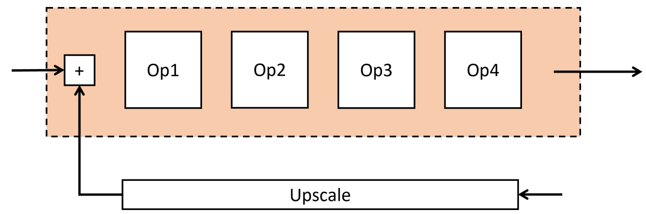

Figure 3: Base decoder block (DB) phenotype.

3.2 Encoder and Decoder Blocks

Each encoder block (*EB*) and decoder block (*DB*) is re-imagined for the purpose of GP as a sequence of 4 arbitrary connected neurons, with *EB* also containing a memory module, see Figs. 2, 3. Each *EB* and *DB* can be evolved to contain some of the following operations:

- 3×3, 5×5 or 7×7 Convolution. All with the option to skip an activation function.
- Linear layer, to allow for intermediary result scaling.
- Identity operation with no trainable parameters, effectively disabling the neuron.

Furthermore:

- The memory neuron inside each encoder block starts as a double convolution operation, but can also be disabled by the evolution.
- Downscaling operation is evolved globally and will remain the same for the entire net. The function is picked from a set holding Average Pooling, Max Pooling or Convolutional downscaling.
- Activation function is also selected globally for the entire network. The function set consists of (P)ReLU, GeLU, Tahn or Mish.

3.3 Neuron Connections

Connections between neurons inside each block are also subjected to the evolution process. The connections are encoded as binary strings, with each bit representing an input from the previous neuron, starting with the input to the block itself. With this approach, we can create a simple linear progression throughout the layer, as well as a skip connection, such as the ones known from ResNets. An example of a connection string can be seen in Figure 4.

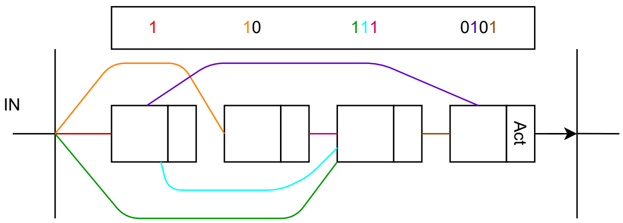

Figure 4: An example of a connection binary string and its resulting phenotype inside an evolved block.

4 RESULTS AND CONCLUSSION

With the setup of both experiments outlined, this section presents the results. The original HelmNet solution was encoded and used as a seed solution for the first generations of our experiments. Both experiments were repeated 5 times, and most figures presented are showing the aggregate results of those runs.

Each candidate net was trained for 20 epochs on the entire training set outlined previously, performing 250 iteration per sample. The performance on the validation set was then considered as the fitness value of the candidate.

4.1 Best Solution

The Self-Adaptation experiment managed to find a better solution than the original (see Fig. 5) - when we took this candidate and trained it as the original solution was trained, the resulting net showed approximately 1 order of magnitude better RMSE (10^{-4} vs 10^{-5}) than the original UNet. However, our solution is using 12, 000 more parameters. Figure 6 shows the parameters spread of the best individuals taken from the exploitative Self-Adaptation runs. We can see that evolution tends to increase the number of parameters to produce individuals with better fitness.

Figure 6 shows the parameters spread of the best individuals taken from the exploitative Self-Adaptation runs. We can see that evolution tends to increase the number of parameters to produce individuals with better fitness. Figure 8 shows the comparison between our evolution-optimized solution and the original Helmnet time to solution. It is clear that the increase in precision comes at the cost of increased inference time and considering that the speed was not a part of the fitness evaluation, it is to be expected. Undeniably, the Helmnet is faster, however, if the requirements for the focused ultrasound planning calls for real-time live updates and navigation, neither of the solutions is fast enough. On the other hand, if some planning time can be afforded, the difference in inference time is not meaningful enough to cause issues. Admittedly, both solutions will need to transition into a 3D domain before a final decision can be made.

4.2 Used Blocks

The crowding approach allowed us to take a look at how disabling different resolution layers influence the quality of a generated individual. Measurably, crowd 7 - the original design of UNet, with all resolution layers enabled - generated the best candidate solutions. As suspected, this crowd is understandably using the most parameters for its candidates on average, indicating once again that an

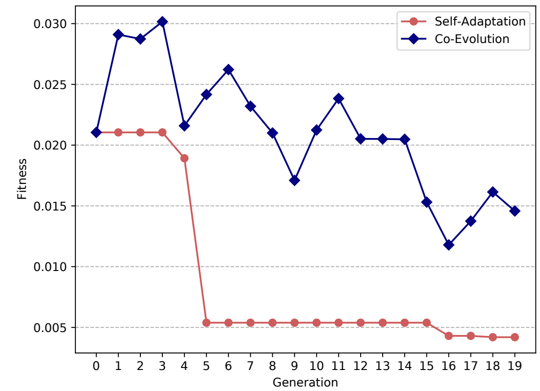

Figure 5: Plots showing the progression of the best individuals found by the crowd self-adaptation and the heuristics-guided evolution experiments.

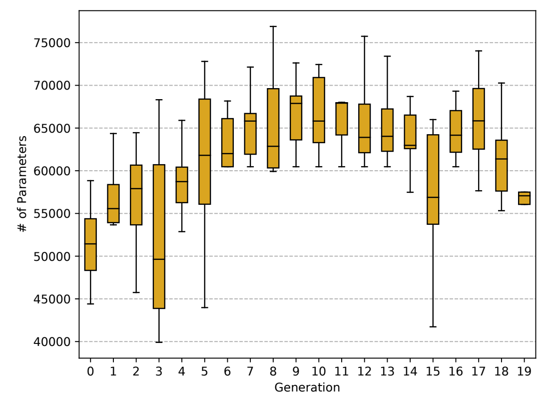

Figure 6: Plots showing the parameters progression of the best individuals during the best crowd self-adaptation experiment.

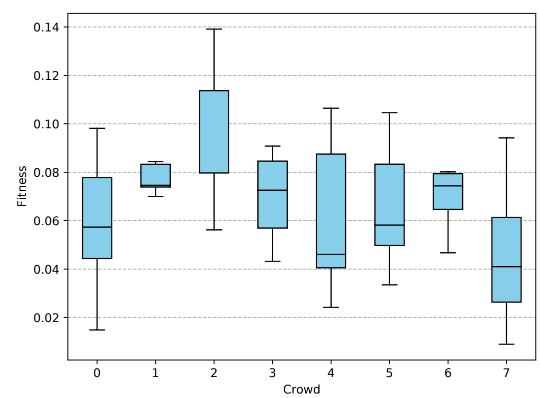

Figure 7: Fitness values of individual crowds, taken as an aggregate from the last generations of the heuristics-guided evolution experiments.

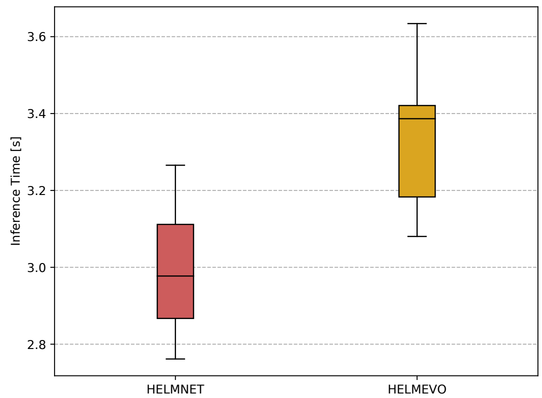

Figure 8: Comparison of inference times of the original Helmnet and the evolution optimized Helmevo.

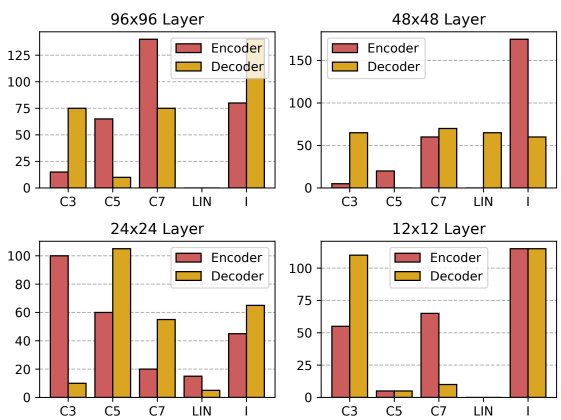

Figure 9: Histograms of blocks used inside the encoding and decoding stages during the last generations, separated by resolution layers of the UNet. Cx stands for convolution with the kernel size of x, LIN is a linear scaling block and I is an identity block. Identity blocks are never followed by an activation function, and could therefore be considered disabled.

increase in parameters seems to be a good way forward if we wish to increase the accuracy of the network further.

Furthermore, figure 9 shows a very interesting spread of convolution window sizes, with the encoder blocks preferring different kernel sizes than the decoders. Additionally, the search algorithm disabled almost the entirety of some encoding or decoding blocks at different resolution layers, thus creating structures usually found in ResNets.

4.3 Memory Blocks and Future Research

While the sample size is too small to make resolute statements, the biggest jump in fitness occurred with the removal of a memory block at the main resolution layer in almost all experiments and this trait is present in all the best solutions. Disabling this block corresponds with the big fitness drops in figure 5. We believe this pattern is worth investigating in the future.

ACKNOWLEDGMENTS

This work was supported by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90140). This work was also supported by the Czech Science Foundation project 21-13001S and by Brno University of Technology project IGA FIT-S-23-8141.

REFERENCES

[1] Agessandro Abrahao, Ying Meng, Maheleth Llinas, Yuexi Huang, Clement Hamani, Todd Mainprize, Isabelle Aubert, Chinthaka Heyn, Sandra E Black, Kullervo Hynynen, Nir Lipsman, and Lorne Zinman. 2019. First-in-human trial of blood-brain barrier opening in amyotrophic lateral sclerosis using MR-guided focused ultrasound. *Nat. Commun.* 10, 1 (Sept. 2019), 4373.

[2] Hongchae Baek, Ki Joo Pahk, and Hyungmin Kim. 2017. A review of low-intensity focused ultrasound for neuromodulation. *Biomedical Engineering Letters* 7, 2 (2017), 135–142. https://doi.org/10.1007/s13534-016-0007-y

[3] Jean-Pierre Berenger. 1994. A perfectly matched layer for the absorption of electromagnetic waves. *J. Comput. Phys.* 114, 2 (1994), 185 – 200. https://doi.org/10.1006/jcph.1994.1159

[4] Marta Cudova, Bradley E. Treeby, and Jiri Jaros. 2018. Design of HIFU Treatment Plans Using an Evolutionary Strategy. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (Kyoto, Japan) *(GECCO '18)*. Association for Computing Machinery, New York, NY, USA, 1568–1575. https://doi.org/10.1145/3205651.3208268

[5] Wonhye Lee, Hyun-Chul Kim, Yujin Jung, Yong An Chung, In-Uk Song, Jong-Hwan Lee, and Seung-Schik Yoo. 2016. Transcranial focused ultrasound stimulation of human primary visual cortex. *Sci. Rep.* 6, 1 (Sept. 2016), 34026.

[6] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. 2018. Regularized Evolution for Image Classifier Architecture Search. *CoRR* abs/1802.01548 (2018). arXiv:1802.01548 http://arxiv.org/abs/1802.01548

[7] Ehsan Rezayat and Iman Toostani. 2016. Review Paper: A Review on Brain Stimulation Using Low Intensity Focused Ultrasound. *Basic and Clinical Neuroscience Journal* 7 (07 2016). https://doi.org/10.15412/J.BCN.03070303

[8] James L. B. Robertson, Ben T. Cox, Jiří Jaroš, and Bradley E. Treeby. 2017. Accurate simulation of transcranial ultrasound propagation for ultrasonic neuromodulation and stimulation. *The Journal of the Acoustical Society of America* 141, 3 (mar 2017), 1726–1738. https://doi.org/10.1121/1.4976339

[9] Kenneth Stanley, Jeff Clune, Joel Lehman, and Risto Miikkulainen. 2019. Designing neural networks through neuroevolution. *Nature Machine Intelligence* 1 (01 2019). https://doi.org/10.1038/s42256-018-0006-z

[10] Antonio Stanziola, Simon R. Arridge, Ben T. Cox, and Bradley E. Treeby. 2021. A Helmholtz equation solver using unsupervised learning: Application to transcranial ultrasound. *J. Comput. Phys.* 441 (sep 2021), 110430. https://doi.org/10.1016/j.jcp.2021.110430

[11] Visa Suomi, Jiri Jaros, et al. 2018. Full Modeling of High-Intensity Focused Ultrasound and Thermal Heating in the Kidney Using Realistic Patient Models. *IEEE Transactions on Biomedical Engineering* PP (09 2018), 1–1. https://doi.org/10.1109/TBME.2018.2870064

[12] Yufeng Zhou. 2011. High intensity focused ultrasound in clinical tumor ablation. *World journal of clinical oncology* 2 (01 2011), 8–27. https://doi.org/10.5306/wjco.v2.i1.8

[13] Hamit Taner Ünal and Fatih Basciftci. 2021. Evolutionary design of neural network architectures: a review of three decades of research. *Artificial Intelligence Review* (07 2021). https://doi.org/10.1007/s10462-021-10049-5