

Using deep learning models to accelerate the design of soft robots with genetic algorithms

Loïc Mosser¹, Laurent Barbé¹, Lennart Rubbert¹ and Pierre Renaud¹

ICube, Université de Strasbourg - CNRS - INSA Strasbourg, France

mosser1@unistra.fr

Abstract. The motion of soft robots is intrinsically linked to their shape. Design of soft robots is then still a challenge, with a very large design space to explore in terms of possible shapes. Generative methods can be of interest, but they require intensive use of robots motion prediction. We assess the interest of using deep learning models to accelerate the synthesis. The case of pneumatically-actuated structures is considered. We show first that a Resnet model can accurately describe the structure motion after learning on a dataset based on finite element simulations. Second, we show that the model accuracy can be maintained during a synthesis, outside the initial dataset, using transfer learning.

Keywords: Soft Robot, Resnet, Transfer learning, Genetic Algorithm

1 Introduction

Soft robots are structures made of flexible material, for which motion is obtained by various means such as cable or pneumatic actuation [1]. The latter is largely considered with so-called soft pneumatic actuators (SPA) [2], the elementary components of a soft robot. Pneumatic chambers are then distributed in the SPA body and they allow deformation of the structure. Design methods for these actuators still need to be developed [3, 4]. With the freedom of shape allowed by additive manufacturing techniques [5], the design space is in particular growing dramatically and efficient design methods are yet to be proposed.

The gold standard to assess a SPA motion is to use finite element analysis (FEA) [6, 7]. Designing SPA using FEA and evolutionary algorithms such as genetic algorithms could be interesting to get a generic design method [8], but the computational cost of FEA limits the feasibility. In the literature, deep learning (DL) models have been proven to be relevant for the prediction of soft structure deformation under various loadings [9–11]. During a SPA design, the problem of prediction using DL is however quite different: the model has to estimate the displacements of a structure, while the latter is modified by the design process. At the same time, the loading is also modified as it is related to the pressurization of the SPA.

In this paper, we thus assess the adequacy for such use of a CNN model. First, we describe the SPAs to be modeled. A simplified design problem is considered to study the limitations of a CNN model when the training dataset is of limited size compared to the full design space. We then study the impact of bias in the training dataset in a situation where the design space can be exhaustively explored. Finally, we assess the model capacity to remain accurate while being used in other areas of the large design space, using transfer learning (TL).

2 SPA definitions and CNN construction

2.1 SPA under consideration

Our assessment is based on the design of SPAs defined by a discrete distribution of soft material through a $25 \times 25 \times 75$ matrix of 0.4 mm voxels, to be compatible with experimental assessment in the future. The prismatic outer shape is considered fixed, considering it is imposed by size requirements. The air supply is fixed at the center of the bottom surface. A functional SPA is composed of a hollow structure connected to the air supply within the prismatic outer shape. Examples of three-dimensional SPAs generated randomly accordingly are given in Fig. 1-a.




Fig. 1: In (a), examples of SPA randomly generated for the learning dataset (blue : void, red : filled), in (b) the points of interest used to define the SPA behavior with, from left to right, 1, 4, 12 and 20 points.

Our goal is to estimate the 3D displacements at n points of interest, located on the external border of the SPA. The output of the CNN model is thus a vector of size $3n$. In the following, we estimate the prediction accuracy in 4 situations, *i.e.* with $n=1, 4, 12, 20$ as represented in Fig. 1-b. As a reference and for CNN training, FEA is performed using Comsol. The SPA body is considered as composed of a soft linear elastic material ($E = 2$ MPa, $\nu = 0.3$).




Fig. 2: Schematic representation of the reduced problem with the point whose displacement is being tracked (red point), the outer edge (crossed-out voxels), the pneumatic input (violet point), the fixed surface at the bottom of the design (blue surface) and the pneumatic chamber (yellow) with full (grey) and empty (white) voxels.

2.2 Simplified design of SPA

Implementing a deep learning model for a large design problem makes it difficult to identify the influence of biases during learning. Here, with a structure composed of $25 \times 25 \times 75$ voxels, there are 10^{14096} possible designs of SPA. Thus, a simplified problem of SPA design is considered for initial evaluation of CNN behavior. The design is 2D, and described by 5×5 matrix. SPA material is the same. There are then 1,338,341 possible SPAs, all of which can be simulated using FEA: the whole design space can therefore be explored. One example is depicted in Fig. 2.

2.3 CNN model construction

In [12], a Resnet model was successfully used to estimate the behavior of composite structures under static loading. The composite structure is then defined as an array, with 2 possible materials for each cell. The model is then employed to predict the structure mechanical behavior, e.g. displacement at a specific point. This is closely related to our problem, so we decided to assess a similar approach. The implementation and training of the Resnet is carried out using Tensorflow and Keras on the same computer.

To determine the number of parameters required to predict the mechanical behavior of SPA, the number of residual convolution layers (RCL) is increased gradually, until an RMSE of the order of 10^{-2} mm is reached. RCL are defined using full pre-activation (Fig. 5) as it has shown great generalization performances. The Resnet takes as input the matter distribution matrix defined as a matrix boolean values and outputs the $3n$ vector representing the displacements of n points defined in Fig. 1-b and expressed in millimeters.

3 Initial assessment of CNN model prediction accuracy

3.1 Creation of the dataset

For this initial assessment, the simplified design is being used. There is one point of interest ($n = 1$) on the outer edge of the SPA (Fig. 2, red point). The SPA structure is defined (Fig. 2) by a 5×5 matrix. The dataset is generated by considering 3 rules: 1) the

outer strip of material needs to remain present to keep the SPA sealed ;2) the pneumatic input is set at the center of the lower surface of the design (Fig. 2, violet point) ; 3) the lower surface is attached to the base, so it has no displacement. During the generation of designs, all voxels composing the inner structure of the SPA, that is submitted to the internal pressure, must be connected to the pneumatic supply, as shown in Fig 2. Designs that do not follow this rule are not considered.

With a reduced problem formulated in this way, we use a greedy algorithm to determine the set of possible SPAs. The dataset contains about 1.4×10^6 SPAs. From this dataset, we analyze the training results of a Resnet network. The dataset is first ranked using the value of displacement at the point of interest. This creates on purpose a bias in the training data, which impact is analysed in the following.

3.2 CNN model construction

The constructed Resnet architecture is composed of an input convolution layer (CL), 11 residual convolution layers, a last CL followed by 2 dense layers of 128 neurons each. Each CL have $64 \ 3 \times 3$ kernels. RCL are defined using full pre-activation as it has shown great generalization performances. The Resnet takes as input the matter distribution matrix defined as a matrix of 5×5 boolean values and outputs the displacement of the monitored displacement along \vec{y} . In order to facilitate the representation of results, we limit ourselves to the study of vertical displacement coordinates.

An initial training of the Resnet is done using 100,000 SPA that randomly picked in the dataset. 90 % are used as a training set and 10 % as a validation set. RMSProp is being used for the training as it here provides better performances than Adam. The remaining individuals are used to test the model performance. The training is done through 300 epochs with a learning rate of 10^{-4} and early-stopping. The performance of this initial training, evaluated on the test set, is shown in Fig. 3. The results show that the Resnet is able to give an estimate of displacement with an accuracy of less than $10 \ \mu\text{m}$ for the entire dataset. The R^2 coefficient is then of 0.99.




Fig. 3: Learning results on the whole dataset without the introduction of bias with training data (yellow) and test data (black).

3.3 CNN behavior

We seek to quantify the loss of performance of the network’s estimation when its learning is conducted on a non-representative set of all SPAs. To this end, the complete dataset obtained after the initial sorting operation is segmented in 2 domains, using a threshold on the achievable displacement at the point of interest. The domain with the smallest displacements is used to build the training dataset and the test dataset. The other domain is used as the bias control dataset. 4 situations are considered, with 20%, 40%, 60% or 80% of designs with smallest displacements. The training dataset contains 100,000 SPAs for each assessment.

The training results are available on Fig. 4 and on Table 1.

considered displacement	training RMSE	training R^2	Test RMSE	Test R^2	bias control RMSE	bias control R^2
80% selection	$9.0 \cdot 10^{-3}$	0.999	$9.0 \cdot 10^{-3}$	0.999	$1.14 \cdot 10^{-1}$	0.350
60% selection	$8.0 \cdot 10^{-3}$	0.999	$8.0 \cdot 10^{-3}$	0.999	$2.54 \cdot 10^{-1}$	-1.745
40% selection	$7.0 \cdot 10^{-3}$	0.999	$7.0 \cdot 10^{-3}$	0.999	$3.54 \cdot 10^{-1}$	-3.687
20% selection	$6.0 \cdot 10^{-3}$	0.998	$6.0 \cdot 10^{-3}$	0.998	$8.43 \cdot 10^{-1}$	-38.752

Table 1: Presentation of training results on 100,000 data points taken from a set of chambers with displacement associated with the best 80, 60, 40 and 20 percent of the data in the dataset.

When we evaluate the network’s performance on the training domain, we find that it does not differ from the initial training, for both training (Fig. 4, yellow) and test (Fig. 4, black) datasets. The CNN is then able to provide a satisfactory estimate of vertical displacement at the monitored point. The CNN performance is, however, insufficient on the bias control dataset (Fig. 4, blue). We note that this estimator loses precision when the maximum displacement of SPAs considered in the training is diminished.

The bias introduced in this learning process allows us to understand the possible issues during the design of SPAs. During the initial generation of SPAs, it is difficult to propose a dataset that offers an exhaustive representation of all possible displacements. If the initial SPA generation process is based on unguided random generation, we can anticipate that a bias will be introduced, that will impact the model built by learning. We can also anticipate that other biases, such as a bias linked to the size of the pneumatic chambers, will impact our initial dataset. So, rather than focusing on the initial generation method, in the following we investigate the efficiency of correcting this bias through learning transfer steps.

4 CNN for SPA design with transfer learning

4.1 CNN model learning

We now focus on the SPA under consideration initially, as described in 2.1. The constructed Resnet architecture (Fig. 5) is composed of an input convolution layer (CL), 16 residual convolution layers, a last CL followed by an average pooling layer and 2 dense layers of 128 neurons each. Each CL have $64 \ 3 \times 3 \times 3$ kernels. RCL are defined using full pre-activation (Fig. 5) as it has shown great generalization performances. The Resnet takes as input the matter distribution matrix defined as a matrix of $25 \times 25 \times 75$ boolean

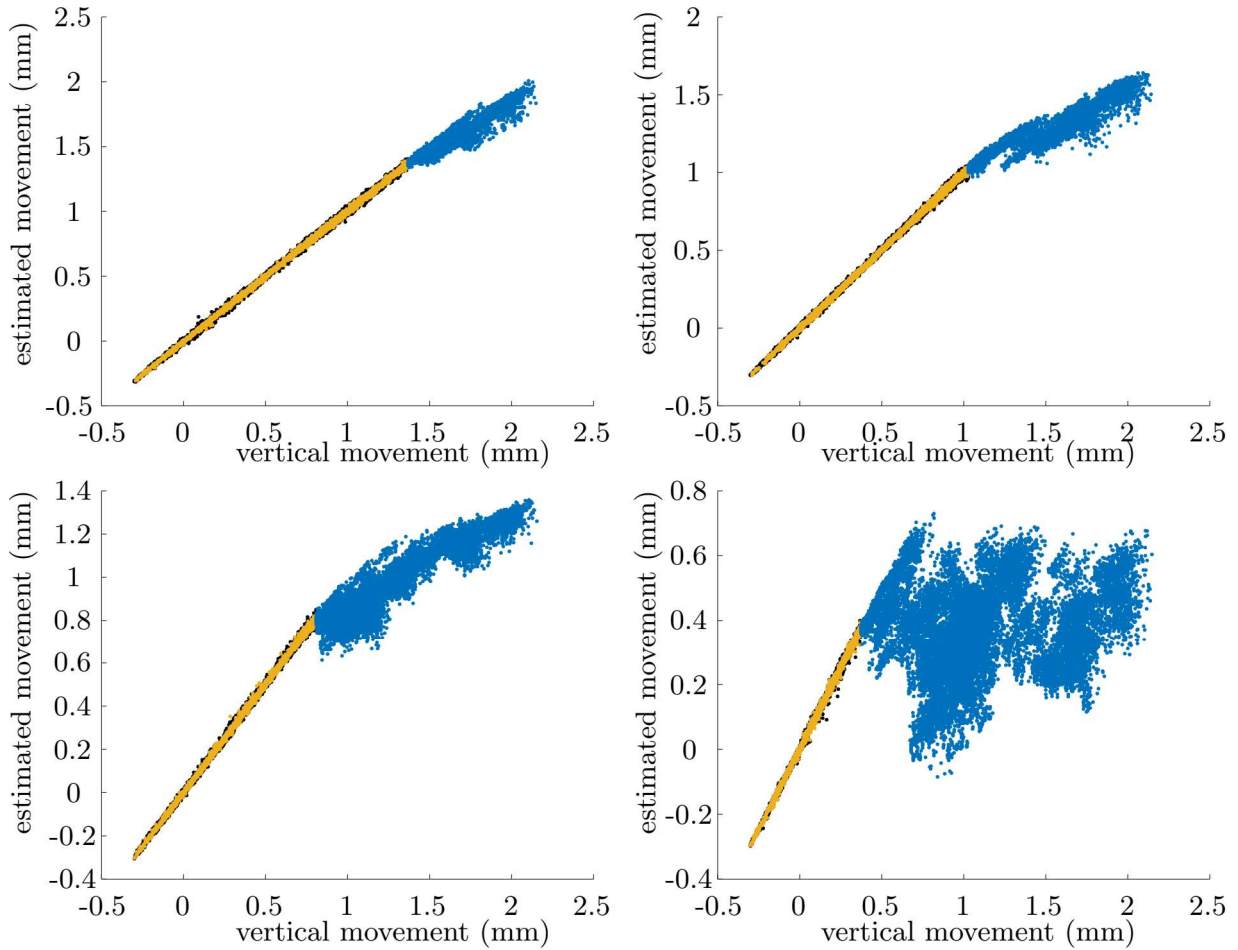


Fig. 4: Plot of estimated displacements versus simulated displacements for the different drives performed on the lowest 80% (top-left), 60% (top-right), 40% (bottom-left) and 20% (bottom-right) displacements.

values and outputs the $3n$ vector representing the displacements of n points defined in Fig. 1-b and expressed in millimeters.

For its training, an initial dataset of 100,000 designs is first randomly generated. The simulation time is about 8 days using 1 PC (Intel i9-10900KF, 64 GB of RAM, NVIDIA RTX-3090).

The training of the Resnet is done using 60,000 SPA randomly drawn in the initial dataset (90 % as training set and 10 % as validation set) using RMSProp as it here provides better performances than Adam. The remaining 40,000 SPA are used to test the model performance. The training is done through 20 epochs with a learning rate of 10^{-4} and early-stopping. The performance of this initial training, evaluated on the test set, is shown in the table 2. The results show that, with initial training, the Resnet is able to give an estimate of displacement with an accuracy of less than $10 \mu\text{m}$ for the 4 situations under consideration. In addition, one estimation of SPA performance requires a computational time of less than 1 ms, compared with 30 seconds for FEA.

4.2 CNN accuracy with transfer learning

During the SPA design, the design space is being explored to gradually tend to obtain larger displacements, meaning better performances out of the initial dataset domain.

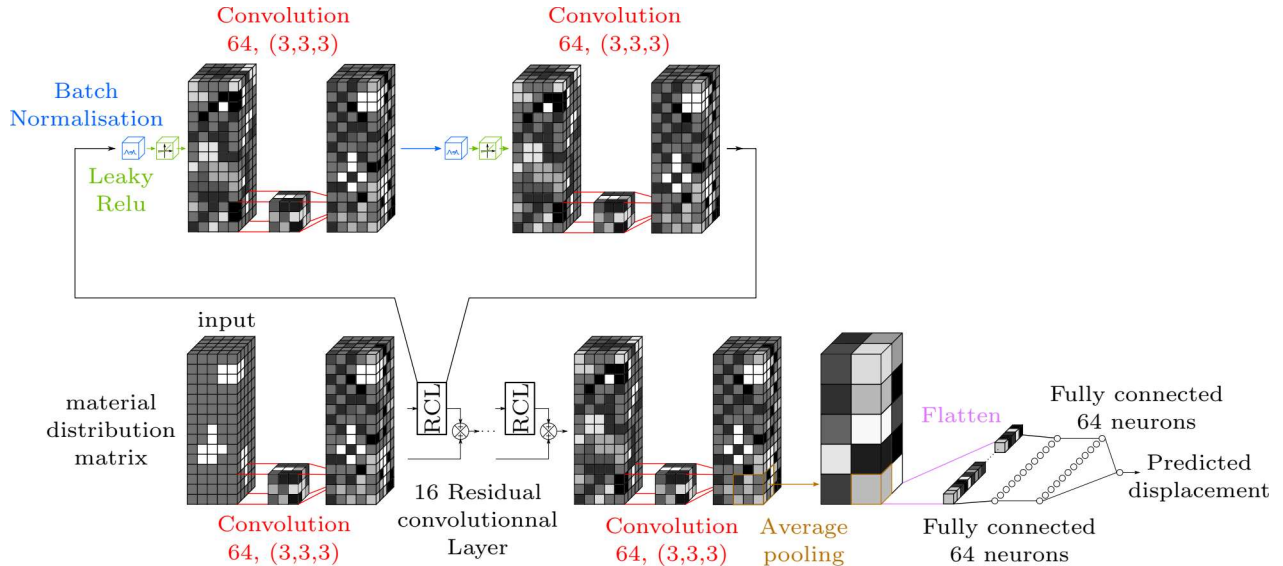


Fig. 5: Resnet representation with its input as the material distribution matrix and its output as a dense layer of $3n$ neurons.

When the Resnet is being used in a domain out of the initial dataset, prediction accuracy may decrease significantly as we showed earlier on the simplified problem. To assess this, a specific dataset of 830 SPA designs with greater displacements and far from the initial dataset has been generated specifically (Fig. 6, blue and orange dots). We define our test dataset with 800 SPA. The initial performance of the CNN is indicated in the second column of Table 2. For example, precision is almost decreased by a factor 20 when 1 point of interest is considered, which means a strong loss of accuracy.

A transfer learning (TL) step is then considered. It is achieved through 10 epochs with a learning rate of 10^{-5} with 3000 SPA. 1500 SPA are randomly drawn in the training dataset of 60,000 SPA and the 30 extra SPA geometry are duplicated 50 times to have an equal representation in the TL dataset. Performances of the test dataset are indicated in the last column of Table 2. It highlights the relevance of TL in design space exploration out of the initial domain. The loss of precision is then only by a factor 2, which could be acceptable as it is only used to identify best design subspaces.

Considered output	Metrics	Initial training	Before TL	After TL
1 point	R^2 RMSE	0.99 $7.7 \cdot 10^{-3}$	0.90 $1.4 \cdot 10^{-1}$	0.96 $1.6 \cdot 10^{-2}$
4 points	R^2 RMSE	0.99 $9.0 \cdot 10^{-3}$	0.92 $1.8 \cdot 10^{-1}$	0.94 $1.8 \cdot 10^{-2}$
12 points	R^2 RMSE	0.99 $6.8 \cdot 10^{-3}$	0.69 $1.3 \cdot 10^{-1}$	0.92 $1.3 \cdot 10^{-2}$
20 points	R^2 RMSE	0.98 $7.3 \cdot 10^{-3}$	0.81 $6.5 \cdot 10^{-2}$	0.89 $1.1 \cdot 10^{-2}$

Table 2: Resnet performance as a function of the number of points of interest, the use after initial training, without and with transfer learning.




Fig. 6: Displacement of 1 point of interest of the SPA. Black/Blue/Orange crosses represent the initial dataset/the test dataset for TL/ the train dataset for TL

5 Conclusion

We investigated the use of a Resnet model to estimate the behavior of SPA during its design. First, it reduces the computation time from 30 s to less than 1 ms in comparison with FEA. The Resnet accuracy on a randomly generated population of designs is in the order of 10 μm . This performance is obtained when focusing on 1 or multiple, here up to 20, points of interest. We have identified that learning the CNN on an initial population of SPAs may be subject to certain biases. We then proposed to carry out a transfer learning step on a reduced number of new data to correct the estimate proposed by the CNN. The capacity to use transfer learning to improve the Resnet accuracy when it is used during the synthesis, far from the initial dataset, was also evaluated. The loss of precision can be reduced by a factor of 20 to 2 thanks to the TL.

The overall impact of computational time reduction obviously depends on the evolutionary process, which is another aspect to investigate. A perspective will then to include more advanced material models, to refine the design process, for instance including material non-linearities that can be present with very soft materials.

References

1. Rus, D., Tolley, M.T.: Design, fabrication and control of soft robots. *Nature* **521**(7553) (may 2015) 467–475
2. Walker, J., Zidek, T., Harbel, C., Yoon, S., Strickland, F.S., Kumar, S., Shin, M.: Soft Robotics: A Review of Recent Developments of Pneumatic Soft Actuators. *Actuators* **9**(1) (jan 2020) 3
3. Chen, F., Wang, M.Y.: Design Optimization of Soft Robots: A Review of the State of the Art. *IEEE Robotics & Automation Magazine* **27**(4) (dec 2020) 27–43
4. Lee, C., Kim, M., Kim, Y.J., Hong, N., Ryu, S., Kim, H.J., Kim, S.: Soft robot review. *International Journal of Control, Automation and Systems* **15**(1) (feb 2017) 3–15
5. Stano, G., Percoco, G.: Additive manufacturing aimed to soft robots fabrication: A review. *Extreme Mechanics Letters* **42** (jan 2021) 101079

6. Pagoli, A.: Innovative Design of a Soft Robotic Gripper for In-hand Amir Pagoli. PhD thesis, Université Clermont Auvergne (2021)
7. Coevoet, E., Morales-Bieze, T., Largilliere, F., Zhang, Z., Thieffry, M., Sanz-Lopez, M., Carrez, B., Marchal, D., Goury, O., Dequidt, J., Duriez, C.: Software toolkit for modeling, simulation, and control of soft robots. *Advanced Robotics* **31**(22) (nov 2017) 1208–1224
8. Hiller, J., Lipson, H.: Automatic Design and Manufacture of Soft Robots. *IEEE Transactions on Robotics* **28**(2) (apr 2012) 457–466
9. Odot, A., Haferssas, R., Cotin, S.: DeepPhysics: a physics aware deep learning framework for real-time simulation. *International Journal for Numerical Methods in Engineering* **123**(10) (sep 2021) 2381–2398
10. Deshpande, S., Lengiewicz, J., Bordas, S.P.A.: Probabilistic Deep Learning for Real-Time Large Deformation Simulations. *Computer Methods in Applied Mechanics and Engineering* **398** (nov 2021) 115307
11. Mendizabal, A., Márquez-Neila, P., Cotin, S.: Simulation of hyperelastic materials in real-time using Deep Learning. *Medical Image Analysis* **59** (apr 2019) 101569
12. Kim, Y., Kim, Y., Yang, C., Park, K., Gu, G.X., Ryu, S.: Deep learning framework for material design space exploration using active transfer learning and data augmentation. *npj Computational Materials* **7**(1) (sep 2021) 140