

Aarhus School of Architecture // Design School Kolding // Royal Danish Academy

Localised and Learnt Applications of Machine Learning for Robotic Incremental Sheet Forming

Nicholas, Paul; Zwierzycki, Mateusz; Ramsgaard Thomsen, Mette

Published in:
Humanizing Digital Reality

DOI:
[10.1007/978-981-10-6611-5_32](https://doi.org/10.1007/978-981-10-6611-5_32)

Publication date:
2017

Document Version:
Publisher's PDF, also known as Version of record

Document License:
Unspecified

[Link to publication](#)

Citation for published version (APA):

Nicholas, P., Zwierzycki, M., & Ramsgaard Thomsen, M. (2017). Localised and Learnt Applications of Machine Learning for Robotic Incremental Sheet Forming. In *Humanizing Digital Reality: Design Modelling Symposium Paris 2017* (pp. 373-382). Springer. https://doi.org/10.1007/978-981-10-6611-5_32

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Localised and Learnt Applications of Machine Learning for Robotic Incremental Sheet Forming

Mateusz Zwierzycki, Paul Nicholas^(✉), and Mette Ramsgaard Thomsen

Centre for IT and Architecture, Royal Danish Academy of Fine Art, School of Architecture, Copenhagen, Denmark
paul.nicholas@kadm.dk

Abstract. While fabrication is becoming a well-established field for architectural robotics, new possibilities for modelling and control situate feedback, modelling methods and adaptation as key concerns. In this paper we detail two methods for implementing adaptation, in the context of Robotic Incremental Sheet Forming (ISF) and exemplified in the fabrication of a bridge structure. The methods we describe compensate for springback and improve forming tolerance by using localised in-process distance sensing to adapt tool-paths, and by using pre-process supervised machine learning to predict stringback and generate corrected fabrication models.

Keywords: Machine learning · Incremental sheet forming · Robotic fabrication

Introduction

Incremental Sheet Forming (ISF) is a formative fabrication process, in which mechanical forces are applied to a flat material so as to form it into a desired 3D shape. A particular problem associated with ISF is the springback of the metal, and the development of methods to counteract it. In this paper we counteract springback in the fabrication of architectural panels by creating two different material behaviour models—the first based on on-line distance measurement and linear regression, and the second based on a neural network. This paper particularly describes the application, architecture and accuracy improvement of the artificial neural network. The trained model is based on 102-dimensional input data, which the artificial neural network learns and interprets, so that it is able to predict the shape of the panel after the fabrication process. This work is a continuation of the research described in Nicholas. The limitations and problems encountered in the approach presented in that paper are addressed here.

Background

The structure A Bridge Too Far spans 4 m, and is made from 1 mm aluminium panels rigidized using a process of robotic incremental sheet forming (ISF). It was exhibited

together with two further projects—*Inflated Restraint & Lace Wall*—as part of the research exhibition *Complex Modelling* held at Meldahls Smedie, The Royal Danish Academy of Fine Arts (KADK), over the period September–December 2016.

In the ISF process, a rounded head tool is applied from either one or two sides of a flat sheet, and moves over the surface of that sheet to cause localized plastic deformation out of plane (Jeswiet et al. 2005). In a double-sided approach two tools, one working as a forming tool and one as a support, enable forming out of plane in opposing directions, and significant freedom and complexity in the formed geometry. The double-sided setup used to fabricate *A Bridge Too Far* incorporates two industrial robots working on each side of a moment frame that allows for a working area on approximately 1000×1000 mm.

The primary advantage of ISF is to remove the need for complex moulding and dies, which only become economically feasible with large quantities. However, this means that positional accuracy is lowered compared to other approaches, as well as more highly dependent upon a combination of material behaviour and forming parameters. Research into resultant incrementally formed geometries has shown significant deviations from planned geometries (Bambach et al. 2009), which is recognised as a key deterrent from the widespread take up of the process (Jeswiet et al. 2005). Our testing has shown that parameters including the forming velocity, the tool path, the size of the sheet and distance to supports all affect springback of the sheet and the geometric accuracy of its forming. It would be extremely difficult to develop a linear, analytical model for springback control including all of these factors, on account of their complex interdependency.

Machine Learning. As an alternative, machine learning approaches can be considered. Data mining techniques have been previously applied to prediction and minimisation of springback in sheet metal forming, including statistical regression (Behera et al. 2013) and neural networks (Khan et al. 2014). These approaches have combined Finite Element simulation with data mining techniques (Ruffini and Cao 1998). Finite Element simulation can also be combined with point cloud data, obtained via scanning of formed results, with the difference between simulation and scan taken as training data. While Khan et al. (2014) extends the prediction of springback errors to new shapes and the generation of corrective geometries, this is restricted to a simple pyramidal geometry. We go beyond state of the art by applying a deep learning based approach to significantly more complex geometries.

A second trajectory has been the use of dynamic tool path correction, in conjunction with online monitoring, to predict and correct distortion due to springback. While adaption of toolpaths, for example to change tool speed or ‘jump’ zones of material fracture has been proposed (Paniti and Rauschecker 2012), this trajectory has not yet been explored in ISF. A significant limit is that previous limitations on fabrication technology have required that the tool path be specified in advance, rather than as the process develops. For this reason, toolpath correction has been limited to application on simulation results, however recently developed control and communication tools for robotic control—i.e. HAL (2017), Robots (2017), KUKA|prc (2017)—together with a shift to industrial robots, rather than purpose built machines, provide new opportunities for adaptive control as a means to manage forming tolerance.

Methods

Localised Method

A Bridge Too Far (Fig. 1) relies on several hundred point connections being made precisely between upper and lower skins. In the first method, rather than rely on model-based predictions of forming, an in-process measurement-based approach was developed to achieve the required tolerances. This automated process involved measurement of formed geometry at these points, calculation of the local deviation, and automatic selection of an appropriate toolpath for a cone that iteratively corrects these deviations to make a precise connection. A single point laser distance measure mounted to the robot arm instead (Fig. 2) was used to measure, at each cone centre-point, the local deviation between actual formed depth and ideal geometry. Some knowledge of springback could also be gathered, modelled and used during this process.



Fig. 1. A Bridge Too Far exhibited at the research exhibition *Complex Modelling* KADK, 2016



Fig. 2. Single point distance sensor mounted to the robot arm

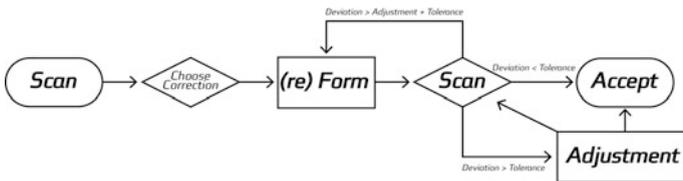


Fig. 3. Workflow diagram of the localised method

To determine the required amount of over-forming for each cone, the relationship between target depth and forming depth was modelled during the entire cone forming process using a linear regression. After forming, the resultant depth was scanned, and this data was used to refine the curve fitting model. This allowed a continued improvement in accuracy across the course of fabrication. Scanning occurred several times in the forming of each cone, and could trigger two automated correction methods. If the formed cone geometry was greater than 5 mm off the target geometry, the cone was reformed. If between 5 and 2 mm off the target geometry, the tip of the cone was extended by 2 mm; tolerances below 2 mm were considered acceptable. This procedure is illustrated as a diagram in Fig. 3.

Learnt Method

Contrary to the Localized method which acts during the fabrication process, the Learnt method is utilizing supervised machine learning to construct a material model a priori the fabrication. Therefore the adaptability competence of the overall workflow is shifted to the earlier design phase. This simplifies the fabrication process (no need for on-the-fly toolpath generation), but requires to collect the material behaviour samples

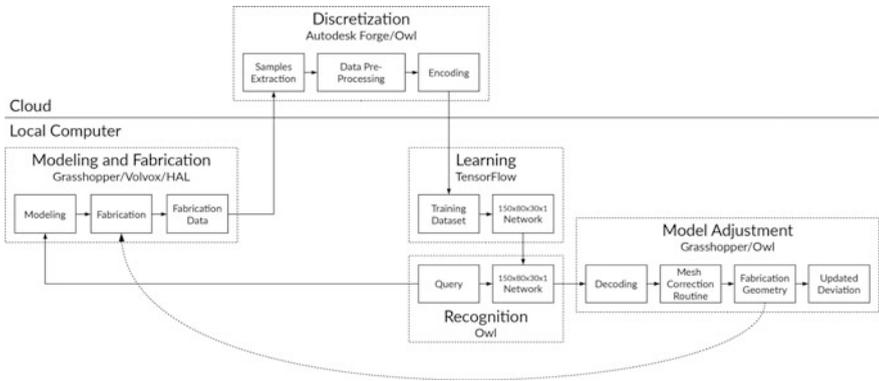


Fig. 4. Workflow diagram of the learnt method

beforehand. While this can be a time consuming venture, there is an advantage of recycling the acquired database for use in other projects or workflows. Another advantage of developing a material model compared with the toolpath adaptation seen in the Localized method, is the decoupling of the model from the fabrication technique.

In the most general terms, the Learning method is a mapping between 102-dimensional input and 1-dimensional output. The accuracy of the map is a subject of investigation, and the means required to achieve a satisfactory results are described hereafter. This section is divided into 4 subsections focusing on the different aspects of the developed workflow and the challenges encountered during the development process: *Acquisition*, *Clean-up*, *Pre-processing* and finally *Learning*. The presented workflow is a continuation and elaboration of the previously published work in Nicholas et al. (2017), which should be seen as a test bench for the methods described in this publication (Fig. 4).

Acquisition. The acquisition phase was the first problem to be addressed, as it influences directly the quality of the constructed material behaviour database. The highly reflective surface of the aluminium panels made the 3d scanning problematic because the laser-based scanner is prone to reflections, resulting with points scattered in space. While it wasn't a matter of concern with the dataset collected for the first experiment (due to its highly prototypical nature), in the presented workflow this problem was addressed to minimise the data loss or noise-induced training inaccuracy.

A chalk-based, water-soluble spray paint was used to effectively (and temporarily) cover the metal surface with a thin layer of a non-reflective coating. The scattering incidents effectively remove the points from the scope of the scan, creating "holes" in the point cloud—making the resulting outliers greatly stand-out and guide the learning process away from the most accurate state. The one issue which wasn't resolved at the hardware level of the acquisition process is the inherent noise of the 3d scanner—which is up to ± 2 mm (Faro Focus3D 2017) (Fig. 5).

Clean-up. The time consuming process of 3d scanning is succeeded with manual process of further point cloud cleaning. With the toolset based on Grasshopper (2017) and Volvox (2017), the clean-up task became more intuitive and efficient, as it was



Fig. 5. The reflective surface of the panel compared with the coated surface

possible to automate the general cropping of the scan data. Removing the points is required to prepare the data for the automatized pre-processing methods.

Pre-processing. In the presented workflow context means preparation of the input Tensors for the learning process. As described above, the input is a 102 Tensor with 100 dimensions to represent the local shape of the surface as a heightfield (of size 5×5 cm), and 2 additional values: distance to the edge of the panel and distance to the edge of the fabrication rig frame. This is a slight difference introduced in the new approach compared with the one from [blinded reference], which was using a 9×9 heightfield + distance to the panel edge value. The significance of this change is further investigated in the evaluation section.

The sample extraction process, gathers up to 3000 samples per one panel. Each sample is mirrored and rotated 180° to effectively obtain 4 samples. While for this workflow the procedure takes only up to a minute for each panel, the reusability of the toolset and its scalability was a matter of concern. The Autodesk Forge platform (2017) was used to parallelize the pre-processing using a cloud-based solution. Two applications and a common library was developed to make the software setup and usage more agile. This enabled to experiment with different extraction parameters, such as the number of points per sample, the size of the sampling frame etc. In total there was over 45,000 samples extracted from 9 panels during the first training session and almost 100,000 for the reversed mapping session (Fig. 6).

Learning. The architecture of the artificial neural network is slightly expanded in comparison with the one used in the previous workflow (Nicholas et al. 2017). With the learning process based on the TensorFlow framework (2017) running on GPU, it was possible to expand the network size to 4 layers. The layers have 150, 80, 30 and 1 neurons, with 102 inputs. A 3-layered network was also a subject of tests, yet it didn't reach a comparable accuracy with the 4-layered one (Fig. 7).

The training process was performed on a single PC equipped with a GPU card, and few sessions were performed to reach the final accuracy. Each session initialized the

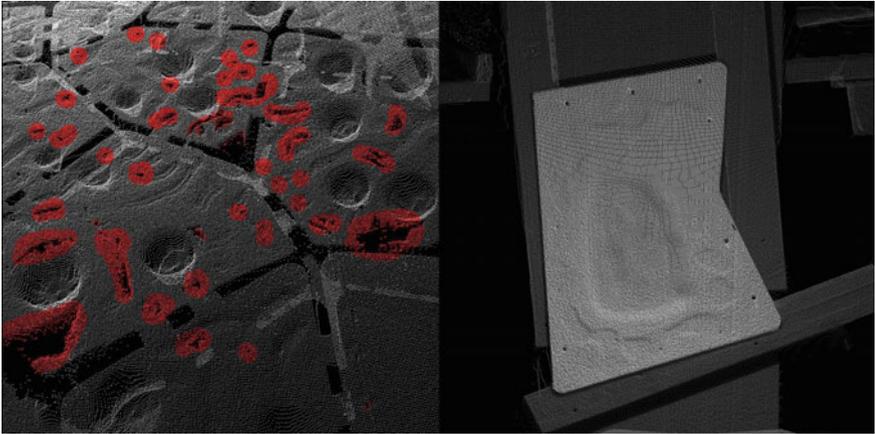


Fig. 6. From *left* an example of a point cloud with missing points (marked in *red*) which is a result of the point scattering problem; The anti-reflective coating reduces the noise and the point scattering to minimum

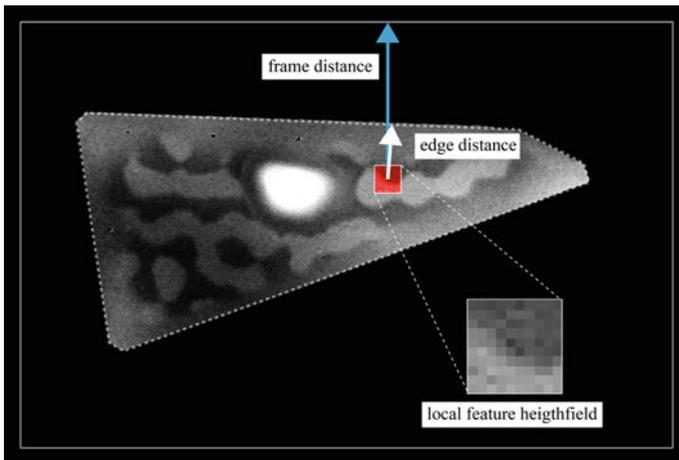


Fig. 7. Data encoded within the 102-dimensional input of each of the training samples

network with randomised values and used different training parameters (learning rate, optimizer type, batch size), taking up to 15 min to train down to the point of optimizer equilibrium.

Results

The evaluation of the neural network prediction error was conducted on $\sim 10,000$ samples from which the output mesh was reconstructed. An overlay of the input mesh (the desired geometry), the result of fabrication and the prediction are shown in Fig. 8.

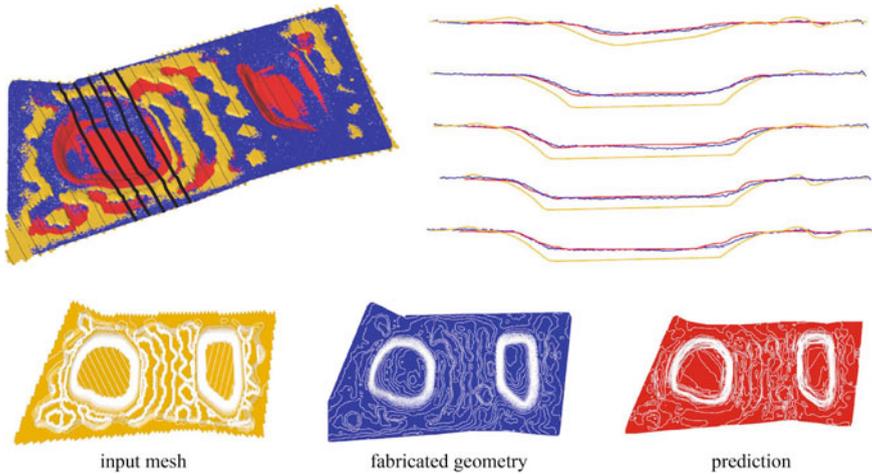


Fig. 8. The prediction closely follows the fabrication results, up to the point where the scanner induced noise (up to ± 2 mm) becomes an obstacle in reaching a higher prediction accuracy. The standard deviation of the prediction error is 1.075 mm

While there are areas with a larger error (of up to 7 mm), the majority of predictions (>70%) are within ± 2 mm deviation (which also incorporates the hardware error of up to ± 2 mm), and 90% of samples are of error less than 3 mm.

Certainly the areas with a large difference in the slope gradient yield predictions with lesser accuracy, as shown in the cross-sections in Fig. 8. Arguably the source of this problem lies in their sparsity which translates directly to the relatively small amount of the representative samples in the training dataset

Reversed prediction. The presented experiments were conveyed with a plan of predicting the fabricated geometry shape based on the input mesh model. Given the network is able to learn that relationship, it is in principle possible to train the network in the opposite direction: predicting the input mesh geometry based on the desired output. This experiment yielded a network which is able to generate the geometry with lesser accuracy. Arguably this accuracy drop is caused by the amount of the hardware-induced noise in the training set inputs. A comparison of the results obtained from the 2 networks (forward and reversed prediction) is compiled in Fig. 9.

Future Work

While for the presented fabrication method the prediction accuracy is satisfactory, there is room for improvement. An identified obstacle to achieving better results is the 3d scanning hardware used for the data acquisition. This problem will be addressed in the future by using a digitizer or a blue light laser scanner rather than a regular mid-to-high range 3d scanner.

The greatest prediction deviation in both of the models (“forward” and “backward” predicting) occurs in the areas with a steep slope. The reason of that error is considered

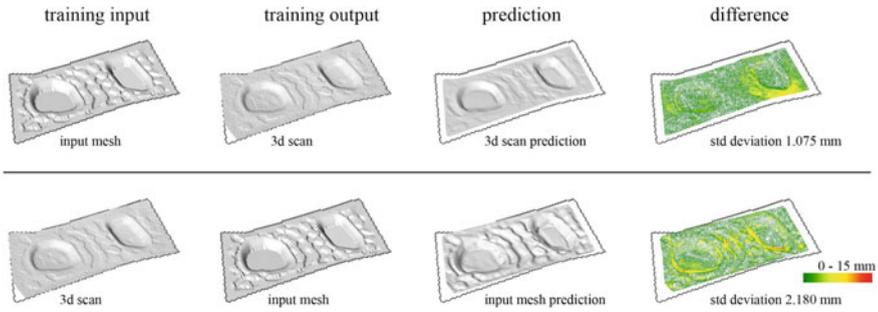
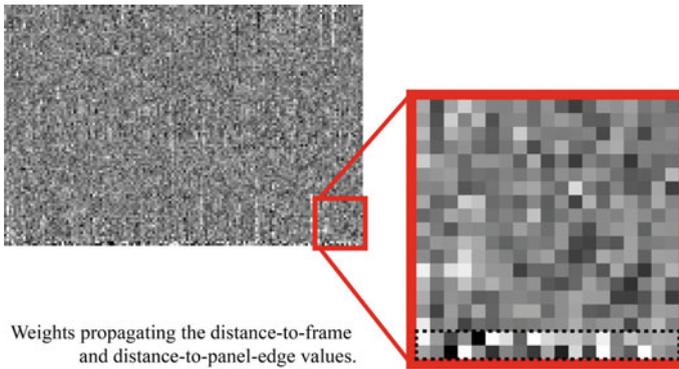


Fig. 9. The comparison of different input-output training sets and the achieved accuracy. *Top row* “forward” prediction, *bottom row* “reversed” prediction



Weights propagating the distance-to-frame and distance-to-panel-edge values.

Fig. 10. A plot of the network’s first layer weight values

to be the sparsity of those regions in the training dataset in comparison with the rest of the samples. This could possibly be solved by collecting more samples from a greater variety of panels.

An interesting research direction is seen in the analysis of the trained network. The plot of the first layer’s weights shows a pattern formed by the last 2 rows of the weight matrix (Fig. 10). Those values are trained on the data encoding the distance to the edge of the panel and the distance to the fabrication rig frame. This pattern indicates that there is some importance of these data in the trained model, although we don’t know if it is possible to quantify.

Conclusion

This paper demonstrates the use of sensing and learning to manage springback, for the purpose of reducing geometric inaccuracies within the incremental sheet forming process. Two different methods have been presented, the first based on in-process localised adaptation, and the second based on learnt prediction. Each introduces new

potentials for design decision making, extending specification into an ongoing contact with the process of fabrication. However, while both methods negotiate between the design model and the fabrication process, the first method is limited to a simple process of continual correction, as it does not develop a model of material behaviour, while the second method does construct such a model. We believe that the application of machine learning in this project indicates their wider feasibility in architectural fabrication.

Acknowledgements. This project was undertaken as part of the Sapere Aude Advanced Grant research project Complex Modelling, supported by The Danish Council for Independent Research (DFF). The authors want to acknowledge the collaboration of Autodesk Forge, SICK Sensor Intelligence Denmark, Monash University Materials Science and Engineering, Bollinger + Grohmann, and the robot command and control software HAL.

References

- Autodesk Forge: <https://forge.autodesk.com/>. Accessed 19 May 2017
- Bambach, M., Taleb Araghi, B., Hirt, G.: Strategies to improve the geometric accuracy in asymmetric single point incremental forming. *Prod Eng* **3**, 145–156 (2009). doi:10.1007/s11740-009-0150-8
- Behera, A.K., Verbert, J., Lauwers, B., Duflou, J.R.: Tool path compensation strategies for single point incremental sheet forming using multivariate adaptive regression splines. *CAD Comput. Aided Des.* **45**, 575–590 (2013). doi:10.1016/j.cad.2012.10.045
- Faro Focus3D 120: <http://www.faro.com/en-us/products/3d-surveying/faro-focus3d>. Accessed 19 May 2017
- Grasshopper 3D: <http://www.grasshopper3d.com/>. Accessed 19 May 2017
- HAL: <http://www.food4rhino.com/app/hal-robot-programming-control>. Accessed 19 May 2017
- Jeswiet, J., Bramley, A.N., Micari, F., et al.: Asymmetric single point incremental forming of sheet metal. *CIRP Ann. Manuf. Technol.* **54**, 623–649 (2005)
- Khan, M.S., Coenen, F., Dixon, C., et al.: An intelligent process model: predicting springback in single point incremental forming. *Int. J. Adv. Manuf. Technol.* **76**, 2071–2082 (2014). doi:10.1007/s00170-014-6431-1
- KUKA|prc: <http://www.robotsinarchitecture.org/kuka-prc>. Accessed 19 May 2017
- Nicholas, P., Zwierzycki, M., Nørgaard, E., Stasiuk, D., Ramsgaard Thomsen, M.: Adaptive robotic fabrication for conditions of material inconsistency: increasing the geometric accuracy of incrementally formed metal panels. In: *Proceedings of Fabricate 2017*, pp. 114–121 (2017)
- Owl Framework: <https://github.com/mateuszzwierzycki/Owl>. Accessed 19 May 2017
- Paniti, I., Rauschecker, U.: Integration of incremental sheet forming with an adaptive control into cloud manufacturing. In: *MITIP* (2012)
- Ruffini, R., Cao, J.: Using neural network for springback minimization in a channel forming process. *SAE Tech. Pap.* (1998). doi:10.4271/980082
- Robots IO: <https://robots.io/wp/>. Accessed 19 May 2017
- TensorFlow Framework: <https://www.tensorflow.org/>. Accessed 19 May 2017
- Volvox: <https://github.com/CITA-cph/Volvox>. Accessed 19 May 2017