DRAWING RECOGNITION

Integrating Machine Learning Systems into Architectural Design Workflows

L. BROWN,
University of New South Wales, Sydney, Australia
Lachlan.benjamin.brown@gmail.com

Abstract. Machine Learning (ML) has valuable applications that are yet to be proliferated in the AEC industry, currently there exists only handful of meaningful research examples. ML offers significant new ways to produce and assist design; neural networks possess the unprecedented power of independent decision-making based solely off interpreting massive quantities of data. However, these tools are largely out of the reach of designers, severely limiting opportunities to improve the methods by which designers design. To optimise the practices of designers, this paper aims to create a ML tool that can be integrated into architectural design workflows. This research investigates how ML can be used to universally move BIM data across various design platforms through the development of a convolutional neural network (CNN) for the recognition and labelling of rooms within floor plan images of multi-residential apartments. This task currently has no standard means of operation and is laborious, subject to human-error and data loss/discardment. CNNs extract patterns from spatial data (e.g. visual data) and are thus suitable for this image recognition task. Ultimately, A fully trained CNN will obtain deep understandings of design, common patterns and the relationships between components, this information when interpolated to architects may be used to infer how and why we design. A further step in the development of ML will be the qualification and generation of designs based off the knowledge the CNN has already extracted. The long-term aim of this work is to open the door to further practical implementations of ML applications, providing designers with accessibility to these sophisticated tools. The effects of this computation and thinking shift will have meaningful impacts on future practices enveloping all major aspects of our built environment from designing, to construction to management

Keywords. machine learning; convolutional neural networks; labelling and classification; design recognition

1. Introduction: (Research context and motivations)

If you can understand a floor plan, you can design a floor plan, and to understand you must learn. With ML, provided enough training, a program can become the master of nearly anything. Whether it be strategy such as DeepMind's AlphaGo Zero which, after three days of self-teaching defeated world no.1 Go player Ke Jie (Hassabis, Silver, 2017), or creative endeavours observed in *The Next Rembrandt project, that generated a new* Rembrandt artwork after study the artist's body of work (Nudd, 2016). Considering the versatility, and rampant accessibility of modern ML algorithms, this research is provoked into exploring ML within architecture, a field sparsely populated with meaningful applied results (Khean, 2017). Directing ML efforts towards architectural pursuits will enable algorithms to understand floor plans and thus become the masters of designing them, leading to a plethora of design implications from optimising the practices within architecture to providing tools by which those underprivileged can access design.

Furthermore, the AEC industry is often perceived quite negatively when discussing the integration of modern technologies into their workflows. There currently exists a visible reluctance to adopt emerging technologies, which only serves to hinder the means for architects to improve their designs towards more economic, efficient, environmental and aesthetic structures. BIM data's significance to architectural practices continues to grow exponentially, most firms have become reliant on data to the extent that the mastery and fluency of creating, organising, analysing and extracting BIM data has become a necessity to futureproof themselves (Davis, 2019). However, we find that commonly AEC architects are not mastering BIM data. Data and time is frequently wasted when projects are transferred across design platforms (e.g. exporting a special geometry from rhino to revit or revit to rhino to conduct solar analysis), due to a lack of interoperability, data is inevitably purposely discarded or lost. More time is wasted either re-creating lost data or finding inefficient workarounds in destination software to combat incompatibilities. These issues drain energy from core design work and compound in a compromised end product resulting in the end users experiencing the built environment negatively. Specifically, for PTW Architects whom this research was conduct with there existed the need to improve solar analysis compliance checking, only living rooms and balconies needed to be assessed, however the BIM data defining these rooms failed to be transferred from revit to rhino. They desired a means for which ML could extract relevant information from one place and interpolate it in another.

Accordingly, it is the aim of this research to combat this issue of interoperability and introduce future-proof practices into design workflows through the development of a ML algorithm in the form of a CNN that recognises architectural drawings and is thus able to automatically classify and label elements within floor plans such as individual room properties. It will be the role of the CNN to be given a floor plan as a visual input and output an instance segmentation function on said image, where each room, is classified and the region it occupies is masked (i.e. the pixels inside the region are labelled corresponding to the classification). The CNN accomplishing this aim by automatically producing universal BIM data capable of being delineated across various design platforms. This research additionally contributes knowledge towards architectural understandings of ML, the data created through this research can be extrapolated to assist future machine generation of floor plans.

2. Research Aims

Despite the nearly ubiquitous presence of ML in research fields and even our popular zeitgeist, there continues to be an underwhelming representation of ML systems within architectural practices. It appears theoretical developments produced in research scenarios fail to be converted into tools implemented in practical scenarios. The result of this is somewhat of a frustrating situation, where there exists technology that can improve not only the practices but the architecture in which we inhabit, yet it lies dormant and unused.

Considering this context, the aims of this research are twofold, to investigate how ML systems can be implemented into architectural design workflows, and if ML systems can develop some form of understanding of design to draw nearer to a goal of machine designed architecture. To satisfy these aims, this research details the development of a CNN that can interpret architectural drawings to classify and label the rooms of residential floor plan images. The CNN can inform the development of an automated room labelling application that can be implemented into architectural design workflows and contribute knowledge to ML understandings.

3. Research Questions

From these aims the research is provoked into exploring the following key questions:

1. How can convolutional neural networks be utilised within AEC design workflows to optimise or automate the recognition, classification and extraction of architectural elements such as rooms?

2. How can machine learning algorithms synthesise the elements, topologies, values, etc. of floor plans to identify, distinguish and separate between the spaces within?
3. To what extent can machine learning algorithms understand and interpolate architectural design? And to what extent of sophistication must a ML algorithm have to generate design?

4. Methodology

This project adopts the design action research methodology to inform and guide the course of work undertaken. Action design research as proposed by Maung K. Sein, is an approach whereby the generation of knowledge is motivated by the discovery of problem situations in a specific organisational setting and through an iterative cycle produces an evaluated IT artifact that addresses said problem (Sein, et al, 2011). Likewise, this research is motivated by present issues observed within AEC design workflows, and through the construction of a ML algorithm produces knowledge which will be used to inform an IT artifact that will automate, and thus solve/alleviate the issue. Following the structure of action design research this research is divided into three separate stages which operate in this manner: investigate, act, observe.

In the first stage (problem formulation), architectural design workflows are observed and investigated and after exploring ML a specific problem is formulated and the strategy by which to solve it. In the second stage (building intervention and evaluation), the NN is constructed and the accuracy of its predictions and the robustness of the system is assessed, since this is an iterative approach this step may repeated a number of times before a satisfactory result is attained. Finally, in the third stage (reflection and learning), moves thinking into a conceptual frame where the knowledge, applications and implications obtained from producing the NN is ascertained.

5. Background Research/Literature review

ML is by no means a recent technology, examples of applied ML can be traced back to Arthur Samuel's checker player for the IBM 701 in 1952 (Barto, Sutton, 1992). However, today the public has access to unprecedented computational powers and open source tools such as TensorFlow and PyTorch that have enabled us to openly explore ML outside of traditional computer science fields. Mario Carpo predicts that the effects of this computational shift will see design informed by the mass retrieval of data and information, an environment where ML technologies may thrive and potentially create a new form of artificial intelligence (Carpo, 2017). ML enables a program to perform given tasks without explicit instruction through a process of training. Professor Tom M. Mitchel of Carnegie Mellon University states that “A program is said to learn from experience with respect to some class of tasks, and performance measure, if its performance at tasks, as measured by performance measure, improves with experience.”. (Mitchell, 1997).

A particular field of ML comes under the umbrella of deep learning, commonly associated with the implementation of artificial neural network (NN) architectures. NNs probe mass quantities of data to find statistical correlations, patterns, trends etc. and extract this metadata to inform an output such as some form of prediction, classification or suggestion. As their name suggests NN draw loose inspiration from biological models of neural networks like the brain, a NN is a series of interconnected neurons that receive, affect and send data (Miller, 2015). The process by which a NN learns is surprisingly brute force by nature, using forward and backpropagation and NN receives an input and is told to produce an output, the data it is fed has labels associated with it (supervised learning) that inform the NN of the desired (correct) outputs. The NN iteratively receives a piece of training data, produces and output, finds out its wrong, fine tunes itself, repeats, and now is slightly less wrong. The is a process referred to as gradient decent and can be likened to climbing down a mountain blindfolded, and after a long enough cycle the NN has a series of values and hyperparameters that consistently produce (mostly) correct outputs when it receives a certain class of input data.

Of particular interest are convolutional neural networks (CNN), they possess the ability to extract patterns from spatial data types and are thus frequently used in image recognition tasks. With accurate training methods a CNN architecture have achieved high accuracy rates in image classification tasks (Hinton, et al, 2017). CNNs operate in two stages, firstly they detect features by applying convolutional operations over images on a pixel level, and secondly, they classify features in later layers by developing understandings of detected features. In an architectural context Jennifer Ng employed CNNs to distinguish between sections and plans (Ng, 2018).

There currently exists only a handful of meaningful research examples that explore the capacity of ML system to be applied to architectural endeavours, leaving many questions concerning the legitimacy of ML in this field open. Stanislas Chaillou’s *AI + Architecture* demonstrates ML’s ability to classify, validate and generate architectural drawings (Chaillou, 2019), furthermore, ML can comprehend design to the extent that researchers from MIT managed to differentiate designs between architects using a NN (Ratti, et al, 2018). During these investigations specific research was discovered that helped inform the course of this research. In one case a group used a generative adversarial network for the recognition and generation of floor plan images (Huang, Zheng, 2018), they achieved interesting results, however, the decision to generate a training dataset from a single data source (that being marketing floor plans from a Chinese website) resulted in an overfitted NN where it was unable to correctly classify non-orthogonal plans or plans that did not pertain to the specific design style the NN was trained on, thus rendering itself useless in all other situations. This situation highlights the incredible importance of the data to determine the success of ML applications. Another group investigated the CNNs ability to distinguish between building typologies of monetarises and mosques, training their CNN on floor plans of these typologies (Ferrnado, et al, 2019). Again, it was the process of data collection and pre-processing that proved most crucial and laborious/time consuming, this research proved to be quite insular whilst training was, in most respects successful, their does not appear to be application produced as a result of these research examples.

6. Case Study

This research is a collaboration with PTW Architects that aims to investigate ML technology towards an architecturally oriented goal. Existing research demonstrates that ML can be used to classify architectural topologies, assess designs through a certain metric such as compliance regulations, make predictions and assessments on projects or even directly influence design decisions. But this research project specifically aims to explore drawing recognition and classification of spatial types in multi-scale residential floor plans to inform the development of an automated room labelling workflow for use on residential architecture projects. The project was divided into four key stages that involved ML processes research, data collection and processing, script development, testing, reflection and the production of a ML floor plan recognition model.

The first stage of the project involved researching machine learning and existing applications in the AEC industry. These findings informed the appropriate ML methodology to adopt, the requirements of this methodology and the feasibility to perform all requirements to contribute new findings. This research's aims necessitate the creation of a CNN, train it on labelled floor plan images, and a program that takes its predications and produce classification instance segmentation labels. If this project work is completed successfully and all components work harmoniously with each other then we will have produced a ML algorithm whose outputs will inform an automated room labelling workflow and ML understandings towards design.

Research concludes using the python programming language with the open source ML library TensorFlow and the Keras API in the Visual Studio Code IDE is the optimal approach to take to create the CNN. Due to the extensive documentation, learning resources, compatible tools and community support provided by using these popular tools enables a higher chance of successfully producing an optimal CNN that will accurately perform its desired task. The development of a CNN can be synthesised into three main components; assembling layers, compiling model and fitting model. The process of ML in CNNs has previously been delineated in the background research section of this paper, therefore we will describe what these three components are with expectation of that knowledge. Assembling layers defines the structure of the CNN; the number of layers and the purpose they have (whether they perform a convolutional or pooling operation for instance), the number of neurons and their activations (sigmoid, Relu, softmax, etc.) and, what amount of potential outputs it can produce (size of the output layer). Essentially, we are defining how data is fed into the CNN, interpreted and converted into an output. Compiling the model defines ML metrics for the processes of learning in forward and back propagation, the loss and optimiser functions are used to analyse the performance of the CNN during training and inform how the weights and biases are altered to improve output accuracy. Fitting the model segments the dataset into training, and validation and gives them to the CNN, the training set is the largest (approximately 90% of the images) and is used to inform the learning, the images are accompanied by labels which are the image's desired outputs if it were to be fed to the CNN, within the training cycle the CNN wants to produces outputs that match these labels to the greatest degree of accuracy possible. The validation dataset used almost as a test; it withholds a small section (10%) of the data to show after training to assess the model's performance and ensure that the CNN has not been overfitted with the training data (when the CNN is too closely trained on the specific examples in the training data and thus becomes useless in any practical scenario).

The second stage involved gathering floor plan images, labelling images and pre-processing images into a dataset to feed the CNN to enable training to commence. Typically, ML algorithms require massive quantities of varied data which posed a problem to this research in collecting enough to form a substantial dataset to enable any meaningful ML developments. Initially it was thought that marketing floor plans could easily be sourced from real-estate companies, this proved false. Fortunately, PTW Architects provided a dataset of 454 floor plan images from their previous projects, these consisted of CAD drawings, marketing images, and hand-drawn floor plans from a variety of multi-scale residential buildings. Despite, still the limited size of this dataset, it was believed that quality of this data would make up for this limitation, furthermore, the feasibility of manually labelling 454 within the span of this research made it quickly apparent that whether the dataset size was satisfactory or not, didn't matter but whether or the required work could be completed.

ML is mostly about the data, this stage of the project surprised us by being by far the most extensive and time consuming. This is logical, ML is entirely dependent on the data it is provided to learn, if there exists flaws within this data the ML product is made redundant. The supervised learning approach employed by this researched required all the images collected to be accompanied by a label in the form of a .JSON file that defined the regions in which certain room classes were occupying. To simplify this task and focus the efforts of the ML algorithm, it was decided to make six room classes for the CNN to search for, these being: bedroom, bathroom (includes ensuite), living room, dining room (if connected to the living room, counted as part of the living room), kitchen and balcony (includes terraces, patios, etc.). The images were labelled using the VGG image annotator (Dutta, Zisserman, 2019), this involved drawing the regions that rooms occupied and defining said region as whatever room it was. The output of this process was a massive .JSON that contained every label for all the images and would be referenced for the training cycle.

The third stage involved developing a program that could perform instance segmentation on test data. Instance segmentation is the classification of multiple class regions within a single image and produce a mask that defines the space the classes occupy within the image. This task proved itself to be too computational advanced for this research to build from the ground up. That is why the Mask R-CNN framework developed by the Facebook AI Research Group was employed (He, et al, 2018). Using the pre-defined functions from Mask R-CNN and altering them to fit our data and project's desired outcomes within our own scripts enabled us to create a CNN capable of meeting the aims of this research. With a CNN, a dataset and a framework for instance segmentation, training commenced over a period of 1000 epochs. Initial training cycles failed due to various memory and iterative errors, after a process of debugging and re-writing to more efficiently pass data (alleviating the stresses on the system), a full training cycle was completed. The finalised CNN model was then exported as a .h5 file containing the finalised hyperparameters to an inference script where testing was conducted.

The fourth and final stage involved the testing and evaluation of the CNN. With a completed CNN running instance segmentation on floor plan images, new test images were inputted, overall the results were satisfactory and demonstrated that CNN was beginning to develop a fundamental understanding of design. The test images were purposely varied in design typologies, visual style, and size, they even included floor plans from projects outside of Australia and plans where features such as furnishing were absent, through this the CNN was validated as tool of wide capabilities. The CNN began with randomised values and through the training cycle developed its own methods to detect and interpret floor plans. Through testing it is evident that a multitude of factors informs its decision-making processes these are; visual features including symbols, icons, linework, patterns, etc. the shape and form of elements within the image and relational data, for instance the layout of floors and general design principals that dictate how we design the spaces we inhabit. The ability for the CNN to rely on multiple features is good because it means it makes it a more robust system capable of dealing with more floor plans that may be unique in their features and styles. However, the CNN is far from perfect, it is evident that from viewing its outputs that there is yet to a be a single 'perfect' output where rooms are labelled exactly and correctly, these issues will be further discussed in the next section of this paper.

Figure 1. Room Labelling Results.

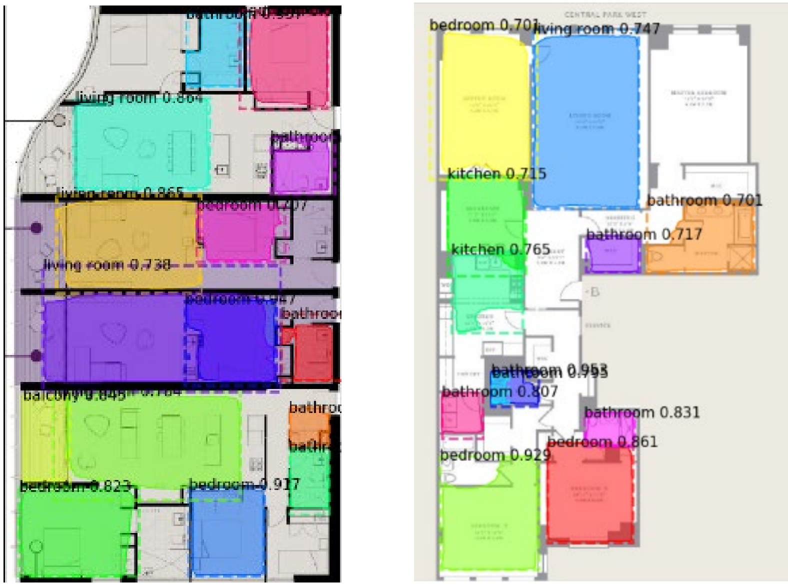

Figure 2. Room Labelling Results.

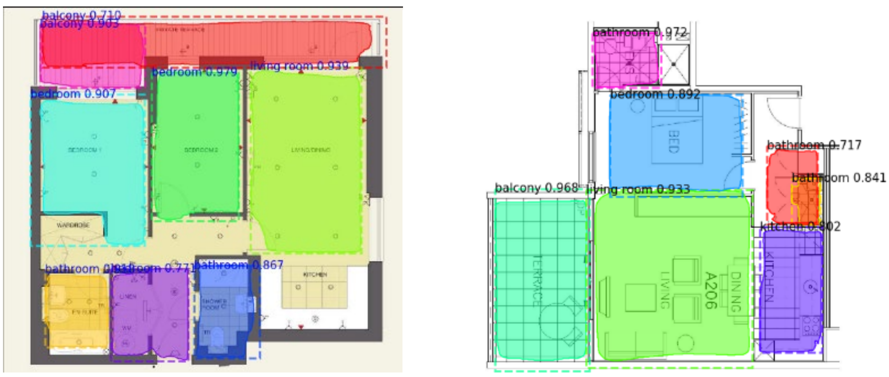

Figure 3. Room Labelling Results.

7. Discussion (evaluation and significance)

The aim of this research was to develop a ML algorithm that could understand design a subsequently be used to inform the creation of an automated room labelling application. This research documented the development of the CNN for this purpose, it has shown an initial success to reading and labelling architectural drawings, that with further refinement and training with additional data will be sophisticated, reliable tool, thus demonstrating how ML systems and thinking may be applied and integrated into architectural design workflows to optimise practices within design workflows.

CNNs do possess the ability to, if directed and trained, understand and interpret design within an architectural framework. The CNN possesses an elementary ability to classify rooms and an advanced ability to distinguish between spaces that with additional training will only improve. The CNN initially randomises the variables by which it detects features and only through trial and error during training discovers what is of significance to complete its task successfully, so there exists a layer of mystic concerning how ML algorithms truly understand their subjects. However, from analysing several of the CNNs outputs it is evident that a number of factors informs its predictions. These factors include features such as the addition of text, symbols and icons (e.g. the word bedroom and a bed symbol), linework such as room boundaries (walls) and patterns (floors), and most interesting to this research, relational data. It is possible that the CNN is unintentionally learning design rules and compliances through its brute force training approach to learning, such as room sizes (e.g. the largest room is the living room), room layout (e.g. the balcony is separated from the rest of the floor plan, often accessible only through another room) and room shapes (e.g. the bedrooms often have wardrobes which extrude out from their otherwise thick rectangular typology).

In a real example for instance, the CNN may find that bathrooms typically contain a number of visually distinct graphics (e.g. the symbol for a toilet, shower, bath and sink) within a narrow much smaller space in relation to the rest of the floor plan, confined by thick lines and chequered pattern within the bounds of said lines and contextually close to bedroom regions. So, when the CNN is viewing an image containing a bathroom/s it may rely on some or all these things and potentially more to formulate a bathroom prediction.

There are, however, aspects that continue to confound the CNNs prediction making process, larger images with too many class instances appear to overwhelm the program, this however may just be a limitation of the hardware available to this research. More serious issues pertain to aspects of classifying and distinguishing. Popular trends in modern design facilitate open space living, thus blurring the separation between kitchen living room and dining room, in many tests cases there was an observable uncertainty and hesitation in the CNNs predictions concerning the distinguishing typically were unsatisfactory either over or under compensating. Rooms that are visually similar on floor plans such as a laundry and bathroom occasionally result in incorrect classifications. With more extensive training these nuanced problems can be overcome, it is clear that these problems are not fundamental issues within CNNs but are a result of research limitations, despite this, there presently would be a risk if designers were to use this tool for their projects that they may in fact simplify or alter aspects of their design to accommodate for the limitations of the program, which is worrying and creates a clear motivation to quickly overcome these limitations.

A further point of contention to draw out comes from the selection of training data, the issue of size is an obvious one, ideally the size of the dataset would have been in the thousands, despite this limitation, this research has already proven with this dataset that a CNN can understand design, a larger dataset would simply make it more sophisticated. However, a more important, less visible limitation was data availability, this research used exclusively floor plans that PTW Architects designed, this is significant because it may mean that the CNN is learning the biases and design preferences of PTW, resulting in a layer of subjectivity in a tool that is desired to be objective. Another fear sparked by this revelation is that of overfitting the CNN, where the data is too specialized to a specific context, that being PTW's work that the CNN fails to be a universal tool that other firms can benefit from. It is unclear whether simply brute forcing this issue with more training data will overcome these concerns and is a more philosophical question that will be investigated with further inquiry beyond this research.

8. Conclusion

Currently we stand at a crucial time for architecture, the decisions made by firms now, will dictate whether they will survive the oncoming decades. It is clear that the need to adapt and become masters of BIM data and emerging technologies is a necessity and, if architects are still forced to deal with the interoperability issue described in this research and other similar issues then they will fail to accommodate their design/output to the growing needs of market. That is why investing in further ML investigations and applications in architecture is of significance, the knowledge collated in this research has proved that even with a several limitations a CNN can be produced to optimise, automate and improve several processes that previously hindered design workflows.

This research is definitely preliminary, addressing the statement made in the introduction alluding to floor plans designed by machines, the technology produced by this research is not capable of such a feet, the level to which the CNN understands design is yet to surpass or even attain to a human's cognitive ability. In the spectrum of design oriented applications, however, initial applications where the ML algorithms work in tandem with architects to suggest and inform design decisions are within reach. The applications that can be produced as a direct consequence are quite staggering, the results of this research possess a multitude of direct and indirect implications. Directly, we know that an application where floor plans are read and automatically labelled can be produced after another, more in depth training cycle and indirectly we can see that ML opens the door to the concept of automatically generated, universal BIM data, where this thinking can be applied to a plethora of other scenarios.

With further developments of ML in architecture, and if we divulge more of our data and general architectural information to ML algorithms, then the concept of machine designed architecture will not be one of insane over-ambition. The effects of this kind of computation shift will be staggering, ML will divulge information concerning design that humans are cognitively incapable of perceiving, informing better, optimised and automated design practices. The consequence of this computational shift will ultimately augment the perceptions we as designers have towards our built environment and the means by which we conceptualise, design and create it.

Acknowledgements

Nicole is amazing, the best lecturer evah. Hank, Yannis, and Cristina are also pretty awesome.

Thanks to my industry partner PTW Architects, Michael and Brandon from PTW and to Nariddh Khean for your crucial expertise and mentorship, whom without your help this research would not be what it is.

References

Abdelfattah, A.: 2017, “Image Classification using Deep Neural Networks – A beginner friendly approach to using TensorFlow”, Available from: Medium <https://medium.com/@tifa2up/image-classification-using-deep-neural-networks-a-beginner-friendly-approach-using-tensorflow-94b0a090ccd4>, (accessed 22 October 2019).

Barto, A, Sutton, R, 1992, Reinforcement Learning: An Introduction, The MIT Press, London

Hinton, G, Krizhevsky, A, Sutskever, I, 2017, ‘ImageNet Classification with Deep Convolutional Neural Networks’, University of Toronto, Toronto.

Carpo, M.: 2017. “The Second Digital Turn: Design Beyond Intelligence”. Available from: The MIT Press <https://mitpress.mit.edu/books/second-digital-turn> (accessed 27 October 2019).

Chaillou, S, 2019, ‘AI + Architecture | Towards a New Approach’, Harvard University

Davis, D.: 2019, “What ‘The Future of The Professions’ Reveals About the Future of Architecture” Available from: Daniel Davis <https://www.danieldavis.com/future-of-the-professions/> (accessed 14 November 2019).

Dutta, A., Zisserman, A: 2019, *The VIA Annotation Software for Images, Audio and Video,* ACM International Conference on Multimedia, Nice.

Ferrando, C, Dalmasso, N, Mai, J, Llach, D, 2019, ‘Architectural Distant Reading, Using Machine Learning to Identify Typological Traits Across Multiple Building’, CAADFutures19

Haung, W, Zheng, H 2018, ‘Architectural Drawings Recognition and Generation through Machine Learning’, ACADIA

He, K, Gkioxari, G, Dollar, P, Girshick, R ,2018, ‘Mask R-CNN’, Facebook AI Research (FAIR)

Isola, P. Zhu, J. Zhou, T. Efros, A 2018, ‘Image-to-Image Translation with Conditional Adversarial Networks’, Berkeley AI Research (BAIR), Laboratory, UC Berkeley

Kalervo, A, Ylioinas, J, Haikio, M, Karhu, A, Kannala, J, 2019, ‘CubiCasa5K: A Dataset and an Improved’, Department of Computer Science, Aalto University, Espoo

Karn, U.: 2016, “An Intuitive Explanation of Convolutional Neural Networks”, Available from: The Data Science Blog <https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/> (accessed 12 October 2019).

Karpathy, A.: 2017, “Convolutional Neural Networks (CNNs / ConvNets)”, Available from: GitHub <http://cs231n.github.io/convolutional-networks/>, (accessed 11 October 2019).

Khean, N., Kim, L., Martinez, J., Doherty, B., Fabbri, A., Gardner, N. and Haeusler, M. H.: 2018, The Introspection of Deep Neural Networks - Towards Illuminating the Black Box, *Proceedings of CAADRIA 2018*, Beijing, 237 - 246.

Miller, S.: 2015, “Mind: How to Build a Neural Network”, Available from: GitHub <https://stevenmiller888.github.io/mind-how-to-build-a-neural-network/> (accessed 4 October 2019).

Mitchell, M 1997, Machine Learning, McGrawHill

Ng, J, 2018, ‘Optimising Image Classification’, University of New South Wales, Sydney

Nudd, T.: 2016, “Inside ‘The Next Rembrandt’: How JWT Got a Computer to Paint Like the Old Master” Available from: Adweek <https://www.adweek.com/brand-marketing/inside-next-rembrandt-how-jwt-got-computer-paint-old-master-172257/> (accessed 22 November).

Santos, L.: 2018, “Object Localization and Detection”, Available from: GitBooks <https://leonardoaraujosantos.gitbooks.io/artificial-inteligence/content/object_localization_and_detection.html> (accessed 26 October 2019).

Sein, M., Henfridsson, O., Purao, S., Rossi, M., Lindgren, R., 2011, *Action Design Research, MIS Quarterly Vol. 35 No. 1 pp. 37-56/March*.

Silver, D., Hassabis, D.: 2017, “AlphaGoZero: Starting from scratch”, Available from: DeepMind <https://deepmind.com/blog/article/alphago-zero-starting-scratch> (accessed 19 November).

Touretzky, D., Pomerleau, D.: 1989, *What’s Hidden in the Hidden Layers?*

Yoshimura, Y, Cai B, Wang, Z, Ratti, C, 2018, ‘Deep Learning Architect: Classification for Architectural Design through the Eye of Artificial Intelligence’, Massachusetts Institute of Technology.