Drawing Recognition

Integrating Machine Learning Systems into Architectural Design Workflows

Lachlan Brown[1], Michael Yip[2], Nicole Gardner[3], M. Hank Haeusler[4], Nariddh Khean[5], Yannis Zavoleas[6], Cristina Ramos[7]
[1,3,4] UNSW / Computational Design [2] PTW Sydney [5] Austrian Institute of Technology
[6,7] UNSW / Computational Design
[1] lachlan.brown@student.unsw.edu.au [2] m.yip@ptw.com.au [3,4,6,7] {n.gardner|m.haeusler|y.zavoleas|c.ramos}@unsw.edu.au
[5] Nariddh.Khean@ait.ac.at

Machine Learning (ML) has valuable applications that are yet to be proliferated in the AEC industry. Yet, ML offers arguably significant new ways to produce and assist design. However, ML tools are too often out of the reach of designers, severely limiting opportunities to improve the methods by which designers design. To address this and to optimise the practices of designers, the research aims to create a ML tool that can be integrated into architectural design workflows. Thus, 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. 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*

INTRODUCTION, CONTEXT, MOTIVATION

If you can understand a floor plan, you can design a floor plan, and to understand you must learn. While this statement might simplify learning one can argue that with Machine Learning (ML), provided enough training, a program can currently become the master of an increased number of things. Popular examples might include strategy such as DeepMind's AlphaGo Zero which, after three days of self-teaching defeated world no.1 Go player Ke Jie [1], or creative endeavours observed in *The Next Rembrandt project*, that generated a new Rembrandt artwork after studying the artist's body of work [2]. Considering the versatility, and rampant accessibility of modern ML algorithms, this research is provoked into exploring ML within architecture, a field with an increasing number of meaningful applied results and attempts (Khean, 2017, 2019).

Yet there is evidence that the Architecture, Engineering and Construction (AEC) industry often resists the integration of modern technologies into their workflows [3]. As outlined by McKinsey in [3] 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. Consequently 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 and compromising the productivity of the AEC sector [4]. When analysing the productivity gain in other sectors enabled through the introduction of ML [5] (McAfee, Brynjolfsson, 2017) one can clearly argue for the need to introduce ML into the AEC sector as it serves the creation of a nation's critical infrastructure and building stock with currently AU$200B in the pipeline [6] needed to keep pace with the nations population growth [7]. It is an industry crucial to the national economy responsible for 8% of GDP or $134B P.A. [8]. It has, however, been slow to adopt new technology (Deutsch 2019), digitally enabled workflow processes that unlock productivity [4] (Susskind & Susskind 2016; McAfee et al. 2017). We argue therefore for research and investigations into ML's ability to automate *'mundane'* tasks that are essential but not *'very sexy'* in order to increase the productivity of the AEC industry. Items that, amongst others can fall under the category of mundane tasks are compliance checks and the question of how to perform compliance checks (Knodel, Naab, 2016). Specifically, for [NAME WITHHOLD] Architects whom this research was conduct with identified a need to improve solar analysis compliance checking according to the local building code. The code required that only living rooms and balconies needed to be assessed, however the BIM data defining these rooms failed to be transferred from Revit to Rhino. Consequently the [NAME WITHHOLD] Architects questioned if ML could extract relevant information from one place and interpolate it in another.

RESEARCH AIMS, OBJECTIVES AND OUTCOMES

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 Convolutional Neural Network (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.

With this aim in mind the research has the following objectives:

- 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. This research's objective is to contribute to how ML can be used in the AEC sector using a pragmatist approach (Gardner et al 2020)
- 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. A further objective of this research is using the industry partnership and its applied context as vehicle to identify how practice might use ML.

Considering this context, the outcomes based on the aims and objectives 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.

Hence as an outcome 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.

RESEARCH QUESTIONS

From these aims, objectives and outcomes the research is provoked into exploring the following key questions:

- *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?*
- *How can machine learning algorithms synthesise the elements, topologies, values, etc. of floor plans to identify, distinguish and separate between the spaces within?*
- *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?*

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 (2011), 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. Likewise, this research is motivated by present issues observed within AEC design workflows (through the research partnership with [NAME WITHHOLD] Architects, 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 Neural Network (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.

BACKGROUND RESEARCH AND 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 with access to unprecedented computational powers and open source tools such as TensorFlow and PyTorch have enabled us to openly explore ML outside of traditional computer science fields. Mario Carpo (2017) 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. Certain types of ML enables a program to perform given tasks without explicit instruction through a process of training. Tom M. Mitchel of Carnegie Mellon University (1997) 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."*

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 [9]. 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 prior mentioned 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. But what is the current state of play in the AEC industry?

There currently exists only a handful but growing number of 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. To provide a few examples; Ng (2018) employed CNNs to distinguish between sections and plans (Ng, 2018). Chaillou's (2019) AI + Architecture demonstrates ML's ability to classify, validate and generate architectural drawings. Yoshimura, Cai, Wang & Ratti's (2018) research suggests that ML can comprehend design to the extent that researchers from MIT managed to differentiate designs between architects using a NN [10].

Through the literature review in the above, specific research findings were discovered that helped inform the course of this research. In one case Huang, Zheng (2018) used a generative adversarial network for the recognition and generation of floor plan images, and achieving 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 which 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 example highlights the importance of the data to determine the success of ML applications.

Ferrando et al. (2019) investigated the CNNs ability to distinguish between building typologies of monasteries and mosques, training their CNN on floor plans of these typologies. 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 in its final application in practice as it concentrated on monasteries and mosques whilst training of the CNN was, in most respects successful.

Equipped with these insights, the research team started its collaboration with [NAME WITH HOLD] Architects with the aim to investigate ML technology towards the architecturally oriented goal of recognising architectural drawings and thus being able to automatically classify and label elements within floor plans such as individual room properties.

CASE STUDY

The existing research above 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.

Researching Machine Learning and existing applications in the AEC industry.

The first stage of the project involved researching machine learning and existing applications in the AEC industry. In combination with the above listed literature review 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 stage necessitate the creation of a CNN, trains it on labelled floor plan images, and generate a program that takes its predications and produce classification instance segmentation labels. When completed successfully and all components work harmoniously with each other one has produced a ML algorithm whose outputs will inform an automated room labelling workflow and ML understandings towards design.

Research specific to this stage concludes that 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 create the CNN. This step is based on existence of the extensive documentation, learning resources, compatible tools and community support provided by users of these popular tools. Consequently it 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: (a) assembling layers, (b) compiling model and (c) fitting model. The process of ML in CNNs has previously been delineated in the background research section of this paper, therefore the paper will describe these three components in detail.

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, one defines how data is fed into the CNN, interpreted and converted into an output. Compiling the model defines the ML metrics for the processes of learning in forward and back propagation, the loss and optimiser functions that 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 validating and giving 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 are 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).

Gathering floor plan images, labelling images and pre-processing images into a dataset to feed the CNN

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, [NAME WITH HOLD] 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 images within the span of this research made it quickly apparent that whether the dataset size was satisfactory or not was of lesser importance to the feasibility of compiling and processing data.

As argued earlier, ML is mostly about the data. Consequently this stage of the project proofed to be 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 research 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
- 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 large .JSON file that contained every label for all the images and would be referenced for the training cycle.

Develop program that could perform instance segmentation on test data

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. Consequently the Mask R-CNN framework developed by the Facebook AI Research Group was employed [11]. 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.

Testing and evaluation of the CNN

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: (1) visual features including symbols, icons, line work, patterns, etc and (2) 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.

DISCUSSION, EVALUATION AND SIGNIFICANCE

The aim of this research was to develop a ML algorithm that could understand designs of floor plans and 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 in reading and labelling architectural drawings, that with further refinement and training with additional data will be a sophisticated, reliable tool. Hence the research was able to demonstrate 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 understand and interpret design within an architectural framework if directed and trained correctly. A CNN possesses an elementary ability to classify rooms and an advanced ability to distinguish between spaces. Due to only being able to train the CNN with a dataset of 454 floor plan images we believe that the results with additional training data sets will 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. 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), line work 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).

To elaborate on this point further. In an 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 CNN's 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 (standard university lap-

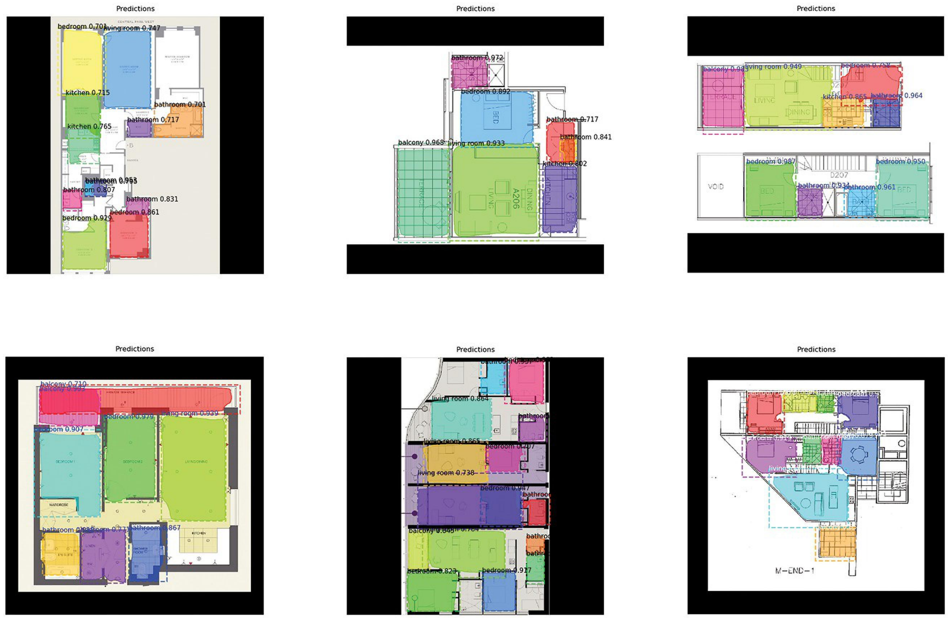

Figure 1
Various Room
Labelling Results

top - Microsoft Surface Pro) 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 the kitchen, living room, and dining room. In many tests cases the CNN's distinction between between rooms connected in open space living were unsatisfactory, and often displayed an uncertainty. 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. Furthermore, it is clear that these problems are not fundamental issues within CNN's but are a result of research limitations.

A further point of contention comes from the selection of training data. Here the issue of data size (454 floor plan images) is an obvious one, ideally the size of the dataset would have been in the thousands. Despite this limitation, this research has already proven that within this limited dataset a CNN can understand design. Consequently a larger dataset would simply make it more sophisticated. However, a more important, less visible limitation was data variety availability, this research used exclusively floor plans that [WITH HOLD] Architects designed. This is significant because it may mean that the CNN is learning the biases and design preferences of [WITH HOLD] Architects, 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 [WITH HOLD] Architects'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.

CONCLUSION

Currently we stand at a crucial time for architecture (Deutsch, 2019) [12] [13], the decisions made by firms now, will dictate their prosperity for the next decade. It is clear that it is a necessity for the AEC sector to adapt and open up to emerging technologies. Proponents of technology have identified a lack of automation and adoption of technology as the primary reasons for the AEC industry's poor performance (Chapman 2005; Deutsch 2019; Khean et al. 2019). We argue amongst others that the AEC industry needs reform if it is to sustain and grow in the new economy (Brynjolfsson et al. 2014; Susskind, Susskind 2015; McAfee 2017; Parker et al. 2016, van Rijmenam 2019). Automation technologies have the potential to help AEC firms sustain or even grow as they experience inevitable digital transformation. 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 and 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 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 a Synthetic Design method combining ML with computational design for an optimised design workflow consumable for humans via a web-browser in and for Architecture, Engineering and Construction (AEC) disciplines becomes feasible (Haeusler, 2019). 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.

ACKNOWLEDGEMENT

The authors would like to acknowledge the contribution of [NAME WITH HOLD] Architects for providing us with the research challenge, the data sets and an ongoing conversation about improving the research during the duration of this project. The project is part of the ongoing research in AI Enabled Workflow Automation for the AEC Industry an [NAME WITH HOLD] research group to grow and improve competitiveness in the Australian AEC industry by accelerating AI-enabled workflow automation.

REFERENCES

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

Brynjolfson, E and McAfee, A 2014, *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*, W. W. Norton&Company

Carpo, M 2017, *The Second Digital Turn: Design Beyond Intelligence*, The MIT Press, London

Chaillou, S 2019, *AI + Architecture | Towards a New Approach*, Master's Thesis, Harvard University
Chapman, R 2005 'Inadequate Interoperability: A Closer Look at the Costs', *22nd International Symposium on Automation and Robotics in Construction ISARC 2005 - September 11-14*, Ferrara, Italy, pp. 1-6
Deutsch, R 2019, *Superusers: Design Technology Specialists and the Future of Practice*, Routledge
Dutta, A and Zisserman, A 2019 'The VIA Annotation Software for Images, Audio and Video', *27th ACM International Conference on Multimedia*, Nice, France, pp. 2276-2279
Ferrando, C, Dalmasso, N, Mai, J and Llach, D 2019 'Architectural Distant Reading, Using Machine Learning to Identify Typological Traits Across Multiple Building', *18th International Conference, CAAD Futures 2019, Proceedings*, Daejeon, Korea, pp. 114-127
Haeusler, M H 2019, 'Theory (Methods) and Design in the Second Machine Age', in Gardner, N, Haeusler, M H and Zavoleas, Y (eds) 2019, *Computational Design: From Promise to Practice*, av edition , Ludwigsburg, Germany, pp. 56-67
Hinton, G, Krizhevsky, A and Sutskever, I 2012, 'ImageNet Classification with Deep Convolutional Neural Networks', *Advances in neural information processing systems*, 25(2), pp. 1-9
Huang, W and Zheng, H 2018 'Architectural Drawings Recognition and Generation through Machine Learning', *Proceedings of ACADIA 2018*, Mexico City, Mexico, pp. 156-165
Khean, N, Gerber, D, Fabbri, A and Haeusler, M H 2019 'Examining Potential Socio-economic Factors that Affect Machine Learning Research in the AEC Industry', *18th International Conference, CAAD Futures 2019, Proceedings*, Daejeon, Korea, p. 254
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, pp. 237-246
Knodel, J and Naab, M 2016, 'How to Perform the Architecture Compliance Check (ACC)?', in The Fraunhofer IESE Series on Software, and Systems Engineering (eds) 2016, *Pragmatic Evaluation of Software Architectures.*, Springer, pp. 83-94
McAfee, A and Brynolfsson, E 2017, *Machine, Platform, Crowd: Harnessing Our Digital Future';*, W. W. Norton&Company
Mitchell, T M 1997, *Machine Learning*, McGraw Hill
Ng, J M Y, Khean, N, Madden, D, Fabbri, A, Gardner, N, Haeusler, M H and Zavoleas, Y 2019 'Optimising Image Classification - Implementation of Convolutional Neural Network Algorithms to Distinguish Between Plans and Sections within the Architectural, Engineering and Construction (AEC) Industry', *Intelligent & Informed - Proceedings of the 24th CAADRIA Conference - Volume 2,*, Wellington, New Zealand, pp. 795-804
Parker, G G, Van Alstyne, M W and Sangeet, P 2016, *Platform Revolution: How Networked Markets Are Transforming the Economy - and How to Make Them Work for You*, Harvard Business Review
van Rijmenam, M 2019, *The Organisation of Tomorrow – How AI, blockchain and analytics turn your business into a data organisation*, Routledge
Susskind, R and Susskind, D 2016, *The Future of the Professions – How technology will transform the work of human experts*, Oxford University Press
[1] https://deepmind.com/blog/article/alphago-zero-starting-scratch
[2] https://www.adweek.com/brand-marketing/inside-next-rembrandt-how-jwt-got-computer-paint-old-master-172257/
[3] https://www.mckinsey.com/industries/capital-projects-and-infrastructure/our-insights/decoding-digital-transformation-in-construction
[4] https://nvlpubs.nist.gov/nistpubs/gcr/2004/NIST.GCR.04-867.pdf
[5] https://www.mckinsey.com/industries/capital-projects-and-infrastructure/our-insights/reinventing-construction-through-a-productivity-revolution
[6] https://www.infrastructureaustralia.gov.au/publications/australian-infrastructure-audit-2019
[7] https://population.gov.au/
[8] https://www.buildsoft.com.au/blog/10-statistics-defining-the-australian-construction-industry
[9] https://stevenmiller888.github.io/mind-how-to-build-a-neural-network/
[10] https://www.researchgate.net/publication/329441652_Deep_Learning_Architect_Classification_for_Architectural_Design_through_the_Eye_of_Artificial_Intelligence/stats
[11] https://arxiv.org/pdf/1703.06870.pdf
[12] https://www.danieldavis.com/future-of-the-professions/
[13] https://www.weforum.org/agenda/2018/06/construction-industry-future-scenarios-labour-technology/