Machine Learning as Enabler of Design-to-Robotic-Operation

Liu Cheng, Alexander

Publication date
2018

Document Version
Final published version

Published in
Archidoct

archi|DOCT

The e-journal for the dissemination of doctoral research in architecture.

Supported by the ENHSA Network | Fueled by the ENHSA Observatory

11

July 2018
www.enhsa.net/archidoct
ISSN 2309-0103

INTELLIGENCE

Lifelong Learning Programme

enhsa

European Observatory of Doctoral Research in Architecture

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

Machine Learning as enabler of Design -to-Robotic -Operation

Alexander Liu Cheng

Abstract

This essay promotes Artificial Intelligence (AI) via Machine Learning (ML) as a fundamental enabler of technically intelligent built-environments. It does this by detailing ML's successful application within three deployment domains: (1) Human Activity Recognition, (2) Object as well as Facial-Identity and -Expression Recognition, and (3) Speech and Voice-Command Recognition. With respect to the first, the essay details previously developed ML mechanisms implemented via Support Vector Machine and k-Nearest Neighbor classifiers capable of recognizing a variety of physical human activities, which enables the built-environment to engage with the occupant(s) in a highly informed manner. With respect to the second, it details three previously developed ML mechanisms implemented individually via (i) BerryNet—for Object Recognition; (ii) TensorFlow—for Facial-Identity Recognition; and (3) Cloud Vision API—for Facial-Expression Recognition; all of which enable the built-environment to identify and to differentiate between non-human and human objects as well as to ascertain the latter's corresponding identities and possible mood-states. Finally, and with respect to the third, it details a presently developed ML mechanism implemented via Cloud Speech-to-Text that enables the transcription of spoken speech—in several languages—into string text used to trigger pertinent events within the built-environment. The sophistication of said ML mechanisms collectively imbues the intelligent built-environment with a continuously and dynamically adaptive character that is central to Design-to-Robotic-Operation (D2RO), which is the Architecture-informed and Information and Communication Technologies (ICTs)-based component of a Design-to-Robotic-Production & -Operation (D2RP&O) framework that represents an alternative to existing intelligent built-environment paradigms.

Keywords

Design-to-Robotic-Operation, Machine Learning, Human Activity Recognition, Computer Vision, Voice Recognition

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

Introduction

Intelligence in the built-environment as a discourse began in the late 60s and early 70s (Cook, 1970, 1972; Eastman, 1972; Negroponte, 1969, 1975; Pask, 1975a, 1975b). Due to the rudimentary state and forbidding costs of Information and Communication Technologies (ICTs) during this period, explorations were principally limited to theoretical and/or hypothetical. But over the next two decades, and driven by increasingly sophisticated and accessible ICTs, explorations gradually produced physical implementations. From said nascent period throughout early physical implementations, two main emphases emerged within the same discourse: one centered around the technical context and the other around the architectural.

With respect to the technical, Ambient Intelligence (AmI) was coined in the late 90s to describe a vision of a future digital living room, a built-environment whose ICTs imbued its dwelling space with serviceable intelligence to the benefit of its occupant(s) (Zelkha et al., 1998). Within AmI a further specialized domain developed, i.e., that of Ambient Assisted Living—or Active and Assisted Living—(AAL), which framed its inquiry around the promotion of quality of life as well as the prolongation of independence with respect to Activities of Daily Living (ADLs) among the elderly via technical assistance. By the first decade of the 21st century, AmI and AAL were established and proliferating topics within the fields of Computer Science and related Engineerings (Lindgren et al., 2016; Paz Santana et al., 2017), Architectural Engineering (Bock et al., 2015; Georgoulas et al., 2014), and—indirectly—in the Medical Sciences (Acampora et al., 2013).

With respect to the architectural, and beginning with Cedric Price's pioneering Generator Project and corresponding programs by John and Julia Frazer (Steenson, 2014) in the late 70s, notions of interaction between non-human and human agents in the built-environment began to be envisioned. For example, in Price's project, architecture was conceived as a set of interchangeable subsystems integrated into a unifying computer system, which enabled a reconfigurability sensitive to function. More importantly, both Price and the Frazers intended for the system itself to suggest its own reconfigurations, denoting non-human agency in the built-environment. Although the Generator Project was never realized, it became the de facto first instance of a subset field in Architecture concerned with bi-directional communication and interaction between non-human and human agents in the built-environment, viz., Interactive Architecture (IA) (Fox, 2010; Oosterhuis, 2012) first and Adaptive Architecture (AA) (Jaskiewicz, 2013; Kolarevic, 2014) later, which—like AmI—have also proliferated in the 21st century.

The proliferation of intelligence in the built-environment with respect to AmI/AAL surpasses that of IA/AA in terms of technical complexity, reliability, and performance. This is due to their differing emphases, with the technical focusing on ICTs and corresponding services and the architectural on spatial experience, materiality, function, and form. That is, the technical proliferated alongside sustained development of ICTs over decades in ways that the architectural could not, at least not with the same affinity and immediacy. Nevertheless, technical sophistication or lack thereof alone has not necessarily guaranteed or disqualified contributions in the discourse. Indeed, principally technical as well as principally architectural explorations have both independently identified key effective as well as affective desiderata common to built-environments—intelligent or otherwise—construed as successful with respect to function as well as to spatial experience. This consideration includes a caveat: while both the technical as well as the architectural have yielded independent contributions, these have been otherwise limited by the lack of mutually provided input and/or feedback. However, the promise of solutions yielded by both principally technical AmI/AAL and principally architectural

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

IA/AA explorations have been unwittingly and invariably limited by the rigid and increasingly outdated character of their complementing frameworks. This is because the sophistication of a system depends on that of its mutually complementing subsystems; and two or more subsystems may not mutually complement, sustain, and/or support one another adequately if their levels of development and sophistication do not correspond (Milgrom, 1990). More succinctly expressed: at present, the architectural does not correspond to the technically superior AmI/AAL, while the technical does not correspond to the architecturally superior IA/AA. Consequently, a different design paradigm / framework is required in order to enable comprehensively and cohesively intelligent built-environments with corresponding levels of technical and architectural sophistication.
The work detailed in this essay belongs to the Design-to-Robotic-Operation (D2RO) component of the Design-to-Robotic-Production & -Operation (D2RP&O) framework. D2RP&O considers the technical as well as the architectural in conjunction from the early stages of the design and development processes, where the built-environment is construed as a highly sophisticated and integrated Cyber-Physical System (CPS) (Rajkumar et al., 2010) consisting of mutually informing computational and physical mechanisms that operate cooperatively and continuously via a highly heterogeneous, partially meshed, and self-healing Wireless Sensor and Actuator Network (WSAN) (Yang, 2014). Via a series of limited and progressively complex proof-of-concept implementations (Liu Cheng, 2016; Liu Cheng et al., 2018; Liu Cheng and Bier, 2016a, 2016b, 2018; Liu Cheng, Bier, Latorre et al., 2017; Liu Cheng, Bier, Mostafavi, 2017), the feasibility and promise of D2RP&O in general and D2RPO in particular have been demonstrated. In this essay, two previously and one presently developed core Machine Learning (ML) mechanisms are detailed in order to assert the promise of Artificial Intelligence (AI) as a fundamental enabler of intelligent built-environments: (I) ML and Human Activity Recognition (HAR), (II) ML and Object as well as Facial-Identity and -Expression Recognition, and (III) ML and Speech and Voice-Command Recognition. These mechanisms are part of a prescriptive System Architecture intended to serve as the technical backbone of highly sophisticated (i.e., artificially intelligent, intuitively adaptable, and continuously evolving, etc.) intelligent built-environments.

Machine Learning and Human Activity Recognition

HAR enhances the built-environment's ability to respond adequately to the daily habits of the occupant(s). It enables said environment to build an accurate activity profile that informs proactive intervention routines intended to promote well-being. For example, via HAR a built-environment may prompt the occupant(s) to engage in physical activity when prolonged periods of inactivity have been detected. Furthermore, ventilation systems may be engaged whenever HAR and temperature / humidity sensors integrated in the built-environment detect an increase of interior temperature correlated with high physical activity. As with all other mechanisms within the System Architecture, HAR increases the resolution of the information that the built-environment receives as sensed input, which is directly correlated with the quality and pertinence of the actuated output.

In this section, previously developed (see Liu Cheng, Bier, Latorre et al., 2017) HAR mechanisms are detailed. These mechanisms integrate both cloud-based as well as localized ML capabilities in order to ascertain robustness and resilience. Whenever possible, ML processes are locally and dynamically executed via ad hoc node-clustering. But should this prove impossible either due to failure or unavailability of adequate resources, cloud-based ML services are used. More specifically, two ML mechanisms are integrated into the prescribed System Architecture: (1) a localized ad hoc cluster system based on open-source and purpose-written Python scripts, and (2) a simulated cloud-based analytics service using MathWorks® MATLAB™. Both mechanisms in this system use accelerometer data streamed from a smartphone and each uses polynomial programming of Support Vector

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

Machine (SVM) and k-Nearest Neighbor (k-NN) classifiers.

Due to their evolving and resilient characters, ML classifiers have been implemented in a variety of applications built on WSANs (Alsheikh et al., 2014). HAR, as one such application, has successfully exploited said classifiers in the last five years (see, for example, Andreu and Angelov, 2013; Villa, 2012; Xiao and Lu, 2015). However, due to the cost-effective and low energy-consumption character typical of WSAN nodes, computational processing with respect to feature extraction has been considerably limited (Salomons et al., 2016). The implementation in question overcomes this limitation by instantiating ad hoc clusters consisting of a variety of high-performance nodes. Furthermore, several clusters may be instantiated simultaneously in order to enable parallel high-performance information processing activities. The system's clustering mechanism uses the Message Passing Interface (MPI) standard via MPI for Python (mpi4py) (Dalcin et al., 2011). Another way to overcome this limitation—and one also implemented—is to avoid it altogether by outsourcing all high-performance information processing to cloud-based ML services. But there are a number of limitations with this approach. The first, and perhaps the most salient, is the cost incurred by including proprietary services in any proposed intelligent built-environment solution. A second yet no less important limitation may be the impact to the solution's resilience. That is to say, should said built-environment lose access to the Internet, it would be incapable of generating classification models.

In the local mechanism, a script based on pyOSC (V2_Lab, 2008) is first written to receive OSC data from any device and application capable of broadcasting in said protocol. While all the WiFi-enabled nodes in the system's WSAN have the capacity to receive this data-streaming, only one of the nodes of the cluster instantiated to generate classification models stores it locally and streams it to a cloud-based data visualization service. Should the receiving node fail, another high-performance node replaces it automatically. The proposed solution uses a smartphone (ML for HAR has typically used gyroscopic / accelerometer data collected via portable devices—see Anguita et al., 2013; Ortiz, 2015—or via sensor-fusion—see Palumbo et al., 2016), and the script in question proceeds to parse and to reduce the noise in the received data in order to generate a robust dataset. At this point the dataset is processed through two ML scripts based on scikit-learn (Buitinck et al., 2013; Pedregosa et al., 2017), one for SVM and another for k-NN classification models. In this particular implementation, the SVM model attained a 95.71% HAR prediction success rate and the k-NN model a 97.85%. The success rate attained by the local k-NN model was 2% higher than its counterpart model generated externally with MATLAB (i.e., 95.9%), while the local SVM model was 2% less successful than its MATLAB counterpart (see figure 1). This is a strong argument to prefer the open-source alternative.

Machine Learning and Object as well as Facial-Identity and -Expression Recognition

Computer Vision enables the built-environment to recognize the object and persons within it, which is a pertinent prerequisite for actuations that involve interaction with them. For example, if an object has collapsed within the built-environment, whether emergency intervention and notification protocols be initiated or not would depend on whether said object was a person or not. Object Recognition enables the built-environment to do this. Moreover, if the collapsed object has been detected to be a person, perhaps the identity and facial expression of the person would serve as indicators of the nature (e.g., intentional, accidental, etc.) of this fall. This is where Facial-Identity and -Expressions recognition plays a crucial role. These three Computer Vision features combined enable the built-environment so see its context and to corroborate phenomena as perceived by

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

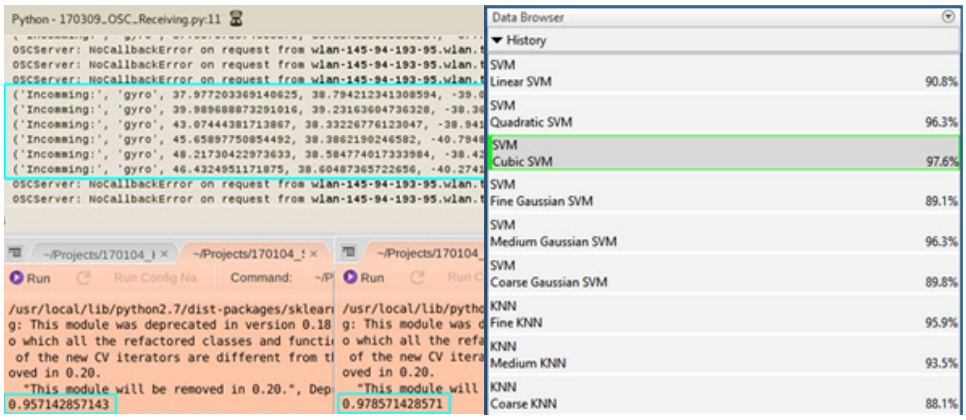

Figure 2. Left: Receiving OSC-data; bottom: 95.71% prediction success with respect to HAR via SVM (left) vs. 97.85% via k-NN (right). Right: MATLAB ML results

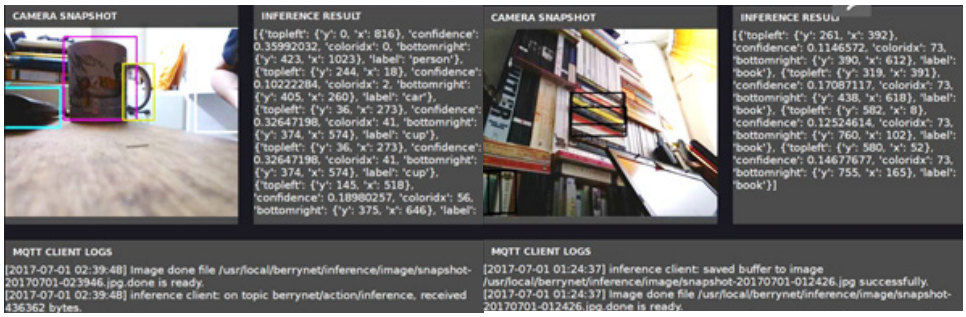

Figure 2. Multiple-object detection via BerryNet (DT42©, Ltd., 2017); Left: 'person', 'cup'; Right: 'books'

other sensing mechanisms.

In this section, three previously developed mechanisms are detailed. The first pertains strictly to Object Recognition (see Liu Cheng, Bier, Mostafavi, 2017); the second and third to Facial-Identity and -Expression Recognition (forthcoming publication—see Notes), respectively. The first mechanism is implemented with open-source BerryNet® (DT42©, Ltd., 2017), which is built with a classification model—viz., Inception® ver. 3 (Szegedy et al., 2015)—as well as a detection model –viz., TinyYOLO® (Redmon and Farhadi, 2016). The classification model uses Convolutional Neural Networks (CNNs), which are at the forefront of ML research (Szegedy et al., 2015). An advantage of BerryNet® is that it is a fully implementable gateway on a cluster of RPi3s. On an individual RPi3, the inference process is slow, requiring a delay between object-recognition sessions. This situation is ameliorated by the dynamic clustering feature of the WSAN. Another benefit-cum-limitation is that BerryNet®'s classification and detection models are pretrained, which avoids the need to generate said models locally.

The Object Recognition mechanism (see figure 2) in the D2RO System Architecture is intended to be deployed across a variety of cameras in the overall built-environment, and that instances of detection were to be cross-referenced to minimize false positives. In order to implement this setup, each RPi3 node in the WSAN is equipped with a low-cost Raspberry Pi Camera® V2.1, then BerryNet® is installed in every node and the inference mechanism tested individually. The next step is to enable the nodes to share their detection results, which could be done via WiFi. Nevertheless, in order to reduce energy-consumption for every object-detection cross-referencing instance, ZigBee is preferred. In order to enable ZigBee on BerryNet®'s detection_server.py and classify_server.py

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

were modified and made compliant with python-xbee (n.io Innovation©, LLC, 2017).
The second and third mechanisms—i.e., Facial-identity and -Expression Recognition—are implemented via two independent yet interrelated components. The first is implemented locally via Google Brain®'s TensorFlow™ (TensorFlow™, 2018): while the second via Google Cloud Platform®'s Cloud Vision API (Google Cloud Platform®, 2018b). In the implementation of the first component, TensorFlow™ is installed on a Linux (Ubuntu) virtual environment and executed in Python. During execution of its Multi-Task Convolutional Neural Network (MTCNN) face detection model, TensorFlow™ requests the user to capture images of his/her face from a variety of positions, orientations, and angles. After completing this phase, facial identity recognition is successfully tested real-time (see figure 3, Top). In the implementation of the second component, Python is used to integrate the services of Cloud Vision API into the inherited WSAN. The same visual input is provided to both components to yield a correlated recognition of an identity as well as of a facial expression (see figure 3).

Machine Learning and Speech and Voice-Command Recognition

Speech and Voice-Command Recognition enable the built-environment to listen to the occupant(s). Perceived speech and subsequent processing of command may serve to override and/or to adjustment automatic actuations effected by the built-environment according to the preferences suggested by occupant-profiles. They may also serve to explicitly engage an actuation or to feed information to the system. For example, should the mechanism that ascertains comfortable temperature and humidity conditions within the built-environment actuate against the occupants wishes for that particular moment, he/she could verbally command the built-environment to stop. In a different scenario, one where the occupant is in a state of emergency, he/she could verbally ask the built-environment to call for help (see figure 4, Top).

In this section, a purpose-built implementation of this mechanism is detailed. This mechanism is designed to work in tandem with but independent of a previously implemented Alexa Voice Service (Amazon®, 2017b) (AVS) mechanism (see Liu Cheng and Bier, 2018). The AVS mechanism enabled the built-environment to access an array of preset voice commands made available by Amazon®, and to connect the former's services to the Internet. However, the usefulness of AVS centered around consumer-based services online, not within the local built-environment. Admittedly, AVS may be extended to work with customized commands within local built-environments via Alexa Skills Kit (Amazon®, 2017a), but these must rely on Amazon®'s developer and cloud services. Although AVS does provide advantages to the services provided by the local built-environment, a more flexible and easy to customize Speech and Voice-Command Recognition mechanism is preferred for the control of local actuations. Via a Python script, this mechanism first uses PyAudio (Pham, 2017) to listen to an initial key trigger command and to process following spoken speech locally (compare to AVS's remote processing), and then sends the result to Google Cloud Platform®'s Cloud Speech-to-Text (Google Cloud Platform®, 2018a) to generate a string text in return. This text—now effectively a local variable—is then used to trigger particular events in the local built-environment. Since the trigger mechanism is locally programmed, there is no limit—beyond that of the system's storage capacity—as to how many new speech-to-actuation correlations may be configured (see figure 4).

Conclusions

The purpose of this essay is to promote AI's role in the realization of highly sophisticated intelligent

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

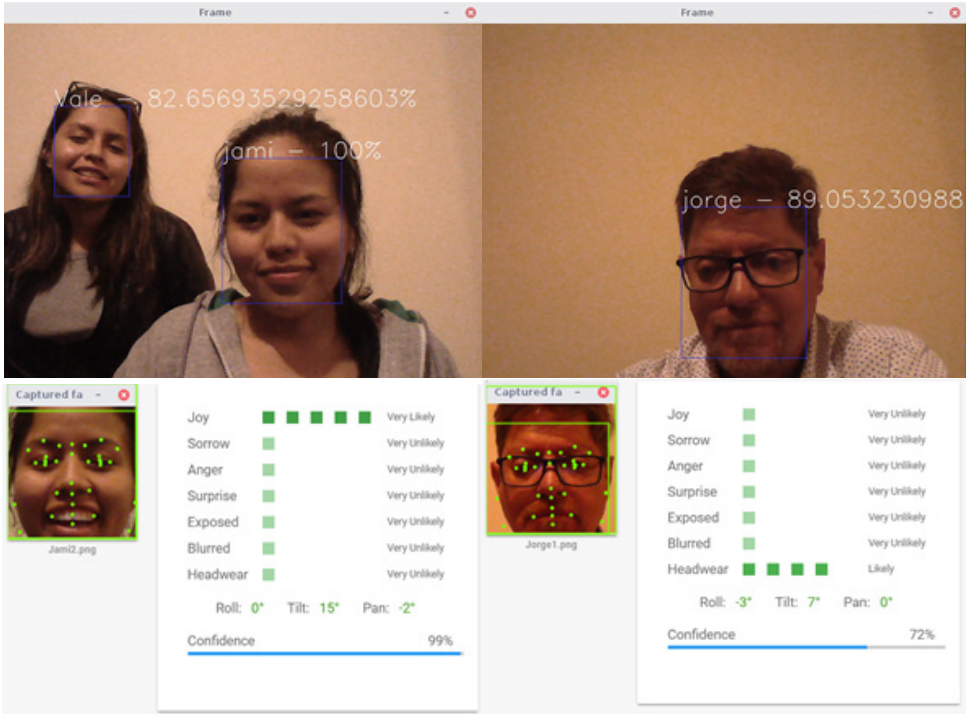

Figure 3.
Top: Facial-Identity Recognition via TensorFlow (TensorFlow™, 2018). Bottom: Facial-Expression Recognition via Cloud Vision API's (Google Cloud Platform®, 2018b).

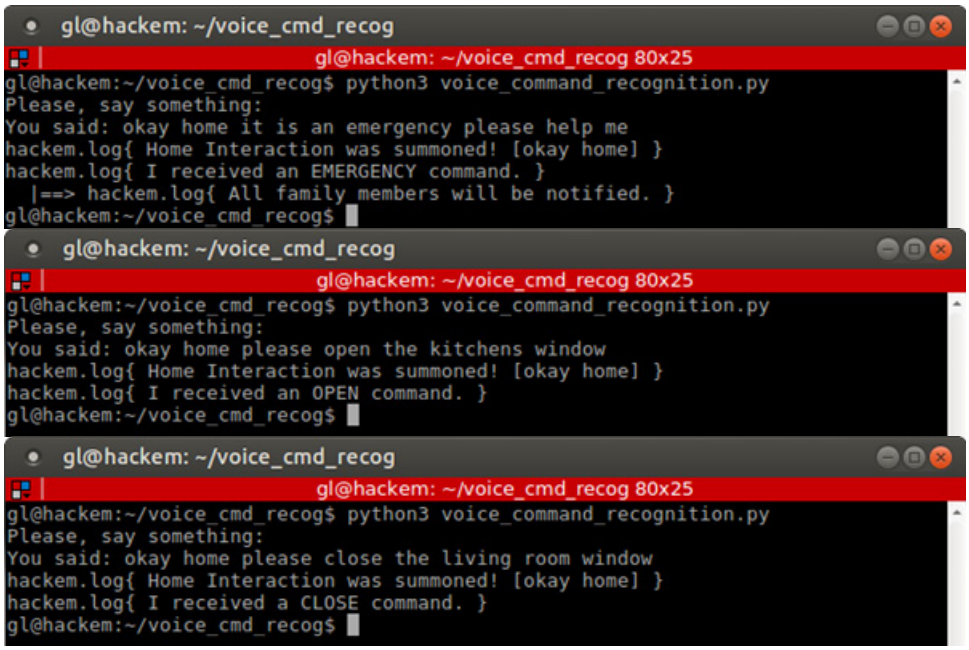

Figure 4.
Speech and Voice-Command Recognition via Cloud Speech-to-Text (Google Cloud Platform®, 2018a)

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

built-environments by illustrating three fundamental ML mechanisms in D2RO's prescriptive System Architecture. Each of the described mechanisms highlights the sophisticated way via which ML processes seemingly random high-volume data to yield meaningful results. These mechanisms are also highlighted because no viable non-ML counterparts exist, at least not ones capable of inherent evolution and increase in precision over time. AI via ML enables the built-environment to detect patterns otherwise undetectable, patterns that mean the difference between an intuitive solution and a cumbersome imposition. Especially in the context of intelligent built-environments, this difference and the likes determine user acceptability as well as system effectiveness with respect to promotion of occupant well-being. The intelligent built-environment without AI is simply not intelligent enough.

Acknowledgements

This paper has profited from the invaluable feedback provided by my doctoral supervisor, Dr.-Ing. Henriette Bier, Assistant Professor at Robotic Building, Faculty of Architecture and the Built Environment, Technische Universiteit Delft (Delft, The Netherlands). Moreover, I would like to acknowledge Galoget Latorre, Researcher at Facultad de Ingeniería de Sistemas, Escuela Politécnica Nacional (Quito, Ecuador), for his assistance in the testing and development of the Speech and Voice-Command Recognition mechanism.

Notes

Parts of the Introduction and Sections 1 and 2 are adapted from or reference a forthcoming chapter in the first volume of Springer Series in Adaptive Environments: Robotic Building (2018).

References

Acampora, G., Cook, D. J., Rashidi, P. and Vasilakos, A. V. (2013) 'A Survey on Ambient Intelligence in Healthcare', Proceedings of the IEEE, vol. 101, no. 12, pp. 2470–2494.

Alsheikh, M. A., Lin, S., Niyato, D. and Tan, H.-P. (2014) 'Machine Learning in Wireless Sensor Networks: Algorithms, Strategies, and Applications', IEEE Communications Surveys & Tutorials, vol. 16, no. 4, pp. 1996–2018.

Amazon® (2017a) Alexa Skills Kit: Teach Alexa a new skill and reach millions of customers [Online]. Available at https://developer.amazon.com/alexa-skills-kit (Accessed 19 April 2017).

Amazon® (2017b) Alexa Voice Service: Bring Alexa to your connected products with Amazon's intelligent cloud-based voice service [Online]. Available at https://developer.amazon.com/alexa-voice-service (Accessed 19 April 2017).

Andreu, J. and Angelov, P. (2013) 'An evolving machine learning method for human activity recognition systems', Journal of Ambient Intelligence and Humanized Computing, vol. 4, no. 2, pp. 195–206 [Online]. DOI: 10.1007/s12652-011-0068-9.

Anguita, D., Ghio, A., Oneto, L., Parra, X. and Reyes-Ortiz, J. L. (2013) 'A Public Domain Dataset for Human Activity Recognition Using Smartphones', Proceedings of the 21th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2013.

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

Bruges, Belgium, 24-26 April.
Bock, T., Güttler, J., Georgoulas, C. and Linner, T. (2015) 'The Development of Intra-House Mobility, Logistics and Transfer Solutions in PASSAge', Journal of Robotics and Mechatronics, vol. 27, no. 1, p. 108.

Buitinck, L., Louppe, G., Blondel, M., Pedregosa, F., Mueller, A., Grisel, O., Niculae, V., Prettenhofer, P., Gramfort, A., Grobler, J., Layton, R., Vanderplas, J., Joly, A., Holt, B. and Varoquaux, G. (2013/09/01) API design for machine learning software: experiences from the scikit-learn project [Online]. Available at http://arxiv.org/pdf/1309.0238.

Cook, P. (1970) Experimental architecture, New York, Universe Books.

Cook, P. (1972) Archigram, London, Studio Vista.

Dalcin, L. D., Paz, R. R., Kler, P. A. and Cosimo, A. (2011) 'Parallel distributed computing using Python', Advances in Water Resources, vol. 34, no. 9, pp. 1124–1139.

DT42©, Ltd. (2017) BerryNet®: Deep learning gateway on Raspberry Pi [Online]. Available at https://github.com/DT42/BerryNet (Accessed 21 June 2017).

Eastman, C. M. (1972) Adaptive conditional architecture, Pittsburgh, Institute of Physical Planning, School of Urban and Public Affairs, Carnegie-Mellon University.

Fox, M. (2010) 'Catching up with the Past: A Small Contribution to a Long History of Interactive Environments', 5-18 Pages / FOOTPRINT, Issue # 6 | Spring 2010 | Digitally-Driven Architecture.

Georgoulas, C., Raza, A., Güttler, J., Linner, T. and Bock, T. (2014) 'Home Environment Interaction via Service Robots and the Leap Motion Controller', Proceedings of the 31st International Symposium on Automation and Robotics in Construction (ISARC 2014).

Google Cloud Platform® (2018a) Cloud Speech-to-Text: Speech to text conversion powered by machine learning and available for short or long-form audio [Online], Google®. Available at https://cloud.google.com/speech-to-text/ (Accessed 25 April 2018).

Google Cloud Platform® (2018b) Cloud Vision API: Derive insight from images with our powerful Cloud Vision API [Online], Google®. Available at https://cloud.google.com/vision/ (Accessed 20 April 2018).

Jaskiewicz, T. (2013) Towards a Methodology for Complex Adaptive Interactive Architecture.

Kolarevic, B. (2014) 'Outlook: Adaptive Architecture: Low-Tech, High-Tech, or Both?', in Kretzer, M. and Hovestadt, L. (eds) ALIVE: Advancements in adaptive architecture, Basel/Berlin/Boston, Birkhäuser, pp. 148–157.

Lindgren, H., Paz Santana, J. F. d., Novais, P., Fernández-Caballero, A., Yoe, H., Ramirez, A. J. and Villarrubia, G. (2016) Ambient intelligence: Software and applications -- 7th International Symposium on Ambient Intelligence (ISAmI 2016), Switzerland, Springer.

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 201

Liu Cheng, A. (2016) 'Towards embedding high-resolution intelligence into the built-environment', Archidoct, vol. 4, no. 1, pp. 29–40 [Online]. Available at http://www.enhsa.net/archidoct/Issues/ArchiDoct_vol4_iss1.pdf (Accessed 15 September 2016).

Liu Cheng, A. and Bier, H. H. (2016a) 'Adaptive Building-Skin Components as Context-Aware Nodes in an Extended Cyber-Physical Network', Proceedings of the 3rd IEEE World Forum on Internet of Things. Reston, VA, USA, 12-14 December, IEEE, pp. 257–262.

Liu Cheng, A. and Bier, H. H. (2016b) 'An Extended Ambient Intelligence Implementation for Enhanced Human-Space Interaction', Proceedings of the 33rd International Symposium on Automation and Robotics in Construction (ISARC 2016). Auburn, Alabama, 18-21 July, pp. 778–786.

Liu Cheng, A. and Bier, H. H. (2018) 'Extension of a High-Resolution Intelligence Implementation via Design-to-Robotic-Production and -Operation strategies', Proceedings of the 35th International Symposium on Automation and Robotics in Construction (ISARC 2018). Berlin, Germany, July 20-25. Berlin, Germany.

Liu Cheng, A., Bier, H. H., Latorre, G., Kemper, B. and Fischer, D. (2017) 'A High-Resolution Intelligence Implementation based on Design-to-Robotic-Production and -Operation strategies', Proceedings of the 34th International Symposium on Automation and Robotics in Construction (ISARC 2017). Taipei, Taiwan (R.O.C.), 28 June - 1 July.

Liu Cheng, A., Bier, H. H. and Mostafavi, S. (2017) 'Deep Learning Object-Recognition in a Design-to-Robotic-Production and -Operation Implementation', Proceedings of the 2nd IEEE Ecuador Technical Chapters Meeting 2017. Salinas, Ecuador, 18-20 October. Guayaquil, Ecuador.

Liu Cheng, A., Bier, H. H. and Mostafavi, S. (2018) 'Integration of a Wearable Interface in a Design-to-Robotic-Production and -Operation Development', Proceedings of the 35th International Symposium on Automation and Robotics in Construction (ISARC 2018). Berlin, Germany, July 20-25. Berlin, Germany.

Milgrom, P. R. (1990) 'The economics of modern manufacturing: technology, strategy, and organization', The American Economic Review, vol. 80, no. 3, pp. 511–528.

n.io Innovation©, LLC (2017) python-xbee: Python tools for working with XBee radios [Online]. Available at https://github.com/nioinnovation/python-xbee (Accessed 1 June 2017).

Negroponte, N. (1969) 'Toward a Theory of Architecture Machines', Journal of Architectural Education (1947-1974), vol. 23, no. 2, p. 9.

Negroponte, N. (1975) 'The architecture machine', Computer-Aided Design, vol. 7, no. 3, pp. 190–195.

Oosterhuis, K. (2012) Hyperbody: First decade of interactive architecture, Heijningen, Jap Sam Books.

Ortiz, J. L. R. (2015) Smartphone-based human activity recognition, Cham, Springer.

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

Palumbo, F., Gallicchio, C., Pucci, R. and Micheli, A. (2016) 'Human activity recognition using multisensor data fusion based on Reservoir Computing', Journal of Ambient Intelligence and Smart Environments, vol. 8, no. 2, pp. 87–107.

Pask, G. (1975a) Conversation, cognition and learning: A cybernetic theory and methodology, Amsterdam, New York, Elsevier.

Pask, G. (1975b) The cybernetics of human learning and performance: A guide to theory and research, London, Hutchinson Educational.

Paz Santana, J. F. d., Julian, V., Villarrubia, G., Marreiros, G. and Novais, P. (2017) Ambient intelligence: Software and applications : 8th International Symposium on Ambient Intelligence (ISAmI 2017), Cham, Switzerland, Springer.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Louppe, G., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, É. (2017/01/10) Scikit-learn: Machine Learning in Python [Online]. Available at http://arxiv.org/pdf/1201.0490.

Pham, H. (2017) PyAudio [Online], Massachusetts Institute of Technology. Available at http://people.csail.mit.edu/hubert/pyaudio/ (Accessed 29 May 2018).

Rajkumar, R., Lee, I., Sha, L. and Stankovic, J., eds. (2010) Cyber-physical systems: the next computing revolution, ACM.

Redmon, J. and Farhadi, A. (2016/12/25) YOLO9000: Better, Faster, Stronger [Online]. Available at http://arxiv.org/pdf/1612.08242.

Salomons, E. L., Havinga, P. J. M. and van Leeuwen, H. (2016) 'Inferring Human Activity Recognition with Ambient Sound on Wireless Sensor Nodes', Sensors (Basel, Switzerland), vol. 16, no. 10.

Steenson, M. W. (2014) Architectures of information: Christopher Alexander, Cedric Price, and Nicholas Negroponte and MIT's Architecture Machine Group.

Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J. and Wojna, Z. (2015/12/11) Rethinking the Inception Architecture for Computer Vision [Online]. Available at http://arxiv.org/pdf/1512.00567.

TensorFlow™ (2018) An open source machine learning framework for everyone [Online], Google Brain®. Available at https://www.tensorflow.org/ (Accessed 20 April 2018).

V2_Lab (2008) pyOSC [Online]. Available at https://github.com/ptone/pyosc (Accessed 10 June 2016).

Villa, A. E. (2012) Artificial neural networks and machine learning-- ICANN 2012: 22nd International Conference on Artificial Neural Networks, Lausanne, Switzerland, September 11-14, 2012, Proceedings, Berlin, New York, Springer.

Xiao, W. and Lu, Y. (2015) 'Daily Human Physical Activity Recognition Based on Kernel Discriminant

ISSN 2309-0103
www.enhsa.net/archidoct
Vol. 6 (1) / July 2018

Analysis and Extreme Learning Machine', Mathematical Problems in Engineering, vol. 2015, no. 1, pp. 1–8.
Yang, S.-H. (2014) Wireless Sensor Networks Principles, Design and Applications, London, Springer London.
Zelkha, E., Epstein, B., Birrell, S. and Dodsworth, C. (1998) 'From Devices to 'Ambient Intelligence': The Transformation of Consumer Electronics (Conference Keynote)', Digital Living Room Conference.