Jan Willmann
Philippe Block
Marco Hutter
Kendra Byrne
Tim Schork
Editors

Robotic Fabrication in Architecture, Art and Design 2018

Springer

Robotic Fabrication in Architecture, Art and Design 2018

Jan Willmann · Philippe Block
Marco Hutter · Kendra Byrne
Tim Schork
Editors

Robotic Fabrication in Architecture, Art and Design 2018

Foreword by Sigrid Brell-Çokcan and Johannes Braumann,
Association for Robots in Architecture

Springer

Editors
Jan Willmann
Faculty of Art and Design
Bauhaus-Universität Weimar
Weimar, Germany

Philippe Block
Department of Architecture
Swiss Federal Institute of Technology
Zurich, Switzerland

Marco Hutter
Department of Mechanical and Process Engineering
Swiss Federal Institute of Technology
Zurich, Switzerland

Kendra Byrne
San Francisco, CA, USA

Tim Schork
Faculty of Design, Architecture and Building
University of Technology, Sydney
Sydney, NSW, Australia

Funded by KUKA Robotics Germany and the Association for Robots in Architecture

ISBN 978-3-319-92293-5 ISBN 978-3-319-92294-2 (eBook)
https://doi.org/10.1007/978-3-319-92294-2

Library of Congress Control Number: 2018950095

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Foreword by the Association for Robots in Architecture

Since the beginning of the Association for Robots in Architecture, it has been our goal to promote collaboration within the community, but also with partners from industry and cross-disciplines. Robots in Architecture is proud to be associated with leading research organizations within architecture, such as ACADIA, or eCAADe but the scope of work that is being done by the community is starting to exceed the field of architecture into many other new domains, and we feel that it is important to accompany such steps. One measure to ensure the exchange across disciplines and other domains is to open the community to disciplines such as engineering and robotics by establishing common platforms where people can meet and exchange their ideas and research. Together with Springer, the Association for Robots in Architecture has therefore established a new Journal for Construction Robotics with the first published issue at the end of 2017 to foster collaborative papers and high-quality research in architecture.

Another goal for the Association was to join forces with associations in Robotics. In 2016, Robots in Architecture joined euRobotics AISBL, the largest public–private partnership involving robotics in Europe. Sigrid Brell-Çokcan co-established a new Topic Group on Construction Robotics, acting alongside the other established 30 topic groups within euRobotics, ranging from wearables, bio-inspired robotics, health care, mining to infrastructure. In 2018, Sigrid has joined the board of directors.

Through networks such as euRobotics, it is not only possible to promote a field of research, but also to actively shape policy, so that the importance of Construction Robotics is recognized, and appropriate research funding is allocated to relevant topics. Through the Multi-Annual Roadmap in research for the EU and its Horizon 2020 programme, these initiatives not only reach a chosen few researchers in academia, but also a wide range of commercial and non-commercial research institutions, robotics developers and users alike in Europe and beyond.

This year, we recognize the importance of such public–private partnerships and euRobotics by presenting euRobotics chairman Bernd Liepert with the Rob|Arch Community Contribution award. It is our goal that more researchers within the Robots in Architecture community will reach out to large-scale research and believe strongly that euRobotics AISBL is a prime example on how to combine economic with academic and political interests, fostering the common goal of creating robotic innovation.

Innovation is also one of the core qualities we are looking for when selecting awardees. This year's Pioneering Achievement and Pioneering Industry Award goes to two very different architectural companies, who have greatly facilitated the potential of robotic processes in their work.

The internationally highly reputed architectural office Snøhetta was one of the first architectural companies to invest in robotic arms, joining us for our first KUKA|prc workshop in 2010 at the Advances in Architectural Geometry Conference in Vienna. For them, the robot has been an important tool for prototyping new approaches and design, placed closely to the open office in Oslo.

The second awardee is Branch Technology from Tennessee in the USA, who have gone even further by not just using robots as CNC machines, but by developing their own robotic processes for large-scale robotic 3D printing. What we feel is special about Branch is that it is a company by architects who develop products for architects. As such, they did not stem from academia but from practice and therefore had to find investors to fund their ideas. Today they have a team with a wide variety of backgrounds and have realized a number of large-scale projects, collaborated with companies such as Foster + Partners and even won an award for their joint vision of future construction on Mars by the NASA. We believe that this drive embodies the true spirit of the community. We see in Branch a perfect role model where robotics lies at the core to enable technology-driven creativity in new business models.

By Yu Lei's research at Tsinghua University in China and his own professional workshop, we do not just honour a single person, who has made significant contributions to architectural robotics in China, but also the entire Chinese community of companies and researchers, where the past years have seen great developments and a surge of new ideas and initiatives, as was demonstrated at this year's CAADRIA conference organized by Prof. Xu Weiguo and the DADA community. While the potential for Construction Robotics in China is huge, there is also a great need for education and research and thus educators like Yu Lei are important trailblazers by sharing and starting new business models in architecture.

We also believe that it is important to recognize the researchers, without whose work into robotics we would not have today's sophisticated methods and machines at hand, and who create tomorrow's tools and processes today. Jonas Buchli, director of the ETH Zurich Agile & Dexterous Robotics Lab, is one such pioneer whose work spans across many disciplines—from neurobiology and human locomotion to service robotics and also architecture, where has been a Principal Investigator of the NCCR Digital Fabrication, developing the In situ Fabricator (IF), with the goal of enabling the machine to autonomously perform precise mobile manipulation tasks in unstructured and unpredictable environments. We recognize his work with a Pioneering Research Award.

In 2018, the Rob|Arch conference is back in Europe for the first time since Vienna in 2012 and is being hosted at ETH Zurich. ETH Zurich, in particular through the work of Gramazio Kohler Research, has been a central part of the research community and further solidified that status with the creation of the NCCR Digital Fabrication, a multi-disciplinary research cluster with more than 14 professors. We see ETH Zurich and the NCCR Digital Fabrication as a prime example how research can happen on a very large scale, with highly interdisciplinary and diverse research teams. At the same time, we believe that innovative research can also be done at a smaller scale, as it is demonstrated by this year's 15 Rob|Arch conference workshops, involving tutors from more than 25 institutions, that are hosted centrally in ETH Zurich's Robotic Fabrication Laboratory.

An effort like Rob|Arch 2018 is only possible when many people work together towards a common goal. We would like to thank our local hosts at ETH Zurich and the NCCR Digital Fabrication, the Scientific Board led by Jan Willmann and all minds and hands involved in setting up such a big event.

We are grateful to Matt Jezyk (Autodesk) and Alois Buchstab (KUKA) for their continuing and enthusiastic support of the community through their respective companies, and we would like to thank all sponsors of this conference – Arup, BCG, Sika, Erne, Moog, and Bachmann Engineering – for this year's collaboration in making Rob|Arch 2018 a success.

Sigrid Brell-Çokcan
Association for Robots in Architecture

Johannes Braumann
Association for Robots in Architecture

Preface

New Scientific Frontiers

The emergence of robotics with the creative sectors has led to an entirely new epistemology of collective making that is inextricably open and future-oriented. Challenged by increasingly complex technological and environmental problems, architects, designers, civil and process engineers, and roboticists are seeking novel practices of collaboration and exchange that deliberately overcome and dissolve traditional disciplinary boundaries. This collective approach to working with robots is not only revolutionizing how things are designed and made, but is fundamentally transforming the culture, politics and economics of the creative industries as a whole.

What distinguishes contemporary industrial robots from their industrial predecessors—and indeed from other contemporary computer-controlled devices—is their versatility. Like computers, today's robotic arms are suitable for a wide variety of tasks: they are "generic", open-ended, adaptable and not restricted to any particular application or disciplinary focus. This versatility allows them to be readily customized and programmed to suit a wide range of specific intentions, both at the material and conceptual levels. It has also allowed us to shift our perception of robots as mechanistic, utilitarian devices suited to standardized serial production, towards understanding them as creative tools for exploring, designing and realizing physical objects and the built environment. If the first robotic age—the age of industrial automation—vastly improved our physical productivity, the second robotic age will surely come to distinguish itself as a driver of creative capacity.

The present moment is ripe for connecting robot technology with imagination and materialization, inspiring new fundamental discoveries and opening new scientific frontiers. In fact, we have within reach access to volumes of information and centuries of knowledge about how to design and realize the material world. Aided by global digital connectedness, open-source ideals and collective encounters, robotics rejuvenates traditional disciplinary wisdom with entirely new practices of scientific collaboration and knowledge transfer. Now, more than ever, we are coming to understand that robotics research should not be bound by constricting disciplinary standards, constraints or ideologies lest we limit its potential. Yet to explore this unprecedented potential requires not only a technical grasp of robots' capabilities and limitations, but also an in-depth understanding of the disciplinary consequences of robotics research. With its theme of "Radical Cross-disciplinarity", Rob|Arch 2018 facilitates this understanding by encouraging novel scientific approaches, applications and collaborations, not just in robotics, but beyond.

Closing the Loop

The Rob|Arch conference series was first launched in Vienna, Austria, in 2012 by Sigrid Brell-Çokcan and Johannes Braumann, the founders of the Association for Robots in Architecture. Their purpose was to make industrial robots more accessible to the creative industries—including art, design and architecture—by sharing ideas, research results and technological developments. The series has since become a biannual tradition in the international community (travelling to Michigan, US, in 2014 and to Sydney, Australia, in 2016) and has decisively boosted the exchange and dynamics within.

In 2018, Rob|Arch lands at the Swiss Federal Institute of Technology in Zurich (ETH Zurich), marking an important milestone for the digital fabrication community: ETH Zurich is not only one of the leading international universities for technology and science, it is also the institution where the first industrial robotic fabrication laboratory for non-standard architectural fabrication processes was installed in 2005. Closing this loop gives us the opportunity to foster novel explorations and state-of-the-art knowledge, techniques and methods, while consolidating and advancing our collective understanding of the evolution and impact of robotics in art, design and architecture.

It is no coincidence that Rob|Arch 2018 is also co-hosted by the Swiss National Centre of Competence in Research (NCCR) in Digital Fabrication. Launched in 2014, the NCCR Digital Fabrication is itself a truly cross-disciplinary research platform meant to foster the seamless combination of digital technologies and physical building processes through cooperation and exchange beyond disciplinary boundaries.

Content and Contributions

The Rob|Arch 2018 publication features the most important contributions to the conference. Rather than featuring merely formalist or technicist robotic adventures, this publication goes beyond pure built outcome to forward fresh approaches to scientific innovation, knowledge exchange and cross-disciplinary collaboration.

This includes designers, artists and architects, and also—and increasingly—computation and robotics experts and builders, materials scientists and engineers, process and systems specialists and manufacturers, to name just a few. As a consequence, this book gathers exceptional, scientifically rigorous projects that not only transform the way we design and make, but which also build collaborative capacity in the field of robotic fabrication.

The structure of this publication addresses this "new territory" of collaborative research. Stepping beyond theoretical observation, it outlines five distinct epicentres of practical research, which range from design and simulation research to automated assembly and real-world applications. Robotics and material and structural engineering play an integral role in each of these five areas.

Chapter 1 ("Design and Simulation") discusses new computational approaches to image classification using neural networks, stochastic assembly and deep learning for robotic construction; it also presents procedural fabrication workflows and haptic programming techniques, automatic path planning methods, visual feedback techniques, and function representation models.

Novel materials and material processes for robotic fabrication are introduced in Chapter 2 ("Material and Processes"), including thermally tuned concrete panel printing, time-based material deposition, and digitally controlled concrete injection processes. This is complemented by research into the robotic manipulation of filament material and the automated control of material behaviour for spatial extrusion processes.

In Chapter 3 ("Construction and Structure"), the emphasis is on new robotic construction processes and structural applications, for example bespoke concrete reinforcement, highly versatile wood processing, automated band-saw cutting for complex timber structures, fabrication-aware methods for the realization of non-standard timber shells, and an advanced hybrid subtractive-additive approach to robotically construct double-curved concrete shells. Finally, the chapter presents a novel approach to the construction of jammed architectural structures.

Robotic control, machinery, tooling and fabrication are discussed in Chapter 4 ("Control and Fabrication"), involving tubular composite fabrication with the aid of robotic swarms, automated manufacturing of natural composites, 3D printing with clay on freeform moulds, choreographic robotic wood manipulation, aerial construction using a cyber-physical macro-material system, as well as adaptive robotic carving. Also outlined in this chapter are approaches for multi-mode hybrid fabrication, robotic extrusion of functionally graded building components, as well as of elements with non-standard topology, on-site robotic construction and additive manufacturing techniques for non-woven textiles.

The transfer to larger scales of real-world applications and practices is addressed in Chapter 5 ("Application and Practice"). Here we present automated slipforming for façade elements, robotic brick printing and stacking, robotic sewing of wooden shells, additive manufacturing of truss-shaped concrete pillars, and the realization of topology-optimized concrete structures using abrasive techniques. Large-scale bespoke timber frame construction and cooperative robotic brick assembly are also discussed.

Workshop Activities

Rob|Arch 2018 features a variety of formats and sessions to encourage creative dabbling and encounters with different research topics, practices and field-wide issues. Led by experts from academia, practice and industry, Rob|Arch 2018 workshops empower participants to learn and practise hands-on skills, and discuss cutting-edge fabrication techniques and trends with their peers in a collaborative environment.

This year's workshops offer a broad range of topics, including multiple robotic fabrication, industry-grade robotic programming using HAL, robotic real-time control using Grasshopper, robotic fabrication through the COMPAS framework, chainsawed wood joinery, cooperative robotic assembly of spatial timber structures, large-scale robotic construction, hybrid robotic 3D printing of concrete shell structures, autonomous robotic swarm systems, adaptive spatial 3D printing of space frame structures, automated assembly in constrained sites, mixed reality environments for complex steel structures, mixed reality simulation for collaborative design exploration, as well as an introduction to KUKA|prc for Dynamo.

Beyond Boundaries

Rob|Arch 2018 aims to bring the community ground-breaking approaches to robotic fabrication from the most innovative research laboratories in the world, all while illuminating alternative pathways to boosting cross-disciplinary research and exchange. This publication therefore highlights contributions that not only substantially advance the state-of-the-art in robotic fabrication, but also challenge the reputedly clear division between research, practice and industry.

It is our belief that effective knowledge transfer and exchange between different disciplines is crucial for the development of truly innovative and high-impact research in robotics, a priori, rather than a posteriori. Specifically, Rob|Arch 2018 looks at new paradigms of scientific collaboration, along with the challenges, risks and dynamics within this process. Given that our collective expertise includes autonomous control systems, advanced construction, collaborative design tools, computerized materials and structures, adaptive sensing and actuation, on-site and cooperative robotics, machine-learning, human–machine interaction, large-scale robotic fabrication and networked workflows (the list goes on), we can no longer discuss cross-disciplinarity, cooperation and collaboration in abstract terms. Doing so would be utterly inadequate to address the manifold cultures and practices of robotics that have emerged to master the increasingly complex technological and environmental challenges we face today.

While we have observed a growing capacity for knowledge transfer and exchange in Rob|Arch submissions with each subsequent edition of the conference, this year the blurring of disciplinary boundaries between creative-, scientific- and practice-based domains is particularly significant. We view this as a sign that complex problems cannot be dealt with from a single disciplinary perspective alone.

Yet, while this blurring has yielded many new robotic explorations and real-world applications, these have not taken place uniformly. For example, the fields of intelligent computational design and simulation systems are particularly benefiting from an expanded set of collaborations and exchange between researchers and industry practitioners. Other areas that have especially benefitted from collaborative exchange include: advanced robotic control systems, and feedback processes that enable robots to adapt to different material conditions and changing environments. In all these cases, constant interaction and knowledge transfer between architects, designers, engineers and roboticists are pivotal, both as a result and as a catalysing instrument.

The fast pace of creative and scientific research documented by Rob|Arch is no doubt a result of the bringing-together of diverse disciplines, competences and cultures. Perhaps the emerging cross-disciplinary culture of robotic fabrication research will, through the collaboratively built future environment, one day yield a generational change in how we view the collaborative creative process more broadly. As Richard Sennett once described it: it stimulates a gathering of creative explorations similar to collective encounters that in the pre-machinic age used to be related with, and venerated for, all things man-made.

Acknowledgments

The Scientific Chairs would like to express their gratitude to the Conference Chairs, Fabio Gramazio and Matthias Kohler, for entrusting us with the development of Rob|Arch 2018. We would like to extend our gratitude to the Association for Robots in Architecture, namely Sigrid Brell-Çokcan and Johannes Braumann, for their invaluable support and commitment, and, above all, for the forming of a global (and cross-disciplinary) creative robotics community through the development and promotion of Rob|Arch. In addition, we would also like to thank Autodesk, KUKA, ARUP, Boston Consulting Group, Sika, ERNE, Moog and Bachmann Engineering who financially supported Rob|Arch 2018. Our sincere appreciation goes out as well to the Paper Committee; this conference and publication would not have been possible without their timeless effort and support.

The Scientific Chairs also wish to thank the National Centre of Competence in Research (NCCR) Digital Fabrication for co-hosting and supporting Rob|Arch 2018. The engagement of the NCCR Digital Fabrication, including its management staff, technicians and researchers, has been decisive in making this conference and publication possible. As such, a special thanks goes to Russell Loveridge, Orkun Kasap and Kaitlin McNally for their extraordinary commitment and work in coordinating and pushing Rob|Arch 2018 forward. We would also like to thank our Workshop Chair, Romana Rust, and all our workshop partners for their exceptional engagement. And, we would also like to extend our gratitude to ETH Zurich and the

Department of Architecture for the generous opportunity to pursue Rob|Arch 2018 in Zurich at the Hoenggerberg Campus.

Last but not least, we would like to thank all our research partners and peer institutions, our local supporters and colleagues at ETH Zurich, University of Technology, Sydney, and Bauhaus-Universität Weimar. Finally, we would also like to thank Springer Engineering for their kind support in editing and publishing this scientific publication.

June 2018

Jan Willmann
Philippe Block
Kendra Byrne
Marco Hutter
Tim Schork

Paper Committee

Sigrid Adriaenssens	The Department of Civil and Environmental Engineering, Princeton University, USA
Mania Aghaei Meibodi	Chair for Digital Building Technologies, ETH Zurich, Switzerland
Francis Aish	Foster + Partners, UK
Shajay Bhooshan	Block Research Group, ETH Zurich, Switzerland/CODE, Zaha Hadid Architects, UK
Tobias Bonwetsch	ROB Technologies AG, Switzerland
Johannes Braumann	Association for Robots in Architecture, Austria
Jonas Buchli	Agile & Dexterous Robotics Lab, ETH Zurich, Switzerland
Michael Budig	Architecture and Sustainable Design, Singapore University of Technology and Design, Singapore
Jane Burry	School of Design in the Faculty of Health Arts and Design, Swinburne University, Australia
Xavier De Kestellier	HASSELL, UK
Karola Dierichs	Institute for Computational Design and Construction, University of Stuttgart, Germany
Benjamin Dillenburger	Chair for Digital Building Technologies, ETH Zurich, Switzerland
Leda Dimitriadi	Department of Digital Knowledge, ENSA Paris-Malaquais, France
Thomas Feix	Adidas FUTURE, Germany
Jelle Feringa	Aectual, The Netherlands
Mary Franck	ESI Design, USA
Fadri Furrer	Autononous Systems Lab, ETH Zurich, Switzerland

Norman Hack	Institute of Structural Design, Technical University of Braunschweig, Germany
Volker Helm	Department of Architecture, University of Applied Sciences and Arts Dortmund, Germany
Mats Isaksson	Department of Mechanical Engineering and Product Design Engineering, Swinburne University, Australia
Roderick Jackson	National Renewable Energy Laboratory, USA
Jason Kelly Johnson	Architecture Division, California College of Arts, USA
Hanif Kara	AKT II, UK
Steve Keating	Mediated Matter Group, Massachusetts Institute of Technology, USA
Axel Kilian	School of Architecture, Princeton University, USA
Toni Kotnik	Department of Architecture, Aalto University, Finland
Torsten Kroeger	Intelligent Process Automation and Robotics Lab, Karlsruhe Institute of Technology, Germany
Nitish Kumar	Computational Robotics Laboratory, ETH Zurich, Switzerland
George Legendre	Graduate School of Design, Harvard University, USA
Katharina Lehmann	Blumer-Lehmann AG, Switzerland
Christiane Luible-Bär	Department of Fashion & Technology, University of Art and Design Linz, Austria
Areti Markopoulou	Institute for Advanced Architecture of Catalonia, Spain
Iain Maxwell	School of Architecture, University of Technology, Sydney, Australia
Wes McGee	Taubmann College of Architecture and Urban Planning, University of Michigan, USA
Marek Michalowski	X: The Moonshot Factory, USA
Ammar Mirjan	Chair of Architecture and Digital Fabrication, ETH Zurich, Switzerland
Stefanie Mueller	Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA
Paul Nicholas	Schools of Architecture, Design and Conservation, The Royal Danish Academy of Fine Arts, Denmark
Andy Payne	Autodesk, USA
Marshall Prado	College of Architecture and Design, The University of Tennessee, Knoxville, USA

Mette Ramsgaard Thomsen	Schools of Architecture, Design and Conservation, The Royal Danish Academy of Fine Arts, Denmark
Dagmar Reinhardt	Sydney School of Architecture, Design & Planning, The University of Sydney, Australia
Matthias Rippmann	Block Research Group, ETH Zurich, Switzerland
Christopher Robeller	Digital Timber Construction, Technical University of Kaiserslautern, Germany
Romana Rust	Chair of Architecture and Digital Fabrication, ETH Zurich, Switzerland
Fabian Scheurer	Design-to-Production, Switzerland
Jonatan Schumacher	Konstru, USA
Tobias Schwinn	Institute for Computational Design and Construction, University of Stuttgart, Germany
Claire Sheridan	Brick, USA
Asbjørn Søndergaard	Aarhus School of Architecture, Denmark
Hanno Stehling	Design-to-Production, Switzerland
Bratislav Svetozarevic	Chair of Architecture and Building Systems, ETH Zurich, Switzerland
Paul Tierman	Morris Adjmi Architects, USA
Brian Trump	Gehry Technologies, USA
Jaime Valls Miro	Centre for Autonomous Systems, University of Technology, Sydney, Australia
Teresa Vidal-Calleja	School of Mechanical and Mechatronic Engineering, University of Technology, Sydney, Australia
Timothy Wangler	Institute for Building Materials, ETH Zurich, Switzerland
Aaron Willette	WeWork, USA

Contents

Material and Processes

Construction and Structure

Design and Simulation

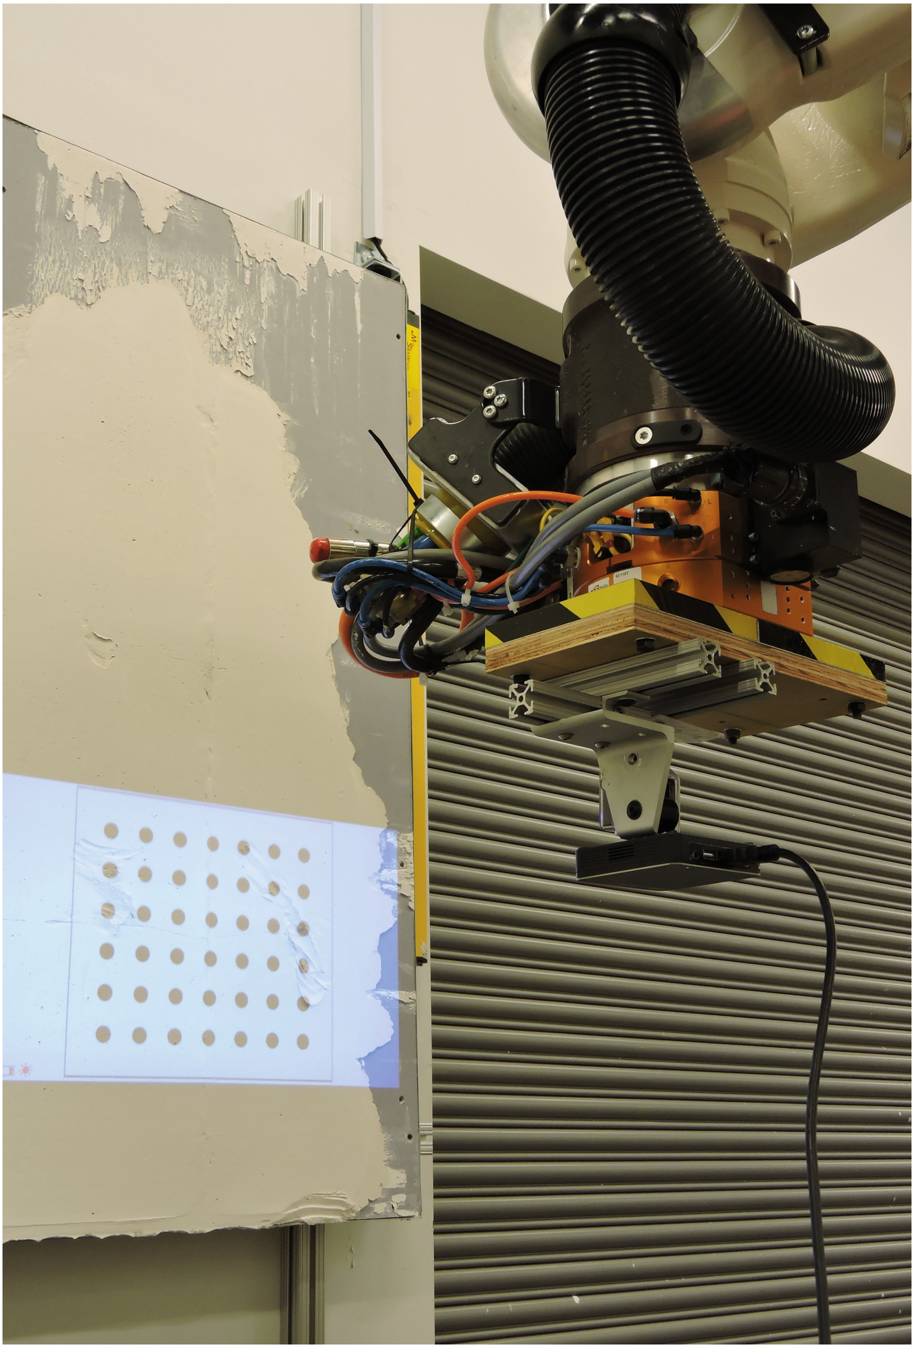

Image Classification for Robotic Plastering with Convolutional Neural Network

Joshua Bard(✉), Ardavan Bidgoli, and Wei Wei Chi

Carnegie Mellon University, Pittsburgh, PA 15213, USA
jdbard@cmu.edu, {abidgoli,weiweic}@andrew.cmu.edu

Abstract. Inspecting robotically fabricated objects to detect and classify discrepancies between virtual target models and as-built realities is one of the challenges that faces robotic fabrication. Industrial-grade computer vision methods have been widely used to detect manufacturing flaws in mass production lines. However, in mass-customization, a versatile and robust method should be flexible enough to ignore construction tolerances while detecting specified flaws in varied parts. This study aims to leverage recent developments in machine learning and convolutional neural networks to improve the resiliency and accuracy of surface inspections in architectural robotics. Under a supervised learning scenario, the authors compared two approaches: (1) transfer learning on a general purpose Convolutional Neural Network (CNN) image classifier, and (2) design and train a CNN from scratch to detect and categorize flaws in a robotic plastering workflow. Both CNNs were combined with conventional search methods to improve the accuracy and efficiency of the system. A web-based graphical user interface and a real-time video projection method were also developed to facilitate user interactions and control over the workflow.

Keywords: Architectural robotics · Machine learning
Convolutional neural networks · Image classification

1 Motivation

Surface finishing is an essential domain in the architectural construction practice, which requires high-skilled workers and demand accurate quality control procedures. By way of example, the authors have developed a robotic workflow to use industrial robots for decorative plastering techniques (Bard et al. 2016a, b). One of the remaining challenges in this workflow is to implement an automated, precise, and reliable quality control pipeline to guarantee satisfying results through a touch-up scenario. The touch-up procedure would let the user automatically inspect the surface and detect any unwanted fabrication artifact and command the robot to correct it.

Researchers have developed a wide range of scanning systems using different combination of sensory data, including, but not limited to, multiple RGB cameras/view (Vasey et al. 2014), RGB cameras combined with pattern projection (Rocchini et al. 2001; Zhang et al. 2002), RGB-D sensory data (Amtsberg et al. 2015), and depth data (Bard et al. 2016a, b) to reconstruct a digital representation of the physical models that can be used in the feedback loop.

J. Willmann et al. (Eds.): ROBARCH 2018, *Robotic Fabrication in Architecture, Art and Design 2018*, pp. 3–15, 2019.
https://doi.org/10.1007/978-3-319-92294-2_1

Construction tolerances standards may vary from rough to finish application resulting in different level of accuracy in each of these phases (Bard et al. 2016a, b). The achievable level of accuracy by the above-mentioned techniques with respect to the construction tolerances might not be desirable for such a delicate task as surface finishing and touch-up tasks.

Our approach requires a vision-based solution to detect texture flaws (i.e., scratches, bubbles, …) and small-scale 3D finishing issues (i.e., holes, unfinished patches). It proposes a single-camera solution without 3D reconstruction as the main input for the quality check workflow. This will result in simpler hardware setup, faster workflow, and lower costs. This approach can also be useful for other fabrication workflows, for example subtractive and deforming manufacturing.

The proposed system takes advantage of a state-of-the-art computer vision method based on *Convolutional Neural Network* (CNNs or *ConvNets*) for image classification and object detection.

Recently CNNs have dominated the image processing field, outperforming other image processing and computer vision methods by a large margin. Since 1990's CNNs have been used for different applications, including, but not limited to optical character recognition, medical image processing, feature extraction, object recognition, image understanding, and optimization (Egmont-Petersen et al. 2002; LeCun et al. 1989), thanks to their robust and real-time performance even in noisy spaces (Pal and Pal 1993).

The breakthrough advancements in computational hardware (efficient GPU architectures, possibility of distributed/multi-core/cloud-based/parallel processing, and dramatic cut in the hardware and service prices), alongside the open-source and widely accessible software platforms for machine learning, simplified and enhanced the implementation of CNNs in different contexts. CNNs can now be deployed with a reasonable budget and fewer technical challenges. These factors render CNNs as the method of choice for image classification and object detection in the past few years.

2 Methodology

The authors tested two approaches to implement CNN for this task, (1) transfer learning on *Google*'s *Inception v3* model and (2) design and training a CNN from scratch. The resulting CNN were tested for accuracy and efficiency in a robotic plastering workflow as a vision-based feedback loop for surface touch-ups.

Since its public release, Inception v3 has been used as an almost off-the-shelf image classifier for different use-case scenario. Several research teams leveraged this architecture for scientific studies, possibly most notably, for example detecting skin cancer classification (Esteva et al. 2017). In its purest form, it only requires users to organize the training dataset in a folder structure and run the provided script for a desired number of epochs. This doesn't necessitate any substantial machine learning knowledge or advanced programming skills. On the other hand, design and training a model from scratch demands for both but may provide simple and optimized results.

2.1 Apparatus Setup

Hardware. The hardware setup consists of two different end-of-arm-tool assembly (EOAT) and the connected computer unit. The first EOAT was designed to apply plaster and the second one was dedicated to the quality control and feedback loop. It includes a camera and a video projector to capture images and project results on the probing surface (see Fig. 1).
Despite the resiliency and scalability, CNN algorithms are computationally complex and expensive. While embedding computing solutions from *nVIDIA*, i.e. *Jetson* series, are capable of running such CNNs, authors decided to centralize all the computational process on a connected computer, leveraging GPU acceleration.

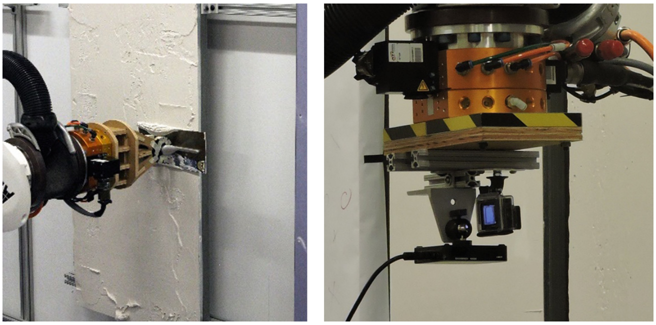

Fig. 1. End-of-arm-tools, left: plastering tool, right: camera and projector

Software.[1] A user interface was developed as a web-based application using *Django* framework. This architecture makes it possible to use *HTML/JavaScript* interactions at the front-end while leveraging *Python* scripts on the back-end. User can interact with the robot for motion commands, triggering the image processing workflow, and monitor the results.

To control the cameras, an in-house library was developed leveraging GoPro's built-in Wi-Fi protocol that provides full control over the camera functionalities.

The robotic control module was developed based on project *Open-ABB* (Dawson-Haggerty, n.d.). It communicates with the *IRC5* controller to transfer motion commands and inquire robot's status (see Fig. 2).

2.2 Data Set

The data set consists of images taken from a series of plaster finishes applied by a robot on drywall test panels. To collect the training samples, the GoPro camera was used to take 5 mega-pixel images of available plastered panels. Due to the GoPro camera significant lens distortion an image calibration method was applied using *OpenCV*, and

[1] Code for this project are available on GitHub (https://github.com/Ardibid/RoboticPlasteringCNN).

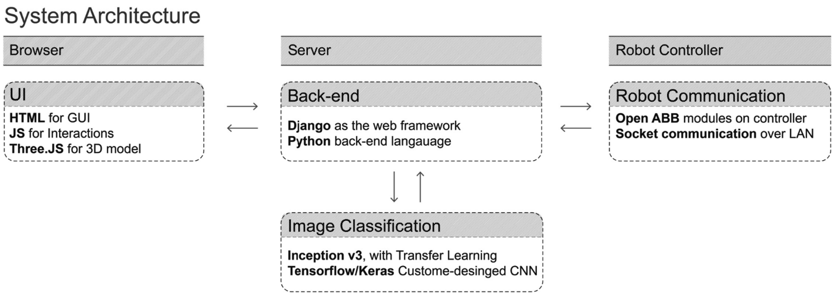

Fig. 2. System architecture

only the central 1024 px × 1024 px region of each image was used. Images were manually cropped and labeled into smaller sections as one of the three main classes: (1) pass, (2) fail, (3) markup (see Table 1). Then the same data set was categorized in five classes; perfect or near-perfect plaster regions were labeled as (1) pass, while images containing fabrication flaws including: (3) holes, (4) scratches, and (5) unfinished surfaces were categorized as fail. The markup class was left intact (see Table 2).

The markup class was dedicated to hand drawn characters that users could sketch on the work surface to communicate with the robot. Markup training samples were taken from hand drawn marks on a white surface in the same lighting condition as the plastered panels.

Table 1. Training data set for three classes

Class	Training samples	Validation samples	Test samples
1 Pass	77	23	26
2 Fail	135	42	52
3 User markup	91	28	36
Total = 510	303 (~%60)	93(~%18)	114(~%22)

Table 2. Training data set for five classes

Class	Training samples	Validation samples	Test samples
1 Bad finish	26	9	10
2 Holes	29	8	11
3 Rough finish	78	22	29
4 User markup	87	29	36
5 Pass	68	23	27
Total: 492	288 (~%59)	91(~%18)	113(~%23)

2.3 Image Classification Methods

In addition to designing and training a CNN from scratch, it is a well-established practice to repurpose currently available CNN architectures and their pre-trained models for a new task. Shin et al. categorized three methods to repurpose CNNs for image detection as: (1) training CNN from scratch, (2) using available CNNs without training the network, and (3) using unsupervised pre-training and fine-tuning (Shin et al. 2016).

Transfer learning on Inception v3. In 2014, *Google*'s entry for the *Large-Scale Visual Recognition Challenge* (ILSVRC2014), titled *GoogleNet* demonstrated astonishing performance (Russakovsky et al. 2015; Szegedy et al. 2015; Szegedy et al. 2016). The core model behind *GoogleNet* was called *Inception* (Szegedy et al. 2016) (Fig. 3) which adopted a relatively simpler architecture compared with other competitors and was computationally less expensive. Since then, the inception model has been used in cutting-edge research for object classification in different contexts, notably medical image processing (Esteva et al. 2017).

What makes the *Inception* model an ideal platform for object categorization in different contexts is its flexibility to be retrained for relatively similar tasks. This approach, called *transfer learning*, is a well-practiced method to fine-tune and repurpose CNN models for new tasks with very small training data set to train a deep CNN (Donahue et al. 2014).[2]

By only modifying the second to last layer of the model, transfer learning on inception v3 eliminates the need for training a whole new model from scratch for every new set of classes. Transfer learning could be a proper choice for this use case scenario since (1) it has already been trained to detect a wide range of features and there is no need to train it from scratch, theoretically this will save substantial amount of time; (2) it performs well when the size of training dataset is relatively small; (3) it requires a

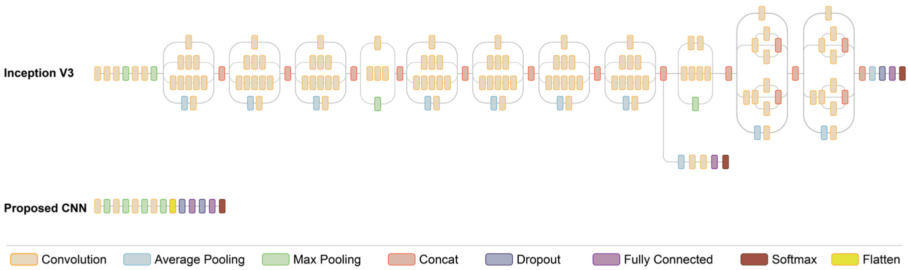

Fig. 3. Inception v3 architecture compared with the proposed architecture (Inception architecture diagram is reproduced from ("Inception v3," 2018))

[2] In the same year, at CVPR2014, Oquab et al. and Sharif et al. also addressed transfer learning and representation. For further information please look at (Oquab et al. 2014; Sharif Razavian et al. 2014).

simple workflow to repurpose the model. In light of these three advantages, we decided to exploit transfer learning technique as a possible back-end methods for our feedback loop system.

Proposed CNN. A significant trade-off of using transfer learning is the heavy model that it entails. Trained to classify one thousand classes of objects, the CNN trained model occupies hundreds of megabytes on the system storage and requires expensive computation to process a single image.

However, in our case, most of the captured images are of low contrast with primarily white backgrounds and subtle changes in color. This color space requires different feature layers for an efficient classification. Accordingly, the authors designed and trained a sequential multi-layer CNN. This architecture has already been proved its performance in several state-of-the-art models, including *AlexNet* (Krizhevsky et al. 2012) and later *VGGNet* (Simonyan and Zisserman 2014).[3] The proposed architecture is significantly simpler than of the *Inception*, resulting in a speed boost.

The authors designed and tested a series of CNNs using *Keras* with *Tensorflow* back-end to find an optimum architecture. Several combinations of convolutional, dropouts, and fully connected layers have been tested. In each architecture, all models have been trained for a fixed number of epochs and the model with the highest f1 score were selected. The results from each architecture were then compared with each other to select the optimum architecture. The selected architecture demonstrated the highest f1 score on both 5 and 3-class classification, while the others failed to demonstrate same f1 score or took longer epochs to converge to the same score.

The proposed architecture consists of four convolutional layers (3×3 kernel) paired with *Relu* activation function, and *maxPooling* (2×2), followed by three fully connected layers and *softmax* at the end. To reduce the effects of overfitting, it also leverages *dropout* to prevent inter-dependencies between hidden layer nodes (see Fig. 4 and Table 3).

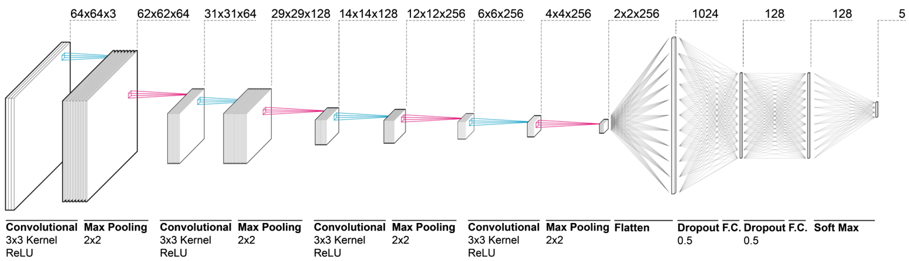

Fig. 4. Our CNN architecture for 5-class model

[3] AlexNet architecture might be confusing at the first sight since it has two parallel pipelines. However, the reason behind this dual pipeline is to train the model on two separate GPU simultaneously.

Table 3. Proposed CNN layers

Layer (type)	Output shape	Param #
conv2d_1 (Conv2D)	(None, 62, 62, 64)	1792
max_pooling2d_1 (MaxPooling2)	(None, 31, 31, 64)	0
conv2d_2 (Conv2D)	(None, 29, 29, 128)	73856
max_pooling2d_2 (MaxPooling2)	(None, 14, 14, 128)	0
conv2d_3 (Conv2D)	(None, 12, 12, 256)	295168
max_pooling2d_3 (MaxPooling2)	(None, 6, 6, 256)	0
conv2d_4 (Conv2D)	(None, 4, 4, 256)	590080
max_pooling2d_4 (MaxPooling2)	(None, 2, 2, 256)	0
flatten_1 (Flatten)	(None, 1024)	0
dropout_1 (Dropout)	(None, 1024)	0
dense_1 (Dense)	(None, 128)	131200
dropout_2 (Dropout)	(None, 128)	0
dense_2 (Dense)	(None, 128)	16512
dropout_3 (Dropout)	(None, 128)	0
dense_3 (Dense)	(None, 5)	645

3 Training Results

3.1 Transfer Learning

We used the model from the Tensorflow GitHub repository and followed the steps described in Tensorflow documentation ("Image retraining Tutorial," n.d.). The result after 4000 epochs are reported in Table 4.

Table 4. Training results, transfer learning (The training results in this chart directly result from the script provided on TensorFlow's GitHub repository after introducing new training samples. No fine-tuning, modification, or custom loss function has been applied to the re-training process. F1 score is measured on a test data set after the training phase.)

Three-class model	Ave. time to process batch of 114 patches of 64 × 64
Train acc.: 1.000 | Val. acc.: 1.00 (N = 100) Test. acc.: 0.964 (N = 112) Test F1 score: 0.9469	10–11 s
Five-class model	
Train acc.: 0.990 | Val. acc.: 0.79 (N = 100) Test. acc.: 0.904 (N = 104) Test F1 score: 0.8938	12–13 s

3.2 Our Model

In this model, the authors leveraged data augmentation to increase the data set size and improve the model's resiliency against small variations of the input data. Training and test samples where reshaped to the same size (28 × 28 × 3) beforehand. The model

Table 5. Training results, our model

Three-class model | F1 Score	Ave. time to process batch of 114 patches of 64 × 64
Train: 0.9826 | Validation: 0.9892 | Test: 0.9736	19 ms
Five-class model | F1 Score	
Train: 0.8559 | Validation: 0.9333 | Test: 0.9292	19 ms

was trained in two scenarios, one with (1) pass, (2) fail, and (3) markup labels[4] and the second one trained to define different types of fail including (1) bad finish, (2) hole, (3) rough finish (see Table 5 and Fig. 5).

3.3 Method Comparison

Comparing the speed and accuracy of the two approaches, signifies that the lighter and less complex architecture of the proposed CNN is on par and even better than what we could obtain using transfer learning. Transfer learning is significantly less complicated method, from the user point of view, that doesn't require significant understanding of machine learning. However, it is not an efficient method for fast image classification which is essential in this use-case-scenario.

On the other hand, designing, fine-tuning, and testing a CNN from scratch requires additional skills and substantial amounts of time in advance. But it pays off with the accuracy and efficiency it brings to the inspection process. Accordingly, we decided to continue with this proposed CNN.

4 Testing Implementations Scenarios

The proposed CNN was used as the image-classification back-end for a robotic plastering feedback loop, which consist of classification tool, user interface, and user interactions.

4.1 Image Surveying Methods

To survey each image, the algorithm divides it into a gird of 64 × 64 px patches that could be fed into the classifier in batches.[5] With the hardware setup described above, it took 50 ms on average to process a 256 × 64 × 64 batch of data, equivalent of a 1024 × 1024 image.

---

4 Markups are simple user-defined drawings, i.e. circles and crosses, that can be used to communicate with the system.

5 Although the authors first implemented Quad Tree search algorithm to compensate for the possible slow classification pipeline, the final model performance was good enough to provide near real-time experience. Accordingly, we opted for a grid search algorithm and avoided potential challenges that a Quad Tree search would introduce. The biggest drawback being its tendency to ignore small features in the initial steps of the search process when surveying large areas of the given image.