Loci Communes 2022, vol. 1 (2), pp. 1–16
ISSN 2720-6122
https://doi.org/10.31261/LC.2022.02.06

Through a Machine Lens: Case Studies of Computer Vision and Machine Learning in Design Methodology

Andrew Witt

Harvard University, United States of America
https://orcid.org/0000-0003-1049-2945

Abstract

Machine vision (MV) and artificial intelligence (AI) offer strange new augmentations and transformations of how architects perceive and conceptualise both the creation of buildings and the analysis of existing architecture. Seeing through machinic eyes allows architects to amplify their intuitions and engage in a flood of digitally sensed imagery in more quantifiable and extensible ways. Through a series of case studies of projects developed through the design office *Certain Measures*, this article argues for the potential of machine vision and artificial intelligence in the creative practice of design while situating these new developments in the history of mathematical ways of seeing and conceptualising architecture. These case studies, across large and small scales, combine ideas from human and machine perception and mathematical geometry to create new architectural approaches.

Keywords

architecture, machine vision, artificial intelligence, design, methodology

LOCI COMMUNES
International Journal of Studies on Spaces in Arts and Humanities, Anthropology and Architecture

Introduction: Architectural Senses, Mathematical Seeing

Vision is an essential faculty of the designer, and media of visual perception shape how architects perceive and conceptualise. Over the last twenty years, machine vision – the computational capacity to recognise and analyse images and video – has advanced, offering architects new media through which to understand design. Feature recognition and deep learning processes have allowed machine vision to exceed the normal capabilities of human perception. Today, human vision, machine vision, and augmented combinations of the two shape the perception and planning of the built environment in hybridised and technosocial ways. For designers, this symbiosis presents new opportunities for creative design. This paper argues for the new potential of intersecting human and machine vision through key design projects by *Certain Measures*, an office for design science founded by Tobias Nolte and myself. *Certain Measures* bridges the creative space of design, art and technology, with a particular interest in the co-evolution of human and machine experience and perception. This more-than-human perspective attends to the reciprocal ways in which humans and machines see, think, and ultimately create. We apply mathematical techniques (including bespoke computational geometry[1] and morphological analysis[2], data science[3] and AI[4]) to a range of projects – from the material to planetary scales.

The office *Certain Measures*

Projects designed by *Certain Measures* foreground perceptual quantification as a creative act and as a way of revealing more rigorous and technically synthetic modes of design. Though our work spans across scales, each project tries to take an architectural approach and experiments with ongoing mutations of design processes introduced by digital sensing. A key part of this experimental practice is the development of bespoke software and hardware to achieve our creative aims. One central example is the collection of "seeing machines" – software tools used to process imagery and video – that we build to organise large sets of digitised visual information.

[1] Computational geometry being the application of digital automation techniques toward generative design.
[2] Morphological analysis being the application of mathematical geometry toward the quantitative comparison of shapes.
[3] Data science being the use of statistics to derive insights from large data sets.
[4] AI being artificial intelligence generally and more specifically the application of neural networks and other machine-learning techniques.

These machines integrate open-source codebases like Open Computer Vision (OpenCV) or neural classifiers such as Visual Geometry Group (VGG) with our own data science methods of graph segmentation, mesh fitting, and shape classification. In this way, we apply bespoke Machine Learning (ML) and AI methods to the perception and organisation of reality and, in turn, to the shaping of architecture.

Impossible Objects and Artificial Intelligence

An early inspiration for the "seeing machines" developed by our office was the class of optical illusions known as impossible objects. Impossible objects are two-dimensional drawings which appear to be projections of three-dimensional objects which could either not be constructed spatially or could only be constructed with odd anamorphic distortions. These curious drawings emerged in optical psychology in the 1960s and 1970s as a perceptual test for humans, particularly in the experimental work of psychologist Lionel Sharples Penrose (Penrose, L. S., & Penrose, R., 1958). Many impossible objects seem to have qualities of a Möbius strip, folding back on themselves in contradictory ways. Penrose challenged human subjects to describe what the drawings represented spatially. Interestingly, impossible objects were also among the first challenges posed to machine vision systems and precursors of artificial intelligence. How could a machine make sense of these very strange forms?

Figure 1

An example of a knot-like surface generated from an impossible object diagram. © Certain Measures.

In our early research, we explored how we could make computational sense of these impossible objects by taking these 2D enigmas and making them into 3D architectural forms. Using methods from computational topology, such as spanning and Seifert surfaces, we began developing computational methods that allowed us to construct three-dimensional non-self-intersecting surfaces from those two-dimensional self-intersecting visual conundrums. The resulting surfaces twisted like spatial knots, warping back on themselves in unpredictable ways. We began to develop entire combinatorial catalogues of these complex surfaces, each having strange and illusional properties (see Figure 1).

First, with smaller mock-ups and then with larger ones, we ultimately developed a highly specific way of building these illusional objects while minimising torsion and curvature of construction materials. Using mathematical processes of graph segmentation and discrete differential geometry, we developed highly efficient algorithms to slice these impossible objects up and give them a constructable manifestation in the physical world.

Of course, building an object is different from merely visualising it, so we needed to adapt ML processes to enable more complex constructions. We used an ML method of object segmentation that is similar to AI image recognition which identified critical thresholds in curvature difference between distinct areas of a continuous surface and partitioned the surface along these thresholds. We applied this technique to complex generative meshes through our custom-built software to analyse and slice these impossible objects into pieces. These pieces are materially optimised: they are designed to minimise torsion and twist as little as possible in space. They can thus be constructed with essentially flat material and assembled in the span of a few hours. This differs from the mass customisation process, in which a product can be endlessly varied for the maximum difference because we were trying to minimise the underlying number of pieces. This way, we can build highly complex forms with a few calculated fragments.

Mine the Scrap

Our attempts to build impossible objects and complex geometries in highly efficient ways were successful to a degree. Yet even as we tried to build these structures as cheaply as possible, we realised that we were still wasting 30% to 35% of the original material from offcuts of very irregular shapes. So we began to ask: is there a material we could use to build even more cheaply? In fact, we realised that there is nothing cheaper than garbage. So instead of taking raw new material as the basis for architecture, what if we were to take construction waste? Construction waste is about 30-35%

of the volume of landfills in Europe and much of the developed world (Sáez, 2019). We were interested in taking this raw material and repurposing it for architecture in the act of radical reuse. Of course, there were key challenges. One of the reasons why this material is not used is that when it is demolished, it breaks into all sorts of irregular shapes. Classifying and organising those irregular shapes, those thousands of broken fragments, became a new geometric problem that held the key to radical material reuse.

Could we turn our machine learning and geometric expertise toward this problem of making unclassifiable scraps, unclassifiable waste, usable? Our answer to this was a project called *Mine the Scrap*, a data-driven process designers can use to develop new structures from old construction waste. We developed a forty-dimensional machine-learning and pattern recognition metric that allowed us to classify these very irregular shapes and cluster them into similar groups. This metric leveraged existing shape-comparison techniques such as Hu invariants and Shape Context measures, as well as new bespoke processes that we developed around a spectral comparison of skeletal graphs. We developed a search engine for these irregular scrap pieces, almost like assembling a new puzzle from a set of unrelated and random pieces. The idea is to find the unique best use of each piece in a new structure through the processes of scanning and classification. With this scanning, organisation, and fitting process, we can match what we have with what we want to build and find beauty in the intricacy of the neglected waste. Using computer vision – the application of digital algorithms to process and organise visual imagery – and machine learning, we can invert the typical process of design. Instead of material selection being an afterthought, we can begin with this raw material and ask what we can design with it. *Mine the Scra*p negotiates between the designer's desires and the material at hand. *Mine the Scrap* ultimately uses big data to tackle big waste.

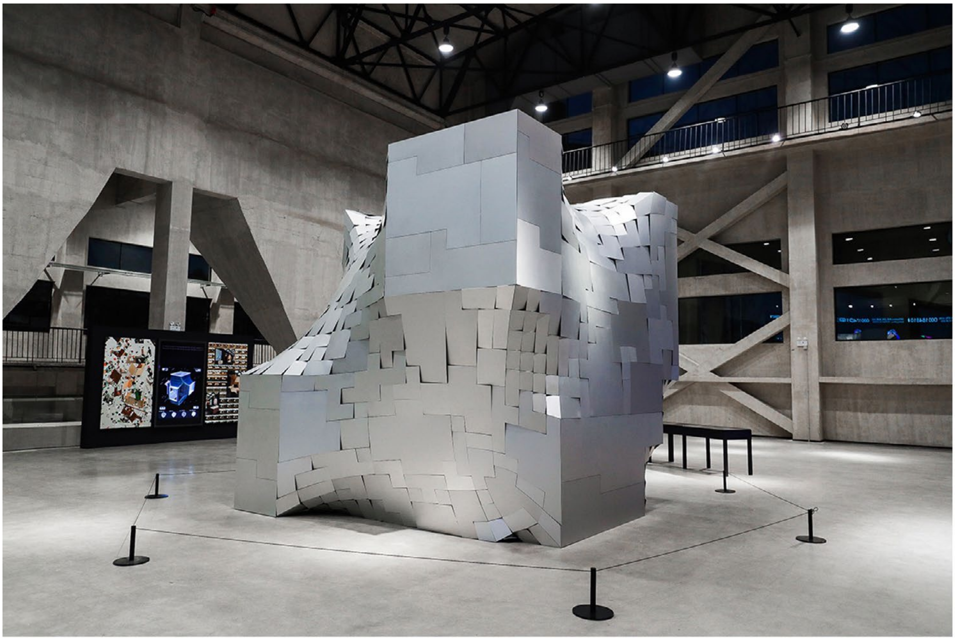

Photo 1
A Mine the Scrap installation at Aiiiii Art Center, Shanghai. © Certain Measures.

We have recently scaled up the geometric methods of the *Mine the Scrap* and presented an installation of that geometric approach in Shanghai at the new *Aiiiii Arts Center*, a museum dedicated entirely to AI and machine learning in design and art (Photo 1). *Mine the Scrap* makes tangible this process of radical reuse and uses AI to classify, sort and see these eccentric geometries in a new way. With this piece, there is an interesting echo of the ideas of the impossible objects that were our first inspiration.

We have begun to also apply the *Mine the Scrap* process of radical waste reuse in other contexts, including with industrial partners. In collaboration with the Austrian fibre-reinforced concrete company Rieder, we have used the *Mine the Scrap* processes to design intricate wall assemblies from waste panels of their fabrication process. Fibre-reinforced concrete is a flat sheet material usually cut with CNC water jets. The cutting process produces offcuts – residual pieces that are not a part of the final intended product. Rieder had retained many of these scrap offcuts and the original digital cut patterns. We compiled these cut patterns into an extensive database from which one could generate larger patchwork assemblies. The resulting panels have a distinct quality similar to the original *Mine the Scrap* structures. In this and most other projects, we were interested in creating these physical artefacts, but we were also interested in developing digital tools and processes that would allow those artefacts to be not just a single piece of architecture but a system. Our work always develops particular design proposals in the context of a wider systems design.

Sculptural Scale: *Kintsugi++*

After our research into construction waste, we began to see the affinities of our work with other precedents of scrap reuse, particularly the Japanese craft of Kintsugi, an ancient technique of resurrecting shattered ceramic vessels (Roma, 2013, p. 63). When a vessel accidentally shatters, the kintsugi technique stitches it back together with sutures and seams of gold. The original vessel is recovered from its fractured state, but it is also elevated, ennobled by its reassembly with precious metals. We were fascinated by this practice, particularly in relation to industrial processes. In high-end ceramic manufactures, if a ceramic piece is defective, it is deliberately shattered. Sometimes the manufacturers would keep these shattered vessels in a separate room as if they were unsure what to do with them. We took fragments of shattered vessels, and 3D scanned them using our ML process to imagine new designs or assemblages of these pieces. We called the project *Kintsugi++*, to evoke C++, the coding language used to achieve the results. We collect these remnants and use our computational process of nonstandard fragment assembly to reclaim them and ennoble them as elements of complex and intricate new forms. When visualised in an immersive digital space, the process creates a galaxy of shards arranged and related by ML engines. The physical artefacts of this process are a series of custom, 3D-printed frames into which these shards are carefully inlaid. Each joint is very finely calibrated to every ceramic piece within a larger 3D-printed frame. Just as with *Mine the Scrap*, we produce both physical objects and digital experiences from this process (Photo 2). The digital experience of *Kuntsugi++* is an immersive dual projection that allows the viewer to enter the dream state of the algorithm and to see the animated galaxy of fragments and shards as a kaleidoscope of artificially imagined vessels. Visitors could also engage this project as a VR experience, seeing this galaxy as a panorama vertiginously creating endless new forms (Figure 2).

Photo 2

Examples of "Kintsugi++" vessels generated from fragments of ceramic vessels and a generative ML process. © Certain Measures.

Figure 2

A panoramic view of the VR experience of "Kintsugi++". © Certain Measures.

Machine Mediated Processes: *Cloudfill*

We see buildings as temporary states of matter, which can be designed as transitions between previous and next states. As designers, our responsibility is to make this transition between states as seamless as possible. A critical aspect of our *Mine the Scrap* process is that it alters the typical design cycle of the architect. Instead of the design process being entirely human iterations between "seeing" and "making", machine vision and learning transform this into a computationally mediated process of scanning and generating. Machinic scanning allows us to take our human intuitions around specific sorts of forms and scale them up massively, and a computational generating process allows us to combinatorially explore a much more comprehensive range of possibilities for that scanned material.

Photos 3, 4

Installation view of "Cloudfill", an application of generative reuse processes applied to a disused datcha. © Certain Measures.

A further application of the *Mine the Scrap* process comes from our *Cloudfill* project, which began with the active disassembly of a disused datcha, a summer home from the outskirts of Berlin. After deconstructing this structure, we exhaustingly catalogued the pieces through a detailed scanning process. This forensic scanning process allows us to consider not only the shape but also the colour, texture, and weathering of these pieces. It is even possible to use tomographic x-ray scanning to understand the internal structural qualities of these members. Thus, we can take into account not only the geometry but a whole range of other qualities of scrap pieces. Since we divert this scrap from landfill and scan them into the digital cloud, we called this project *Cloudfill*. This common scrap database creates a shared resource for reuse. The visualisation of *Cloudfill* also evokes the quality of this cloud-like construction, dynamically recombining these pieces into new datchas or houses of the future. The typically linear cladding pieces of the deconstructed datcha induce a very particular ruled surface geometry on the resulting architecture, which differs from *Mine the Scrap* and its more irregular two-dimensional shapes. This ruled geometry produces a series of extrusion intersections of developable surfaces, while the frames act almost like combinatorial drivers for Rubix cube-like transformations.

A full-scale reconstruction of one corner of a Cloudfill-generated datcha was recently hosted at *Futurium*, Berlin's new museum of the future. The project playfully reconfigures not only the exterior of the datcha but the interior as well. Inside, the visitor sees an artfully reassembled patchwork of 1970s-era East-German wallpaper collaged into this new assembly. Both exterior and interior participate in this extreme reuse process. Ultimately projects like *Cloudfill* combine data science with architecture to propose a future for computational sustainability (Photos 3 and 4).

Machine Vision and the City Scale

Beyond the scale of architecture or products, we were also curious how machine vision classification methods could help us understand urban form. How could machine vision organise the world's cities? We began to apply our *Mine the Scrap* algorithm to building outlines and structures of the cities to categorise them morphologically. We developed a series of video interventions called *Machine View of the City*, which made this scanning and morphological classification process tangible. Each video projection shows the progressive scanning of a target city and a gradual remapping of its buildings based not on geographic location but morphological similarity. The first of these interventions was a *Machine View of London*, which presented a bot that scans, categorises, and maps the shapes of over one million buildings in central London. We extended the project further with other scans and remappings, including Berlin, Boston, Liverpool, Tokyo, Shanghai, and Guangzhou. The idea is to transform the relationships of those buildings from geographic to geometric, revealing formal affinities and similarities across the corpus of built architecture.

By creating a dynamic and machine-generated catalogue of architectural form, we are engaging with a longstanding interest in architecture and art in the catalogue as a creative medium. As we scan and geometrically organise these buildings, we are beginning to understand what the built space of architecture is in a city. A wall-sized printout of *A Machine View of Boston* reveals organised clusters of similarly shaped buildings (Photo 5). In some cases, they are very large clusters, indicating many similar buildings, and in others, they are very sporadic, suggesting unusual or even unique architecture. As the *Machine View of the City* compiles libraries of style, it begins to map out the space of architecture itself. The *Machine View of the City* is a chance to see the city through machinic eyes and to collage these precisely analytic and geometric views with something more atmospheric.

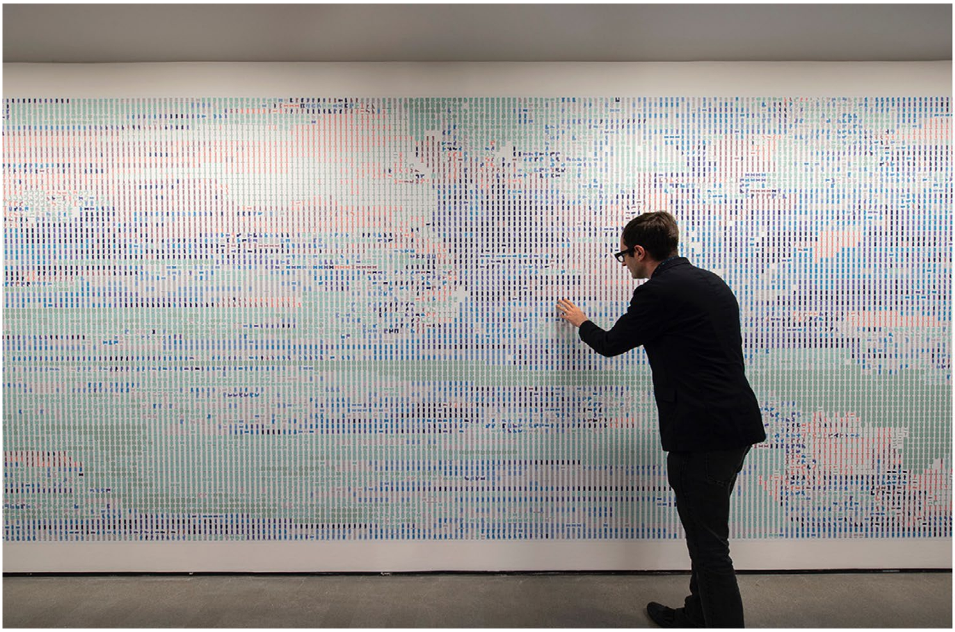

Photo 5

A portion of a large mural of buildings in central Boston, organised by shape. © Certain Measures.

Historic Preservation: The *Neo-Classifier*

The process of machine scanning and classification used in *A Machine View of the City* can be applied to inventory historical building types and chart historic preservation. One example is our recent collaboration with the Office for Urbanization at Harvard to catalogue regional Chinese building typologies using machine vision and satellite imagery. The aim was to identify certain very unusual building forms that were characteristic of particular Chinese areas, train a neural network to identify them, and scan vast areas to detect and inventory their instances. This process could give new insight into the prevalence and distribution of unusual historical building types.

We applied this method to 18 different regional typologies, spanning a variety of geographic regions and building forms. A separate neural network was trained for each building type, which was applied to satellite images to extract outlines of detected buildings. There was considerable variety in how the different typologies cluster and aggregate. Yet this kind of survey can be beneficial for the future of cultural preservation since it gives a comprehensive view of the variety of these unusual building types. We can also apply generative neural networks to speculate on the interior plans of those scanned buildings. When viewed collectively, we can create a self-organising map not only of building outlines but of plans themselves, revealing the intricate interrelationships of spatial structure.

We can apply machine vision processes to data sets of imagery, such as satellite photographs, to identify buildings which might be of historical or cultural interest. But we can also use similar processes to build catalogues of specific architectural styles from elevation imagery of the city. Our *Neo-Classifier* project took the same kinds of ideas that we used to analyse planimetric imagery and applied them to building elevations in the city. We were particularly interested in identifying and tracking architectural styles, so the *Neo-Classifier* was developed to identify the elements of neo-classical style across urban buildings and to extract and compile those elements into exhaustive catalogues of neo-classical style. In the past, when architects talked about style, it was often based on generalisations from a few specific buildings. But with current data and ML methods, we can have a total view of all buildings and understand precisely what is going on with style across the entire range of built architecture.

Beyond the *Neo-Classifier*, we have also looked at other kinds of imagery, such as digitised historical drawings. In one experimental project, in particular, we focused on architectural drawings, such as plans, sections, and elevations, from the École des Beaux-Arts in Paris, during the XIXth century. We compiled a database of 13 000 plan and section drawings of that style and used them to train new neural networks, which could then, in turn, generate new Beaux-arts plans. In effect, we created a machine to produce Beaux-Arts "deepfakes." This is an interesting way in which the AI training process can help us tease out new possibilities that may have been overlooked in historical styles and idioms.

Horizons

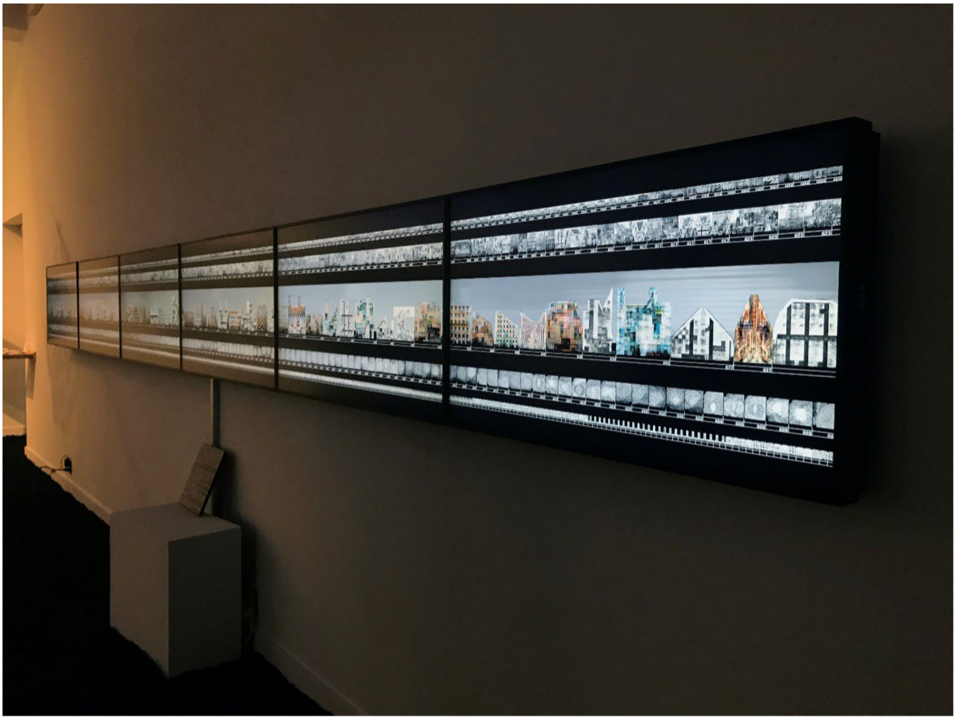

Photo 6
Installation view of Horizons, an endless AI-generated façade. © Certain Measures.

After analysing and classifying architectural forms from imagery, we can also rethink the process of generating new kinds of urban conditions. We can take those fragments and use them to train a new neural network to imagine new city elevations, which is precisely what we did with our project *Horizons*. *Horizons* is a journey into an endless elevation, dreamed by a neural network. We were inspired by the 1966 book of artist Ed Ruscha *Every Building on the Sunset Strip*, a pre-Google Streetview perspective of the city. To create the book, Ruscha mounted a camera on the back of a truck and drove along Sunset Boulevard, a very long and storied street in central Los Angeles. The result is a fold-out book that shows the complete elevation of that strip. We were interested in using our inventoried historical styles and using them to generate an endless elevation produced by an AI. *Horizons* has five related infinite horizontal views. The top is the catalogue of images we take as a training set which is the input for the neural network to imagine these new forms. The second is the outlines of many possible generated buildings. In the centre is the imagined city, the endless elevation created by AI. Finally, you see a collection of machine vision metrics on the bottom strip, exposing the scanning process. *Horizons* invites one to enter the city that bridges human and machine imagination (Photo 6).

Ecology, Pets and Designing Quality

The last project described in the article conceptualises AI not as a generative tool for designers but rather as a new quality of objects that must be designed. In other words, if AI becomes a property of the objects around us, including furniture and products, what does that mean for how we design buildings? AI is not just a set of form-making techniques, but it's a capacity of physical machines as well. We thus began to think about the status of pets – millions of us have pets, quasi-intelligent organisms that we share our homes with. Those pets are not furniture, they have a kind of perception, understanding, or intelligence. Political scientist Heather Roff, a former animal trainer and current AI researcher, has written extensively about AI's large-scale geopolitical implications (Roff, 2017). But at the small scale, she observes that "Animals and animal training can teach us quite a lot about how we ought to think about, approach and interact with artificial intelligence, both now and in the future" (2017, online). That observation was very compelling for us. It led us to a diagram in which we see humans on a spectrum of other intelligences, including animals and AI. If we consider the range of entities that we might design architecture for in the future, we can take this whole ecology into account. When we design, we are designing in the ecosystem, of which humans are just one part.

An example of a project that embraces an ecology of humans and machines is our *Berlin Buoy*. This project proposes a house that is composed almost entirely of robots. Every human activity aligns with certain machine services. There is an entourage of drones, droids, robots, and other smart objects, digitally augmented devices that are matched to specific times of the day and experiences. Your day is a dance between yourself and the machines you share this environment with. We develop an hour-by-hour diagram of the pulse of life and the machines with which you are sharing that moment that becomes tangible architecture in the *Buoy* – it becomes a framework in which these autonomous objects and robots reconfigure themselves in relation to the human occupant. In fact, the *Buoy* itself is a kind of drone – it can move around and explore the world. The idea is that it could even be a literal buoy at some point, just floating on the water, but it can also attach wheels or rotor blades, and it can be an autonomous free agent. The *Buoy* presents an autonomous house, not just as one object but as a constellation of objects that are defining an architecture. That means the objects within this house can become more animal-or pet-like, and our relationships with them can become very familiar in a certain way. In this project, you can see a vision of how the machines would see the environment they would live in and how they might interact with it. Those machines would have their own desires, interests, and attractions – ultimately, they would take part with us in a co-creation of these environments and spaces. The key is that these objects are not hard robots – they are robotic but have soft skins or fur, like animals. We found a future in which these machines could have qualities that we don't normally associate with machines or furniture to be fascinating and suggestive of new ways of living.

Conclusions

Technologies of machine vision enable the designer to process visual information at a vastly expanded scale. With this expanded scale of visual analysis, new modes of design are possible that organise visual data, including material elements and the buildings' forms, in new ways. The idea developed in this article was to present the connections between architecture, machine vision, and artificial intelligence through a range of projects that show these various potentials. From scanning historical building types to radical waste reuse, from autonomous furniture to morphological maps of the city, the potential for AI and machine vision to transform architectural processes is significant. New formal methods enable novel generative techniques for architectural projects, but also allow for fresh approaches to categorising visual historical data. In this way, the projects combine generative approaches with digital humanities and cultural analytics methods. These techniques allow unstructured visual data to be algorithmically organised and thus made much more operative within design. More generally, they allow a normalised view of the pattern and morphological analysis that furnishes designers with more objective techniques for engaging visual data that could previously only be understood qualitatively. By quantifying qualities, AI is opening new windows not only for what we design but also for how we design, prompting questions about the very nature of architectural knowledge. With such tools, we are re-exploring what design is and what the discipline of architecture could be.

References

Penrose, L. S., & Penrose, R. (1958). Impossible objects: A special type of visual illusion. *British Journal of Psychology, 49*, 31–33.

Roma, C. (2013). Kintsugi. *Ceramic Review, 260*, April, 63.

Ruscha, E. (1966). *Every Building on the Sunset Strip*. Self-published book.

Roff, H. (2017). Advancing Human Security Through Artificial Intelligence. *In Emerging Technologies and Human Security*. Chatham House, 3-10.

Roff, H. (2017). *How understanding animals can help us make the most of artificial intelligence, The Conversatioin* (March 30) https://theconversation.com/how-understanding-animals-can-help-us-make-the-most-of-artificial-intelligence-74742

Sáez, V. P., Osmani, M. (2019). A diagnosis of construction and demolition waste generation and recovery practice in the European Union. *Journal of Cleaner Production, 241*, 1-2.

Authors' note

Correspondence concerning this article is to be addressed to Andrew Witt, email: awitt@gsd.harvard.edu