Chapter 7
Cutting-Edge Practical Research on Generative Design, IoT and Digital Twins

Kean Walmsley

Abstract This chapter explores two areas of research undertaken by Autodesk over the last decade and how they both have the potential to impact the construction industry. The first area relates to the Internet of Things and the possibilities around integrating sensor data from smart buildings into a 3D context for exploration and visualization. This has the potential to drive interesting workflows related to understanding and optimizing the performance of the built environment, but more broadly will have the potential to impact design and fabrication processes. The other area of research described in this chapter relates to generative design, particularly in the context of the AEC industry. Autodesk's work in this area started in 2010 with a project named Dreamcatcher. It was further accelerated in 2014 with the acquisition of The Living, an architectural studio which had—under the guidance of its principal David Benjamin—started exploring the application of multi-objective optimization through the use of genetic algorithms for architecture, engineering and construction industry (AEC) workflows. Looking to the future, real-world performance data captured via IoT—and hosted in Digital Twins—will increasingly influence the Generative Design process, as Autodesk and other software providers start to complement algorithmic exploration of the design space with machine learning systems trained with data from prior projects and captured using physical sensors. This has the potential to drive "closed-loop" processes where performance increases with each design iteration or system re-configuration.

Keywords Generative design · Computer vision · Evaluation metrics · Design generation · Data capture

K. Walmsley (✉)

Autodesk Research, Autodesk Development S.à.r.l., Worbstrasse 223, Muri bei Bern, Switzerland CH-3073

e-mail: kean.walmsley@autodesk.com

1   Introduction

This chapter explores two areas of research undertaken by Autodesk over the last decade and how they both have the potential to impact the construction industry. The first area relates to the Internet of Things and the possibilities around integrating sensor data from smart buildings into a 3D context for exploration and visualization. This has the potential to drive interesting workflows related to understanding and optimizing the performance of the built environment, but more broadly will have the potential to impact design and fabrication processes. Autodesk Research started investing in this area in late 2009 and has since been developing Digital Twin technology—with the primary focus of integrating sensor data with BIM—via Project Dasher.

The other area of research described in this chapter relates to Generative Design, particularly in the context of the AEC industry. Autodesk's work in this area started in 2010 with a project named Dreamcatcher. It was further accelerated in 2014 with the acquisition of The Living, an architectural studio which had—under the guidance of its principal David Benjamin—started exploring the application of multi-objective optimization through the use of genetic algorithms for architecture, engineering and construction industry (AEC) workflows. They have since delivered projects at varying scales—from manufactured aerospace components, to office and exhibit hall layouts and even up to the urban scale with studies on the layout of residential neighbourhoods—that demonstrate the potential for Generative Design in the AEC space.

Both Digital Twins and Generative Design are not only useful for architecture: this chapter will show concrete examples of how they are also being applied to both engineering and construction.

Looking to the future, real-world performance data captured via IoT—and hosted in Digital Twins—will increasingly influence the Generative Design process, as Autodesk and other software providers start to complement algorithmic exploration of the design space with machine learning systems trained with data from prior projects and captured using physical sensors. This has the potential to drive "closed-loop" processes where performance increases with each design iteration or system re-configuration.

2   Early Applications of Project Dasher in Building Operations

When Project Dasher (Autodesk, 2021a) was conceived, back in late 2009, the vision was to create a "building debugger" to help people explore data captured from the built environment and allow them to make appropriate changes to drive efficiency, so they might better understand how the building was performing and discover any unexpected ways in which it was not behaving as intended.

The project became possible because of two overall technology trends: the increasing ubiquity of low-cost (and low-power consumption) sensors and the decreasing cost of cloud-based data storage. Proper instrumentation of the built environment will require not only a significant number of sensors but also cheap, accessible storage for their readings: a central thesis was that the value of the data would continue to be high (and potentially increase) as it helps build an overall impression of how performance changes over time. This thinking predated much of the current focus on big data and machine learning, but these are two examples of how this focus on data has proven to be valuable. Another central notion of the data capture was that it be generic: that the sensors be as sensitive as possible and collect higher than average frequencies of data to enable experiments in data mining and sensor fusion that would emerge over time in the style of a living laboratory.

It did indeed turn out that the data collected and made accessible was of significant value in optimizing a building's performance. One example of this was in Autodesk's former Toronto office at 210 King Street East, which was one of the first buildings to be instrumented and feed its data into Project Dasher. It was noticed that 3D prints left to run overnight in the office were often failing, while similar print jobs run during the day would succeed.

The team used Project Dasher to explore the issue and helped uncover the fact that the air conditioning would shut off every day at 6 pm, leading to the temperature in the room housing the 3D printers to spike upwards of 30°C that, coupled with the very low humidity that was exaggerated by the door to this very noisy room always being closed, led to a significant change to the ambient conditions that were at other times favourable to 3D printing with PLA plastic filament.

Beyond just helping people running facilities to optimize performance, Project Dasher always had sustainability as a core goal: as buildings contribute roughly 40% of global CO_2 emissions (UN Environment Programme, 2020), anything that can be done at a wide scale to optimize performance and reduce this impact would have a significant effect on humanity's use of the planet's resources.

2.1 NASA Ames Sustainability Base

When Autodesk Research embarked on a project alongside NASA to instrument their Ames Sustainability Base in California, it was at least partly with this vision in mind. The building was already equipped with an extensive system of sensors, some of which were built into the HVAC system. It also incorporated smart water management both inside and outside the building, with flora that was carefully selected for the local microclimate.

When the project started, there was no holistic view of the various data being captured from the building's disparate systems: a large part of the initial effort was to feed the data from more than a thousand sensors into a centralized, cloud-based, time series database.

This data could then be visualized in a 3D context with the Dasher client—originally a desktop application (Attar et al., 2010) but more recently web-based (Autodesk, 2021b), taking advantage of Autodesk's Forge platform (Autodesk, 2021c)—to allow more intuitive exploration of the building's performance data. The readings for individual sensors could be displayed in graph form, or the data for a particular zone or floor could be displayed as a heat map, showing the variations between multiple sensors. All this data could be animated via a timeline that enabled targeted and interactive exploration of a historical time period.

The graphing system—known as Splash (Glueck et al., 2014)—made use of multiple levels of detail onto the data, allowing the user to see years' worth of data and then quickly zoom into (say) an interval of just minutes or even seconds.

When the visualization shown in Fig. 1 was originally shared with NASA, a member of their team of scientists quickly pointed out an area of the building—highlighted in red—that performed differently from the rest of the floor. He immediately explained that this room was being used to perform a glazing study but that the results had not yet been analysed. In seconds, the results were clearly visible in the contextual 3D visualization. This was an early indication that this kind of explorative visualization tool could provide benefits to people operating a building.
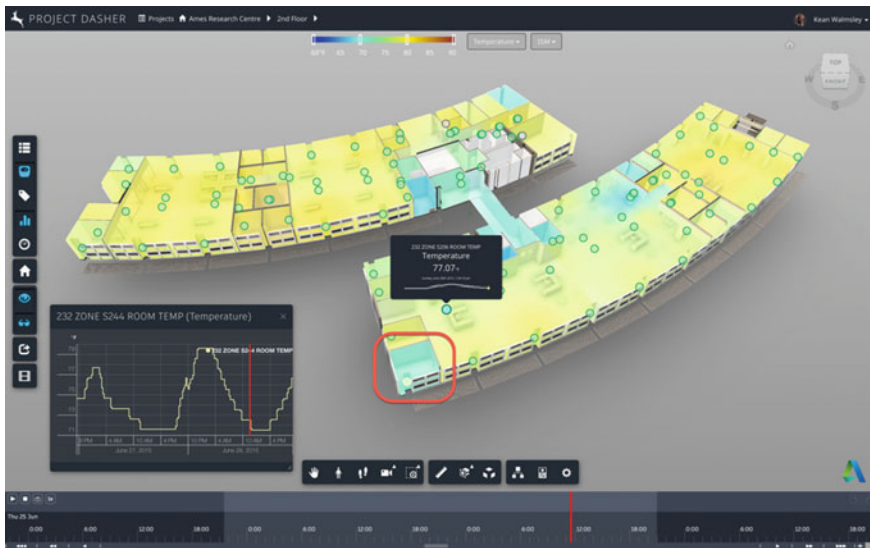

Fig. 1 Project Dasher displaying data captured from NASA's Ames Research Center (author's original)

2.2 Schneider Electric's GreenOValley HQ

Another interesting building-centric pilot project was for Schneider Electric, for a building on their GreenOValley campus in Grenoble, France. Schneider Electric had developed and installed a new occupancy sensor in their building which provided information on the approximate relative (and anonymous) locations of the building's occupants. The data displayed locations of people in the building at moments in time—so Dasher was unfortunately not able to plot the movement of people through a space—but it was certainly enough to be able to plot occupancy against CO_2 levels for a space.

Figure 2 illustrates how Project Dasher allows the correlation of sensor data, highlighting how high occupancy of a particular conference room for an extended period caused the CO_2 levels to spike. Dasher allows quick exploration of such problems, providing the ability to zoom out of the more granular data for the CO_2 sensor for that room, looking at weeks, months or even years of data at a time in order to assess how frequently this situation arises.

Fig. 2 Project Dasher displaying data from a building in Schneider Electric's GreenOValley campus (author's original)

3 From Buildings to Infrastructure: Applying Dasher to Bridge Monitoring for Pier 9 and MX3D

The potential for this technology clearly goes beyond the operation of buildings. The opportunity to explore such possibilities came out of a collaboration with Netherlands-based Joris Laarman Lab.

The laboratory had worked with Autodesk Research via a partnership with the Autodesk Technology Center in San Francisco, where they explored the possibilities for using industrial robots to 3D print large-scale metal objects and structures.

They came up with the idea of using this technique to "print" a generatively designed bridge over a canal in Amsterdam. The bridge would be a metaphor for connecting the city's rich history with its bright technological future.

This was the birth of MX3D, a spin-off from the Joris Laarman Lab that would explore this and other use cases for this innovative technique. Autodesk Research was initially involved in the project to help prototype some initial design options. The final design, shown partially in Fig. 3, was generated and developed by MX3D, but Autodesk remained involved, ultimately shifting focus away from its design to consider possibilities for monitoring the bridge's performance.

Being the first of its kind, constructed via a novel manufacturing technique—shown in Fig. 4—there was clearly no information available on how the bridge would perform when under load or across temperature extremes. A significant amount of effort was put into analysing the performance of the material, of course—primarily by project partners Imperial College, the Turing Institute and Arup—and the recommendation was made to perform structural monitoring of the bridge for a period of time.
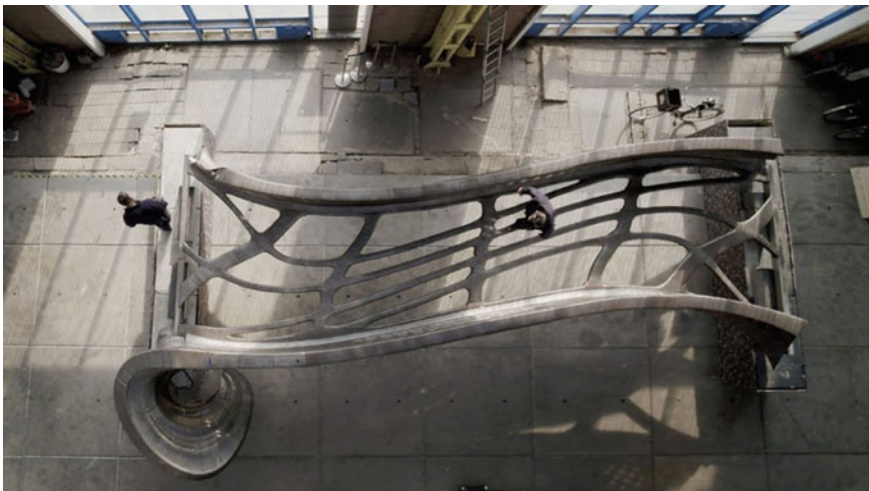

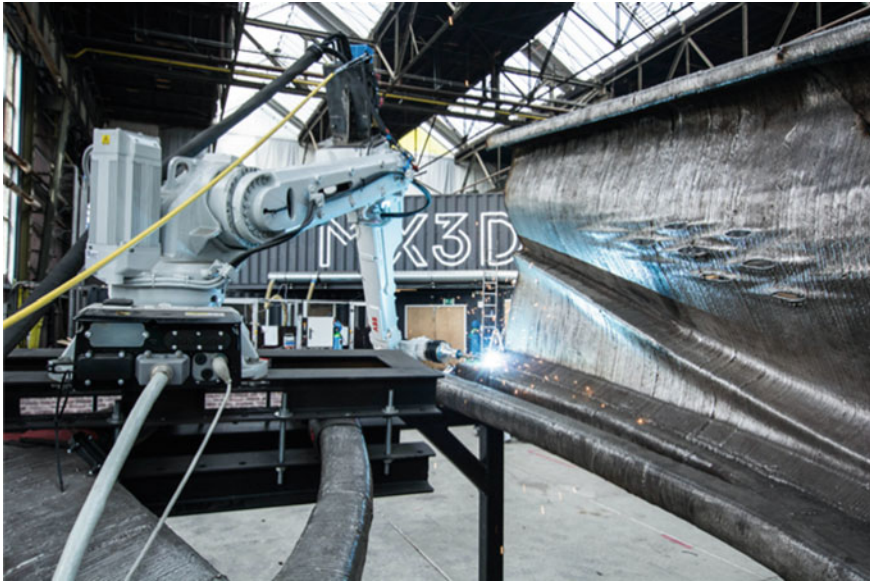**Fig. 3** MX3D smart bridge under construction (photo © Joris Laarman Lab)

Fig. 4 Industrial welding robot depositing material for the MX3D smart bridge (photo © Olivier de Gruijter)

Understanding this as an opportunity for gathering deeper insights from the bridge beyond the structural implications, a group at Autodesk identified the role that smart infrastructure could play within future cities. The idea was to integrate a "nervous system" into the bridge, allowing for constant monitoring, which could unlock the ability to gain insights into how the bridge gets used and performs.

This nervous system would come in the form of sensors on the bridge's surface that would measure temperature, load, strain, acceleration and incline, passing the readings into a cloud-based, time series database for storage and analysis.

To be useful, the data coming from some of these sensors needed to be read at a much higher frequency than in the previous contexts explored by the research team: in a building, one might read the temperature or CO_2 values from a sensor every 5 minutes or so, while an accelerometer on a bridge might be read as often as 1000 times per second (10 Hz is probably more realistic, but that is still significantly more data to marshal and store than had previously been tested with the system).

Rather than waiting for the MX3D bridge to be completed, the decision was made to install sensors on a raised pedestrian walkway connecting areas of Autodesk's Pier 9 office in San Francisco, a bridge of comparable size and function. This would allow the research team to make sure the system was able to deal with the volumes of data generated by a bridge with sensors being sampled at higher frequencies.

The Pier 9 project was an invaluable stepping stone to be able to add a nervous system to the MX3D bridge. A network of 30 sensors was installed on the Pier 9 bridge, as shown in Fig. 5, enabling the collection of data for structural movements
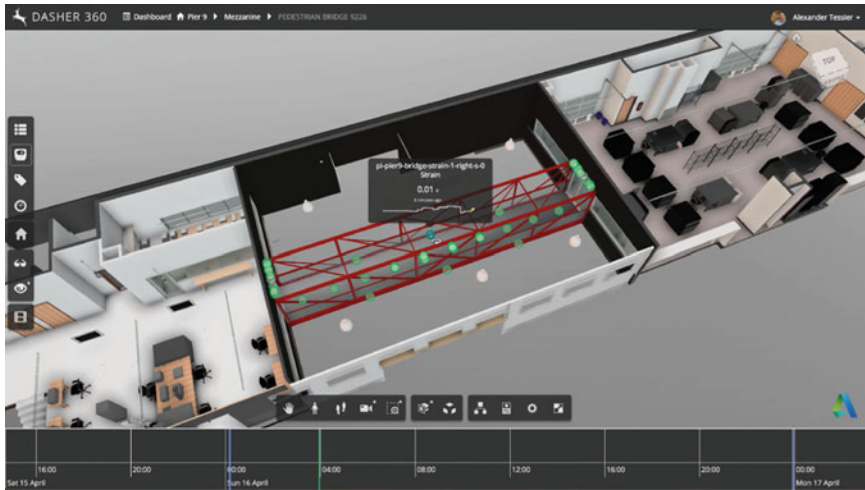
Fig. 5 Raised walkway in Autodesk's Pier 9 office inside Dasher (author's original)

through strain and acceleration, as well as pedestrian motion, sound levels and atmospheric information such as temperature, pressure, humidity and CO_2 levels in the indoor space. Visitors to the space were able to see their impact on the bridge in real time via screens displaying the sensor data inside Dasher.

The Pier 9 project provided Autodesk Research with knowledge that was highly valuable when instrumenting the MX3D bridge with its sensor network. The goal for the smart bridge project was not only to understand the bridge's performance, but also for it to sense its environment beyond the immediate. Alec Shuldiner, who helped initiate the project and drive it to its conclusion said: "I am very interested in what's happening on and around the bridge, and what the bridge can tell us about that. I'm interested in this as a sensor for the neighbourhood.".

A big piece of enabling this was a computer vision system allowing the system to "see" people using the bridge and understand how they engage with the bridge and its environment.

4 Enhancing Privacy: Using Computer Vision to Anonymize Data Capture of Occupants and Passers-By

To "feel" activity on the bridge, Autodesk tried different machine learning approaches. Early results were promising, but the team concluded early on that the best results would be obtained with a properly labelled data set for pedestrians crossing the bridge. In order get as close as possible to "ground truth" with the labelling of the dataset, a key component was the integration of synchronized video
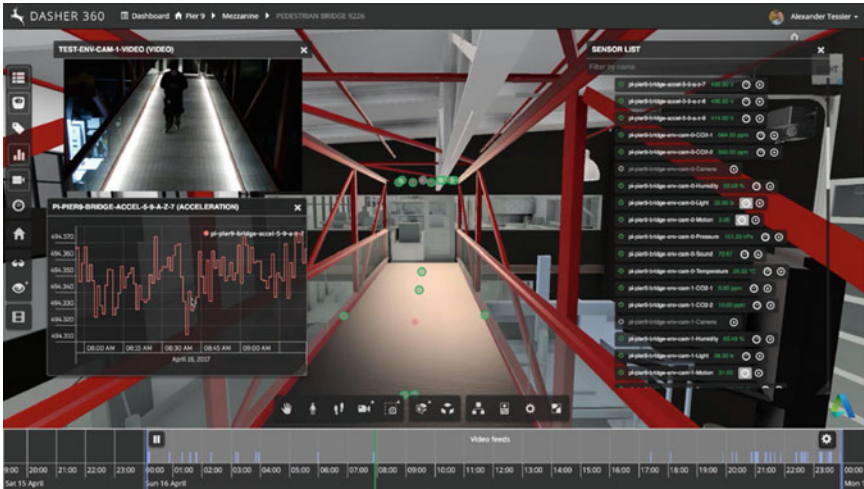
Fig. 6 Pier 9 bridge inside Dasher showing synchronized video content (author's original)

camera footage as shown in Fig. 6, allowing the team to make sense of the performance data relative to what could be seen on and around the walkway. By integrating video into the data set, new possibilities emerged for data correlation and sensor fusion beyond simple annotations.

While it was extremely valuable, the amount of video data was significant but more importantly was also highly sensitive. This led to a project—code-named Ajna—being started to explore the possibilities that modern computer vision algorithms (which effectively allow the computer to "see" real-world objects) could introduce into the context of a 3D visualization system.

Initial phases of the project required reliable pedestrian motion detection to only store video of salient events on the bridge, such as pedestrian crossings, instead of being triggered by the machine movement or lighting changes within the machine shop. These events were integrated into the timeline navigation also shown in Fig. 6. Combined with cloud storage and video compression, the motion detection events proved to be even more useful in understanding where people were on the walkway during traversals and were an integral part of data set labelling for training machine learning human presence, gait and position on the bridge.

The use of integrated video footage quickly became a highly sensitive topic within Autodesk Research. To maintain employee privacy, access to the system was extremely limited. Initial efforts at automatically blurring faces—as shown in Fig. 7—at the hardware level proved insufficient in the perception of anonymization of the data. It was also desirable to derive more complex loading conditions from the data for analysis, such as when people walked in lock step, carried heavy objects or leaned on the handrails along distinct parts of the bridge. To get at the data that was needed, yet remove sensitive individual traits, the research team turned to extracting
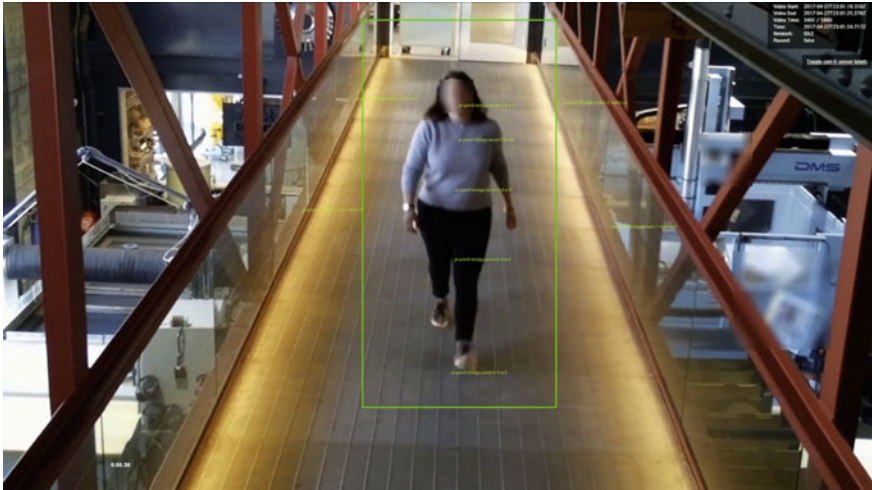
Fig. 7 Person with a blurred face crossing the Pier 9 bridge (author's original)

pedestrian skeletons as shown in Fig. 8, with the potential of negating video storage altogether, storing only the skeleton data.

Part of the project was to map this data into a 3D context to allow visualization of people walking through a space and correlation with other types of data. The initial implementation of this approach allowed mapping of skeletons into a plane that was parallel to a standard (non-depth, non-stereo) camera frustum: essentially
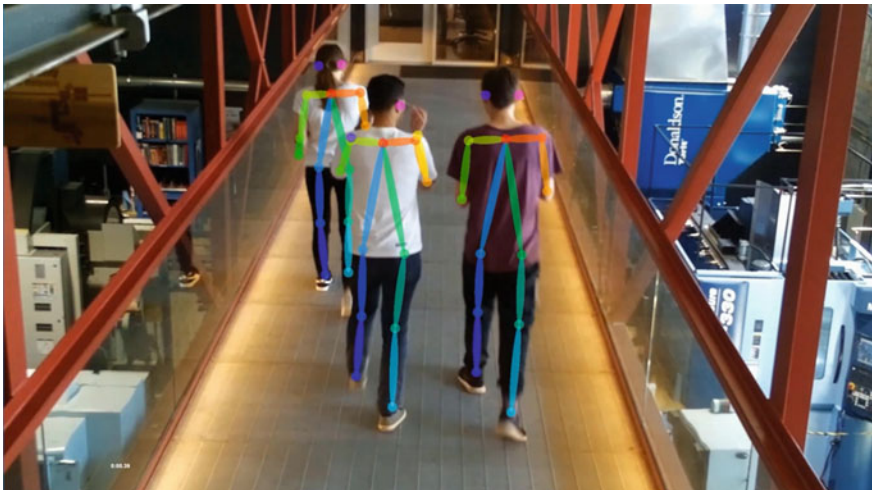

Fig. 8 Extracted skeletons overlaid onto video of people crossing the Pier 9 bridge (author's original)
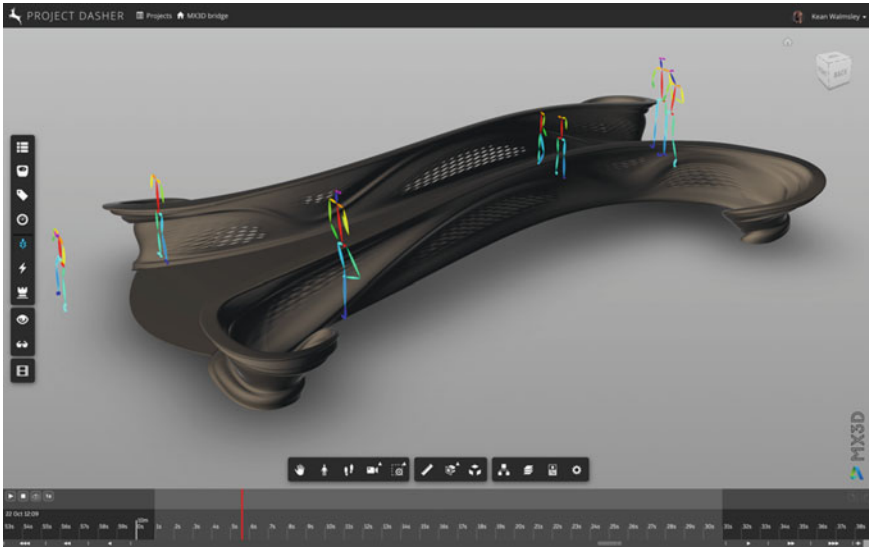
Fig. 9 Displaying 2.5D skeletons extracted from video footage inside Dasher (author's original)

giving 2.5D, in that flat skeletons were positioned correctly in 3D space, as shown in Fig. 9.

A more recent phase of Project Ajna harnessed machine learning—in that it used a library that was trained against various typically body configurations (Kocabas, 2019)—to extract and display full three-dimensional skeletons, an important milestone in the project.

While there was clearly sensitivity about employee privacy in the context of the Pier 9 project, the concerns around privacy of people crossing a bridge placed over a canal in the red light district of Amsterdam were much higher. It was essential that the project complied with GDPR, of course, but also that it met the privacy needs of people who would use the bridge.

The opportunity this kind of positional data provides is significant: firstly, it is possible to compare what is known about people crossing the bridge—as detected by the cameras—with the information from sensors about how the bridge is behaving structurally. Correlating this data inside Dasher, it is possible to understand the specific impact of people crossing the bridge on its performance. See Fig. 10 for an example of how positional information can be correlated with strain information via the combination of skeletons and surface shading.

Secondly, it is now possible to analyse the flow of people across the bridge and the kinds of behaviours they exhibit—anonymously. One can start to reason on how people in Amsterdam move around the city and interact with their infrastructure. The MX3D bridge is ultimately an art exhibit—which does mean that people will engage with it very differently from a more functional piece of infrastructure—but this type
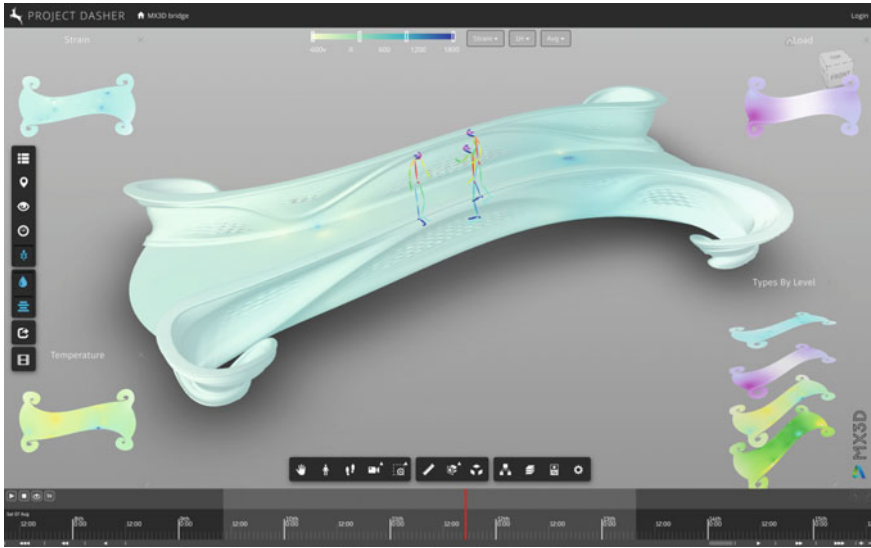
Fig. 10 MX3D bridge displaying 3D skeletons and strain data shaded onto its surface inside Dasher (author's original)

of project paves the way for a future of smart infrastructure that helps the city of the future better meet the needs of its inhabitants and visitors.

5   First Steps with Generative Design: The Airbus Bionic Partition

Autodesk Research started exploring the potential for Generative Design in 2015, shortly after The Living—a New York-based architectural studio headed up by David Benjamin—was acquired by the company.

Generative Design is a methodology where the power of computation is used to explore a solution space for a particular design problem. The performance targets for a design are described in some way, and the system will generate design variants that are measured against these goals.

The project with Airbus was in many ways the beginning of Autodesk's exploration of applications for generative design.

Each A320 airliner currently sports a 65 kg partition at the back—separating the passenger compartment from the galley—that supports two crew seats and can be opened to make space for a stretcher to navigate around the tight corner leading out of the plane in the case of a medical emergency.

The aim of this project—a collaboration between Airbus group's APWorks and Autodesk Research's The Living—was to design a partition that was 50% lighter
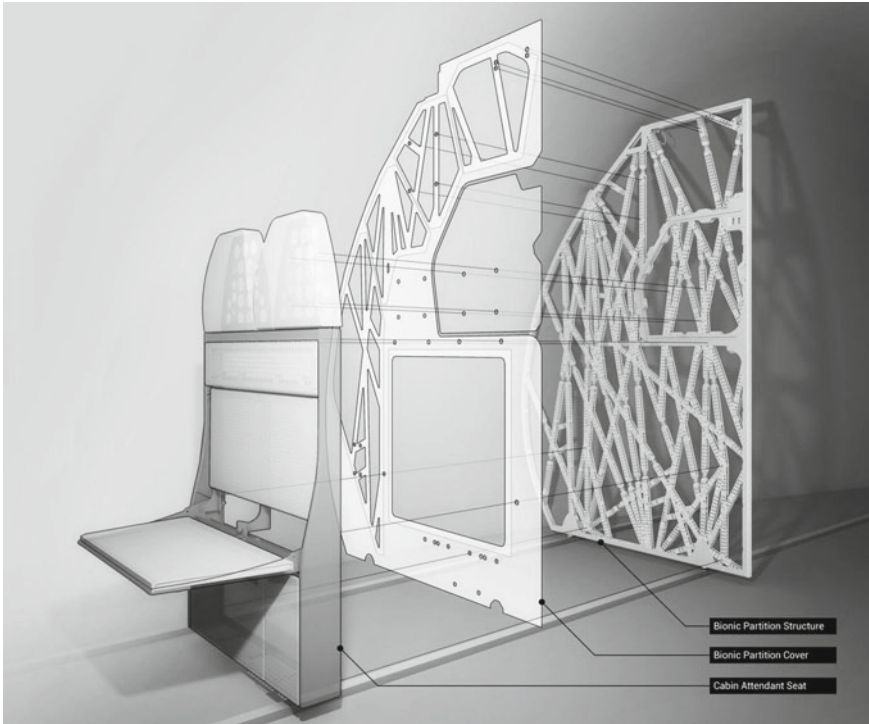
Fig. 11 Airbus A320 Bionic Partition (author's original)

while maintaining or improving its structural properties (Nagy et al., 2017a). The structure and purpose of this partition are shown in Fig. 11.

In October 2020, there were more than 9000 A320s in active service, making it the highest-selling airliner. Reducing the weight of this partition by 50% across all A320s would reduce the amount of fuel needed and—in aggregate—save 1 million tons of CO_2 emissions each year.

For this project, there were two primary metrics used to evaluate each design option: weight and displacement under structural load. The aim was to minimize both of these metrics.

The parametric model used to generate the various options encoded an algorithm inspired by nature: physarum—a type of slime mould—creates efficient, redundant networks while seeking (and linking) its food sources, as shown in Fig. 12.

Physarum sends out tendrils seeking food: those finding it are strengthened, while those that do not wither away. This basic mechanism—coded as a software algorithm—was used to generate many different design alternatives and effectively explore the solution space for this problem. The system was deliberately designed to evaluate options that were outside the space a human designer would typically consider.

Fig. 12 Physarum (slime mould) tendrils connecting food sources. Source: https://en.wikipedia.org/wiki/Physarum_polycephalum#/media/File:Physarum_polycephalum_plasmodium.jpg

The results of this generation process could then be explored by considering the trade-off between weight and displacement: as demonstrated in Fig. 13, it is easy to reduce displacement while increasing weight and vice versa.

It is here that the designer places a key role in the process, once again: using their intuition to guide the appropriate trade-off between how the designs meet the stated objectives.
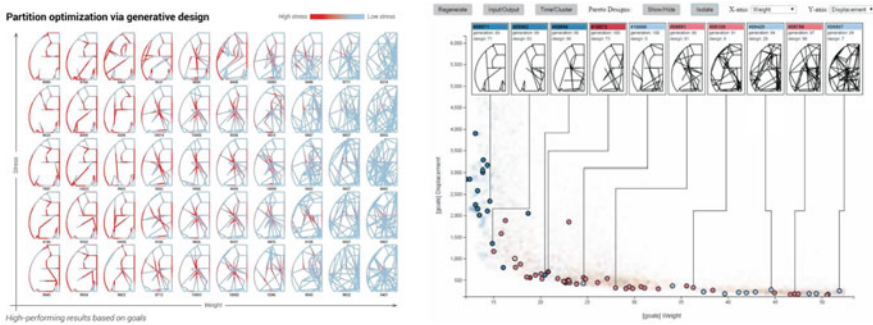

Fig. 13 Generated designs mapped based on their performance (author's original)

When a particular design was selected, a secondary optimization process started: each of the "macro" bars in the design had its topology optimized, replacing each with a set of microbars of varying thickness and strength depending on the load they needed to support.

Of course, extensive testing was performed, both virtual (via finite element analysis) and physical. The results were consistent: the physical objects failed in highly predictable (and predicted) ways.

At the time of the project, the technology did not yet exist to 3D print the entire partition in one piece, so it was divided into 122 parts that were printed in Scalmalloy, a high-performance aluminium–magnesium–scandium alloy designed for metal 3D printing. These parts could be printed in batches, of course, as shown in Fig. 14.

Figure 15 illustrates how these 122 parts were then assembled into a single object using 40 titanium connectors.

These connectors only added 2–3% to the overall weight and allowed the partition to be crated up and transported to Las Vegas for display in the Autodesk University Exhibition Hall, as shown in Fig. 16.

The question remains—longer term—of the appropriate fabrication granularity for this partition: having it all in one piece is likely to be desirable, structurally speaking, but any damage would then mean the full panel would probably need replacing, rather than a smaller component. This is ultimately a manufacturing decision rather than a drawback of the design process, of course (Fig. 17).

The finished partition ended up being 45% lighter—at 35 kg—while being slightly stronger—it had a displacement of 99 m rather than 108 mm—and otherwise
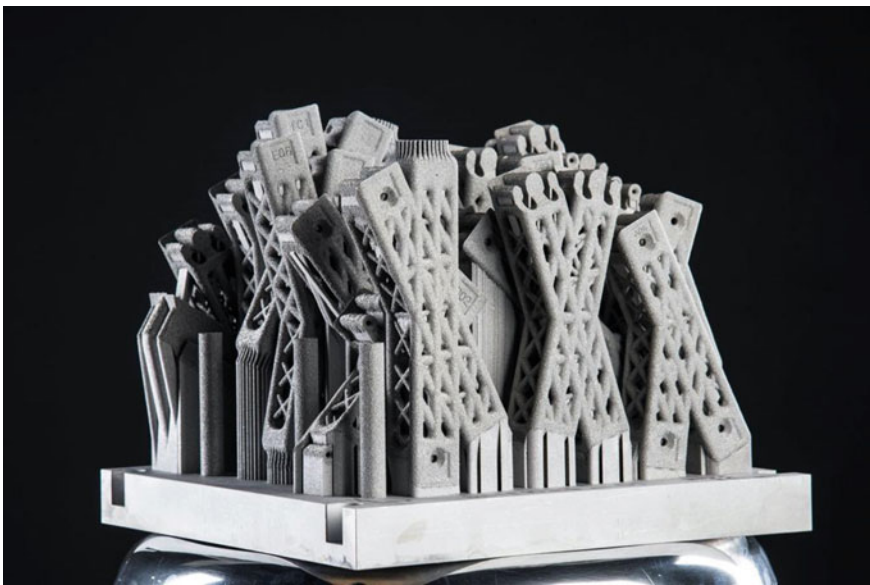

Fig. 14 Parts of the Bionic Partition printed together from Scalmalloy (author's original)
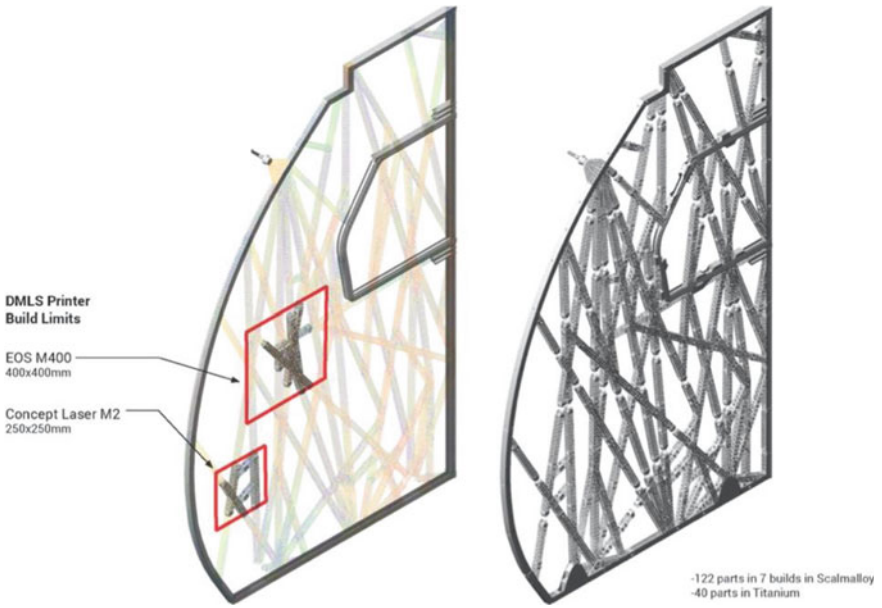
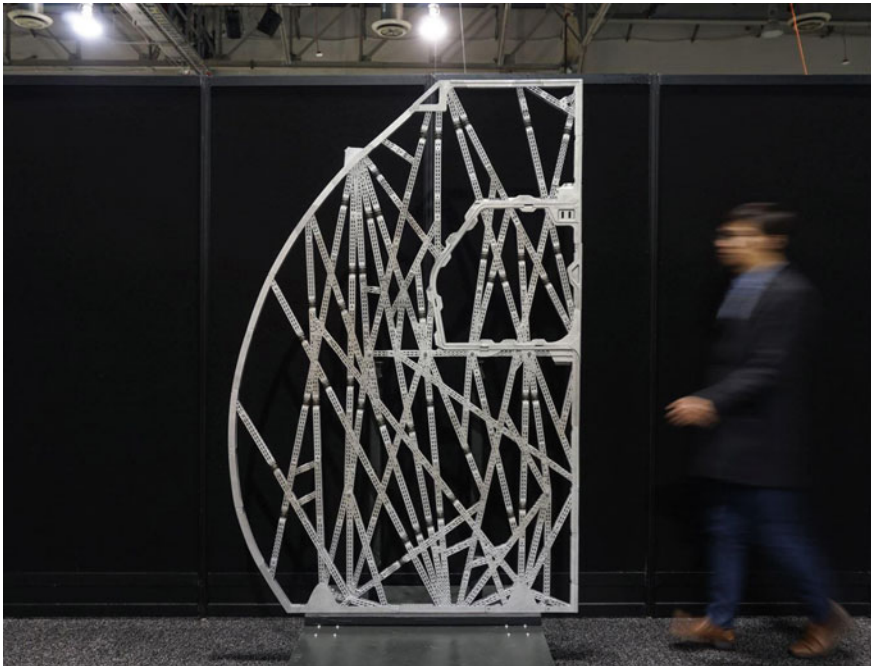**Fig. 15** Titanium connectors between the Scalmalloy parts (author's original)

Fig. 16 Airbus Bionic Partition on display at Autodesk University 2015 (author's original)
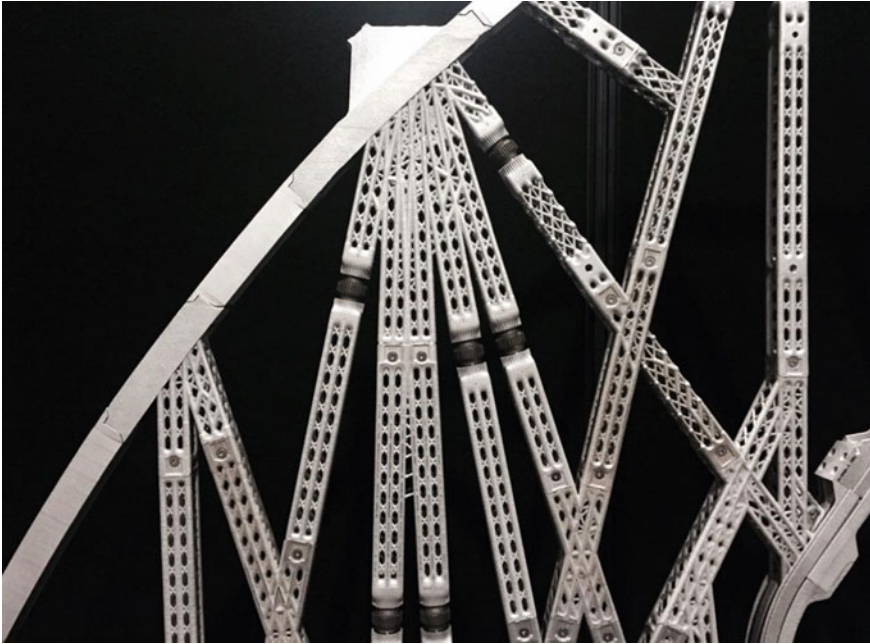
Fig. 17   Closer view of the Bionic Partition (author's original)

performed comparably to the original. Figure 18 illustrates the relative displacement of the existing and new partitions.

The Bionic Partition project was an important step in exploring the possibilities around applying Generative Design for aerospace: it was a relatively simple—when compared with the challenge of designing the airframe—and low-risk way of exploring the technique, while the longer-term opportunity is clearly to use such an approach for more fundamental and radical design work. This will hopefully become feasible as the technology matures.

6   Designing an Office Generatively: Project Discover and Autodesk's Toronto Office

Toronto is an important location for Autodesk Research, with a significant portion of its staff based there. When Autodesk was considering moving its Toronto office to the MaRS district of the city—the largest urban high-tech incubator in North America—there was an opportunity to use the office as a showcase for architectural-scale generative design. The technique had been tested with just two metrics in the Airbus Bionic Partition project, but could it scale effectively to help generate an
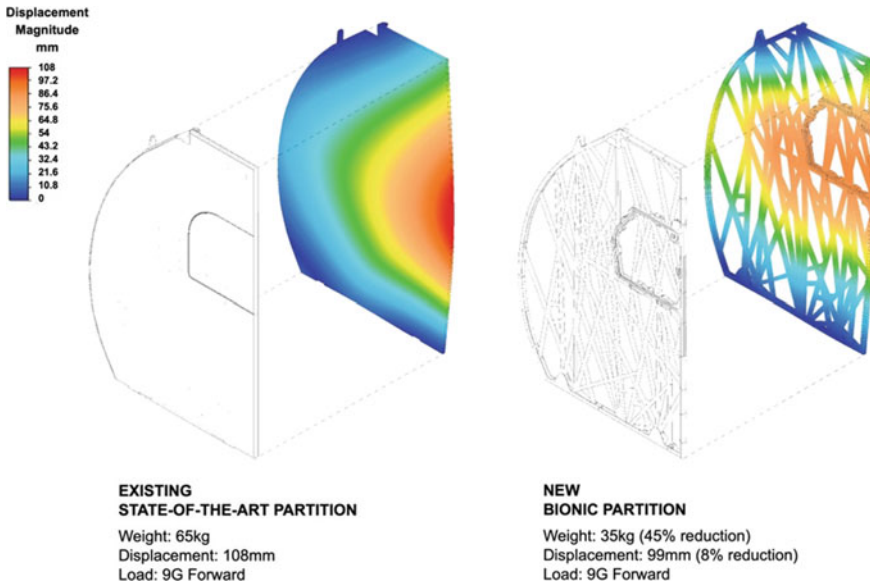
Fig. 18 Comparison of the old and new partitions (author's original)

architectural layout of an office space based on a larger number of performance goals?

Project Discover was designed to answer this very question (Nagy et al., 2017b). The first phase of the project helped establish what Autodesk's Toronto-based employees want from an office space: the employee base was surveyed to understand their requirements and preferences, and—of course—the corporate facilities and HR teams were involved to make sure business needs were considered, too. Based on this input, it was possible to craft a set of metrics that could be used to assess the quality of a particular design.

6.1 Evaluation Metrics

There were six metrics chosen to evaluate potential solutions for this design problem: adjacency, workstyle preference, interconnectivity, distraction, daylight and views to outside, as shown in Fig. 19.

6.1.1 Adjacency

The adjacency metric measured the distance for each employee to travel from their desk to a set of preferred neighbours and amenities. The main algorithm used for this
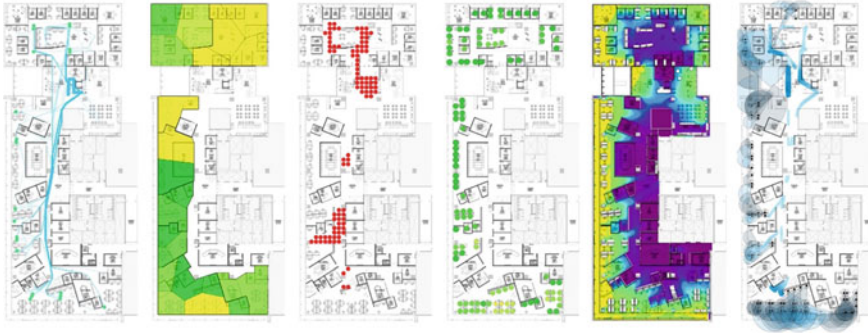
Fig. 19 Evaluation metrics (from left to right: adjacency, workstyle preference, interconnectivity, distraction, daylight, views to outside) (author's original)

was a "shortest path" analysis, to help understand the minimal path of travel from one point in the office to another. The score—between 0 and 10—was an indication of how little individuals need to travel to get to key points in the office (some of which are standard, others are specific to that individual). A score of 0 meant occupants had a high cost of travel, a score of 10 indicated the lowest cost.

6.1.2   Workstyle Preference

The workstyle preference metric measured the suitability of a neighbourhood to the assigned team's preferences. It determined how closely a team's preference and weighting of ambient conditions (light and activity) were met by their assigned neighbourhood. At the extremes, the score indicated whether none of the teams had their preferences met (0) or whether all the teams had them met (10).

6.1.3   Interconnectivity

The interconnectivity metric considered the amount of likely congestion in an office layout based on the paths people will take through the office. Congestion can be considered negative—which is certainly the case when prioritizing office designs that minimize the propagation of viruses through human contact—but in the case of this project, the team was looking to encourage serendipitous encounters and "water cooler" discussions. High interconnectivity scores are also likely to impact occupants looking for quiet work areas, something that will be discussed further for the distraction metric.

6.1.4   Distraction

The distraction metric measured the amount of negative visual and auditory activity from individual workspaces. It counted the number of colleagues who were in an employee's field of view when seated at their desk and considered desks that were close to zones with high auditory activity (as measured by interconnectivity). It scored designs poorly (0) that had all workstations with high visual/auditory distraction and scored them highly (10) if all workstations had no visual/auditory distraction.

6.1.5   Daylight

This was a measurement of the daylight levels in workspaces and amenity spaces. Industry-validated methods were used to calculate light levels via LEED v4 standards, with the score indicating the amount of occupied space that has adequate natural light at both 9 am and 3 pm. A score of 0 indicated 0% of the occupied floor area had adequate lighting while a score of 10 indicated 75% or more.

6.1.6   Views to Outside

This metric indicated whether people at their desks or walking through the office would have a view to the outside. An isovist calculation was performed from each of the sample points, with the overall score indicating how many of these points would have an unobstructed view of a window.

6.2   Design Generation

The next phase was to build a parametric model that could potentially generate thousands of different design variants, each of which could be assessed using the six metrics.

   The inputs to this model would vary the way floorplans get divided into neighbourhoods—a Voronoi pattern was created from a set of variable neighbourhood centres and was then used for the division process—at which point the algorithm could place desks and amenities and then allocate teams to areas of the office.

   While this model could be used to generate many random design options, it was unlikely this would result in finding the highest-performing designs: when a design problem has high dimensionality—and this one has six dimensions—it is (a) unfeasible to perform an exhaustive search by generating every possible design and (b) unlikely that a random search (or even a systematic chopping up of the input parameters to perform "optioneering") would find the most interesting solutions.

   Multi-objective optimization can provide a more intelligent search of a solution space: an optimization engine makes use of a genetic algorithm to maintain a pool

of high-performing designs, and—generation on generation—uses genetic operators such as selection, crossover and mutation to seek even better solutions based on this population. The inputs to the model are tweaked based on the parent solutions' and the child's metrics which are then evaluated to see whether they perform better or worse.

The specific algorithm used during Project Discover was NSGA-II, a genetic algorithm that has been used effectively in a number of different design-related activities such as printed circuit board layout for electronic engineering.

6.3 Design Exploration

The Generative Design process typically does not result in a single best design option: a set of results will be generated, many of which perform well against different goals. Therefore, a key piece of such a Generative Design system needs to be an environment allowing designers to explore the generated results and ideally gain an understanding of trade-offs between different performance metrics.

For the Airbus project—with two metrics—it was a straightforward process to scatter plot the results with one metric on the X axis and the other one on Y, helping us understand the inherent trade-off between weight and displacement (heavier designs typically displace less than lighter ones).

For Project Discover—with six metrics—things were more complicated: while it was possible to create scatter plots with four different metrics (in addition to X and Y, we can use size and colour to convey additional information), it was likely that over time this process would be applied to problems with even higher dimensionality. The future generative designer will need an environment that allows them to assign metrics to the axis of their choice, helping them gain a sense of any trade-offs and to look for interesting solutions.

Figure 20 shows the designs generated during Project Discover plotted based on their interconnectivity (Buzz) and adjacency preference metrics.

This exploration stage can be extremely valuable, in and of itself: it is at this point that the designer is likely to be presented with unexpected results. The generative system—if implemented properly—is not subject to the same biases as a human designer. We are formatted in such a way to typically favour rectilinear layouts, for instance, or at least to have angles that are consistent with an overall style. The generative process is not necessarily limited by these biases and so can often generate results that challenge the designer's natural inclinations.

For instance, Fig. 21 shows interesting results from Project Discover, some of which were unexpected:

1. Multiple types and sizes of amenity spaces surround each neighbourhood, leading to better scoring for productivity.

Fig. 20 Scatter plot of Project Discover designs (author's original)

2. Residual, irregularly shaped areas become semi-private informal social spaces that, while performance neutral, were unexpected, interesting design elements and well-received by the clients.
3. A diagonal line between neighbourhoods allows fitting more meeting rooms while giving each neighbourhood its own character.
4. Unusually shaped room used for open-ended activities.
5. Stepped walls create nooks as a threshold between public and private spaces.
6. Non-orthogonal, non-parallel boundaries obscure sources of distraction (desks in adjacent neighbourhoods and busy corridors) to improve productivity scores.
7. A back-alley connection between neighbourhoods, leading to a better score for adjacency.
8. A neighbourhood expanding out towards the window because the team prefers natural light.
9. A neighbourhood contracting towards the window because the team prefers less distraction from outside.
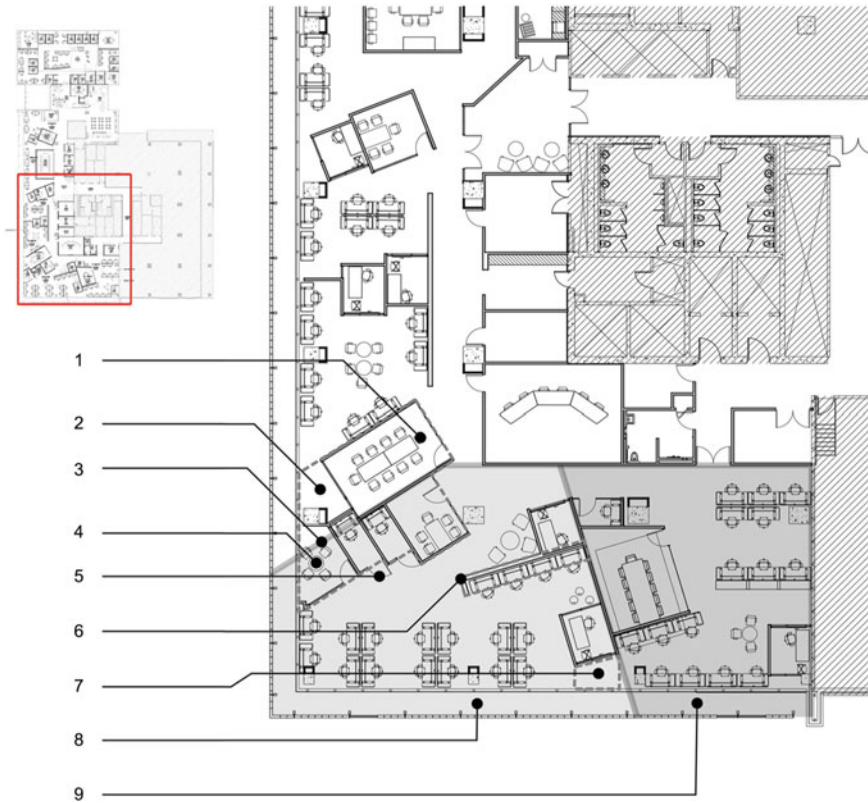
Fig. 21 Interesting outcomes in the design of the MaRS office (author's original)

6.4 Design Elaboration

Once a set of interesting designs has been identified, there is often a review process with project stakeholders. It is typical, at this stage, for a decision to be made on which of the design variants to elaborate further. It is worth noting that the metrics are an extremely valuable tool for this decision, as the selection process becomes much more data centric.

It is also worth noting that the resulting "design" is often far from being a complete blueprint that can immediately drive construction at the current stage of technological maturity: the value is in having important, high-level design decisions suggested by the generative process, while the detailed design work is likely to be performed manually, for now. This will certainly change as the technology matures—as this detailed work is often very time consuming and would also benefit from automation—but with early efforts such as Project Discover, this step was performed in a traditional way that would not impact the high-level evaluation metrics.

So while the generative process does not necessarily result directly in a completed design—today, at least—if successful, it should provide inspiration of interesting ways to attack the design problem, with data to support the inherent logic of the design.

7   Urban Scale: Applying GD for Residential Neighbourhood Layouts with Van Wijnen

Having demonstrated the opportunity for Generative Design to be used for architectural space planning through Project Discover, Autodesk Research was approached by the Netherlands-based development and construction company Van Wijnen to explore the potential of applying this methodology at the urban scale (Nagy et al., 2018).

Van Wijnen builds neighbourhoods of residential homes via an efficient, standardized process. They make heavy use of modern building techniques such as offsite fabrication and modular construction: it is this modularity and standardization that made it feasible to build a Generative Design workflow for the urban scale.

Seven metrics were identified to measure the quality of designs: project cost, profit, solar gain, backyard size, exterior views, programme and variety, as shown in Fig. 22.

The geometry system was driven by inputs indicating the location of streets intersecting the lot, with the various steps shown in Fig. 23.

The first step was to create a boundary-sensitive subdivision mesh that covers the layout's boundary. This will be the same for every design generated for this particular layout, of course, so could be generated just once.

In the second step, streets were placed based on the input parameters to the model, which were then used to subdivide the mesh into discrete lots containing parcels.

In the next two steps, houses and apartment buildings were placed into the allocated parcels.

Finally, the programme for the layout was allocated.

The whole process was driven by a small number of input parameters: the location of streets intersecting the layout. Everything else flowed from this.

Using this parametric model—which encapsulates both the geometry system and an implementation of the metrics that can be used to evaluate each design—the generative process can search the solution space for high-performing designs.

As discussed previously, an exploration step helps the designer to understand the trade-offs between different metrics and identify the most interesting designs to present to stakeholders, as shown in Fig. 24.

Once the various project stakeholders have assessed the selected designs—a process that is greatly facilitated by having data that supports the decision, in the form of the evaluation metrics—a final design can be selected for further refinement, as shown in Fig. 25.
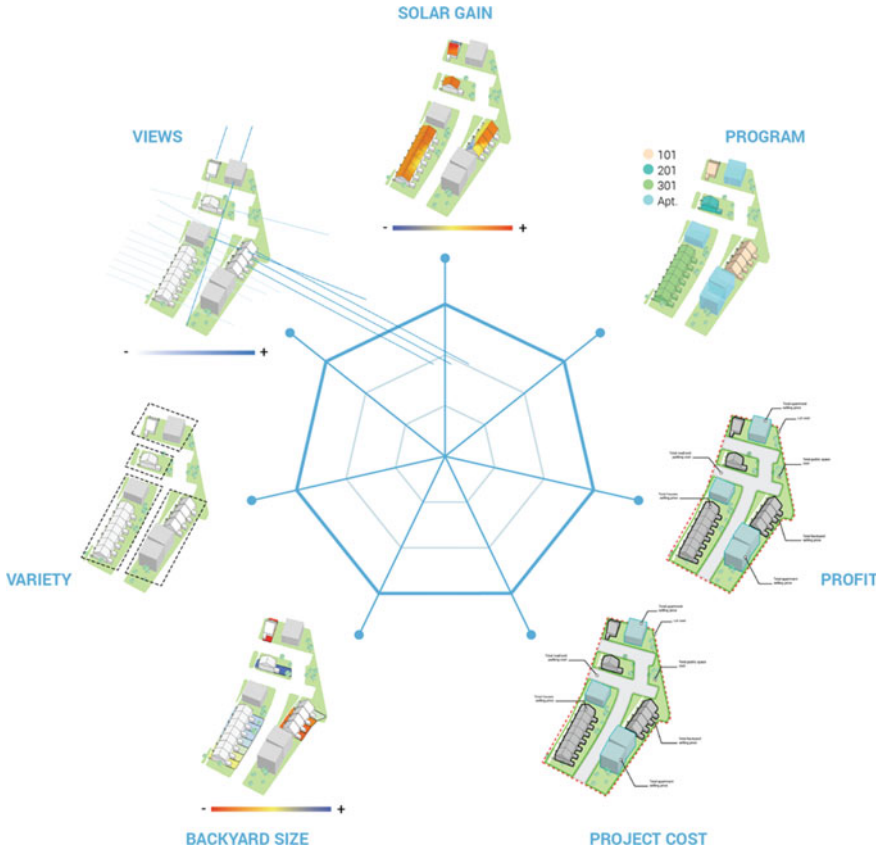
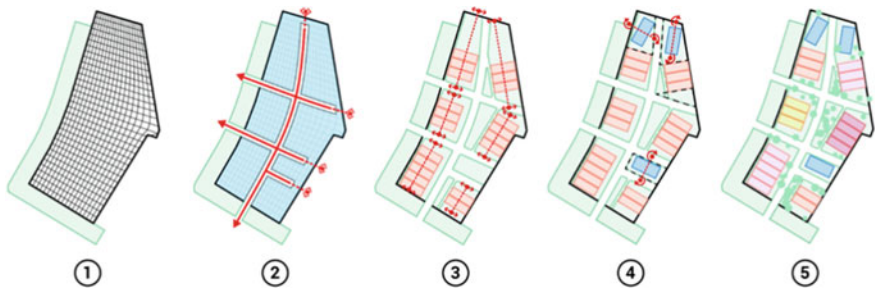**Fig. 22** Evaluation metrics for the project with Van Wijnen (author's original)

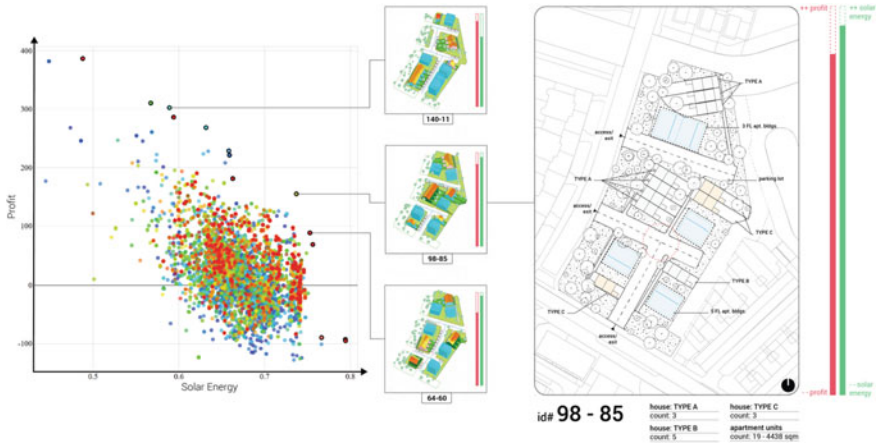**Fig. 23** Steps to define the geometry for a residential layout (author's original)

Fig. 24 Exploration of the generated designs and selection for stakeholder communication (author's original)
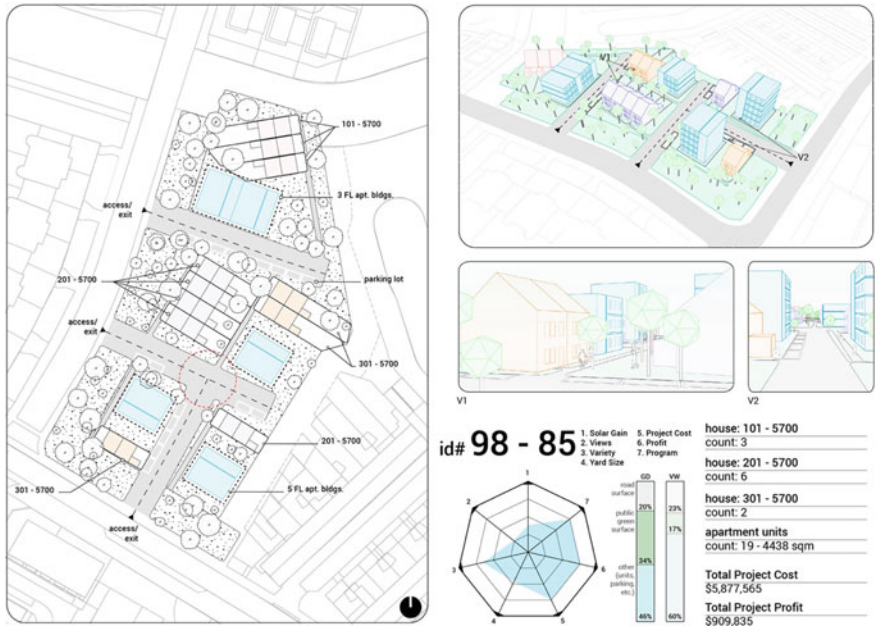

Fig. 25 Developed design (author's original)

This first application for Generative Design at the urban scale has opened the door for further experimentation and development in this space.

When asked about the potential for this technology, Jelmer Frank Wijnia, Generative Design Lead at Van Wijnen, said the following:

In the end, Generative Urban Design will be a big timesaver for Van Wijnen - in seconds functional designs are generated - in a holistic way. It is really easy to see which project scores best on goals set by the user.

Still, lots of functionality needs to be implemented; the more the better. Meeting the right criteria, setting the right goals for a project will take time.

Therefore, at this moment Generative Urban Design feels more like a partner in crime than replacing the job of the designer.

8   The Future: Closing the Loop Between Digital Twins and Generative Design

It is interesting to consider an analogy for the two areas of research highlighted in this chapter: to think of them as being like the two hemispheres of the human brain.

On the one hand, we have the data-centric workflows we have explored via Project Dasher, where the real value is in collecting measurements and taking a highly analytical view of how buildings and infrastructure perform. We can think of this as being a left-brain function.

On the other hand, we have the more creative processes in Generative Design. While not strictly creative, it is sometimes hard to tell the difference between true creativity and something that looks a lot like it. This could be considered more of a right-brain activity.

Things become more interesting when considering opportunities to combine these two types of activity: there is significant potential for the historical data collected when building a Digital Twin to influence a Generative Design process.

For instance, the strain data captured for the MX3D bridge could help engineer the next version (should there be one) more optimally: as the first of its kind, it was always going to be the case that the bridge would be over-engineered until the characteristics of the material in the long-term could be fully understood. Real-world performance data can help future iterations of a design use less material and fewer resources.

Similarly, an office could be laid out based on collected data relating to its average occupancy and the typical behaviours of its occupants.

The data could, of course, inform the design activities performed in a more traditional manner, but the greater opportunity is to have the data drive a generative process, whether influencing it directly or via a surrogate model such as a neural network that gets queried as designs are generated and evaluated.

This could also reduce the complexity of the parametric model that currently needs to contain the full logic to define the geometry of many different design variations. Machine learning could be used to encode a "style" based on prior projects and

allow this to be integrated into the generative process, whether in the creation of the geometry or the evaluation of designs relative to this style.

Projects combining the two disciplines are now starting to emerge, but the journey is just starting.

One early example relates to the use of Generative Design to create the layout for the exhibit hall at Autodesk University 2017 (Nagy & Villaggi, 2021): the design process took into account prior designs—in terms of the underlying logic used to create designs for the space—but there was no data representing the movement of people through an exhibition area. Autodesk Research was able to install cameras that monitored the movement of people through the AU 2017 space—once again using the anonymization technology from Project Ajna—that would allow post-event analyses to validate the quality of the selected design, and whether assumptions about the evaluation metrics were correct.

The longer-term opportunity offered by this data is to have the flow information more directly influence the next iteration of the Generative Design process. The loop is now closed, and over time—and ideally through multiple iterations—the assumptions will be validated and the model improved to reflect reality.

9   Conclusions

Automation is changing the way work is performed across many industries, as "software eats the world" (Andreessen, 2011). As technology advances—whether sensor hardware, cloud-based storage and compute, machine learning or genetic algorithms—there are significant opportunities to make sense of the built environment and to use these insights to improve its operation and to influence the next generation of design. Autodesk Research has been exploring these possibilities for the last decade or more, finding significant potential to apply technology to improve the way things are built and operated. Technologies developed during this time are becoming ready for mainstream adoption, with several vendors providing Digital Twin platforms and others delivering tools enabling Generative Design workflows. Autodesk Research believes that in the longer term, these two areas of technology will converge, closing the loop on the design process and enabling real-world performance to influence many processes that until now have been disconnected from this source of knowledge.

References

Andreessen, M. (2011). Why software is eating the world. *The Wall Street Journal*. https://www.wsj.com/articles/SB10001424053111903480904576512250915629460

Attar, R., Hailemariam, E., Glueck, M., Tessier, A., McCrae, J., & Khan, A. (2010). *BIM-based building performance monitor*. In Proceedings of the 2010 Spring Simulation Multiconference (SpringSim '10).

Autodesk. (2021a). Autodesk Research's project Dasher [WWW Document]. Autodesk Forge. https://www.autodesk.com/research/projects/project-dasher. Accessed on 24 September 2021.

Autodesk. (2021b). Dasher 360 [WWW Document]. https://dasher360.com. Accessed 24 September 2021.

Autodesk. (2021c). Autodesk Forge [WWW Document]. Autodesk Forge. https://forge.autodesk.com. Accessed 31 August 2021.

UN Environment Programme. (2020). Building sector emissions hit record high, but low-carbon pandemic recovery can help transform sector—UN report [WWW Document]. https://www.unep.org/news-and-stories/press-release/building-sector-emissions-hit-record-high-low-carbon-pandemic. Accessed 24 September 2021.

Glueck, M., Khan, A., & Wigdor, D. J. (2014). *Dive in!: Enabling progressive loading for real-time navigation of data visualizations*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14). ACM.

Kocabas, M. (2019). *VIBE: Video inference for human body pose and shape estimation*. https://arxiv.org/abs/1912.05656

Nagy, D., Zhao, Z., & Benjamin, D. (2017a). *Nature-based hybrid computational geometry system for optimizing component structure*. Design Modeling Symposium 2017.

Nagy, D., Lau, D., Locke, J., Stoddart, J., Villaggi, L., Wang, R., Zhao, D., & Benjamin, D. (2017b). *Project discover: An application of generative design for architectural space planning*. In Symposium on Simulation for Architecture and Urban Design 2017.

Nagy, D., Villaggi, L., & Benjamin, D. (2018). *Generative urban design: Integration of financial and energy design goals in a generative design workflow for residential neighborhood layout*. In Symposium on Simulation for Architecture and Urban Design 2018.

Nagy, D., & Villaggi, L. (2021). *Generative design for architectural space planning*. Autodesk University. https://www.autodesk.com/autodesk-university/article/Generative-Design-Architectural-Space-Planning-2020