Form-Making in SIFT Imaged Environments

Matthew Parker
Joshua M. Taron
University of Calgary

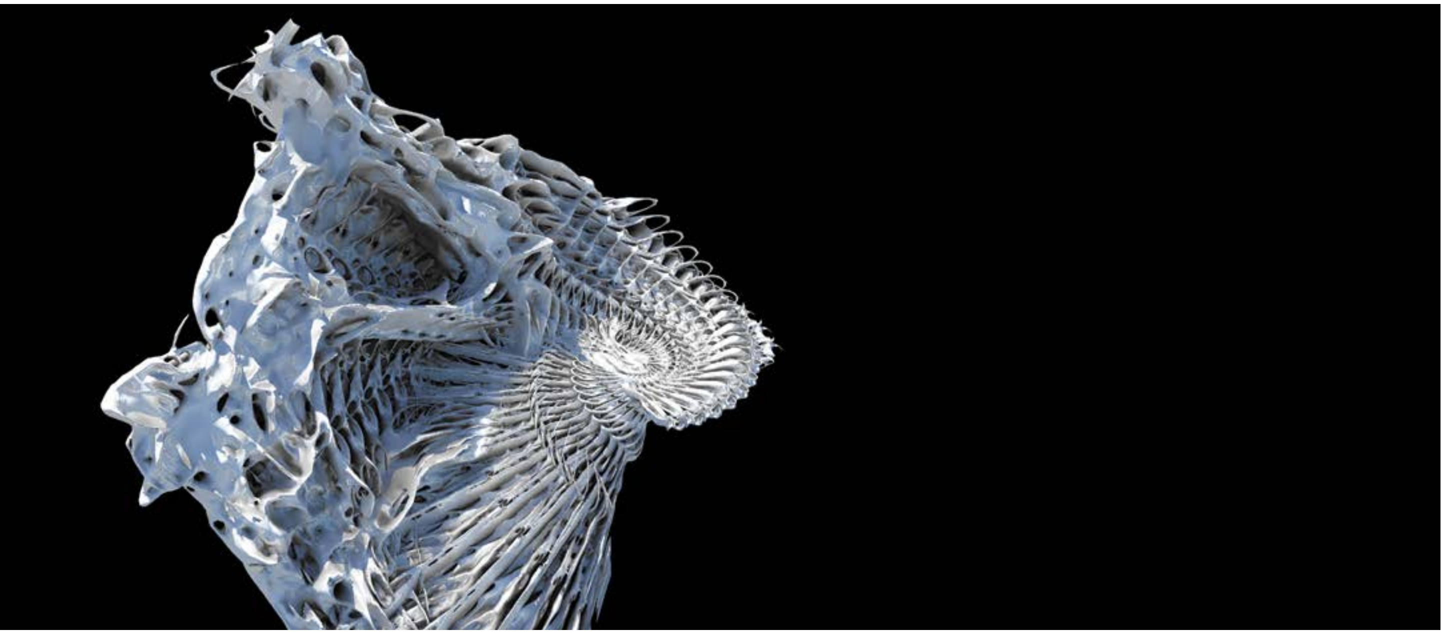

1 Indeterminate n-dimensional form that results from turbulence produced by algorithmic observation.

ABSTRACT

Within the contemporary condition, turbulence that confronts architecture is no longer unpredictable weather patterns or wild beasts, but the unintended forces of a constantly connected digital infrastructure that demands constant attention. If, as Mark Wigley puts it, "architecture is always constructed in and against a storm" it is time for architecture to reevaluate its ability to separate us from a new storm—one that situates technology, global connectivity, human, non-human and composite users, and algorithmic architecture itself as new weather systems. Toward this end, this paper explores architecture's ability to mediate and produce algorithmic turbulence generated through image-based sensing of the built environment. Through a close reading of Le Corbusier's *Urbanisme*, we argue that for much of the 20th and the early part of the 21st century, cities have been designed to produce diagrams of smooth and homogeneous flows. However, distributed personal technologies produce virtual layers that unevenly map onto the city, resulting in turbulent forces that computational platforms aim to conceal behind a visual narrative of accuracy, cohesion, anticipation, and order. By focusing on *SIFT algorithms* and their ability to extract n-dimensional vectors from two-dimensional images, this research explores computational workflows that mobilize turbulence towards the production of indeterminate form. These forms demarcate a new kind of challenge for both architecture and the city, whereby a cultural appetite to deploy algorithms that produce a smooth and seamless image of the world comes hand in hand with the turbulent and disruptive autonomy of those very same algorithms. By revisiting *Urbanisme*, a new set of architectural objectives are established that contextualize SIFTs within an urban agenda.

INTRODUCTION

"Ambiguity is absorbed when it becomes an unambiguous goal. Architects keep turbulence as a secret, secreting it inside to produce the effect of a clear line between the wild and the calm, such that this line, the facade of the building pressed up against the ever-changing exterior and exposed to unpredictable forces, is seen to be unchanging and therefore the calmest line of all." Mark Wigley

Western architectural theory contends that all buildings descend from a primitive hut, a small structure consisting of four poles, four beams, and a roof. The parable of the primitive hut demonstrates a desire to order and stabilize an unruly and chaotic world through the production of interiority that shelters us from a turbulent exterior. It wasn't until Le Corbusier's embrace of technology in the 1920s that an alternative to this framing was produced. For Le Corbusier, shelter from the outside was not only an insufficient task, it was a mode of framing that would have doomed architecture to obsolescence within an increasingly technological city. Le Corbusier recognized the architectural interior not as a shelter from the exterior, but rather as a point of interface and engagement with an ordered city. In *Urbanisme* (1924), the city is framed as a landscape of flows that relies on architecture to mediate the exchange between interior and exterior. But this mode of framing complicates the viability of the architectural interior-exterior distinction and is only exacerbated by the ubiquity of ever-intensifying orders and ever-expanding scales of algorithmic computation. Perhaps the distinction between turbulent and laminar flows—terms borrowed from fluid dynamics—provides a more useful reading of *Urbanisme* that allows algorithmic processes to be layered on top of and throughout Le Corbusier's figure of architecture. In particular, this fluid distinction provides a scalable way to address complex territories that simultaneously blur and articulate interiors.

Within the contemporary condition, turbulence that confronts architecture is no longer unpredictable weather patterns or wild beasts, but the unintended forces of a constantly connected digital infrastructure that demands constant attention. If, as Mark Wigley puts it, "architecture is always constructed in and against a storm" (2007), it is time for architecture to reevaluate its ability to separate us from a new storm—one that situates technology, global connectivity, human, non-human and composite users, and algorithmic architecture itself as new weather systems.

But where exactly do these stormy flows lie? What algorithms do the work of imaging them? How do these algorithms see the world and how do they both smooth and agitate the turbulent and laminar images captured within their gaze? What are the effects of their sensing and making sense of architectural images and images of the environment—and do they even recognize architecture as a thing within the built environment? If so, how might architecture take advantage of its privileged position? If not, does this compromise the architectural object as an agent of environmental mediation; and is it possible to extend the architectural objectives of *Urbanisme* to agents more capable of the task?

This paper explores the ability of Scale-Invariant Feature Transform (SIFT) algorithms to mediate and produce algorithmic turbulence generated through image-based sensing of the built environment while evaluating them in terms specified in *Urbanisme*.

LE CORBUSIER'S TECHNO-DEMATERIALIZATION

Urbanisme expresses a desire for architecture to embrace technologies that flatten spatial and temporal dimensions of the city. 'Intense agitations' of technology would reshape the city by destroying traditional habits of urban planning and force a 'mutation' of both the concept and form of the city itself. Le Corbusier's tactic was to dematerialize architecture such that it would integrate technology into the interior by "domesticat[ing]" all the lines of communication in the city [...] turning the interior into a landscape in which nothing is ever still" (Wigley 2007). He would dissolve walls in order to 'minimize friction' and rooms to give way to 'trajectories.' The transition from a Cartesian system of extensively defined spaces and boundaries to an intensive logic of vectoral trajectories not only marks the obsolescence of the interior-exterior distinction, it also spells out the agenda for a techno-architecture concerned with mediating gradients of informational friction and flow. So for Le Corbusier, the project for architecture was very much one of the city—a utopian agenda of frictionless architecture that would both allow and enable technological integration rather than opposing or resisting it. What Le Corbusier recognized is that for architecture to be most effective, it would have to radically open itself up (thus subjecting its inhabitants) to technological force through the annihilation of interior-exterior disruptions via laminar boundaries.

Corbusier's framing also belied the delicacy of technological performance manifesting in a tabula rasa approach, whereby architecture would strive to produce only the most minor instances of turbulence amidst an otherwise smooth and laminar context—cities and architecture designed free of resistance. He writes, "a city should be treated by its planner as a blank piece of paper, a clean table-cloth, upon which a single, integrated composition is imposed" (Le Corbusier 1971). Within this smooth environment, Le Corbusier called for a centrally located core that was to perform all the 'higher functions' of society; a core consisting of skyscrapers that form the 'brains' not only of the city, but of the entire country. These brains "embody the work of elaboration and command on which all activities depend.

Everything is concentrated there: the tools that conquer time and space; telephones, telegraphs, radios, the banks, trading houses, the organs of decisions for the factories: finance, technology, commerce" (Le Corbusier 1971). Architecture therefore is responsible for turbulent agitations that act as controlling signals. The center of the city does not consult with the citizens of the city, instead, it issues commands that are to be followed towards the production of homogeneity and smoothness—able to register the most insignificant perturbations within the neural network of the technological city. This metaphorical model is as viable now as it was then. However, what Le Corbusier could not possibly have comprehended is the scale and intensity that this model would assume, nor the radically ambiguous and interchangeable positions that urban objects themselves would be able to occupy.

What we arrive with from *Urbanisme* is the privileged territory of the architectural envelope defined as a disruption between some interior and its environment. This territory can be evaluated as architecture by asking the following questions: How is the boundary produced and/or imaged? What is the logic of the envelope? What qualities do its signals produce? What new architectural forms emerge from such an investigation?

The Agitated Relationship Between Technology and the City

Facilitated by advancements in tele-communication, planetary scale computation, and global exchange networks, the city of the early 21st century has become a *Digital City*, a city "composed of multiple 'intelligent' layers, based on 'real-time' interaction, communication and location based content, [existing] beyond the physical buildings and urban environment" (Handlykken 2011). Within the Digital City, architecture is forced to look beyond the physical materiality of buildings and toward software as means of imaging, occupying, and using the city. Furthermore, interior/exterior distinctions also fall away within a digital paradigm. Thus,

2

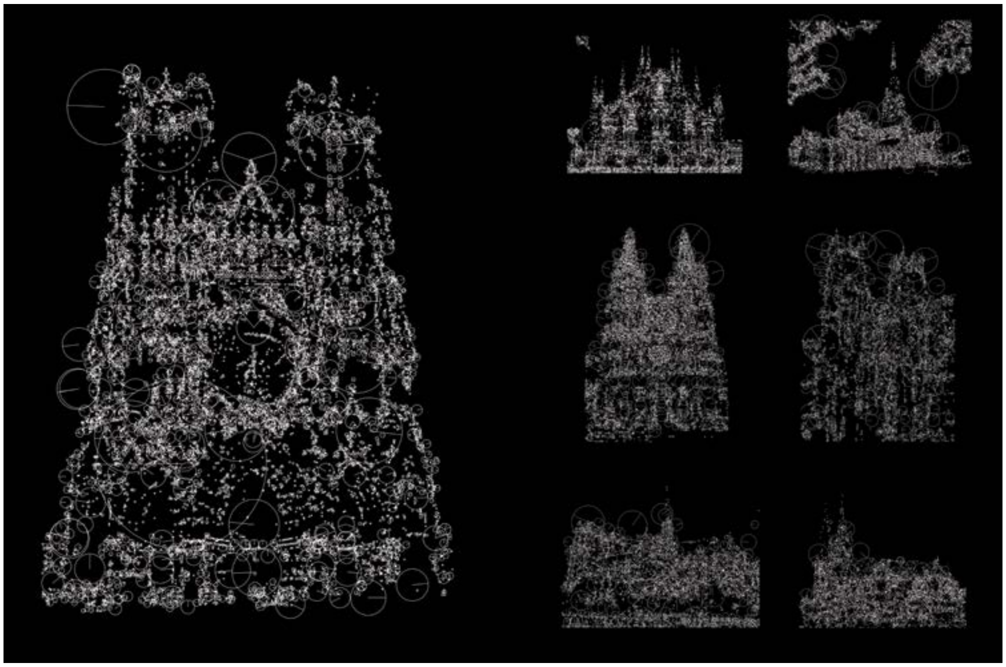

3

2 Visual anomalies of Google Earth: These images illustrate a selection of (mis)representations that result from architecture's agitated relationship to algorithmic observers.

3 SIFT cathedral study: digital images are converted to SIFT keypoints for the purpose of matching, stitching, etc., within the computational protocols of SIFT algorithms.

Le Corbusier's model demonstrates a kind of extension from the modern city into the digital city, where buildings have been replaced by code as the agents of turbulent disruption. But digital layers comprised of virtual mapping platforms and the ubiquitous production and dissemination of images produce a type of virtual tourism that destabilizes and distorts architecture's ability to even render an outside. Instead, what we have is a seamless image of an interior space that scales from our living rooms to the entire surface of the planet. But how is such a turbulent and inaccessible environment rendered in such a flattened and accessible way? And what is architecture to do if it is in the least way concerned with agitating or augmenting the smooth seamlessness of that space? Rather than beginning with the objects captured by this particular algorithmic observer (we will return to that later), the answer to these questions perhaps lies in the algorithms that produce the meta-image of the earth itself.

Despite its best efforts to "iron out the creases," turbulence is no more evident than in the visual anomalies of Google Earth (figure 2). As Johnson and Parker (2014) demonstrate, the visual anomalies of Google Earth are not glitches but the algorithms exposing themselves from behind a veil of anonymity. These digital artifacts are produced through architecture's attempt to present its outside (façade) to a class of *algorithmic observers* that have been tasked with sensing and making sense of the built environment. Algorithmic observers are entities that possess a type of sight despite a lack of eyeballs, rods, cones, and visual cortex. Instead they 'see' through the use of sensors to detect light, heat, motion, and color data, which are computationally processed to produce 'images' of the physical world. To digitally reconstruct the built environment, algorithmic observers reduce images of architectural façades to geometric data that is superimposed, composited, stitched, and texture mapped to produce virtual three-dimensional environments.

Algorithmic observers are interesting precisely because they lack "humanlike or human-level perceptual and aesthetic capacities, but rather [demonstrate] something that is uncanny and interesting because it does not possess those things" (Bratton 2015). It is their lack of humanlike perception and cognition that destabilizes our relationship to the built environment and produces new modes, scales, and patterns of turbulence within systems designed to produce and maintain smooth flows. The (in)ability of algorithmic observers to digitally reconstruct architecture produces turbulent flows previously concealed from human perception. These flows illustrate a new excess that has yet to be folded into the purview of human—or even architectural—perception. This research is not necessarily interested in either absolutely smoothing or agitating these possible flows, but rather focuses on possible ways in which turbulence might be modeled and imaged by misusing a tool that is generally biased toward smooth or laminar flows.

METHODOLOGY: HACKING SIFTS

Algorithmic observers rely on algorithmic protocols of machine vision to deconstruct images into collections of unique Scale-Invariant Feature Transform (SIFT) features that can be identified, organized, and matched across dynamic image sets. Of particular importance within these computational protocols are the algorithms associated with SIFT descriptors. SIFTs, developed by David Lowe (1999), enable algorithmic observers to identify specific image features that are invariant to scaling, rotation, changes in illumination and 3D camera viewpoint (Lowe 1999; 2004). SIFT descriptors are extracted from each pixel of an input image and encoded with contextual information through processes that reduce images to a large number of highly distinctive features, facilitating the filtration of visual clutter/ noise within the image, while providing a high probability of feature-matching and correlation across images (figure 3). SIFTs, with their strong matching capabilities and computational stability, are mobilized for the purposes of image retrieval, image stitching, machine vision, object recognition, gesture recognition, and the digital reconstruction of cities (McClendon 2012; Wu et al. 2013; Yang et al. 2011).

Due to the inherent ability of SIFTs to identify invariant geometric features within an image, an accurate 3D digital model with high-resolution renderings produces nearly identical SIFT data-fields to those of an actual image (figure 4). Additionally, the use of computer modeling and rendered outputs allows for the control of other image characteristics such as camera aperture size, focal length, white balance, global illumination, and lighting levels, along with rigorous control over viewpoint perspective, ambient noise, and visual clutter (trees, clouds, cars, people, garbage, etc.).

The use of 3D modeling and rendering allows for the production of an initial input dataset, where the camera viewpoint with respect to the object is the only dependent variable, with all other variables remaining constant (figure 5). By controlling all variables except for the camera's position to an object, the resulting dataset is defined primarily through an object's ability to produce data through its geometric composition, outside of variable lighting, camera, or environmental impacts. This methodology allows the object of observation to become transferable, as the control over image variables provides a framework for replicable and repeatable investigation with respect to any physical/virtual object.

Earlier research—presented in "This is Not a Glitch: Algorithms and Anomalies in Google Architecture" (Johnson and Parker 2014)—illustrates the internal protocols of SIFT algorithms and a methodology to extract n-dimensional pixel data from two-dimensional images. Whereas this previous research engaged SIFTs and their ability to reveal concealed vectors as a two-dimensional image problem, this paper extends this methodology to explore how these n-dimensional datasets can be mobilized towards the production of 3D form.

Voxels present a unique opportunity for working with turbulent systems, as they possess the ability to contain and respond to a multiplicity of datasets within a single volumetric pixel.

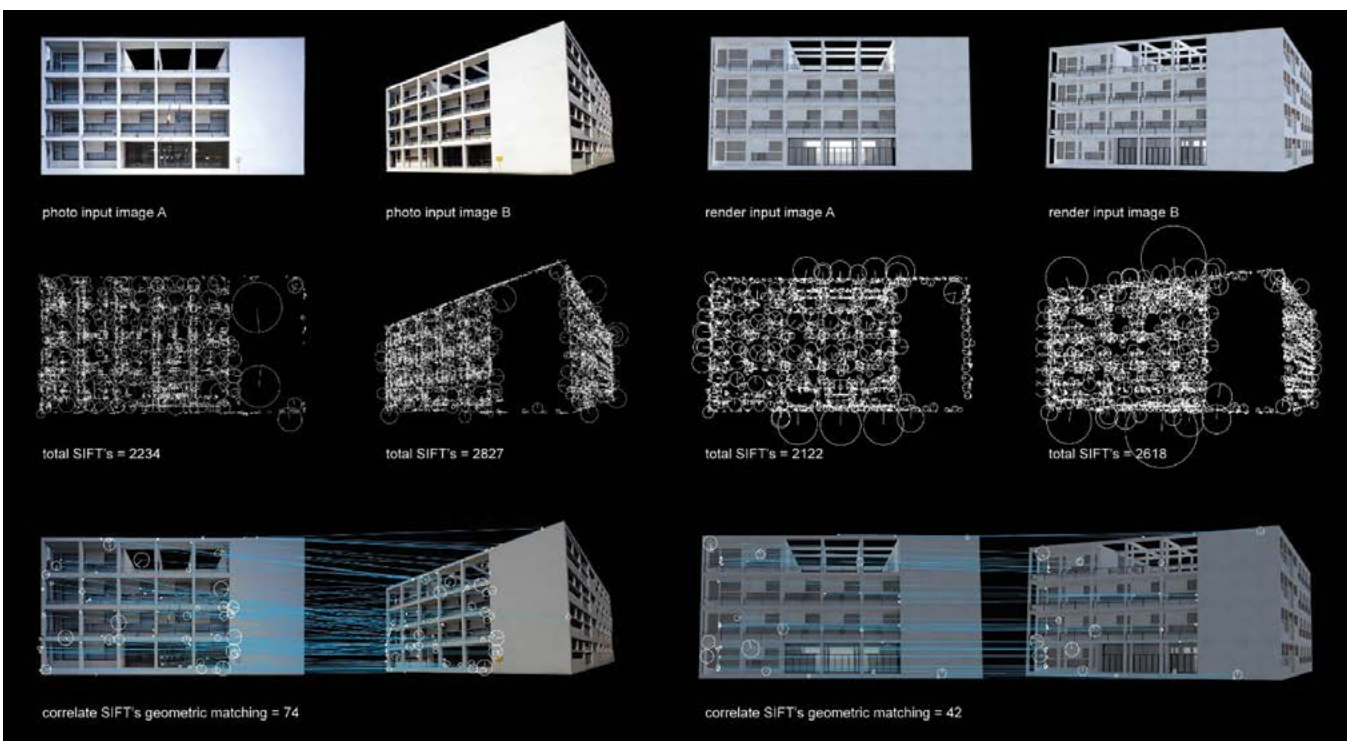

4 SIFTs in Photos vs. Renders: This diagram illustrates how high fidelity renderings produce similar SIFT data fields to those of an actual object.

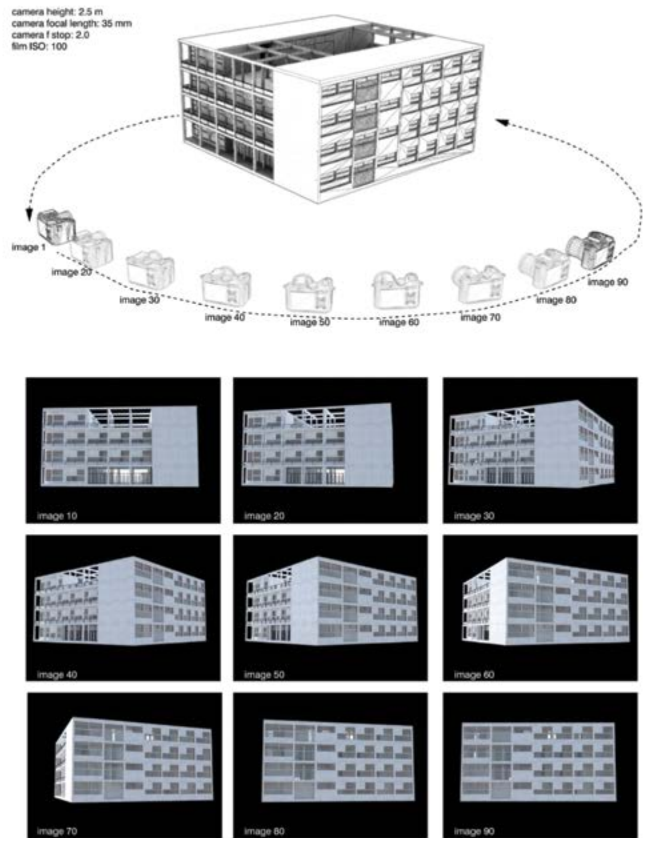

5

5 Virtual Camera Path and Outputs: By using virtual objects to simulate the physical, image variables can be constrained allowing the objects relationship to the observer to be the only dynamic variable.

Due to this ability, voxels are often used within medical image reconstruction, terrain scanning, and computational fluid dynamics. Whereas particles and points are defined by their relationship to coordinate space and their position relative to other particles (thus making them good for swarm logics, etc.), voxels need not respond to their neighbors or community as they can be programmed to operate as independent unique entities. The ability to uniquely codify each voxel is valuable as it pertains to SIFT algorithms, because once a SIFT keypoint is identified and codified, it becomes an autonomous entity, for which SIFT algorithms search for correlate SIFTs regardless of their neighbors and surrounding image context. It is the autonomy of SIFT keypoints that produce the warped images previously discussed in “SIFT Materiality: Indeterminacy and Communication between the Physical and Virtual” (Parker 2014), where correlate SIFTs are superimposed and mapped onto each other, with the surrounding context resulting from degrees of algorithmic interpolation, or what Luciana Parisi refers to as ‘soft thought’ (2013). By mapping the n-dimensional vectors produced through processes of superimposition directly to voxels within programmed gas-solver environments, turbulent systems are mobilized towards the production of dynamic form.

RESULTS AND REFLECTIONS: UNPACKING FLATTENED DIMENSIONS

Early investigations only looked at how voxel fields would deform when the only forces within the system were those internally produced through algorithmic observation. The resultant objects demonstrate the autonomous ability to produce turbulence without context (figures 6–9).

As images are abstracted to SIFTs and composited to produce smooth yet warped images, an excess of n-dimensional vectors are revealed through the production of seamless and ‘accurate’ composite images that have been algorithmically curated to represent a physical object. By mapping these n-dimensional vectors to a corresponding voxel field, voxels assume the attributes of a particular pixel as it transforms within the process of re-composition. By allowing these charged voxels to organize within a computationally turbulent gas-solver, voxels organize, not in relation to their neighbors or a specific coordinate system, but in relation to the dynamic data on a voxel-by-voxel basis. These experiments provide a technique to produce 3D form from 2D images, as each pixel of a warped image generates n-dimensional vectors that are mapped to unique voxels within the simulation (a warped image of 640 X 480 pixels produces vectors mapped to a two-dimensional voxel field of 640 X 480 voxels). Despite the initial field being two-dimensional, the voxels, free to interact within the smooth constraints of a three-dimensional computational environment, result in indeterminate n-dimensional form (figure 1). This is the structural logic of SIFT weather systems.

The workflow allows for images to be mined for their n-dimensional vectors, producing dynamic datasets that aggregate to form increasingly resolved, multi-dimensional ‘images’ of a particular object. Similar to the manner in which algorithmic observers sense an object in the built environment through a series of images or scans from multiple viewpoints over a specified duration, each new set of n-dimensional vectors produces unique magnitudes and directions that are introduced to their corresponding voxel on a per-frame basis. This process creates a turbulent system that negotiates data-intensities over time. By establishing a one-to-one relationship between datasets and frames, the simulation is capable of producing a distinct morphologic manifestation on a per-frame basis, each one a reflection of data excess and the turbulence produced as algorithmic observers attempt to process and composite a multitude of discrete images into a cohesive whole (figure 12).

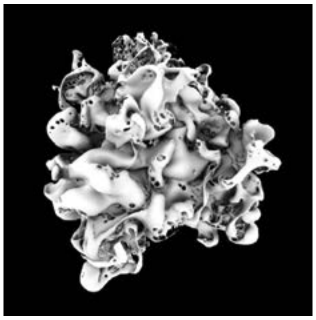

6

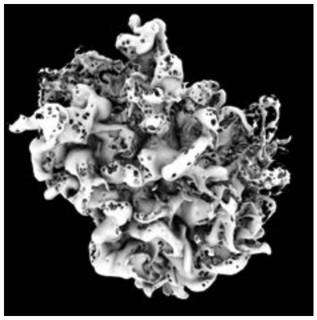

7

8

9

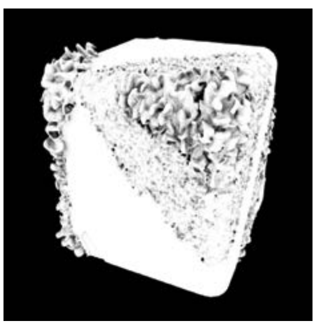

10

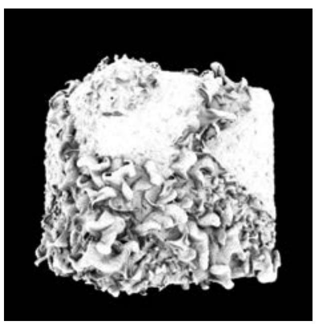

11

Turbulent Façades

From this investigation emerges an interesting discussion about the ability of the familiar to produce the speculative, and a symbiotic relationship between specific smoothing agents and their turbulent byproducts becomes a bit clearer. This set of investigations explores how the previously developed turbulent systems can exert a force back onto and into an existing object towards the production of the *uncanny*. By defamiliarizing the familiar, the uncanny has the ability to mobilize a dialogue around what else an object might be, and what other bodies of information may be concealed from human observation (figures 10 and 11).

Within this research, the process of destabilization relies on computational protocols to expand the intelligence of a mesh, allowing it to operate indeterminately within an environment. To accomplish this, high fidelity computational objects are converted to voxels through a process that reduces an existing poly-mesh into a number of spatial voxels that correspond to the location and number of pixels an object occupies within a digital image. The technical problem emerging from this conversion process is determining exactly how many vectors a 3D mesh is responsible for producing within a 2D image. As the visual noise and clutter are controlled within this workflow through the production of computation and renders, objects exist within the image outside of their context. Whereas an image may be 640 X 480 pixels, the object may only occupy a small number of these pixels, leaving the rest of the image devoid of information and thus incapable of producing n-dimensional vectors. The challenge is to then

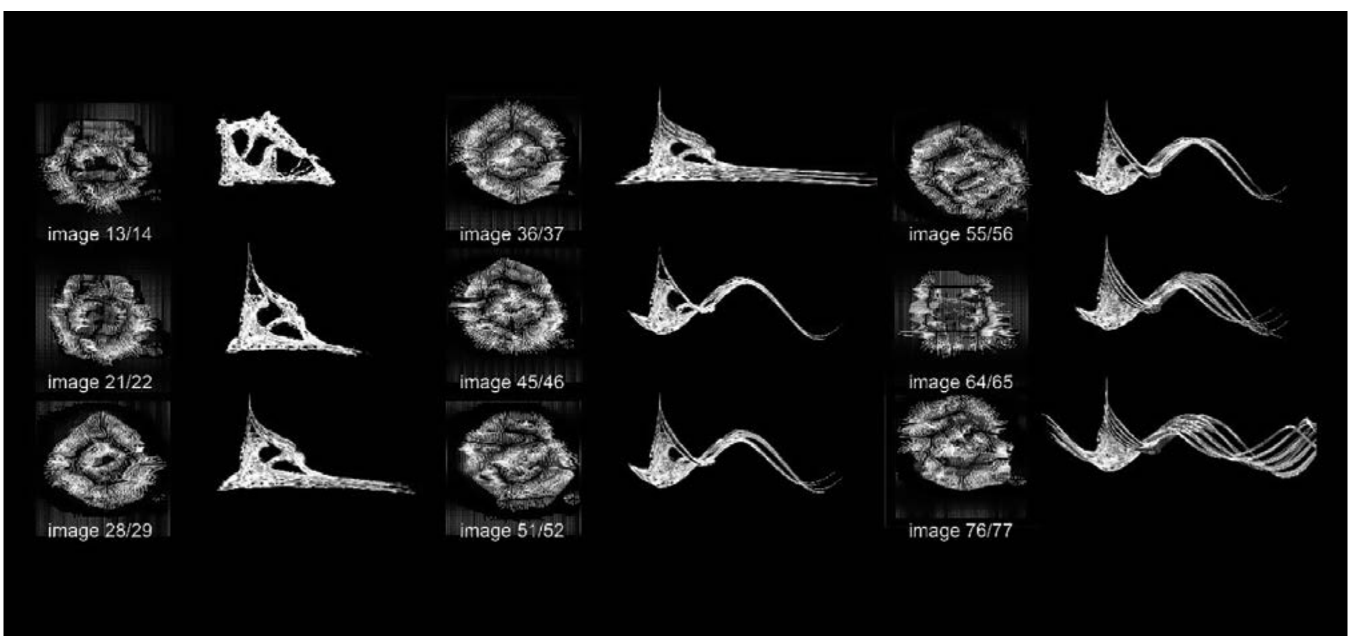

12

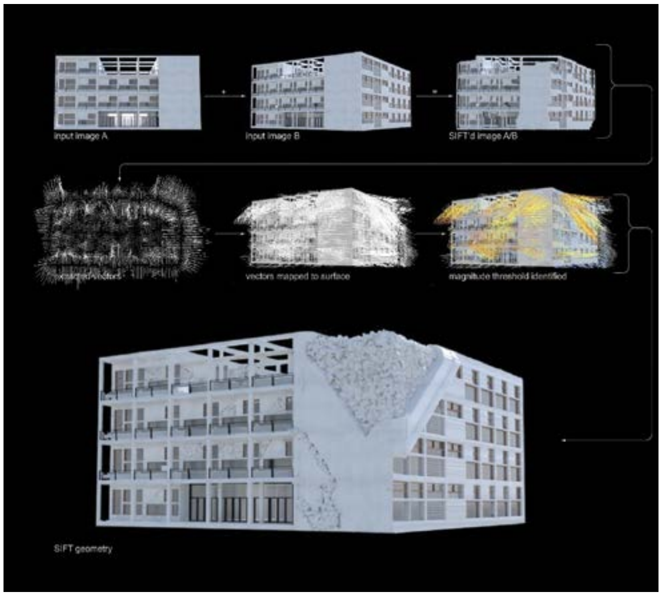

13

6–9 Turbulent Object: produced as vector-flow-fields operate indeterminately across a set of voxels in accordance with the governing logics of a codified computational environment.

10–11 Primitive deformation (Cube A): the faces of a volumetric cube are charged relative to their correlate SIFTs in a 2D image plane. The charged surface exerts a force within a voxel system describing the geometry and the output is deformed relative to the resulting vector-flow-fields produced through algorithmic observation.

12 N-dimensional workflow for extracting and mobilizing the output vector-flow-fields produced through algorithmic observation. Extracted n-dimensional vectors deform voxel based geometry within a computational environment.

13 De-familiar façade: This diagram illustrates two input images warping to produce a composite data-rich territory and resulting n-dimensional vector field. This vector field is extracted from the 2D image and mapped to its 3D form where turbulent thresholds are identified and deformed within the predefined logics of the system.

identify the location and number of pixels an object occupies and convert the computational object to a corresponding number of voxels. To further compound this problem, an object may occupy a different number of pixels within different images, dependent on camera view-point perspective and orientation to the object. To account for these kinds of problems a process of *Taylorian Projection* (Cohen 2001) is utilized to convert a 3D object to its corresponding number of voxels relative to the number of pixels in an image, and to accurately remap the n-dimensional vectors onto their corresponding voxels.

As this process maps a specific vector to each voxel that describes an object, it is possible to deform an object in its entirety; however, this research has explored the potentials of only deforming areas that exhibit charges over a predetermined threshold. By constraining deformation to specified thresholds of intensities, a legibility becomes apparent within the destabilization of an object. This legibility allows for the production of a narrative that supports the uncanny. Through a process of curation that allows the indeterminate nature of the protocols to operate in accordance with predetermined thresholds, the output remains legible. However, the object exhibits a speculative deformation in the areas of high complexity in accordance with the perceptive abilities of algorithmic observers (figure 13). The resulting deformation illustrates the agitated relationship between architecture and its algorithmic observers, and distorts our perception of an object by mobilizing a concealed technologic turbulence.

CONCLUSION: NEW OBJECTIVES

Le Corbusier argues that old Paris subjected its citizens to chaotic, terrifying, and confusing conditions through its 'disturbing,' 'twisted,' 'mis-shapened,' and 'abnormal' patterns of urbanization. To rectify these problems, Le Corbusier proposed an architecture that sheltered its inhabitants from the chaos of the city, producing order by facilitating co-existence. It can be argued that for much of the 20th and the early part of the 21st century, cities have been designed to produce diagrams of smooth and homogeneous flows, with the increasing prevalence of distributed personal technologies allowing us to navigate the urban fabric no matter how twisted, hostile, and misshapen a city's form may be. However, the virtual layers that unevenly map onto the terrestrial constraints of the city produce turbulent forces that computational platforms tasked with sensing and making sense of the built environment aim to conceal behind a visual narrative of accuracy, cohesion, anticipation, and order. What is at stake with the geometries produced within this research is not that the forms are necessarily novel or unique, but that the methods through which they are produced are born from both the image of the world, as well as the algorithms that observe and produce those images. This restlessness demarcates a new kind of challenge for both architecture and the city, whereby a cultural appetite to deploy algorithms that produce a smooth and seamless image of the world comes hand in hand with the turbulent and disruptive autonomy of those very same algorithms.

Whether we accept it or not, algorithmic observation shapes our experiences and relationships to the city. And so looking toward Le Corbusier's *Urbanisme* through a different lens, we are presented with a variety of opportunities to re-frame the problem. Instead of eliminating boundaries, we are now tasked with identifying latent boundaries and querying them for the signals they might be sending. This demands that architecture address more than buildings—allowing Google Earth into our houses, our bedrooms, and our bodies and training them to produce new turbulent patterns—life as art in the eyes of a SIFT observer. In a less conventional sense, it could also mean that architecture treats SIFT networks as architectural sites themselves—exploring the way in which the algorithm can be inhabited, occupied and, instrumentalized for user-based, programmatic purposes. Algorithmic observation challenges architecture to engage a form of practice that no longer situates the physical act of making buildings as the primary role of the architect, but to take part in shaping and designing the technologies that increasingly redefine new ecologies of the city. In order to formally establish algorithmic observation as an architectural thing, it is necessary for architects to address it and the spaces produced as valuable, despite their lack of traditional materiality and familiar formal manifestations. It also demands that we come to understand the nature of architectural boundaries in the context of algorithmic observation as being both laminar and turbulent at the same time. This should not be seen as compromising Corbusier's laminar agenda, but rather as an expression of the ability for architectural flows and objects to simultaneously occupy multiple dimensions within contested fields of algorithmic observation.

This research is merely a first step towards investigating the agitated relationship that exists between architecture and its virtual processes, but it shows that the new forms of urban inhabitation challenge the smoothness of the city and call into question architecture's desire to accommodate heterogeneous forces towards the production of laminar flows. Instead, algorithmic observation challenges architecture to exploit the fluid and vectoral aspects of imaged buildings, to explore turbulence as an agent of potentially catastrophic change, and to investigate the technologies that are obsessed with squeezing new dimensions out of flattened environmental images.

REFERENCES

Bratton, Benjamin. 2015. "Machine Vision: Benjamin Bratton in Conversation with Mike Pepi and Marvin Jordan." *DIS Magazine*. Accessed August 14, 2015. http://dismagazine.com/discussion/73272/benjamin-bratton-machine-vision/

Cohen, Preston S. 2001. *Contested Symmetries and Other Predicaments in Architecture*. New York: Princeton Architectural Press.

Google. 2012. *The Next Dimension of Google Maps*. YouTube video, 51:18. Posted by Google, June 6, 2012. Accessed November 23, 2014 from https://www.youtube.com/watch?v=HMBJ2Hu0NLw

Handlykken, Asne K. 2011. "Digital Cities in the Making: Exploring Perceptions of Space, Agency of Actors and Heterotopia." *Ciber Legenda* 25 (2): 22–37.

Johnson, Jason S. and Matthew Parker. 2014. "This Is Not a Glitch: Algorithms and Anomalies in Google Architecture." In *ACADIA 14: Design Agency—Proceedings of the 34th Annual Conference of the Association for Computer Aided Design in Architecture*, edited by David Gerber, Alvin Huang, and Jose Sanchez. Los Angeles: ACADIA. 389–398.

Le Corbusier. 1971. *The City of To-morrow and its Planning*. Translated by Frederick Etchells. Cambridge, MA: MIT Press. Originally published as *Urbanisme* (Paris: Éditions Crès, 1925).

Lowe, David G. 1999. "Object Recognition from Local Scale-Invariant Features." In *Proceedings of the Seventh IEEE International Conference on Computer Vision* vol. 2. Kerkyra, Greece: ICCV. 1150–1157.

———. 2004. "Distinctive Image Features from Scale-Invariant Keypoints." *International Journal of Computer Vision* 60 (2): 91–110.

Parisi, Luciana. 2013. *Contagious Architecture Computation, Aesthetics, and Space*. Cambridge, MA: MIT Press.

Parker, Matthew. 2014. "SIFT Materiality: Indeterminacy and Communication between the Physical and the Virtual." In *What's the Matter: Materiality and Materialism at the Age of Computation*, edited by Maria Voyatzaki. Barcelona: ENHSA-EAAE. 313–326.

Wigley, Mark. 2007. "Towards Turbulence." *Volume* 10. Accessed April 19, 2016 from http://c-lab.columbia.edu/0135.html

Wu, Jian, Zhiming Cui, Victor S. Sheng, Pengpeng Zhao, Dongliang Su, and Shengrong Gong. 2013. "A Comparative Study of SIFT and its Variants." *Measurement Science Review* 13 (3): 122–131.

Yang, Donglei, Lili Liu, Feiwen Zhu, and Weihua Zhang. 2011. "A Parallel Analysis on Scale Invariant Feature Transform (SIFT) Algorithm." In *Proceedings of the 9th International Conference on Advanced Parallel Processing Technology*, edited by Olivier Temam, Pen-Chung Yew, and Binyu Zang. Shanghai: APPT. 98–111.

IMAGE CREDITS

Figure 1: Parker, 2014
Figure 2: Google Earth, accessed 2014
Figure 3–4, 6–11: Parker, 2015
Figures 5, 12–13: Parker, 2016

Matthew Parker completed his Master of Architecture from the University of Calgary's Faculty of Environmental Design, where he received honors recognition and the AIA Gold Medal. Currently he is completing a Post Professional Masters with his current research focusing on the ability of algorithmic observation to transform, mediate and re-animate architectures' image. Matthew is also a researcher with the Laboratory for Integrative Design (LID), and a studio designer and parametric consultant with Minus Architecture Studio and Synthetiques / Research + Design + Build.

Joshua M. Taron is an Associate Professor of architecture at the University of Calgary Faculty of Environmental Design where he also co-directs the Laboratory for Integrative Design. His current research focuses on designing for divertability as an alternative to conventional wholesale building demolition through the use of advanced fabrication techniques. Taron is also Principal of Synthetiques Research & Design, Inc, a consultancy that specializes in complex architectural form-making. He earned his undergraduate degree in architecture from the University of California, Berkeley and holds a Master of Architecture degree from the Southern California Institute of Architecture.