Architectural Heat Maps: A Workflow for Synthesizing Data

Jason S. Johnson
Matthew Parker
University of Calgary

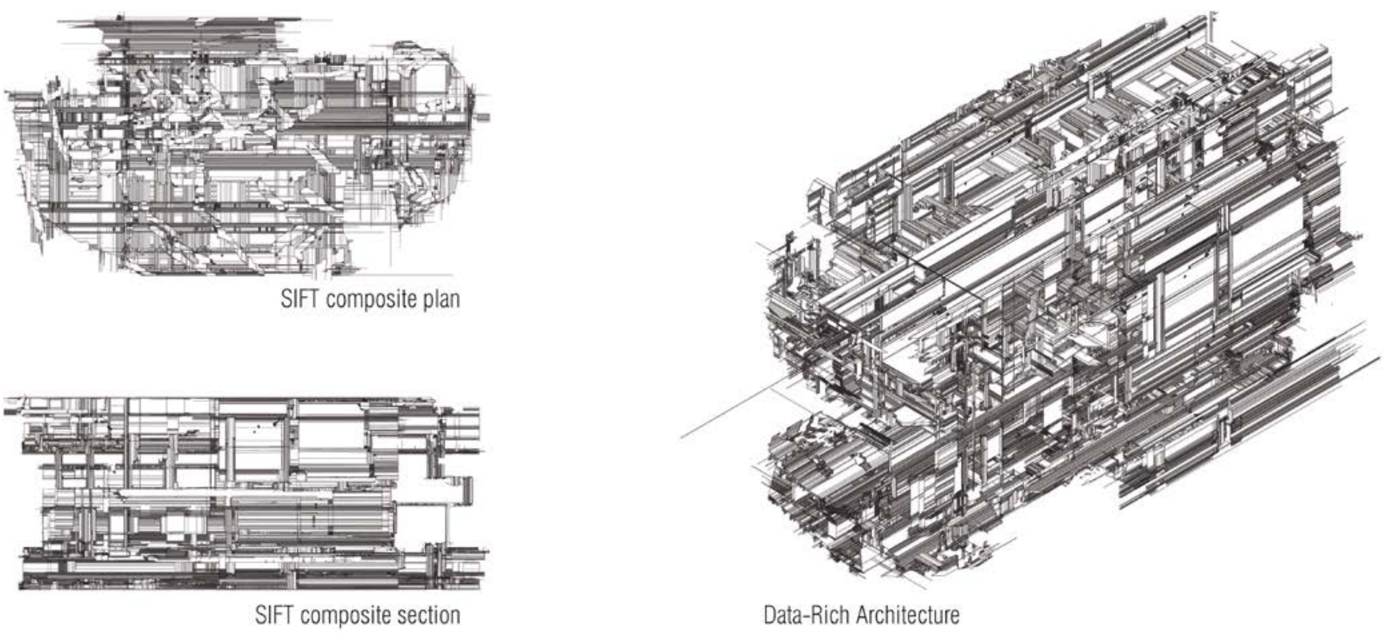

1

1 SIFT composite plans and sections are mobilized within the workflow described in this paper to produce an architecture that operates through a condition we describe as ubiquitous simultaneity.

ABSTRACT

Over the last 5 years, large-scale 'data dumps' of architectural production have been made available online through project-specific websites (mainly competitions) and architectural aggregation/dissemination sites like Architizer, Suckerpunch, and Archinect. This reinforces the broader context of Ubiquitous Simultaneity, in which large amounts of data are continuously updated and easily accessed through a dizzying array of mobile devices. This condition is being exploited by sports leagues and financial speculators through the development of tools that collect, visualize, and analyze historical data for the purpose of producing speculative predictive simulations that could lead to strategies for enhanced performance.

We explore the development of a workflow for deploying computer vision, SIFT algorithms, image aggregation, and heteromorphic deformation as a design strategy. These techniques have all been developed separately for various applications and here we combine them in such a way as to allow for the embedding of the historical and speculative artifacts of architectural production into newly formed three-dimensional architectural bodies. This work builds on past research, which resulted in a more two-dimensional image-based mapping and translation process found in existing imaging protocols for projects like Google Earth, and transitions towards the production of data-rich formal assemblies. Outliers and concentrations of visual data are exploited as a means to encourage innovation within the production of architecture.

INTRODUCTION

Outliers are often the generative fuel for innovation within a wide variety of fields. Something or someone breaks away from the collective pack due to an aberrational quality or glitch that was not fully understood or accepted by the larger system. Innovators seek to both find and describe the aspects of these outliers as a way of exploiting their potential to produce emergent models of differentiated productivity. The cumulative effect of this continual process of the collective either destroying or moving towards these mutations is perhaps best illustrated in the world of sports, where the previously niche field of data analytics has become deeply ingrained across disparate sports at various levels. The success of this model depends on the ability to aggregate and analyze large pools of information in order to understand and leverage the relationships between the individual attributes (singular performance metrics) and the corresponding performance of the team. The National Basketball Association has been using tracking systems to monitor a wide range of statistics relative to a player's positioning, time on the court, and performance. Similar systems that pinpoint the location of players relative to the location of both moving and static objects on the field of play have been deployed by football, baseball, and a number of other sports, and are now being tied to biometric sensors attached to a player's body. These tools of measurement are being used to both analyze past performance and simulate future outcomes in ways that are impacting the strategies deployed in real-time game situations (Goldsberry 2012). The visualizations for these tools as deployed in baseball (Figure 2) and basketball (Figure 3) hint at the potential for overlaying historical data onto a speculative or strategic performance image.

The potential for adapting these approaches for the production of architecture is one that could begin to re-situate the conversation about parametricism and generative design away from a coalescing visual style and towards a more fulsome engagement of historical and speculative formal and functional relationships within designed artifacts of all scales. Schumacher's manifesto for parametricism as a style speculated that the move towards the "inter-articulation of sub-systems" would be critical to the production of differentiated field conditions vs. the production of bounded space (Schumacher 2008). What is perhaps missed in the ways in which generative design projects are currently being developed is the integration of historical and speculative inter-articulation into the parametric models. These models could become more like the data-driven approaches that are redefining the ways in which sports are being visualized in both historical and speculative ways. Here, we propose an approach that leverages computational tools for vision, sensing and recombining the data associated with architectural production, towards a workflow that might begin to engage the age of Ubiquitous Simultaneity (Johnson and Parker 2016), which is allowing for an unprecedented amount of visual, environmental, and regulatory data to be gathered and processed. This research builds on previous explorations into the potential for computer vision and Scale-Invariant Feature Transform (SIFT) algorithms to be deployed for the production of two-dimensional images and introduces into the workflow techniques for three-dimensional form finding and the embedding of bias for the purpose of innovation (Johnson and Parker 2014). It builds out an approach that speculates about the nature of design as a profession that must begin to leverage massive repositories of historical and speculative design strategies and artifacts in order to better understand both past tendencies and future directions. For the purpose of this research we have focused this case study on the design competition held for the Helsinki Guggenheim. With over 1700 entries that share a site as well as programmatic criteria all formatted in two-dimensional images, this dataset allowed us to set up a series of protocols that could form an approach towards a new generative production technique.

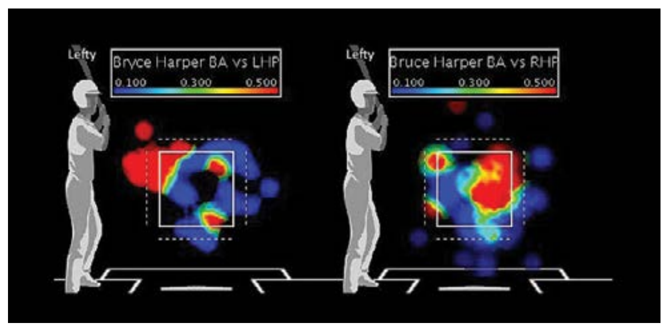

2

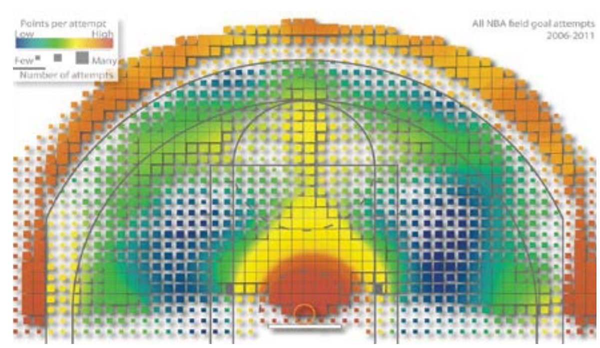

3

2 Heat Map with two variables. Map charts the relative productivity by the specific player when he swings at a pitch in relation to the location of the pitch and handedness of pitcher. Image courtesy of Kirk Goldsberry.

3 This map reveals league-wide tendencies in both shot attempts and points per attempt. Larger squares indicate areas where many field goals were attempted; smaller squares indicate fewer attempts. The color of the squares is determined by a spectral color scheme and indicates the average points per attempt for each location. Image courtesy of Kirk Goldsberry.

BACKGROUND

Computer Vision

Architecture and the city are increasingly experienced through what this paper refers to as algorithmic observation, a type of machine vision that 'sees' despite a lack of eyeballs, rods, cones, and a visual cortex; instead, algorithmic observers perceive the city through the use of sensors capable of detecting light, heat, motion, and color data to produce composite 'images' that describe the physical world. By further abstracting these images into their geometric components, algorithmic observers possess an inhuman ability to identify, sort, and catalog image-based information, making them ideal for deployment within areas of security, surveillance, military observation, feature identification, and a wide array of other uses that are building up a continuously updated catalog of 'seen' data. Within this research, algorithmic observers are utilized for their ability to identify and organize deep catalogs of preexisting speculative projects. These catalogs allow for the indexing of a body of work that not only embeds physical characteristics typical of all buildings, but also the theoretical tendencies that are typically amplified in projects developed for design competitions.

The design competition has long been a vehicle for the avant garde in architecture to push for new formal, social, and functional agendas within the context of a given typology. As such, they typically form a body of work that exists at the edges of architectural practice as a whole. The ability of algorithmic observers to dissect and make sense of the prolific amount of free architectural content disseminated through architectural competitions and the emergence of project aggregation sites like Architizer, Archinect, and Suckerpunch allows for previously 'invisible' insights to be collected and analyzed against each other and other data sets.

The use of computer vision to identify and organize deep catalogs of information originates in the work of Lev Manovich, who produces software in response to a growing desire to track and map the "global digital cultures with their billions of cultural objects and hundreds of millions of contributors." This is the result of the "exponential growth of a number of [...] media producers over the last decade [which] has created a fundamentally new cultural situation and a challenge to our normal ways of tracking and studying culture" (Manovich 2012). Manovich claims that the rise of social media and globalization leaves no other choice than to rely on computational tools for the organization and mapping of the visual artifacts that describe and construct culture. By developing techniques for automatic digital image analysis, Manovich is capable of generating numerical descriptions of various visual characteristics of an image, which allows him to identify tendencies and outliers that aggregate to form a cultural reading (Figure 4). Whereas Manovich's work relies on information codified at the level of a pixel (its hue, saturation, brightness) to catalog images, architecture is faced with a problem of needing to identify specific characteristics within traditional architectural drawings of plans and sections. To accomplish this, we rely on the algorithmic observers' ability to computationally reconstruct images into collections of unique features that can be identified, organized, and matched within and across extensive image sets, through the use of SIFT algorithms.

SIFTs (Scale-Invariant Feature Transform)

SIFTs, first developed by David Lowe, enable algorithmic observers to identify specific images 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. SIFTs, with their strong matching capabilities and computational stability, are mobilized for

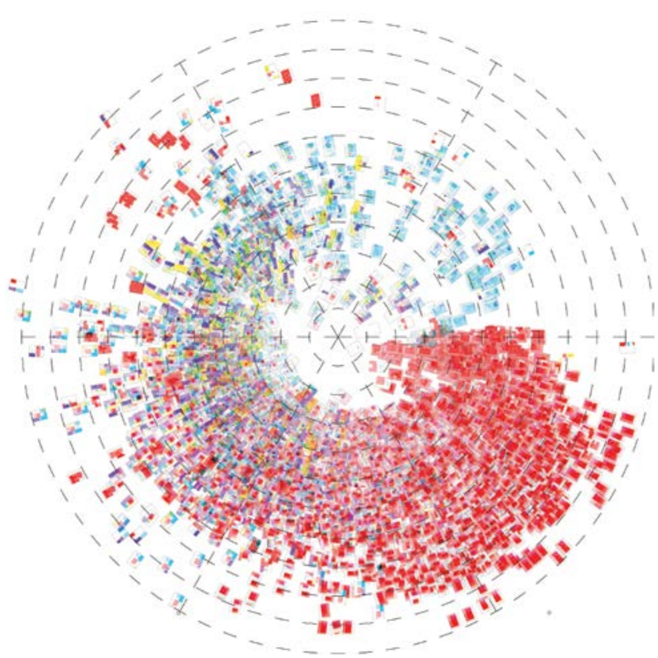

4

4 Identifying outliers: A technique for identifying outliers based off their SIFT-generated heat maps (description to follow in section 3.2), This graph illustrates the outliers within an initial sample of 1700 heat maps of Casa de Fascio when viewed from varying viewpoint perspectives.

5 SIFT production of Warped Images: By identifying correlate SIFTs between 2 input images, the workflow maps and superimposes corresponding data onto other bodies of soft data, tasking SIFT algorithms to exert agency within the production of 'warped images' by speculating towards which data might 'fill-in-the-blanks'.

the purposes of image retrieval, image stitching, machine vision, object recognition, gesture recognition, match moving, and the digital construction of three-dimensional virtual environments (McClendon 2012; Wu et al. 2013; Yang et al. 2011).

Previous research (discussed in Johnson and Parker 2014; Parker 2014) explored the computational protocols of SIFT algorithms and their ability to facilitate computer vision towards the digital (re)construction of the physical environment. These papers draw attention to the ability of algorithmic observation to shape our experiences and relationships with the city, challenging architecture to engage these technologies by intentionally assuming multiple identities within the superimposed virtual and physical layers that compose the contemporary city. The overarching theme of this previous research is that architecture, by and large, retains a level of ambivalence towards algorithmic observation that stems from a traditional desire to cling to certain disciplinary claims, specifically architecture's role in shaping and defining our spatial relationships to the physical materiality of the city. Whereas previous research attempted to hack the virtual layers of the city by augmenting architecture's exterior envelope towards direct interface and communication with algorithmic observers, this research explores a SIFT-based workflow that challenges architects to engage a form of practice that no longer situates the physical act of making buildings as their primary role, but instead challenges them to take part in the shaping and designing of the attributes and actions that redefine the ecologies of the city.

In order to reframe the relationship between algorithmic observers and architecture, it is necessary for architects to address algorithmic observers and the space they produce as valuable, despite their lack of traditional materiality and familiar formal manifestations. This research builds off a previously developed understanding of the computational logics of SIFT algorithms in order to produce a workflow for the production of three-dimensional inter-articulated architecture that can be biased towards historical or speculative tendencies within pre-defined bodies of architectural production. These bodies could be defined as specific sites and programmatic contexts, typologies, or a combination of the two constrained by one another.

Within the context of SIFT algorithms and their role in computer vision, this research claims that at their most basic level, SIFTs convert 2D image data into soft data for the purposes of recombination and superimposition. Soft data possesses the ability to elastically deform in response to external forces while retaining its original topology. With respect to computer vision, SIFT algorithms convert digital images to recognizable geometric descriptors for the production of data-rich territories (DRT) (Figure 5). Initial research into the production of ubiquitous simultaneity utilized a homeomorphic relationship between images, allowing each image to exert an equal force toward the production of DRT. By contrast, this workflow explores techniques that bias particular characteristics within a body of soft data, allowing said characteristics to exert greater influence within the production of DRT.

Importantly, the conversion of images to soft data does not alter the unique characteristics of a particular body of data. To accomplish this, SIFT keypoints—which are unique codified features of an image—are utilized for their capability to transform

5

and receive information from correlate keypoints present across dynamic data sets. Simply put, keypoints allow specific geometric compositions of an image to map and accept multi-dimensional information produced from similar keypoints within a dataset. This process results in an augmentation of the keypoint but not the production of a new keypoint. In this sense, keypoints allow for the codification of multiple features to be contained within a single unique image feature, resulting in what we dub a feature heat map. In the same way that the individual data points within a shot chart (figure 3) both accept a number of characteristics (number of shots, percentages made) and assign a location within a two-dimensional representation of physical space, these feature heat maps allow for a densely calibrated visual dataset to emerge.

By deploying heteromorpholgic deformation within the production of DRT, this research deviates from previous explorations that exploited a one-to-one mapping between bodies of soft data. As such, we are currently invested in developing a workflow for the production of unequal influence in the production of said bodies, allowing for a more speculative generative approach for the establishment of a trajectory towards or away from specific characteristics present in any given image dataset. This notion of bias allows for design decisions to be made within a framework that promotes a gradient approach to projected outcomes.

METHODS AND RESULTS

In their seminal essay "Transparency: Literal and Phenomenal" (1964), Colin Rowe and Robert Slutzky champion a definition of 'transparency' within architecture that supersedes an object's mere ability to transmit light through it without optical destruction. Instead, Rowe and Slutzky suggest 'phenomenal transparency,' the ability to concurrently perceive different spatial locations of an object, allowing space to not only recede but to fluctuate in a continuous activity that allows objects to be understood simultaneously as closer and further away. What this definition does is transform our understanding of the 'transparent' from that which is perfectly clear to that 'which is clearly ambiguous.' This understanding allows for the superimposition of objects/data within a framework that does not necessitate spatial privileging, allowing us to ask if the multiple datasets that (potentially) define a physical object can be compressed into single data-fields for the production of the speculative new. And are there ways in which this transparency between the physical (historical) and the image (speculative) allow for new formulations of space that leverage the latent potentials found within each dataset.

By compositing sets of soft data into a single manifestation of an architectural object, this investigation marks a radical departure from the previous experiments that sought to add information into existing objects. Whereas past research focused on mapping surfaces of existing buildings and producing combinations that resulted in hybrid conditions, this research builds off notions of recycled intentionality or the compression of multiple design-intentions in order to "transpose insignificant singularities into meaningful complexities" (Maholy-Nagy in Rowe and Slutzky 1964). In order to accomplish this, large aggregate collections of architectural drawings (plans, sections, elevations) obtained from the Guggenheim Helsinki Design Competition are reinserted into the workflow to produce a series of volumes that embed the functional aspects of the projects inside of the speculative proposals. These volumes can be read as heat maps of functional and aesthetic intentions. No single instance is as legible as the general tendencies to reinforce the outliers at their perimeters while at the same time producing aggregations that do not lose definition at any one point within the larger assembly.

Producing Bias: Soft Drawings

In order to bias outliers, this workflow relies on protocols that allow a particular instance of soft data to exert varying degrees of influence within the compositional processes involved in the production of DRTs. Within these experiments, we have identified plans that read as clusters and sections that exhibit a relative tallness as outliers within a general population of over 1700 entries. Clustering can be considered an outlier, as the majority of the projects rely on a linear or grid based plan, with tallness being any building that exhibits sectional heights in excess of four stories (Figure 6). This is not to say these are the only two outlier conditions within the submission, or even the most prevalent, but

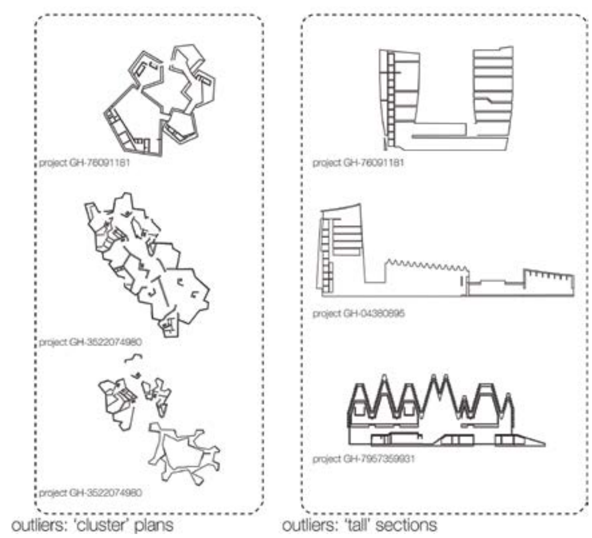

6

6 Outliers: examples of 'cluster' plans and 'tall' sections that we have identified as outliers for the purposes of this investigation.

merely two outliers identified to situate as the biasing characteristics within these investigations.

Bias is produced by converting all available drawings to soft data for the initial purposes of superimposition within a system that allows each body of data to exert an equal force within the process of (re)composition (Figure 7). This process results in a large body of agile data in the form of speculative plans and sections that are ready to be deformed relative to the identified outliers. Bias is introduced once the architectural drawings have been converted to soft data through a computationally codified bias slider that operates similar to an alpha channel or image mask in Photoshop. The bias slider controls how much influence a particular dataset will exert within the composite whole, with a bias value of zero producing no effect within the composite whole and a value of 1.0 exerting maximum influence. If soft dataset A has a bias of 0.8 it will exert a force 200% greater than soft dataset B that exhibits a bias of 0.4, thus the resulting DRT will exhibit a biased towards the characteristics inherent in soft dataset A (figure 8). It follows that if one wishes to bias the characteristics included within the plans identified as ‘clusters,’ one would assign a higher bias value closer to these image-driven datasets within the processes of producing data-rich territories.

Soft-Data: Superimpostion, Composition, and 3D Form

Once SIFT algorithms have converted drawings to soft data and superimposed them towards the production of DRT, the data must be rendered into architectural drawings legible for human observation. These new drawings expose speculative architectural drawing(s) that have elastically deformed in response to a body of soft data’s ascribed bias. Previous research has demonstrated that the process of superimposition that renders soft data as images produces sets of n-dimensional vectors that correspond to the uniquely codified entities of a particular body of soft data. Whereas previous research deployed these vectors towards the production of speculative images, this research remaps these vectors to their correlate SIFT keypoints for extrusion and the production of three-dimensional form (Figure 9),

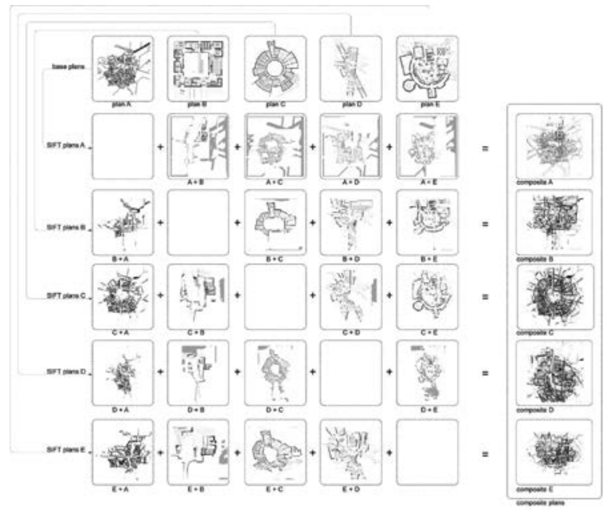

7

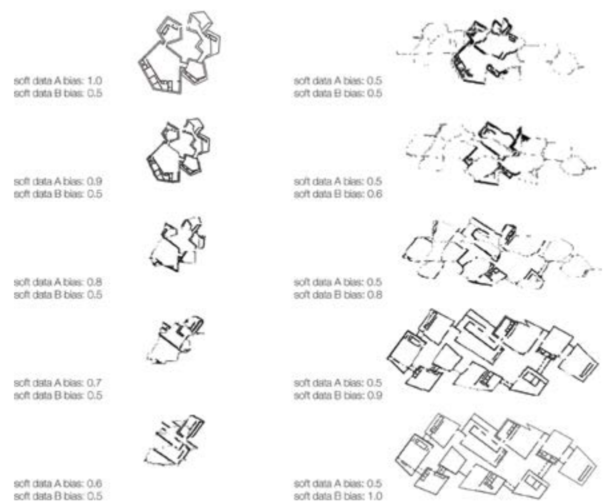

8

7 Drawings are converted to soft data for SIFT processing. Within this initial process, bias is negated and all drawings are a result of equal forces between bodies of soft data.

8 Gradient of biased transformation: Illustration of the bias slider enabling a particular soft dataset to exert more or less influence within the production of data-rich territories.

9 Identified keypoints of SIFTed plans and sections are extruded along their corresponding n-dimensional vectors for intersection and meshing in accordance with the computational protocols of the workflow.

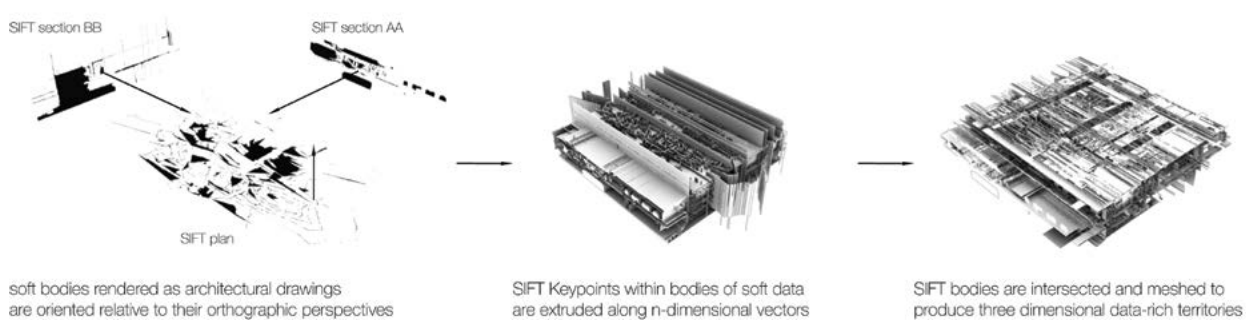

9

producing forms that read as a compression of recycled intentions as opposed to singular ideals.

As this workflow relies on extensive datasets towards the production of extensive collections of data-rich territories, we have started to produce a catalog of deep intentionality that seeks to produce novel forms by intentionally biasing specific characteristics in previously developed architectural artifacts (Figure 10). The characteristics we might choose to bias could be many, of course, but they would share an identification as being instances that fall outside of the formal qualities found in most of the dataset. In the case of the Guggenheim Helsinki Competition, the fact that these projects share a site and space programming allowed us to focus first on the physical characteristics that would be distinct in some projects when compared to the larger group of projects.

Within this workflow, bias is not limited to the production of soft drawings, but is extended to the processes of three-dimensional constructions. Just as any particular body of soft data can exert more influence within the process of image production, the workflow allows for the biasing of a particular plan or section to exert greater force within the resulting formal manifestation. Figure 11 demonstrates how different morphologic manifestations result from the same plan and section simply by biasing the plan over section and vice versa.

Program Identification and Extraction

Just as program is identified through architectural plans and sections this workflow allows for the identification of program within DRTs. By assigning a specific color to designated programmatic areas (public, exhibition, circulation, commercial, service, private), composite image color maps are produced that correlate to a corresponding soft drawing (Figure 12). These color maps produce programmatic heat maps that allow for program identification and organization within an otherwise chaotic environment. Following similar protocols to those developed to produce and extract n-dimensional vectors from soft drawings, programmatic heat maps are SIFTed, producing new sets of vectors that can be mapped to their corresponding image keypoints. These vectors

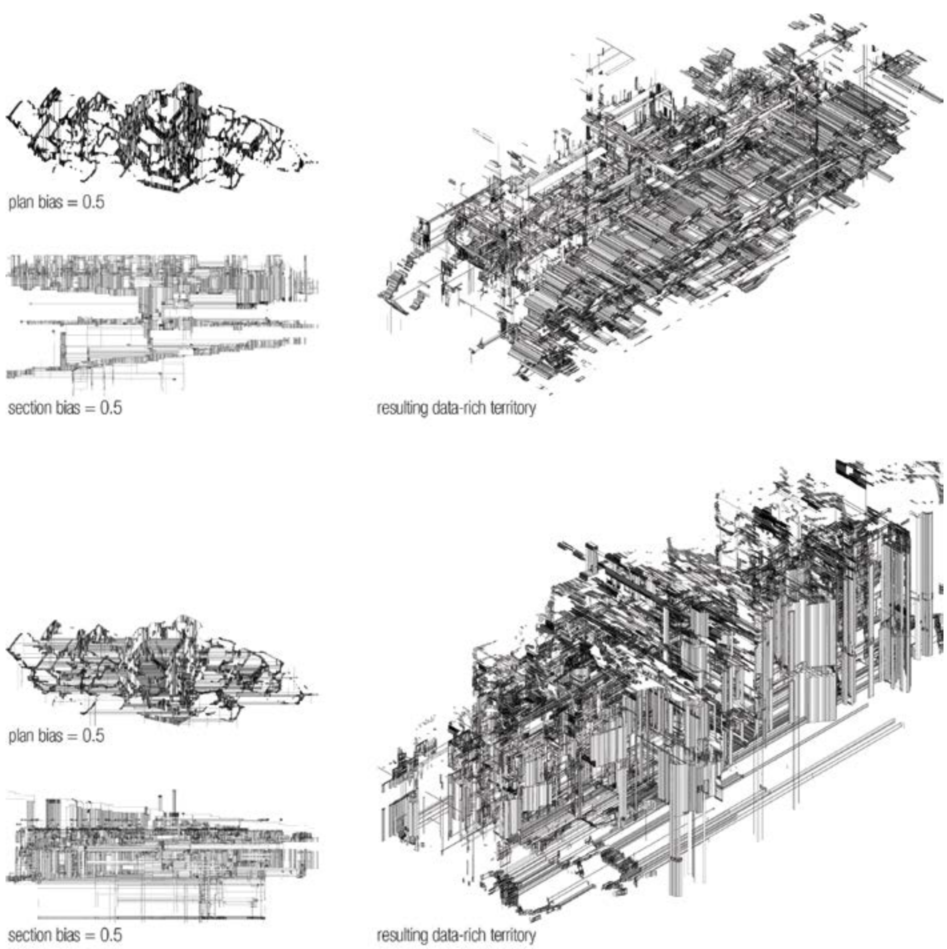

10

10 Selected instances from a growing catalog of deep intentionality. The bias slider allows for the manipulation of certain characteristics that belong to singular plans or sections within the dataset.

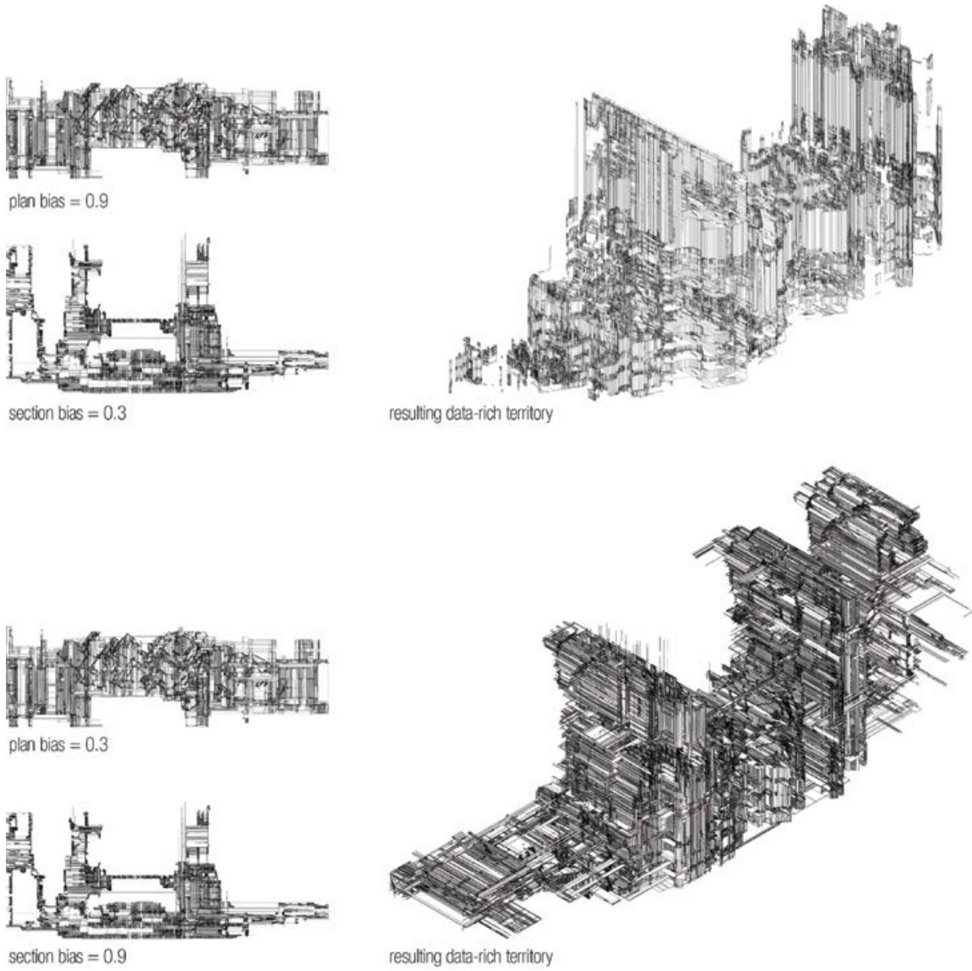

11

11 The initial plan and section that generated these geometries is the same: simply by biasing one drawing set over the other, unique instances can be produced from similar drawings.

are intersected with the DRT, allowing for the identification of program within 3D forms (figure 12). The use of coloration allows for a performative reading of the newly formed territories that is both associative (related to the initial relationships as described in any given project) and hybridized (capable of producing program gradients in three-dimensional space).

CONCLUSION

Generative design strategies have until now predominantly focused on integrating data specific to the environmental responsiveness and formal articulation of buildings within their given regulatory contexts and stylistic/theoretical positions. What we propose here is the exploitation of an ever-increasing catalog of architectural production that is pre-loaded with many of these aspects through a process of computational sensing and evaluation that leads to recombinant data-rich territories. These territories are underutilized resources that represent millions of hours of architectural production and this workflow looks to exploit that resource in ways that are only possible through a computationally mediated framework.

While there has been some success here in creating a workflow that begins to scratch at the potential for mining the available information and processing it iteratively towards new architectural proposals, there are still a number of issues we see as critical to building out a more robust tool. Currently, we are relying on manually identifying outliers for biasing the workflow and an inefficient labor-intensive process for gathering, scaling, and converting architectural image data into a usable format. The ability to automate this function would greatly increase our ability to produce robust catalogs of biased outputs.

Lastly, while the introduction of visual and spatial ‘heat mapping’ programs has led to increased legibility into the potential for innovation and evaluation of outcomes, they would benefit from the introduction of a voiding function which could begin to define spatial relationships within what are currently effectively dense mass structures. This workflow has successfully leveraged existing techniques of data visualization that relate to strategic spatial organization used in sports, and in our view could be used as much as a strategic tool for creating spatial and material

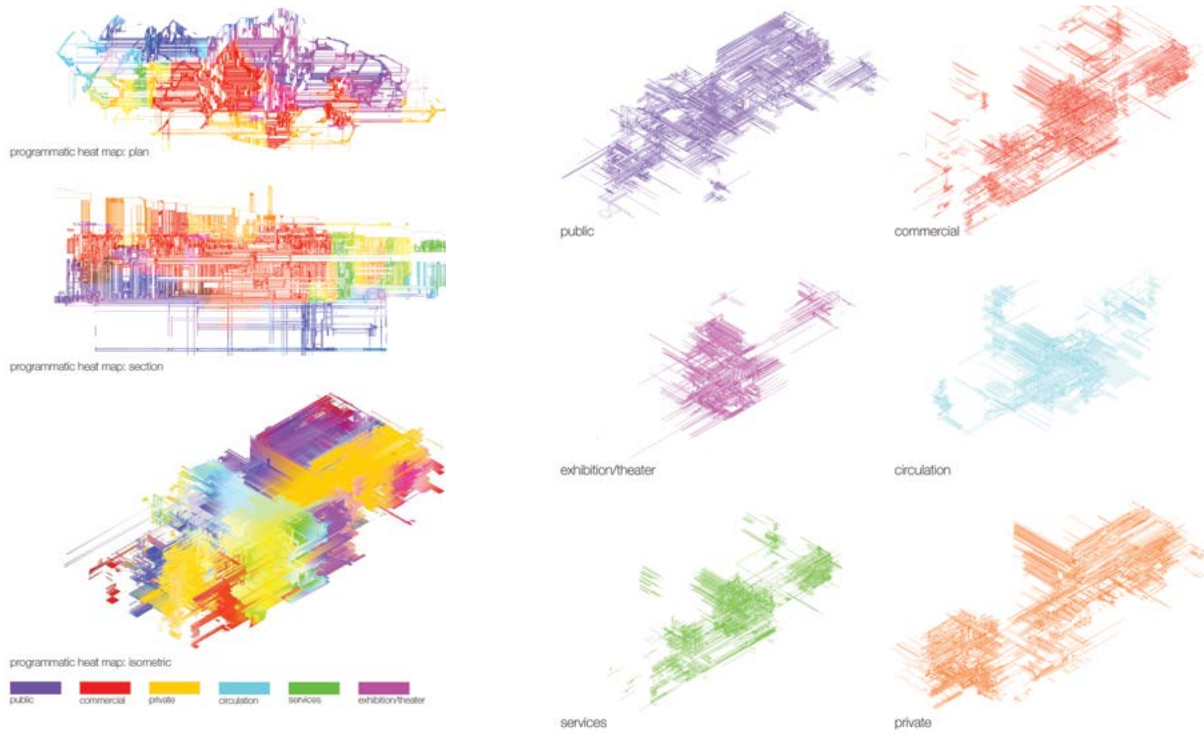

12 Programmatic Heat Maps produced through 2D drawings and mapped to 3D assemblies with the identified programs extracted on the left.

relationships as it could for the production of physical assemblies. The potential exists then for a protocol that is scalable and capable of accessing and processing the ever-increasing body built and unbuilt design proposals whose current value is only that of dead content.

NOTES

1. Ubiquitous Simultaneity is a term coined by the author that identifies a condition which is characterized by the prevalence of massive amounts of data continuously updated and easily accessible through a number of digital interfaces.

REFERENCES

Franks, Alexander, Andrew Miller, Luke Bornn, and Kirk Goldsberry. 2015. "Counterpoints: Advanced Defensive Metrics for NBA Basketball." In *Proceedings of the 9th Annual MIT Sloan Sports Analytics Conference*. Boston: SSAC.

Goldsberry, Kirk. 2012. "CourtVision: New Visual and Spatial Analytics for the NBA." In *Proceedings of the 6th Annual MIT Sloan Sports Analytics Conference*. Boston: SSAC. 1–7.

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

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. doi:10.1109/ICCV.1999.790410.

———. 2004. "Distinctive Image Features from Scale-Invariant Keypoints." *International Journal of Computer Vision* 60 (2): 91–110. doi:10.1023/B:VISI.0000029664.99615.94.

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.

———. 2016. "Ubiquitous Simultaneity: A Design Workflow for an Information Rich Environment." In *CROSS-AMERICAS: Probing Disglobal Networks—The 2016 International Conference of the Association of Collegiate Schools of Architecture.* Santiago, Chile: ACSA.

Manovich, Lev. 2012. "How to Compare One Million Images?" In *Understanding Digital Humanities*, edited by David M. Berry. New York: Palgrave MacMillan. 249–278.

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.

Rowe, Colin, and Robert Slutzky. 1963. "Transparency: Literal and Phenomenal." *Perspecta: The Yale Architectural Journal* 8: 45–54.

Schumacher, Patrik. 2008. "Parametricism as Style - Parametricist Manifesto." Delivered at *Darkside Club, 11th Venice Architecture Biennale*. Venice: Darkside Club: 1–5.

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

Figures 1, 6, 8, 10–12: Johnson and Parker, 2016
Figures 2–3: Goldsberry, 2012
Figures 4, 7, 9: Parker, 2015
Figure 5: Parker, 2014

Jason S. Johnson has practiced and taught architecture in North and South America and Europe. He is a co-director of the Laboratory for Integrative Design, where he holds the position of Associate Professor of architecture. He has exhibited and published his work internationally. His most recent book, co-edited with Joshua Vermillion, *Digital Design Exercises for Architecture Students*, was published in 2016. Mr. Johnson is the founder of Minus Architecture Studio, a design research practice in Calgary, where he lives with his wife and three sons.

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.