On the Architectural Structure of Photographic Space

Mike Christenson[1]

[1]North Dakota State University, Fargo, USA

The ambiguous relationship between photography and architecture is one of constructed and re-constructed identity. As a specific exploration into this relationship, this paper considers the construct of *point-of-vew/field-of-view maps (or POV/FOV maps)*, that is, diagrams which register photographers' positions, fields of view, and directions of view corresponding to a set of photographs of an existing work of architecture. A POV/FOV map can be expected to differ according to whether the set of photographs under consideration is (a) sampled from a image-sharing site such as Flickr; (b) published in an academic monograph; or (c) published in the popular press. This paper tests the extent and significance of these differences through a comparative study of Mies van der Rohe's Crown Hall and Rem Koolhaas's McCormick Tribune Campus Center, both at the Illinois Institute of Technology in Chicago, USA. In both cases, POV/FOV maps are used to compare sets of professional or academic photographs to sets of touristic and popular-press ones. Reflecting the tenuous nature of architectural identity as constructed through photography, the comparison both confirms and denies assumptions concerning differences between professional and amateur approaches. The paper concludes with the speculation that tools like Google Street View are likely to further erode traditional distinctions between modes of identity-construction, in particular, those distinctions which a POV/MAP can register.

Conference theme: Design education
Keywords: Photography, visualization, Mies van der Rohe, Rem Koolhaas

INTRODUCTION

Photography as a practice is uniquely responsible for constructing architecture's identity: knowledge about built works of architecture is largely embodied in and communicated through photographs. Classic examples of the centrality of photography's complicity in constructing architecture's identity include Le Corbusier's appropriation of trade photographs of American and Canadian grain elevators to support his argument about a "New Architecture" (Banham 1989); Bernard Rudofsky's use of apparently anonymous photographs to construct the idea of a "nonpedigreed" architecture (Rudofsky 1964), and more recently, the proliferation of photographs of buildings like Frank Gehry's Bilbao, where the images of the building assume a value which can be understood, in certain contexts, as on par with the building itself (Urry 2002). Yet, architecture simultaneously operates to construct photography's identity as a practice: every building can be understood as a device for selectively revealing and obscuring views. As Beatriz Colomina has succinctly acknowledged, every building is "not simply represented in images but is a mechanism for producing images." (Colomina 1996:97.)

Now, with the advent of technologies like Google Street View and websites like Flickr, new challenges to the mutually constructive relationship between photography and architecture are emerging. Google Street View, or GSV, promises to obliterate the historical idea that architectural photographs in order to be useful must be *of* something: GSV images are of everything, and therefore are of nothing in particular. And yet, their ubiquity makes them uniquely useful and valuable in any effort to "to organise the world's information" concerning architecture (*[Google] Corporate Information* 2010). Flickr promises to obliterate the idea that images must be edited and published by a third party – not the photographer and not the recipient – in order to be widely disseminated. By eliminating the middleman, Flickr and GSV make many more architectural photographs available to everyone anywhere: among these photographs are ones of obvious value and many others whose value is not so obvious. Technologies such as Flickr and GSV, which promise to challenge old ideas about the relationship between photography and architecture, demand new critical tools for assessment. This paper proposes just such a tool in the *point-of-view/field-of-view map.*

1. THE MAPPABLE SPACE OF PHOTOGRAPHIC COLLECTIONS

1.1. Definitions

Given a collection of photographs of a specific building, the photographers' positions, fields of view, and directions of view can be mapped. For a given collection, the resulting map is defined as a *point-of-view/field-of-view map,* or a *POV/FOV map*. In the example shown in Fig. 1, photographers standing at points (a), (b), and (c) aim their cameras

at an existing building to generate specific views.

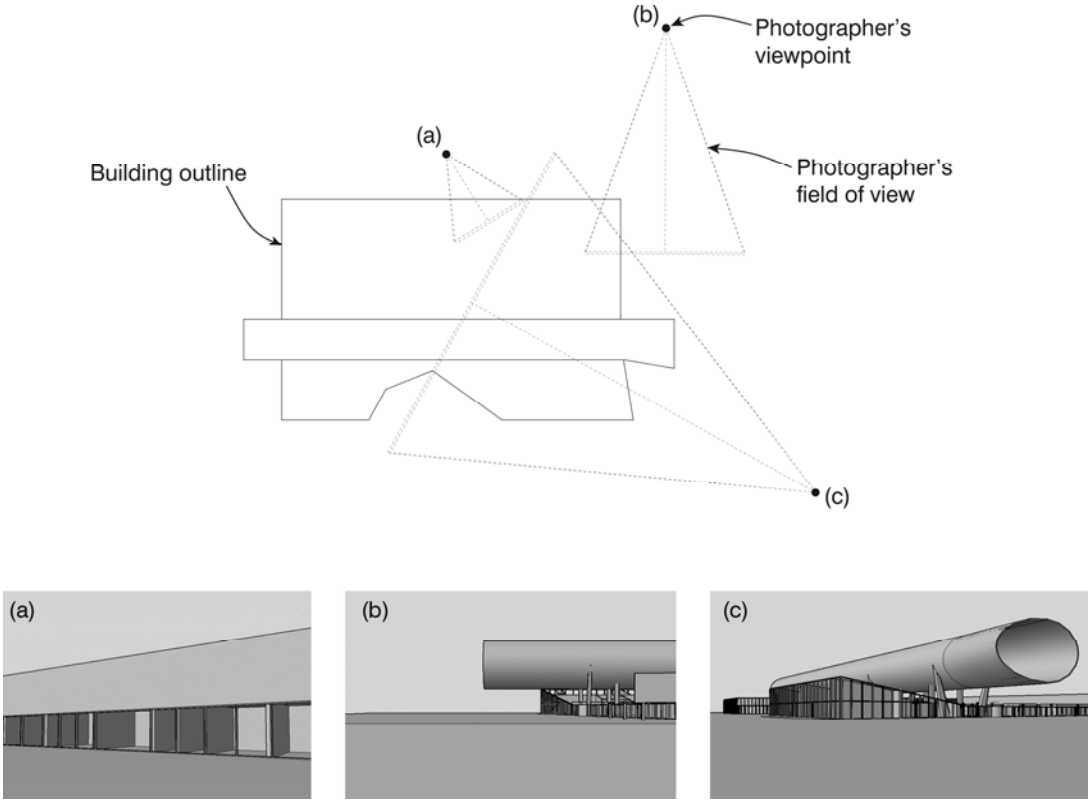

Source: (Author)

Figure 1: Illustration of a POV/FOV map

A diagram which registers multiple points of view and fields of view – a POV/FOV map – depicts what I call the "space of photography," a three-dimensional dataset specific to a given collection of photographs. In the example shown here, the space of photography corresponds to the set of photographs (a), (b), and (c), but POV/FOV maps can be generated from any number of photographic images. My research into the construction of POV/FOV maps, such as the work I describe here and elsewhere (Christenson 2008), differs from related research into photography and architecture in that I focus on two factors – *view geometry* and *viewspace fragmentation* – both of which I see as significant in assessing the mutually constructive relationship between photography and architecture.

First, unlike much closely related research, my research inquires not only into viewpoints and view directions but also into *view geometry:* the shape of individual photographic fields of view. As an illustration of related research, consider the possibility of data-mining information from georeferenced images or from user tags (Quack ***et al*** 2008). Such research assumes the georeferenced images are distributed on a modelable surface, e. g., on geospatial tiles; it makes assertions concerning, for example, the movement of photographers on this surface over a given time period; and it draws conclusions supporting applications such as auto-annotation of otherwise unreferenced images. My own research, by contrast, attempts to say something about the *three-dimensional* space in which photographers operate. My analysis is not limited to points on a surface, but extends to the specific geometries of views from points in three-dimensional space, as occurs for example when cameras are directed upwards or downwards to objects vertically distant from the photographer's vantage point, or when a photograph is taken from the top of a tower or a sunken court, or when a given view is oriented vertically rather than horizontally, or when a photograph is produced with a telephoto or a wide-angle lens.

Viewspace fragmentation is defined here as the failure of a given set of photographs to completely map a given building or site. This issue is of concern in *multi-view 3D-reconstruction research* which attempts to reconstruct three-dimensional models of buildings, sites, and whole cities from collections of photographs (Agarwal ***et al*** 2010). Like 3D-reconstruction research, my research recognises the fragmentary and incomplete nature of photographic collections. However, while 3D-reconstruction research attempts to patch over this incompleteness to provide a complete and accurate "dense" model of a given site, my research attempts to highlight and underscore the differences between idiosyncratic views of the site. 3D-reconstruction work, because it is interested in the object of photography (the building or the site), measures the value of its results by the accuracy of models. By contrast, my research addresses the space created by subjects: because different subjects photograph buildings differently, we can expect that any given "space of photography" will consist only of a fragmentary or incomplete view of the

building. Rather than characterising such fragmentary views as a liability, I am precisely interested in *the mappable differences between "spaces of photography" corresponding to different subjects*.

1.2. Problems with aggregation

Consider the problem of extracting meaningful three-dimensional information from a dataset consisting of the aggregated view content, points of view, and fields of view of a collection of photographs. This problem is faced equally in my research and in related 3D-reconstruction research, and is difficult for two reasons:

a. *Simultaneous unknowns in 3D reconstruction:* specifically, computational difficulties related to constructing a three-dimensional digital model based on information extracted from a set of two-dimensional images, such as photographs taken by a subject in motion around an object. The computational difficulties arise because, in general, both the positions of objects and the positions of photographers are simultaneously unknown. Research into this area attempts to resolve these two simultaneous unknowns through iterative processes seeking optimisation. (Seitz *et al* 2006; Hartley and Zisserman 2003.)

b. *The weak accuracy of georeferencing,* that is, the lack of geographically precise metadata associated with individual photographs. The scale of this problem increases with the explosion of casually-produced photographs available via the Internet. Thus, while it may be possible to identify a subset of images in a photographic collection as belonging to Paris, London, or Bombay, it is not a simple matter to provide more specific georeferencing without access to additional metadata. Moreover, the computational strategies for automating a georeferencing process are not obvious (Crandall *et al* 2009; Serdyukov *et al* 2009).

In my own research, I was able to bypass the simultaneous-unknown difficulties because I constructed highly detailed digital models of the sites prior to conducting image analysis. Thus, it was possible for me to manually calibrate software-simulated photographic images to align with photographs' specific points of view and fields of view. I bypassed the weak-accuracy issue because the collections I analysed were small enough, and sufficiently focused in scope, such that each image in the collection was specifically identifiable to a precise location within well-known sites.

1.3. Model construction

As discussed above, constructing a POV/FOV map of a given "space of photography" depends on having access to an accurate digital site model. For the work described here, although extant digital models of both buildings were available, I chose to construct my own models based on field inspection and document interpretation. Over the course of two visits to the IIT campus in Chicago in 2006 and 2009, I photographed and measured both buildings. With reliance on the field information, precisely-scaled architect's drawings, and images published online at Flickr, I used AutoCAD to construct detailed digital models of each building in 2008-2010. Cross-checking shows the accuracy of the finished models to be within 5% of correspondence to the built structures and thus sufficiently precise to enable generation of accurate perspective images for comparison with subject photographs.

Given a specific photograph to be mapped in a POV/FOV map, I placed a simulated camera in the digital model at a position approximating the apparent point of view in the photograph. I then manually compared the view produced by the simulated camera to the target photograph, making necessary adjustments to camera properties and variables (e. g., station point x-, y-, and z-coordinates; target point x-, y-, and z-coordinates; simulated lens length; and frame orientation), in order to achieve correspondence to the given image. As discussed above, this process differed from the automated processes typical of 3D-reconstruction work because the positions and sizes of objects within the photographs were in this case precisely known and encoded in the digital model.

Following the placement of a complete set of simulated cameras corresponding to a specific collection of photographs, I exported top and side views of the AutoCAD model, including all cameras and the outlines of their cones of vision, to Adobe Illustrator. I used Illustrator to make manual adjustments to lineweight and linetype, supporting legibility and clarity of the map.

1.4. Speculation on difference

As briefly discussed above, a "space of photography" associated with the work of one photographer can be expected to differ from that associated with a different photographer. (Stated more simply, different people behave differently when photographing the same building.) For example, consider what kinds of difference might exist between, on one hand, a collection of photographs produced by a solitary professional or academic photographer engaged in a systematic documentation of an existing building, and on the other hand, the collection of photographs compiled from the aggregated work of several casual photographers and posted to an image-hosting site such as Flickr. We can assume that the POV/FOV maps of these two collections will differ, but to what degree? Would a comparison of POV/FOV maps disclose significant patterns in the photographic practices of different photographers at a single site? Is it necessarily true that collections of professional or academic photographs will differ in a consistent way from collections of tourist photographs of the same building? Conversely, if we know the configuration of a specific POV/FOV map, can we make reasonable assumptions about the work of architecture to which it is likely to correspond, or about the collective biases and practices of the photographers whose work is mapped? And when we consider the act of design, is it possible to make reasonable predictions about architecture's ability to structure a particular kind of photographic collection? Such questions, many of them reflective of long-held assumptions concerning differences between amateur and professional architectural photography (e. g., Rudofsky 1964), are exactly the kinds of questions which my research seeks to address. The POV/FOV map becomes, in my work, the primary tool for testing such assumptions, precisely because it registers in a uniform way the mappable differences

between and among photographic collections.

2. COMPARATIVE POV/FOV MAPS

2.1. Assumptions

As a test case for exploring these questions, I selected two buildings at the Illinois Institute of Technology Campus in Chicago, USA: Mies van der Rohe's Crown Hall and OMA's (Rem Koolhaas's) McCormick Tribune Campus Center. Apart from the obvious and polemical formal and material contrasts between the two buildings (Becker 2003), the buildings can be said to register two distinct philosophies of the mutually constructive relationship between photography and architecture. While Mies, as Colomina has argued, maintained an ambiguous, even "schizophrenic," attitude toward photography's role in the construction of architectural knowledge (Colomina 1996), Koolhaas's work is a forthright acknowledgement of this role and its effect in the shaping of architecture (Koolhaas and Mau 1997).

For what follows, let us assume that there exist distinctly identifiable approaches to the production of architectural photographs. First, suppose that there exists an "academic" approach, requiring a single photographer, working in a fixed, short-duration time period, producing a collection of high-resolution photographs on the assumption that the collection will be professionally edited and published in a topically-specific publication with a limited audience (e. g., an academic or professional one). Next, assume there exists a "touristic" approach to photographing architecture. In this approach, while photographers may or may not expect that their photographs will be widely disseminated via the Internet, they do *not* expect that their photographs will be professionally edited by a third party prior to dissemination. Due to the lack of editorial filtering we can thus reasonably expect touristic images to be less consistent in quality than their academic counterparts. This touristic approach to photography is typical of the photographs which appear on websites like Flickr. (It is obvious that not all photographic collections can be categorised as either academic or touristic; the categories are provisional rather than exhaustive, and arguably, they could overlap in specific cases.) Following these assumptions, I identified two distinct collections of photographs of each subject building for comparison, the first collection representing an academic approach, associated with a single photographer, and the second representing a touristic approach, associated with photographers in the aggregate. I designated these collections *C1* and *C2*, for Crown Hall, and *M1* and *M2*, for the McCormick Tribune Campus Center.

2.2. Comparing Flickr photographs of Crown Hall with Werner Blaser's photographs of Crown Hall

Collection C1 is taken from Werner Blaser's book on Crown Hall (Blaser 2001). C1 contains 28 images, 17 of which are exterior and 11 of which are interior. (An additional 8 photographs from Blaser's book were not included in C1 because they were produced by other photographers and do not necessarily depict Crown Hall.) Collection C2 consists of 25 images downloaded on July 26, 2010 from Flickr.com, using the quoted phrase "crown hall". The total number of hits resulting from this search was 1,850, of which the first several hundred were exclusively of Crown Hall. C2 resulted from the combined efforts of eight photographers. Of the first 26 images which resulted from this search, one was discarded because it did not depict the built structure. C1 represents the "academic" collection while C2 represents an example of a "touristic" collection.

The two collections, when mapped, result in two distinct POV/FOV maps, as diagrammed in Figs. 2 and 3. Each map shows Crown Hall in plan, showing the photographers' points of view and outlines of their respective cones of vision. Each diagram therefore records the visits of one or more people by registering the superimposition of multiple cones of vision.

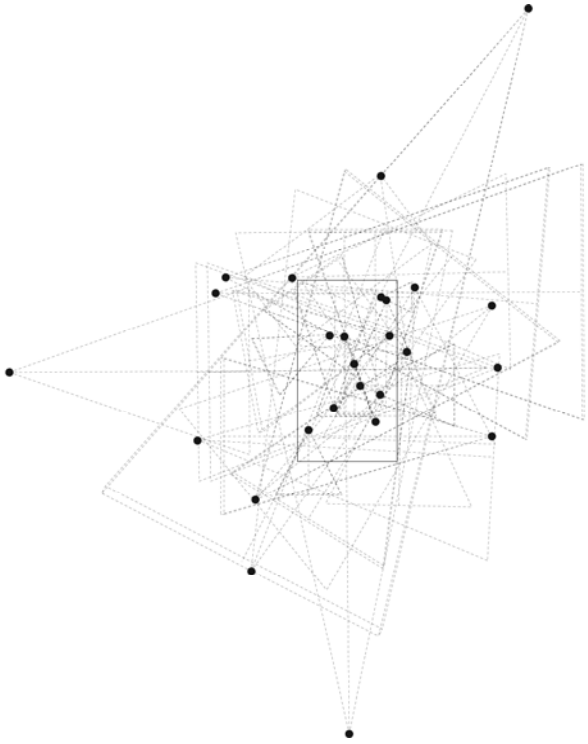

Source: (Author)

Figure 2: POV/FOV map of C1 (Blaser images of Crown Hall)

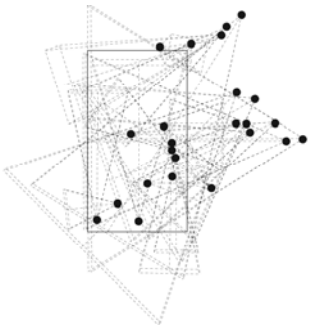

Source: (Author)

Figure 3: POV/FOV map of C2 (Flickr images of Crown Hall)

These superimposed fields register two distinct photographic collections reflecting the divergent but overlapping interests of a dedicated scholar on one hand (C1), and tourists on the other (C2). Comparing the POV/FOV maps of C1 and C2 is to contrast tourists' collective behaviour against the purposeful visit of an individual photographer engaged in the production of a scholarly work attempting to document the building in an academically useful manner. As one might expect to see, the POV/FOV map associated with C1 exhibits by comparison with C2 a relatively uniform distribution of viewpoints and a tendency to pair opposing views, for example, from adjacent corners of the building. This distribution reflects the idea that a deliberate documentation of an existing structure will necessarily proceed in a methodical and consistent way. The POV/FOV map for C2 exhibits a skewed distribution of viewpoints: most of the views are from the corner of the building nearest the local train station – which is not surprising, considering that this is the direction from which visitors are likely to first approach the building.

2.3. Comparing Flickr photographs of the MTCC with Yukio Futagawa's photographs of the MTCC

Collection M1 consists of 24 photographs of the MTCC taken by Yukio Futagawa and published in a photoessay in

GA Document featuring the MTCC (Futagawa 2003). Like Blaser's photographs, Futagawa's photographs were selected because they exemplify a peer-reviewed, single-photographer, editorially-filtered approach to documenting a building. In comparison, collection M2 consists of the first 50 images resulting from a Flickr search, conducted on July 22, 2010, for the unquoted phrase "iit koolhaas." The total number of hits resulting from this search was 2,301, of which the first sixty or so images were exclusively of the MTCC. In the later results, images of other IIT or Chicago-area sites appeared alongside MTCC images. Collection M2 was produced by the combined efforts of 26 photographers, whose Flickr usernames included "!architect4!" and "crazyegg95."

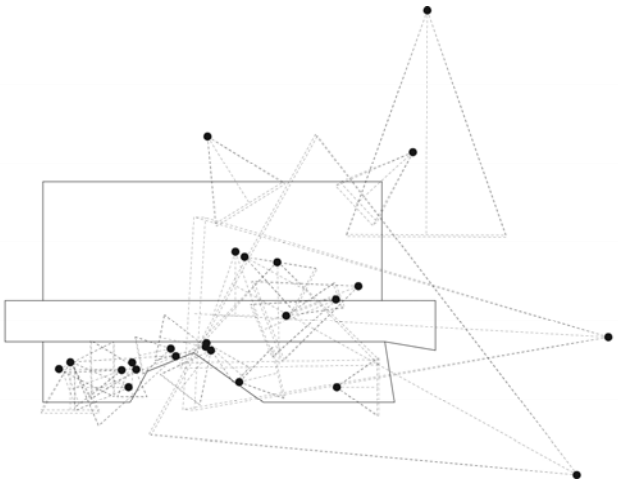

Source: (Author)

Figure 4: POV/FOV map of M1 (Futagawa images of the MTCC)

When M1 and M2 are mapped (Figures 4 and 5), the results differ significantly from the similar comparison of Crown Hall POV/FOV maps. Where C1 and C2 show an obvious, and expected, difference between the so-called "touristic" and "academic" approaches, no such difference is obvious with M1 and M2. both M1 and M2 show the same kind of skewing towards the public spaces, particularly showing a bias toward images which include the orange lenticular glass at the building's west facade. Both M1 and M2 exhibit a tendency for the photographers to avoid the "back" side of the building; unlike Blaser's work at Crown Hall, Futagawa did not make an obvious effort to circumambulate the building.

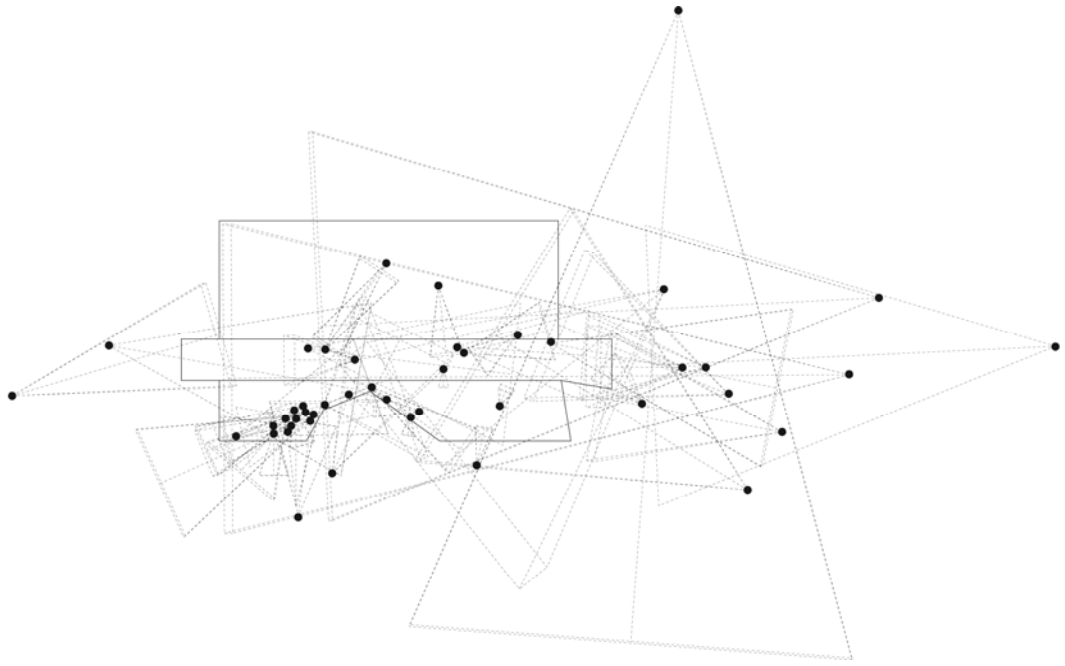

Source: (Author)

Figure 5: POV/FOV map of M2 (Flickr images of the MTCC)

2.4. Discussion

Comparison of the multiple POV/FOV maps does indeed disclose significant differences in the patterns of photographic practices at a single site. Although all of the photographers (Futagawa, Blaser, and Flickr users) are confronted with the problem of constructing architectural knowledge concerning these buildings, their approaches are visibly different. Sets C1 and M1 can be said to represent similar approaches to the production of architectural photography in that both are produced with similar assumptions about dissemination and use. However, a comparison of their associated POV/FOV maps indicates no obvious commonality in approach: C1 indicates Blaser's act of circumambulating the building while M1 has no such indication. But, if in addition to spatially mapping the photographs in M1, we also map *the order in which they appear in their published context,* M1 appears to suggest a "walkthrough" (though not a "walk-around") of the building. By contrast, both C2 and M2 (the "touristic" collections) include several images of essentially identical content taken from vantage points in close proximity. This tendency suggests the effects of a *circle of representation,* wherein people visiting the site attempt to reproduce photographs of the building which they have previously seen (Jenkins 2003), but likely also represents an apparently common practice of photographing a building from the first point where it presents itself to the viewer. Photographs of the "back" of each building are almost completely absent from both C2 and M2, suggesting that visitors tend not to circumambulate the building when they visit.

The evidence of their visits suggests that Blaser and Futagawa, although approaching the problem of architectural photography as scholars, assume different attitudes about how architecture should be understood remotely (i. e., to a viewer of their published photographs). Where Blaser approached the question dogmatically ("neutrally" is not the right word to use), Futagawa appeared more willing to structure his photography according to his movement through the building. Significantly, Futagawa's photographs exhibit a structure very similar to the structure of aggregated tourist views. Flickr users, meanwhile, may be scholars, but they are also tourists; they may be professionally or casually interested in a specific work of architecture. What certainly unites Flickr users is their common interest in "sharing" their work with a global audience in an effectively unedited forum.

Although the work discussed here is a small sample, it is suggestive concerning the possibility of projecting assertions about a work of architecture from a given POV/FOV map. In three of the four cases examined (C2, M1, and M2), the POV/FOV maps showed a distinct bias in image distribution, illustrating a tendency of photographers to concentrate on areas of the building near the entry, or along a popular approach path. If we observe a new POV/FOV map with a similar distribution bias we may be prompted to draw conclusions about the building's immediate surroundings and the location of its main entry relative to its perimeter. Again, allowing for the small sample size, we might also assert that buildings do indeed have the ability to structure particular kinds of photographic collections, such as ones biased toward the building's main entry. The brief evidence presented here suggests that it is also important to understand the building's immediate surroundings – its physical "context" – when attempting to predict the shape a particular collection of photographs will assume. For example, if a designer knows the direction from which visitors will approach the building, the designer already knows something about how the building will be

photographed. To paraphrase Colomina, the building's context becomes a "mechanism for producing images" (Colomina 1996: 97). For any designer interested in managing their work's *mediality* (i. e., reality perceived through media), and particularly in an era when one's access to images of a building increasingly precede one's direct experience of the building (Lasansky and McLaren 2004), the physical limitations of a contextual "image-producing mechanism" in the form of the building's site assume critical importance.

CONCLUSION / FUTURE WORK

This paper has summarised the construction of *POV/FOV maps,* which are diagrams registering a three-dimensional dataset associated with a specific collection of architectural photographs. A test case comparing POV/FOV maps for two existing buildings, in pointing out significant commonalities and differences, demonstrated the importance of analysing POV/FOV maps in an attempt to query the structure of architectural knowledge, or to predict its form. In particular, the test case suggests that an analysis of POV/FOV maps is capable of confirming as well as denying assumptions about differences between professional and amateur architectural photography.

Some possible directions for next steps in this research are *POV/FOV map translation* and *occlusion mapping*. In the construction of a POV/FOV map, the information about photographers' points of view and cones of vision is digitally stored as AutoCAD "cameras." It is possible to copy or translate the cameras as a group from one model to another. Such a translation prompts the question whether a collection of photographs, specifically constituted with respect to one work of architecture, could become an appropriate tool in the knowing or the making-visible of a different work of architecture. Clearly, once a collection of AutoCAD cameras is created, its translation to any other model is technically trivial. But in making the new model visible in a way which was generated by another work, a host of new questions emerge: How can a point of registration, or alignment, between two different models be determined? What are the implications for shifting the point of registration, or the orientation, of a translated collection of cameras?

A possible point of convergence between my research and 3D-reconstruction research arises with the question: What are the features of the built environment which physically restrict or occlude photography? Reconstruction research would seek to develop algorithms to circumvent these features in support of dense model construction, while my research would treat these features as common factors affecting the construction of what might be otherwise structurally divergent photographic collections. It is not obvious whether an efficient methodology for mapping such features could improve the efficiency of either 3D-reconstruction research or of my work.

As a further comparison, I prepared a POV/FOV map associated with Google Street View's images of both buildings discussed in this paper. GSV photographs are produced by an omnidirectional camera mounted atop a moving vehicle. Thus, in general, GSV's points of view are restricted to public streets, but directions of view are not limited. GSV's views, as discussed previously, are not *of* anything in particular, and thus (as should be obvious) the associated POV/FOV map demonstrates no inflection or skew due to the presence of the IIT buildings. To extend this observation past the IIT campus, note that GSV photographs will not "inflect" whether the photographing vehicle is passing by Lincoln Cathedral or Nikolaus Pevsner's bicycle shed (Pevsner 1960). Thus, GSV forces an acknowledgement of a distinction between *intentional* and *unintentional* photography, and the effects this distinction may have on the way knowledge about buildings is constructed. Yet, other effects of GSV (which I discuss elsewhere), tend to *erode* historical distinctions: for example, GSV has the effect of blurring a distinction between on one hand, place-specific photography, and on the other hand, information which is widely disseminated and available everywhere (Christenson 2010). And to return the speculation to POV/FOV maps, note that although GSV does not currently host images of building interiors, in theory there are no technical limitations to doing so. Yet, it is not obvious what configuration interior GSV-like views would take. Would these interior GSV views, in an attempt to be everywhere-applicable like the ubiquitous street-level exterior views, assume a configuration like Blaser's dogmatic approach to Crown Hall? Or, like Futagawa's photographs, would such a collection inflect to something like a casual-walk progression through the building?

Finally, what does this research suggest concerning the mutually constructive relationship between photography and architecture? Because my research questions the degree to which photography should be treated non-critically as a means of constructing architectural knowledge, it seems reasonable to allot this work, and specifically POV/FOV maps, a place in the discussion whenever photographs are used to structure knowledge about architecture.

ACKNOWLEDGEMENT

I gratefully acknowledge the work of this paper's anonymous reviewers. Their comments on an earlier version of this paper reflected their careful attention to detail, substance, and purpose.

REFERENCES

Agarwal, S., Furukawa, Y., Snavely, N., Curless, B., Seitz, S. M. and Szeliski, R. (2010) Reconstructing Rome. *IEEE Computer,* June 2010, 40-47

Banham, R. (1989) *A Concrete Atlantis: U. S. Industrial Building and European Modern Architecture 1900□1925*. The MIT Press: Cambridge, USA and London

Becker, L. (2003) *Of Mies and Rem*. Retrieved from http://www.chicagoreader.com/chicago/of-mies-and-rem/Content?oid=913290

Blaser, W. (2001) *Mies van der Rohe: Crown Hall: Illinois Institute of Technology, Chicago*. Birkhauser: Basel.

Christenson, M. (2008) On dependencies between architecture and media: considering the remote work. *ARCC Journal,* 5 (1), 6-15

Christenson, M. (2010) Google Street View and the transition from the unknown to the known. In *Chang[e]ing Identities: Design Culture Technology (Proceedings of the 2009 ACSA Southwest Regional Conference, Albuquerque, New Mexico),* 76-81.

Colomina, B. (1996) Mies Not. *Columbia documents of architecture and theory: D,* 5, 75-100.

Crandall, D., Backstrom, L., Huttenlocher, D., and Kleinberg, J. (2009) Mapping the world's photos. In *Proceedings of the 18th International World Wide Web Conference,* 761-770

Futagawa, Y. (2003) McCormick Tribune Campus Center, Illinois Institute of Technology, Chicago, Illinois. *GA Document* 76, 8-47

[Google] Corporate Information – Company Overview (2010) Retrieved from http://www.google.com/corporate.

Hartley, R. and Zisserman, A. (2003) *Multiple view geometry in computer vision*. Cambridge University Press: Cambridge, UK

Jenkins, O. H. (2003) Photography and travel brochures: the circle of representation. *Tourism Geographies* 5 (3), 305–328

Koolhaas, R., and Mau, B. (1997) *S, M, L, XL*. Monacelli Press: New York.

Lasansky, D. M., McLaren, B. (2004) *Architecture and tourism: perception, performance, and place*. Berg: Oxford; New York.

Pevsner, N. (1960) *An outline of European architecture*. Penguin Books: Harmondsworth, UK

Quack, T., Leibe, B., Van Gool, L. (2008) World-scale mining of objects and events from community photo collections. In *Proceedings of the 2008 international conference on Content-based image and video retrieval,* 47-56

Rudofsky, B. (1964) *Architecture without architects: an introduction to nonpedigreed architecture*. Museum of Modern Art: New York.

Seitz, S. M., Curless, B., Diebel, J., Scharstein, D., and Szeliski, R. (2006) A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms. In *Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition,* 519-528

Serdyukov, P., Murdock, V., van Zwol, R. (2009) Placing flickr photos on a map. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval,* 484-491

Urry, J. (2002) *The Tourist Gaze*. Sage Publications: London; Thousand Oaks, California, USA.