TOWARD A UNIFIED REPRESENTATION SYSTEM OF PERFORMANCE-RELATED DATA

Ebenezer Hailemariam, Michael Glueck, Ramtin Attar, Alex Tessier, James McCrae, Azam Khan
Autodesk Research
Toronto, Ontario, Canada
first.last@autodesk.com

Figure 1: Visualization techniques combined to simultaneously convey a number of performance-related values.

ABSTRACT

An environmentally-responsible building is not a fixed ideal but a moving target that must be reassessed on an ongoing basis in order to respond to the ever changing patterns of its occupants and its context. Instrumented buildings can generate a tremendous amount of data that can play a key role in gaining insight into their changing behaviour and performance. Although visualization has been used as a technique to interpret this data, representation strategies have been limited to 2D graphs or abstract numerical outputs. In this paper, we integrate a set of well-defined interactive visualization methods into one unified framework where a number of inputs pertaining to building performance data can be simultaneously visualized. Furthermore, we display how this integration allows directly relating both spatial and non-spatial data on top of a high-resolution 3D model in order to provide a rich context for more holistic visualization of interacting performance factors.

Keywords: data visualization, sensor network, existing building.

INTRODUCTION

It is well known that building performance analysis entails a very complex study of many interacting factors. This high level of complexity of inter-related behaviours and properties has given rise to many methods that can be used to determine the impact of a building on the surrounding environment and vice versa. Many of these analysis methods themselves benefit from sophisticated numerical methods, for optimization and efficiency, which represent their outputs as purely analytical data often separated from the visual context of the building.

Historically, visualization has been used as a visual aid to interpret analytical data. A key challenge is to define methods of organizing, studying and communicating data in a manner that promotes a more holistic understanding of building performance in relation to the spatial and contextual configuration of the building. While many visualization techniques have been developed to represent abstract data structures, such as hierarchical or directed graph data, datasets that are inherently spatial can benefit from visualizations that exist within, or overlay, the base geometry. In architecture, although such visualization systems exist that represent the effect of the environment on the building, few systems focus on combining different techniques into one context. Furthermore, current visualization techniques generally do not directly relate spatial and non-spatial data.

Previous research has shown that "perceptualization needs" of performance analysis can be enhanced by overlaying results on 3D objects (Prazeres and Clarke, 2005). The current generation of advanced commercial simulation applications such as Autodesk Ecotect benefit from this approach by providing graphical user interfaces to display simulation results dynamically on top of a 3D model. However, since the design process involves a high degree of abstraction, models are generally presented as light-weight approximations of physical configurations in order to facilitate greater flexibility in exploring alternative configurations.

Therefore, there has been little research in developing novel methods of interaction and visual representation of building performance data within a model with high visual fidelity (Srivastav et al., 2009).

Our research project involves our existing office building located in Toronto, ON, Canada. In this paper, we present our progress toward a unified representation system where various simulated and measured data are visualized within a high-resolution 3D environment to create a much richer context for visualizing various performance data within an existing building.

METHODOLOGY

Overview

Our approach toward a unified representation system is based on a set of well-defined interactive visualization methods that can represent a number of different kinds of input pertaining to building performance data. Our system receives its data either directly through sensor network data collection or by feeding its collected data to a simulator to produce additional data points. In both scenarios, the input to our visualization system is a data field. Next a transfer function is applied to the data. The transfer function quantizes, biases, scales and filters the data according to what is expected by the visualization. In the last stage the visualization is applied to the transformed data and the data is displayed in context (see Figure 2).
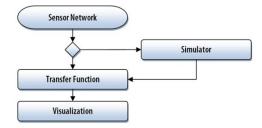

Figure 2: A flow chart representation of our unified representation system.

Sensor-based Input

With the advent of ubiquitous computing, we are witnessing a growing trend in monitoring and evaluation of building performance factors using sensor-based networks (Essa, 2000). Instrumented buildings can generate a massive amount data, which is typically used to reveal patterns of energy consumption or related activities in a building. Collected data is typically aggregated and represented as graphs or numerical tables outside of a 3D environment. We suspect this is largely because existing buildings were conceived using traditional design methods without having a corresponding digital 3D model for

visualization. However, the increasing prevalence of sophisticated commercial 3D applications, which are used to construct and manage the lifecycle of a building, makes it enticing to synchronize a sensor-based monitoring system with a 3D model of a building. As an example, in this paper, using wireless tracking devices we are able to collect and visualize the electricity-usage data of real-world power outlets within the context of the 3D model.

Data-driven Simulation Input

Unlike power usage, some phenomena are more complicated and cannot be accurately measured using sensor data alone. These types of data are still best simulated within the 3D environment, but can be guided by real-world input. For instance, the intensity and temperature of air flow through an office can impact the comfort of the occupants. However, discretely sampling the airflow throughout the office is not an easy task. Thus, we propose that these complex systems can be analysed through simulation. For instance, taking a measurement of air flow intensity and temperature at the vent register can be used to seed a simulation of how that air flow will react within the space of the 3D model. From the standpoint of a unified representation system, the simulation results can then be dynamically visualized in our 3D environment using the techniques presented in the subsequent sections of this paper.

Data Types

Data, acquired from sensors or simulated, can be described as either logical, scalar, vector or semantic.

Logical data. These are discrete data points with two states: true or false. They can be used to describe the presence or absence of a person, object, selection or phenomenon.

Scalar data. These are continuous one-dimensional real values. Data of this type includes temperature and energy usage.

Vector data. These are continuous real values, with direction and magnitude. Data of this type includes air velocity.

Semantic data. These data refer to tags and properties which identify geometry as building elements, such as walls, structure, HVAC, chairs, desks, windows, etc.

All data types are grounded by a contextual 3D position in a continuous volumetric field, so the numerical input to our visualizations can be categorized into *logical fields*, *scalar fields*, *vector fields and semantic fields*. Individual fields contain data types which represent the same instant in time.

VISUALIZATION METHODS

Typically, visualization methods involve mapping particular data values to visual features. In our approach, we map various kinds of input data in the context of a high-fidelity 3D model of the existing building. This plays an important role in reducing the gap between the data and the user's mental model. As noted by Gershon and Eick (1995), since our perceptual capabilities are tuned to the physical world, "it is easier to convey the information to the observer if the information is represented by being mapped to the familiar physical space." Essentially this will decrease a user's cognitive load by harnessing their perceptual capabilities of interpreting visualizations within a 3D space. By representing data visually and within the context of the 3D model, we hope to leverage the affordances provided by the model and also allow users to better understand visualizations in relation to both the space in which they occur as well as other data. Therefore, it is important to be able to clearly and succinctly convey the shape, form, and relative position of geometry in the scene.

3D Model of Existing Building

Our 3D model is basically a highly-detailed representation of the existing condition. The geometry of the 3D scene provides a rich source of contextual information to the user. Whether one is familiar with the specific office space or not, general usage of rooms and spaces can be deduced based on the objects present. For example, a cubicle is commonly understood to be the workspace of a single individual, typically consisting of a chair, a desk and storage compartments. Auxiliary devices typically used in cubicles include computers, desk lamps and telephones. Without explicitly presenting all of these objects, we can safely say that much of this detail is implicitly conveyed to the user.

Incorporating building semantics is another important role of the 3D model in our system. Since our real-time data collection is directly linked to local features in an existing building, our system needs to distinguish between geometries representing specific existing elements. The 3D geometry in the scene was created as a Building Information Model (BIM) using Autodesk Revit Architecture and MEP (Attar et al., 2009). We also simplified the model by exporting a selected area of the building, which we then load into our visualization software. Since our visualization software cannot yet accept semantic building information directly from Revit Architecture, we manually re-authored many of the semantic groups inherent within the original file. These semantics include the hierarchical relationships between geometry, occupancy zonings and interior components. In this way, groupings such as "the research area", "cubicles", and "outlets", as well as hierarchical relationships, such as "the chair at the desk in the cubicle", are maintained in our 3D environment.

When rendering the geometry of the existing building, we avoid realistic texturing and lighting, since it will seem to be a visualization of a lighting system. High visual fidelity does not guarantee that all geometry in the scene is equally emphasized. Instead, we opt for a "less-is-more" approach. We choose to render the scene in a lighting-neutral way with ambient occlusion, a global illumination technique that represents spatial relationships between surfaces (see Figure 3). This technique has been presented as an effective way to not only minimize extraneous visual information, but also provide balance of importance amongst all regions of the scene (Rivotti et al., 2007).
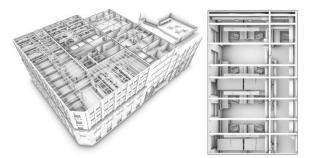

Figure 3: Ambient occlusion shading: the entire office building (right) and "the research area" (left).

Within the context of this 3D scene, we propose three methods of visualizing the captured data fields: direct rendering methods, transient geometry and heads-up displays.

Direct Rendering Methods

Direct rendering methods apply visualizations directly to the scene geometry without greatly disrupting their appearance or affecting visibility.

Geometry of importance can be set apart from the rest of the scene by drawing a distinct outline around its silhouette, or the geometry can be rendered with a coloured hue (see Figure 4). Both of these techniques can be used to indicate the state of logical data, which can either be on or off. For example, the state of a chair occupancy sensor can be indicated by highlighting a chair at a desk when someone is sitting in it, or whether air is currently being forced from a vent register.

It is also possible to apply more complex surface shading to the scene geometry. Given a scalar field, we apply gradient shading to surfaces which attenuate with the distance and intensity of nearby data points. A sparse 3D scalar field can be visualized on surface

geometry by computing an inverse distance weighted-sum between surface points and all the data points in the field. This sum is then mapped to colour and used to shade the corresponding point. Although both highlighting and gradient surface shading are well-known visualization methods in computer graphics, their introduction into a dynamic 3D scene provides generic visualization methods for representing a variety of inputs. Temperature readings, for example, can be taken from different areas in the office and then mapped to provide an approximate visualization of temperature changes across the office (see Figure 5). Power usage can be represented in the same way giving users the ability to quickly determine areas of high power consumption.
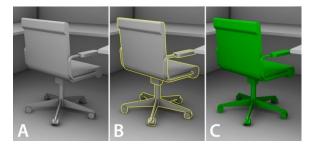

Figure 4: Highlighting important scene geometry. (A) Normal, (B) silhouetted, and (C) shaded hue.

Figure 5: Complex surface shading: warmer areas near walls and windows are differentiated.

Transient Geometry Methods

Transient geometry refers to auxiliary geometry that is not originally present in the scene since it exists only so long as the visualization is presented to the user. The benefit of transient geometry over direct rendering methods is that visualizations are not limited to the surfaces of the existing geometry, thus more complex 3D data can be represented.

The simplest implementation of this group of methods is in the form of glyphs. Glyphs are symbolic representations of single data points, such as an arrow. Arrow glyphs can be used to represent discrete vector data samples making it is possible to visualize complex phenomenon, such as air movement through a space (see Figure 6).

Figure 6: Arrow glyphs can depict air-flow around a cubicle. Larger arrows indicate faster air movement.

Glyphs can also be used as an alternative to direct rendering methods of displaying logical data. For example, point marker glyphs can be used to mark particular points of interest in 3D space. In particular they can be used to mark points whose data originated from a sensor reading. This differentiates them from points whose data values have been interpolated. Similarly, occupancy can also be visualized using glyphs. Using simple motion sensors and a small set of video cameras (Ivanov et al., 2007), building occupancy can be monitored. This data can then be represented using peg-like glyphs, which provide an abstract representation of occupants in a space (Glueck et al., 2009). The combination of base and stem parts ensure that pegs are always visible from a variety of viewing directions, including head-on or from the top. Their abstract nature avoids conveying potentially misleading information about which direction the occupant is facing (this is not the case when representing occupants with human-like billboards). To reduce visual clutter, pegs are aggregated when they are within a certain radius of one another. Aggregated pegs are distinguished from single pegs by using a combination of colouring, text and segmentation of the inner ring (see Figure 7).
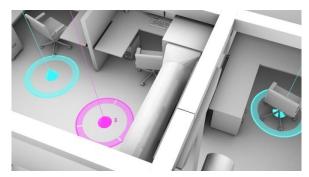

Figure 7: Peg glyphs visualize office occupancy by representing individuals as blue pegs and groups as pink pegs. The number of people in a group is indicated numerically and graphically.

Volume rendering techniques enable high fidelity reconstruction of 3D data fields such as particles or isosurfaces. This technique is ideal for visualization of phenomenon such as smoke or indoor air quality. For instance, we can simulate the diffusion of airborne contaminants using a Computational Fluid Dynamics (CFD) model with boundary conditions that emit into a simulated volume at a rate approximating the diffusion of contaminants. Volumetric rendering when applied to this data enables a comprehensive understanding of the relationship between these airborne particles and adjacent geometry. 3D data can be difficult to inspect visually since it is self occluding (see Figure 8), so user specified cutting planes can be used to isolate specific slices or chunks of the volume. Similarly, to aid in visual analysis, we employ shading to make particular layers of the volume partially or fully transparent (see Figure 1), revealing inner features of the data field. Adding translucency also helps reduce interference with other simultaneous visualizations and helps preserve visibility of scene geometry.
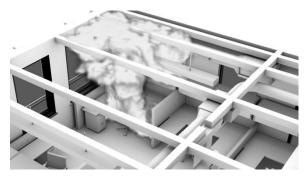

Figure 8: A volumetric rendering of hot buoyant plume.

Heads-Up Displays (HUDs)

Heads-up displays allow us to introduce additional data into the scene that cannot be displayed meaningfully within the original 3D geometry. For example, text and chart information is best viewed orthogonal to the view direction, and if rendered within a scene would require the user to navigate to effectively view it. HUDs are rendered in screen space, but are contextually attached to objects within the 3D scene. They allow presentation of conceptually complex and semantically derived data that would be difficult to display graphically within the scene.

For example, when selecting a power outlet, a HUD appears that not only indicates the current power usage in Watts, but also it presents a bar representation of how much power the outlet is using compared to all other outlets within the office space. This simultaneously provides the user with local and global scope of information, semantically linked to local features (see Figure 9).
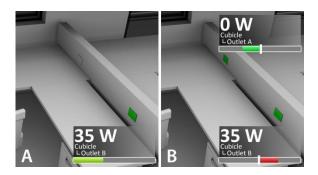

Figure 9: HUDs indicating power usage of outlets. The coloured bar can be used to represent (A) the percentage of overall energy-use of a cubicle, or (B) relative energy-use compared to the average outlet usage.
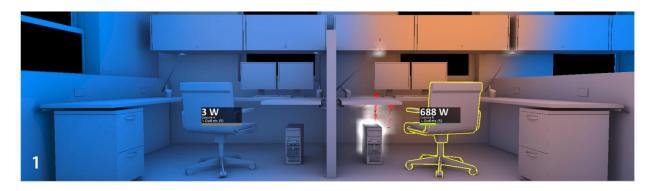
UNIFIED REPRESENTATION SYSTEM IN CONTEXT

The three groups of visualization methodologies presented can be used to display different types of data points. While direct rendering is best to indicate simple logical or scalar values, transient geometry is best suited to display scalar and vector values and heads-up displays are best suited to semantically-derived scalar values (see Table 1).
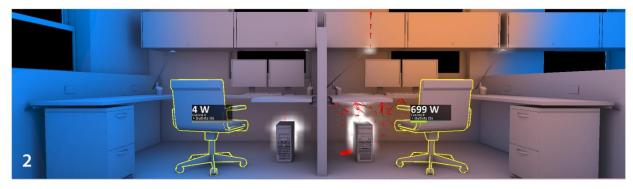
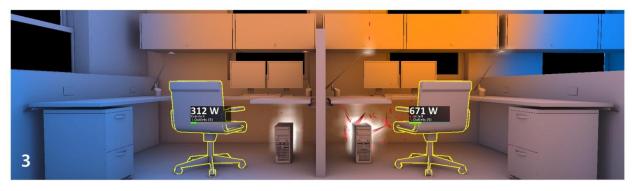|	DATA TYPES			
VISUALIZATION METHODS	Logical	Scalar	Vector	Semantic
Direct Rendering	✓	✓		
Transient Geometry	✓	✓	✓	
Heads-up Display	✓	✓		✓

Table 1: Comparing which data types are best represented by each visualization method.

It is possible to simultaneously display one data set using each of these methods without causing too much confusion, as each resides in a different conceptual space of the model: either on or as part of the geometry, within the 3D space between geometry, or in screen space. Displaying more than one data set using techniques from a single group is potentially confusing, as they would occupy similar space. However, some techniques, such as volume rendering, can benefit from the simultaneous display of glyphs.

As a demonstration, we combine the methods outlined in the previous sections to visualize the performance-data of two adjacent cubicles where usage of personal electrical equipments such as a computer tower and task lights are used as an input for both sensor-based data and simulation visualization (see Figure 10).
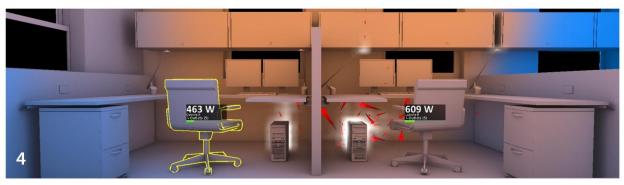
Figure 10: These four sequential frames visualize two adjacent cubicles in the office. The cubicle on the right is already in use, when a co-worker arrives and begins to work at the cubicle on the left. Several visualization methods are simultaneously presented: cubicle occupancy is depicted by highlighted chairs, the heat emitted by the computers and lamps is shown by glyphs and volume rendering, and total energy use of each cubicle is indicated as a HUD. In addition, surface shading on the model represents a combined metric describing a general level of activity across the cubicles based on occupancy, temperature and energy usage.

In this scenario, energy usage of each cubicle is monitored in real-time and is presented as HUDs, which includes semantic data pertaining to the location of each cubicle. In addition to sensor-based data, a CFD simulation model of the cubicle takes heat loads from the various devices into consideration. Radiation from the heat sources is not considered and only convective properties are modeled. A stable Navier-Stokes solver (Stam, 1999) using a 3D regular cubic grid is used to perform the calculations. Only transients are calculated to promote a visually appealing flow-field. Reaching a steady state, although possible, is unnecessary for the purposes of this illustration. All modeling is performed in a separate application which caches the results for each sampled time step. The cache is then interactively read into our system, processed and rendered interactively to visualize different results. Velocity, density and temperature are all computed over multiple solver iterations and sampled at regular intervals to produce visualization data frames. In this example, all values are non-dimensional and are not necessarily physically accurate, but from the visualization standpoint we can observe how all surfaces interact and collide with existing physical elements. Furthermore, performance visualization is represented in relation to cubicle occupancy using direct rendering method of highlighting. Ultimately, surface shading provides an additional metric describing a general distribution level of all activities across the two cubicles.

DISCUSSION AND FUTURE WORK

In recent years we have seen a number of research precedents in presenting a systematic approach toward simulation-based reproduction of performance data from existing buildings (Mahdavi et al., 2007; Wasilowski A. and Reinhart F., 2009). Our current research also stems from our interest to utilize our own office building as a living laboratory to gain insight on how to improve the energy performance of an existing building. Our research aims to capitalize on the insights we can derive from an existing condition where experimental simulation and measured data are mapped onto the same environment for analysis. Within this context, the development of a rich visual environment that can support representation of both measured and simulated datasets is a key step, since traditionally these two methods belong to separate stages of a building's life cycle. Simulation often serves an important role in the design analysis stage whereas energy monitoring systems have been a domain of building maintenance. The integration of these methods through a unified-representation system will foster higher degrees of integration in data representation, which is particularly important for the ongoing analysis of a living building.

This paper does not present its visualization framework as part of the design process, but we believe that a simulation-driven environment can also benefit from effective visualization techniques of performance-related data in-context within a 3D scene. Numerical output, prevalent in standard visualization systems, are crucial for accurate analysis, however research has shown that graphic representation can actually have much richer psychological impact to grasp information (White and Feiner, 2009). This is essentially to assist users with interpreting the underlying message in its appropriate context.

To further enhance our framework as a hybrid visualization system, we have identified number of key issues to address in our future work.

Temporal Visualization and Information Clustering

As we continue to expand our current implementation of a sensor-based monitoring system to collect multiple kinds of data, there are two important visualization factors that we plan to address in our next implementation phase. When dealing with real-time data collection for visual analytics, one challenge concerns the spatiotemporal dimension of time-varying data. Our visualization system provides a method for collecting real-time building performance data representation, however, it is crucial to present this data for different time intervals. Our system should be able to go back and forth in time and view the visualization related to that time interval. Subsequently, these data should be visually available and linked to simulation results to validate and calibrate the output against the real measurements.

Secondly, it is important to be able to visualize current building performance data in comparison to previously recorded data by aggregating across different time steps such as weeks, months or years. This would provide the user with richer diagnostic data, allowing them to better evaluate the ongoing performance of the building.

Building Semantics and interoperability

Synchronization between a modeling program and a free-standing visualization environment still remains a challenge. As stated above, we had to manually incorporate building semantics and intelligence into various geometrical components such as floors, cubicles or power outlets. Ideally, having authored all the components in a BIM environment, we should be able to transport the geometry and its associated semantics easily into our system via the Industry Foundation Class (IFC) file format.

In working toward further integration of additional building performance data in an existing building, there

are important technical and usability challenges. Additional research is required to allow simpler accessibility to visualization of more complex interaction of inter-related factors. The growing demand for environmentally-friendly buildings can certainly benefit from further research into such visualization techniques where building performance should be comprehensively evaluated on an ongoing basis.

CONCLUSION

This paper is described as an interim step in our current research that combines some existing visualization methods within one comprehensive framework. These methods essentially provide a flexible visualization framework for interpreting various kinds of data input within a unified representation system. An important aspect of our strategy is how our system decouples data from visualizations. The software architecture we present is designed to accommodate arbitrary visualizations since the data in this system can be originated from an arbitrary source. We have presented these visualization techniques in the context of an existing building where building performance data is defined as a hybrid of measured and simulated data. While many aspects of building performance in an existing building can be metered through a sensor-network, there are certain factors that cannot be easily measured using this method. For instance, we can measure the temperature of various spots locally, but we cannot necessarily evaluate the behavioural aspects of temperature and air flow using a sensor network. Instead, it is prudent to take advantage of simulation as a powerful method to generate plausible data based on the existing condition.

Another important aspect of our approach is to address the common complexity inherent in a massive dataset generated either through simulation or real-time sensor-based monitoring systems. Although these methods are both well positioned as important research issues, the visualization of their results is often presented as 2D graphs, spreadsheets or numerical outputs, which is typically difficult for users to interpret. Data dimensionality is an important aspect of visualization and by integrating both spatial and non-spatial building performance data into a 3D environment, we can introduce new possibilities in how we understand and analyse a building through its entire life cycle. By mapping the data to local features that correspond to their real configuration in an existing building, we can reduce the complexity of analysis by balancing between "complexity and novelty in images" (Ormerod and Aouad, 1997).

REFERENCES

Attar R., Prabhu V., Glueck M., and Khan A., (2010), '*210 King Street: A Dataset for Integrated Performance Assessment*', Symposium on Simulation for Architecture and Urban Design (In Press).

Essa, I., (2000), '*Ubiquitous Sensing for Smart and Aware Environments*', IEEE Personal Communication, 47-49.

Gershon N. and Eick S.G., (1995), '*Visualization's New Tack: Making Sense of Information*', IEEE Spectrum 32, 11, 38-40.

Glueck M., Crane K., Anderson S., Andres R., and Azam K., (2009), '*Multiscale 3D Reference Visualization*', I3D Conference Proceedings, 255-232.

Ivanov Y.A., Wren C.R., Sorokin A., and Kaur I., (2007), '*Visualizing the History of Living Spaces*', IEEE Transactions on Visualization and Computer Graphics, Vol. 13, Issue 6, 1153-1160.

Mahdavi A., Orehounig K., Mikats N., Lambeva L., and EI-Habashi A., (2007), '*Analysing Traditional Building via Empirically Calibrated Building Performance Models*', Building Simulation, 71-78.

Ormerod M., and Aouad G., (1997), '*The Need for Matching Visualization Techniques to Client Understanding in the UK Construction Industry*', IEEE Conference on Information Visualization, IV, 322-328.

Prazeres L. and Clarke J.A. (2005), '*Qualitative Analysis on the Usefulness of Perceptualization techniques in Communicating Building Simulation Outputs*', Building Simulation, 961-968.

Rivotti V. Proenca J., and Sousa M.C., (2007), '*Composition Principles for Quality Depiction and Aesthetics*', Computational Aesthetics in Graphics, Visualization and Imaging, 37-44.

Stam J., (1999), '*Stable Fluids*', SIGGRAPH Conference Proceedings, Annual Conference Series, 121-128.

Wasilowski H.A. and Reinhart C.F, (2009), '*Modelling an Existing Building in DesignBuilder/Energyplus: Custom versus Default Inputs*', Building Simulation, Glasgow, Scotland, 1252-1259.

White S., and Feiner S, (2009), '*SiteLens: Situated Visualization Techniques for Urban Site Visits*', Human Factors in Computing Systems, 1117-1120.