Street Light View: Enriching Navigable Panoramic Street View Maps with Informative Illumination Thumbnails

Charles Rojo, Wei Xu, and Klaus Mueller

Center for Visual Computing, Computer Science Department, Stony Brook University

ABSTRACT

Google Street View is a technology featured in Google Maps and Google Earth that provides panoramic and immersive views of street scenes in many cities around the world. The imagery is taken from a fleet of specially adapted cars, and advanced computer vision algorithms produce a visual environment that users can interactively navigate in a nearly seamless fashion. Currently, the additional information displayed is confined to address labels, ads, and tourist information. We propose to augment these location-mapped environments with modern data visualization displays, to communicate additional information of interest. To illustrate our framework, we describe the mapping of illumination statistics information onto Google Street View building facades, to aid users in the selection of real estate.

KEYWORDS: panoramic maps, information fusion, illumination

1 INTRODUCTION

To those looking for an understanding of what life in a particular location is like, the utility of Google Street View is undeniable. Anyone can simply only open a web browser, navigate to Google, and enter an address. They are immediately dropped into a visual still-life of the landscape, Street View. This landscape is an interactive environment where one can modify the camera's orientation into new perspectives and change location on the fly. This interactivity provides immersion. It leaves the viewer with a degree of perception of the area so that they no longer need to physically visit it to feel they've understood it. This visual glimpse is useful, but it only begins to scratch the surface at presenting the rich body of data that has been collected for real-estate in the world.

Embedding these data (and the information derived from them) into this large-scale immersive imagery of the real world is essentially an information fusion task. While efforts exist that embed ads and review snippets into these displays, using the available location information as a placement index into the imagery, our work significantly extends these application by the type of information it provides, and its presentation as well.

We shall illustrate our framework by a concrete example. Assume one is in the market for an apartment in New York City. There are a number of online databases available to help in this quest. A popular online database is StreetEasy.com which periodically crawls all real estate broker websites and consolidates the information into a single browsable list, searchable by location, price range, size, building history, and so on. Typically, pictures of the apartment's interior and sometimes the building exterior are also available. However, additional crucial pieces of information are also neighborhood, sun exposure, window views, and the like. Links to Google Street View are usually provided that allow users to gain some insight into these variables. However, Google Street View is necessarily only a snapshot of time. For example, in crowded urban environments one typically wants to gain information about daily and seasonable exterior lighting (illumination) patterns to determine the brightness (and heat) of the apartment. Currently, this information can only be obtained by educated guesses, from the real-estate broker (which assumes a certain level of trust), or by regularly observing the apartment over an extended amount of time. The latter is obviously impractical given the fast pace of the real estate market and the associated time and access overhead. Similar is also true to assess environmental noise, access to taxis, and a myriad of other variable one may have taken for granted at the current living location. We have therefore set out to first compute and then provide these types of information items using Google Street View to enable a direct mental mapping of this information to 3D location. Further, doing so also allows easy comparisons of different apartments along the street or inside the building.

2 COLLECTING THE DATA

We had access to a large GIS building database of Manhattan, but our algorithm only requires a rather coarse level of detail (on the order of floor heights or even less) which can be obtained quite easily from Google Earth. We compute two types of illumination per facet: (1) direct sun exposure taking into account seasonal effects, and (2) indirect light exposure as reflected from nearby (light-colored) building facades. We simulate these contributions on an hourly basis for all days of a year and use publicly available satellite data from NASA [1] to obtain information on sun orbits and cloud covers to determine the incident illumination.

Let us assume a local neighbourhood of buildings. We allocate a GPU floating point texture for each building face to hold the sunlight accumulation data. For each hour, a two-pass rendering of the scene of building geometry is performed from the position of the sun. Orthogonal projection is sufficient since the sun is far away. The first pass removes hidden surfaces, and the second pass renders each face individually, with the illumination accumulating on the texture buffers. This direct light accumulation is further refined by reducing the accumulation value at each hour based on cloud-coverage. NASA's GEOS satellite [1] offers world-located imagery of cloud coverage over time for each day which we can use as an environment map. An occlusion factor derived from the intensity value of the clouds in the area can be used to dampen the effect of the direct lighting for that date and time. The portions of the light texture mapping to windows then adds to the sun exposure (or *sunshine factor*) of the associated apartment. Additionally, higher texture intensity also reflects more direct sunlight illumination, which we use for the indirect illumination stage described next.

Indirect lighting significantly adds to the sun exposure of a given apartment. It is very important in particular for north-facing apartments. We use the following radiosity-like scheme to compute these contributions from the direct illumination textures. For each window, a perspective rendering is performed. The near-plane frustum of the perspective rendering is made to be oriented and skewed to match the window's relevant immediate facade neighbourhood. The field of view is set to be 180 degrees outward from the surface chosen. The scenery around the area is textured with the exact floating-point textures derived from the direct-lighting procedure. The result of the rendering is a new perspective rendering of floating point values that corresponds to the lighting accumulation visible from the area of interest. This correlates with the direct illumination that the selected spot of interest will "see," or be exposed to, over time. This rendering produces an "indirect lighting" texture that is summed into a single “indirect lighting contribution” factor that captures the overall amount of illumination from reflected direct lighting that the window will receive. While more sophisticated indirect illumination techniques [2] could have been implemented, we aimed for the fast computation described here. This approximation gives users an interactive means to compare the impact of indirect lighting from set of nearby constructs.

3 Calculating the Information

Temporal variations in lighting, such as those caused by occlusion from high-rises, are quite common in crowded cityscapes and heavily depend on location. An apartment may lose the bulk of its expected lighting due to a poor location. Moreover, these effects can be seasonal due to the shift of the sun. Fig. 1 illustrates this point. It depicts a typical living situation found in lower Manhattan. A high-rise is situated across a low-rise building with residential apartments. At approximately 2pm in January, sunlight reaches the opposing apartments to the left of the high-rise, while the apartments directly facing it and those to its right have their light occluded at this time. Another, more distant high-rise that will affect lighting in the evening can also be seen in Fig 1b. These complex effects are hard to gauge in a 1-hour showing, and they often strike property buyers unexpectedly once they move in.

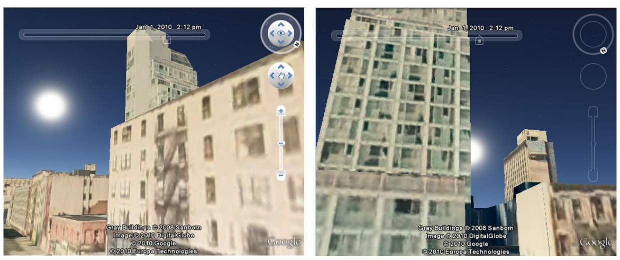

Fig. 1: Lighting in an urban scene at different times of the day.

Our Street *Light* View application offers an understanding of these effects and so makes purchasing (and marketing) efforts more informed. The user first uses Google Street View or Google Earth to navigate to a location of interest and sees a 3D virtual representation of the area. He then uses our tools to define the apartment of interest on the 3D building facade. Next, we specify a half-sphere around this facade to define the buildings participating in the illumination calculations. These setup operations are followed by the process described in Section 2.

4 Displaying the Information

Fig. 2 shows a situation in which the user has selected a date range of March to November, 8am to noon, within Google Earth for a selected window and a set of buildings that will impact the direct and indirect sun light exposure. Dimmer areas on the actual wall faces represent areas that get less direct illumination. For the selected window, the illumination effects over time are computed and displayed in an *illumination thumbnail* that captures the lighting conditions for the given time range on that window. The horizontal axis represents the days, and the vertical axis represents the hours of a given day. Dimmer parts of the thumbnail indicate less light relative to the other times represented by different locations in the thumbnail. We can clearly see the seasonal change in lighting – the summer months enjoy extended periods of sunshine and the apartments appears to be well lit all the time (note that the thumbnails use intensity windowing to bring out variations at higher illumination levels – therefore the brightness levels in the morning appear darker than expected).

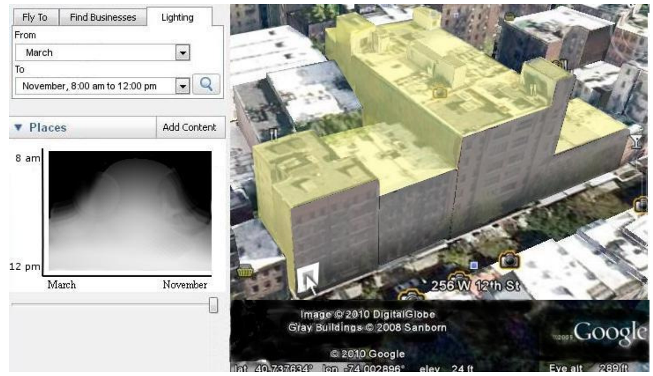

Fig. 2: Street View with thumbnail. The apartment is located on the lower left (white box on facade indicated with arrow).

Fig. 3 compares the lighting in two different apartments up for sale in the block shown in Fig.1. The user has selected two different areas across the street from the blue high-rise. The difference in lighting is apparent on the thumbnail, showing that the first apartment area suffers from light occlusion at around noon, and again toward the evening. The second apartment area is occluded only toward the evening, nearing 3pm. The buyer’s daily routine and preferences will decide which apartment to choose.

5 Future Work

Our current application can only show the thumbnails arranged into a side panel or pasted into the scene. We are currently working on mapping it to the corresponding apartment window directly or onto an apartment’s facade (as an average of all window thumbnails). Further, the thumbnails we compute could also be displayed on the real estate web page as an informative information icon along with other data encoded visually, such as the average neighbourhood income, children per family, etc. Finally, a collection of these thumbnails (of different apartments) could be shown side by side, allowing apartment shoppers to compare the lighting details for a set of apartment candidates.

References

[1] http://weather.msfc.nasa.gov/GOES/

[2] "High-Quality Global Illumination Rendering Using Rasterization”, GPU Gems Volume 2 Chapter 38.

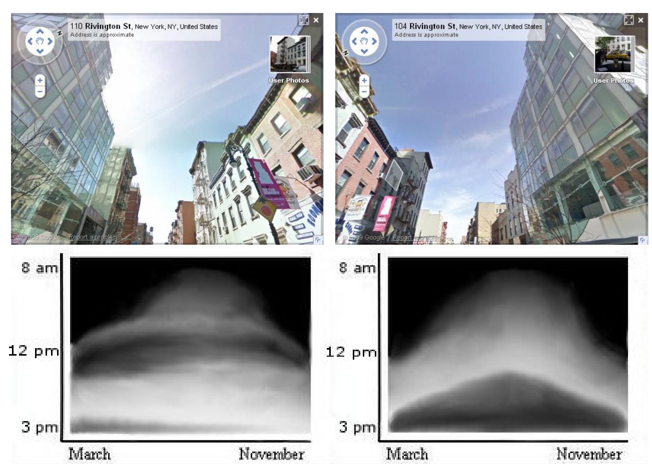

Fig. 3: Illumination of two apartments (above) compared.