

Realtime Estimation of Illumination Direction for Augmented Reality on Mobile Devices

Ibrahim Arief, Simon McCallum, Jon Yngve Hardeberg
Gjøvik University College, Gjøvik, Norway

Abstract

Augmented reality simulations aims to provide realistic blending between real world and virtual objects. One of the important factors for realistic augmented reality is correct illumination simulation. Mobile augmented reality systems is one of the best options for introducing augmented reality to the mass market due to its low production cost and ubiquitousness. In mobile augmented reality systems, the ability to correctly simulate in realtime the illumination direction that matches the illumination direction of the real world is limited. Developing a mobile augmented reality systems with the ability to estimate illumination direction presents a challenge due to low computation power and dynamically changing environment. In this paper, we described a new method that we have developed for realtime illumination direction estimation for mobile augmented reality systems, using analysis of shadow produced by a reference object that doubles as a 3D augmented reality marker. The implementation of the method could estimate the direction of a single strong light source in a controlled environment with a very good degree of accuracy, with angular error averaging lower than 0.038 radians. The current implementation achieved 2.1 FPS performance in a low-end Android mobile device, produced proper estimation within 15 seconds using a uniform surface, and demonstrated scalability potential.

Background

An augmented reality system aims to add simulated virtual objects, for example computer-generated 3D objects, to a real world scene, and presents the combined scene to the user. This combination of the real world and virtual objects augments the user's perception of the real world, and offer a new method of interacting with his or her environment. Augmented reality is used in multiple fields, including medical, manufacturing, entertainment and military [1].

Taken in a larger context, Milgram *et al.* [2] positioned augmented reality as an intermediate between a real world environment and a virtual reality environment. They also define real world environment and virtual reality environment as polar opposites in what they called a reality-virtuality continuum.

Mobile Augmented Reality

One important aspect of an augmented reality system is portability [1]. An augmented reality system with high portability would give the user greater degree of liberty during his or her interaction with the augmented environment. Based on the positioning of its display, two major categories of augmented reality system that have relatively high degrees of portability are head-worn augmented reality and hand-held augmented reality [3]. From

these two categories, the hand-held augmented reality systems are currently considered to be the best for introducing augmented reality to the mass market due to its low production cost and ease of use [3]. The ubiquitous proliferation of smart mobile devices that could be used as hand-held augmented reality systems also helps the ease of adoption by potential users.

Registration and Tracking

To properly align the simulated virtual object with the real world scene and produce a consistent simulation, an augmented reality system needs to be able to detect and track its position and orientation with respect to the real world. This is typically performed using markers with known shapes or textures [4]. Augmented reality markers are typically two-dimensional markers, but some libraries provide the ability to use three-dimensional markers. Recent development using feature extraction and recognition provides a new markerless approach for augmented reality tracking, although this approach is still under development [5].

Reproduction Fidelity

Tracking and alignment is required to provide overlay information. However, if the objective is to make the virtual object appear integrated into the real world scene, we must embed the virtual objects in the real world [2]. To improve the user experience the system must provide more realistic blending between the virtual objects and the real world. The degree of this blending could be represented as reproduction fidelity [2]. Low reproduction fidelity is exemplified by having a simple wireframe simulation of the virtual objects, while higher reproduction fidelity takes into account more complex factors for the simulation such as shading, texture, transparency, and illuminations, with the ultimate target of producing simulations with photorealistic fidelity.

Current state of the reproduction fidelity for mobile augmented reality systems is limited to texture, shading, and transparency simulation of the rendered virtual objects. The ability to properly simulate the illumination condition of the virtual objects to match the illumination condition of the real world is still limited, mainly due to the inability to estimate the illumination conditions of the real world environment. Without this estimate, simulating the illumination condition in an augmented reality systems using abstract parameters would risk producing less realistic simulations.

In order to address this issue, several methods are proposed and developed to estimate the illumination conditions and incorporate the estimate in an augmented reality system [6, 7, 8, 9, 10, 11, 12, 13, 14]. However, none of the explored methods is applied to mobile augmented reality systems, and most of them requires off-line processing that precludes real time application

of the illumination estimation on mobile devices. In this paper, we developed a illumination estimation method that is suitable for simulating real time illumination on mobile augmented reality systems.

Illumination in Augmented Reality

The first issue in understanding illumination in augmented reality is understanding the types and properties of illumination sources that could be simulated in an augmented reality system. Since the simulated part of an augmented reality system is essentially a virtual reality simulation, types of illumination sources that could be simulated in an augmented reality follow the types normally found in a virtual reality. Strauss [15] defines three basic types of illumination sources in a virtual reality simulation:

1. Point illumination source. The illumination originates from a single point in space (e.g. a light bulb).
2. Directional illumination source. The illumination rays run parallel from a certain direction (e.g. the sun). Could be considered similar to a point illumination source, but positioned at a large or infinite distance.
3. Ambient illumination source. The illumination is spread diffusely and evenly (e.g. illumination from a lamp equipped with lampshades).

Strauss [15] states that direction and color of the illumination source are two important properties necessary to simulate proper illumination effect. Moreover, experiment by Slater, Usoh, and Chrysanthou in [16] shows that the existence of shadows could assist the spatial perception of the user. We therefore focused on detecting directional and point illumination sources in order to be able to simulate realistic shadows for a mobile augmented reality system. To achieve this, we require real-time estimation of the illumination condition of the environment surrounding the mobile device.

Methods for Estimating Real World Illumination Conditions

We explored three different methods that could be used to estimate the real world illumination condition. The first method is using a light probe. A light probe is a spherical object coated with reflective materials that could be used to capture the surrounding illumination of an environment [17, 18]. Since the intensity of illumination sources in an environment normally spans a large dynamic range compared to the environment itself, it is common to use HDR imaging to capture the light probe images. The light probe method has been applied to augmented reality simulations to produce photorealistic virtual objects [6, 7, 8, 9, 10].

In general, the hardware setup of augmented reality simulations developed using this method is using two separate cameras [6, 8], one camera dedicated for viewing the light probe and another camera dedicated to viewing and tracking the AR marker. A one-camera setup is possible [9, 6], but necessitate the light probe to be constantly viewable in the same scene with the AR marker. An additional problem exists since the position of the light probe is dynamically changing with respect to the camera, thus introducing the need to detect and track the position of the light probe. Mounting the probe in a fixed platform is one possible solution for this problem [9].

The second method is using a fish-eye lens to directly detect and estimate the illumination sources surrounding the environment [19]. HDR imaging could be done by varying the shutter speed and aperture size of the imaging device. However, since capturing the environment in several different dynamic range requires the time of several seconds, this means that this method is not suitable for real time applications. The fish-eye lens method has been applied to augmented reality simulations [6, 20].

The third method is using analysis of shadows generated by reference objects with known geometry to estimate the direction of the illumination sources [21, 22] and, in an outdoor environment, the strength of sky irradiance [11]. Panagopoulos *et al.* use von Mises-Fisher distribution [21] and Markov Random Field [22] to model the direction of the illumination sources from the shadows detected in the input image. Their method and the resulting illumination models was implemented to simulate shadows of virtual objects for augmented reality, but they noted that their algorithm requires three to five minutes of processing time per image, making it unsuitable for real-time augmented reality. Madsen and Nielsen [11] use contrast obtained by comparing the brightness of the shadowed and unshadowed region of an outdoor image to model the radiance of the sky and sun.

Among the three illumination estimation methods explored above, we consider the shadow analysis method to be the one most suitable for the purpose of real-time illumination estimation on mobile augmented reality systems. The methods that use light probe and fish-eye cameras requires computations that are too costly to be processed in real time using the computational power available in current mobile devices. Those methods also require a specialized hardware setup. Although the shadow analysis method could not be used to produce illumination simulation with photorealistic quality, it is sufficient to be used to estimate the direction of the illumination source. However, the current usage of the shadow analysis methods found in [21, 22, 11] is not optimized for mobile devices, and requires the development of a new method that strikes a balance between good estimation accuracy and lightweight computational requirement. In the next section, we describe the new method that we developed in order to answer this requirement.

Proposed Method for Real Time Estimation of Illumination Direction

The key point of the shadow analysis method that we developed is based on two observations. The first observation is that shadow analyses are typically performed by comparing the shape of the projected shadow with the geometry of the object that generates the shadow. By having the geometry of the shadow-generating object known in advance, the computation required in order to perform the comparison could be greatly reduced. Moreover, another reduction in computational complexity could be achieved by adding a constraint that the shadows are projected to planar surfaces.

The second observation is that, rather than the traditional 2D augmented reality marker, we could merge several 2D markers within a rigidly constrained geometry to form a 3D augmented reality marker. For example, a cube where all its six surfaces are each textured with different patterns could act as a single augmented reality marker, where each surface acts as an independent marker by itself. The accuracy of marker registrations in this case

could also be increased, since we now have six independent markers that could be registered from different viewing angles.

Combining the two observations, we reasoned that a 3D augmented reality marker could also be used at the same time as a reference object for shadow analyses. This reasoning provided several benefits. Since the augmented reality marker would always be within the view of the mobile augmented reality system, it would mean the reference object would also be simultaneously viewed, and we eliminate the need of having two separate objects that need to be constantly within the view for the method to work. Another benefit is automatic geometrical recognition of the reference object, reducing the computation required during performing the comparison between the projected shadow and the geometrical shape of the reference object. Moreover, from the implementation viewpoint, augmented reality SDK such as Vuforia™ from Qualcomm already supports the usage of three-dimensional markers with simple geometrical shapes.

From the observations and reasonings above, we choose to use a cuboid 3D augmented reality marker as a reference object for our shadow analysis method, due to its simple geometrical shape. The basic outline of the shadow analysis method we developed could be divided into three steps. The first step is shadow contour extraction, in which we extract a collection of points that lies in the boundary between shadowed and unshadowed region of the image. From this collection of points, we are able to reconstruct the basic shape of the shadow. An example of this step is shown in Figure 1, where the left side of the figure shows the example of shadow generated by our cuboid reference object and the right side of the figure shows the result of the shadow contour detection and tracing.




Figure 1. Example of the shadow contour extraction step.

The second step in our method is analyzing the shape of the shadow contour extracted from the first step. The goal is to recognize points in the shadow contour that corresponds to points in the geometrical shape of the reference object that generates the shadow in that particular location. In our cuboid reference object, the sharp corners of the shadow contour corresponds to the upper corners of the cuboid. Using this method, we could then obtain several pairs of points between the shadow contour point and geometrical object point. Example of this pairing could be seen in Figure 2, where each corners in the shadow contour are color coded with its corresponding corners in the cuboid reference object.

The last step in our method, after we obtain pairs of corresponding points from the previous step, is forming a set of lines from these pairs of points. We then project these lines and find their intersection. This intersection of lines is interpreted as the estimated position of the illumination source. Example of forming the set of lines from obtained pairs of points could be seen in the left side of Figure 3, and the projection of these lines could be seen in the right side of the Figure. It could be seen that the pro-

jected lines converges at roughly the position of the illumination source used in the example.

In order to have a real time shadow estimation, we perform several levels of optimizations from the steps outlined above. We took advantage of the existing augmented reality registration and tracking system to extract the position of the corners in the image captured by the mobile augmented reality system. Since shadows generated by our reference object always originate from one of the lower corners, we initialize our search space on a small radius surrounding this corner. This drastically reduces the number of computations required to successfully trace the contour of the shadow, as outlined in step one above. We initially tried to adapt the canny edge detector using first order gradient, but we found out that the computational load, multiplied by the number of pixels that need to be processed, is too large to be adequately processed in real-time using mobile CPUs. Therefore, we use a simple high-variance search in determining potential edge candidates, since edges usually have a high variance between the neighboring pixels. We then performed a post-processing step using cluster-based and vector-based analysis in order to reduce false positives generated by noises. An example result of this step could be seen in Figure 4, where the green (left) edge is the recognized shadow edge after our post-processing.

Our second optimization is done during the tracing of the shadow contours. Since reading pixel data is a costly operation, we aim to minimize the number of pixels that need to be read during shadow tracing. We performed the trace using jumps along the current best estimate of the shadow vector. If the jumps failed to land in a shadow edge candidate, we halved the jump distance and re-sampled the pixels. If we succeeded in landing on a shadow




Figure 2. Example of finding corresponding points.




Figure 3. Example of estimating illumination source by correlated line projection.




Figure 4. Detected shadow edge candidate.

edge, we double the next jump distance. Using this approach, we are able to trace most shadow contours using less than 1500 pixel readings, a huge savings compared to the hundreds of thousands of pixels that need to be searched. The result of this binary jumps could be seen on Figure 5.




Figure 5. Result of binary jump shadow contouring. The white crosses shows the sampled pixels.

After obtaining the shadow contour, the second and third step are relatively straightforward. We detect sharp corners from the shadow contours and pair them with upper corners of the cuboid. Since the detected shadow corners is in two-dimensional screen coordinates, and the corners of the cuboid is in three-dimensional world coordinates, we need to perform a single ray tracing from each of the shadow corner to the lower plane of the cuboid in order to obtain the three-dimensional world coordinates of the shadow corners. We then form a line projection from the shadow corner through the cuboid corner. From this, we calculate the intersection between the lines and use it as the estimated position of the light source. We use the midpoint of the closest point between two lines in 3D space as the intersection point. In case we have three projected lines, we calculate the intersection point between the three pairings and use the mean of the coordinate values as the intersection point.

Results

In order to simplify the calibration of the system for testing purposes, we put the reference object and the illumination source in a static position in our testing area. In this way, the position of the illumination source in respect to the marker is fixed and could be used for our ground truth calculation. We put the illumination source, the reference object, and the mobile device in the arrangement shown in Figure 6.

The figure was overlaid by axes used in our coordinate system, which is centered in the reference object as its origin point. The green, blue, and red axes is the X, Y, and Z axes, respectively. The solid lines shows the positive value of the axes, while the dotted lines show the negative values. Using this coordinate system, we then measure the position of the illumination source relative to the reference object. The measured position of the illumination source is $(-54.3, 66.2, 54.3)$, all units in centimeters. This value becomes our ground truth during testing the accuracy of the system. For the illumination source, we use a 60 Watt clear incandescent lamp with luminous power of 710 lumen. We implemented our method in an Android mobile device with a single-core 1 GHz CPU.

Accuracy

Our tests showed that, within our controlled environment, for a successful tracing of the shadow, the directional accuracy for the estimated light source position using our method is very good, with angular error averaging lower than 0.038 radians. The method is not so good in estimating the distance to the light source, with the average estimated distance of 162.83 centimeters is 61% larger than the ground truth distance of 101.39 centimeters. This is because small errors on angular estimation could translate into a large shift in intersection points between the projected shadow lines. Generally, accurate estimation of the direction of the light source could already provide good illumination simulation in augmented reality systems.

Demonstration

After obtaining the estimated illumination source position, we perform a simple demonstration in order to show how our system could enhance the immersion of existing augmented reality systems. We created an augmented reality scene as shown in Figure 7. The augmented object is simulated to be located in front of the reference object, but a lack of visual cues made it harder for the depth to be perceived. The lighting for the scene came from the upper left side of the image.

Using our illumination estimation system, we are able to estimate the position of the illumination source, and use this esti-




Figure 6. Arrangement of components in the experiment setup.




Figure 7. AR scene without illumination estimation.

mation to project a virtual shadow for the augmented object. Figure 8 shows the projection from the estimated illumination source, through the corners of the augmented object, and into the planar surface, marked as circled dots on the surface.




Figure 8. AR scene with shadow projections from estimated illumination source. For clarity, we added an outline to the real world shadow of the reference object.

From the above projection, a simple rendering of a virtual shadow for the augmented object could be done as shown in Figure 9. Although the shape of the virtual shadow is not an exact match with the real world shadow, the direction of the virtual shadow closely matches the real world shadow generated by the reference object. Compared to the original image in Figure 7, the addition of virtual shadows provides visual cues that ease the perception of depth for the augmented object.




Figure 9. AR scene with rendered shadow projections.

Performance

On average, the implementation of our method is able to perform the entire estimation steps under 400 ms in a 1GHz single-core CPU mobile device. In practice, not every frame managed to be properly analyzed and estimated, and successful estimation was done only 12.7% of the time. Most of the failures are caused by the failure in the shadow contour tracing step. Combining the estimation steps and the success rate means that in most cases, a proper estimation of the illumination could be achieved within 15 seconds. Since a typical illumination source does not move considerably in the real world, once a proper estimation managed to be done, no further re-estimation should be needed. This partially

mitigates the low success rate of our implementation. In the worst case that each frame need to do full estimations, our implementation achieves an average of 2.1 FPS. We also identify the shadow contour detection step to be a major performance bottleneck, taking up to 98% of the processing time.

Weaknesses and Constraints

There are several weaknesses in the current implementation of the method. The most prominent one is the need for a sharply contrasted shadow and a uniform surface below the reference object in order to ease the shadow detection and reduce the computational complexity. Another drawback is that the method fails to do proper estimation if the light source lies behind the camera, making the shadow to be fully occluded by the reference object. Also, the current implementation only take into account a single dominant illumination source. In addition, for illumination sources positioned close to the reference object, large variance in distance estimation could considerably affect the length of the virtual shadows. In an outdoor scene, where the sun as the primary illumination source is positioned at a large distance, this effect is less noticeable.

Conclusions and Perspectives

The new method that we developed is able to estimate the direction of illumination in a mobile augmented reality system, despite the challenge of much weaker processing power available for mobile devices compared to desktop computing hardware. The first implementation of the system is able to perform the estimation in under 400 ms, achieving real time performance, but due to the low success rate of the shadow contour tracing, the system took about 15 seconds to properly detect and estimate the direction of illumination under controlled condition.

There are several advantages of the method that we developed compared to existing methods for estimating illumination direction for augmented reality scene. This method does not require specialized hardware setup such as light probes or two-camera setup, making the method suitable for mobile augmented reality systems. Moreover, this method is computationally efficient and implementable in low-end mobile devices, opening the prospect of mass adoption of the method.

Further works could be done in order to reduce the processing time of the method. Preliminary experiments shows that deploying the implementation in a more powerful mobile device, with a quad-core 1.6 GHz CPU compared to the single-core 1 GHz CPU used in most of our experiments, reduced the processing time by a factor of two, demonstrating the scalability of the method. Further optimizations, such as moving part of the steps to the GPU, or using mobile-optimized image processing libraries such as FastCV, could further reduce the processing time of the implementation.

A faster processing time would also mean that we could use a more complex shadow detection and tracing method that provides better accuracy compared to the present one. Since the shadow tracing step is the step with highest failure rates, increasing the accuracy of this step could increase the overall performance dramatically. The geometrical shape of the reference object could also be more complex, opening the possibility of having the markers as tangible objects that the user could directly interact with. Using a more complex reference object also increases the number of

points that could be compared between the object and the shadow contours, potentially increasing the accuracy of the distance estimation between the object and the illumination sources.

In addition, having a larger processing budget would also enable the color of the illumination source to be estimated in real-time. One possible approach for color estimation aligned with this method is to embed color patches to the surface textures of the reference object, and estimate the color of the illumination from the imaged reflectance values.

Acknowledgments

The work and research done in this paper is performed as a part of the Erasmus Mundus Master programme CIMET (Color in Informatics and Media Technology, <http://www.master-erasmusmundus-color.eu>).

References

- [1] R. T. Azuma, "A survey of augmented reality," *Presence: Teleoperators and Virtual Environments*, vol. 6, pp. 355–385, Aug. 1997.
- [2] P. Milgram, H. Takemura, A. Utsumi, and F. Kishino, "Augmented Reality: A Class of Displays on the Reality-Virtuality Continuum," in *Proceedings of the SPIE Conference on Telemicroscopy and Telepresence Technologies*, vol. 2351 of *Proceedings of SPIE*, (Boston, Massachusetts, USA), pp. 282–292, Nov. 1995.
- [3] D. W. F. van Krevelen and R. Poelman, "A Survey of Augmented Reality Technologies, Applications and Limitations," *The International Journal of Virtual Reality*, vol. 9, pp. 1–20, jun 2010.
- [4] X. Zhang, S. Fronz, and N. Navab, "Visual marker detection and decoding in ar systems: a comparative study," in *Mixed and Augmented Reality, 2002. ISMAR 2002. Proceedings. International Symposium on*, pp. 97 – 106, 2002.
- [5] D. Wagner, G. Reitmayr, A. Mulloni, T. Drummond, and D. Schmalstieg, "Real-time detection and tracking for augmented reality on mobile phones," *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, pp. 355–368, May 2010.
- [6] P. Supan, I. Stuppacher, and M. Haller, "Image based shadowing in real-time augmented reality," *IJVR*, vol. 5, no. 3, pp. 1–7, 2006.
- [7] K. Agusanto, L. Li, Z. Chuangui, and N. W. Sing, "Photorealistic rendering for augmented reality using environment illumination," in *Proceedings of the 2nd IEEE/ACM International Symposium on Mixed and Augmented Reality, ISMAR '03*, (Washington, DC, USA), pp. 208–, IEEE Computer Society, 2003.
- [8] P. Supan and I. Stuppacher, "Interactive image based lighting in augmented reality," in *The 10th Central European Seminar on Computer Graphics (CESCG 2006)*, Apr. 2006.
- [9] S. DiVerdi and T. Höllerer, "Combining dynamic physical and virtual illumination in augmented reality," Technical Report UCSB//CSD-04-28, University of California at Santa Barbara, Nov. 2004.
- [10] P. W. Pardel and K. Wojciechowski, "Three cameras method of light sources extraction in augmented reality," in *Proceedings of the 2010 international conference on Computer vision and graphics: Part II, ICCVG'10*, (Berlin, Heidelberg), pp. 183–192, Springer-Verlag, 2010.
- [11] C. B. Madsen and M. Nielsen, "Towards probe-less augmented reality - a position paper," in *GRAPP*, pp. 255–261, 2008.
- [12] K. Jacobs, J.-D. Nahmias, C. Angus, A. Reche, C. Loscos, and A. Steed, "Automatic generation of consistent shadows for augmented reality," in *Proceedings of Graphics Interface 2005, GI '05*, (University of Waterloo, Ontario, Canada), pp. 113–120, Canadian Human-Computer Communications Society, 2005.
- [13] J.-F. Lalonde, A. A. Efros, and S. G. Narasimhan, "Estimating natural illumination from a single outdoor image," in *IEEE International Conference on Computer Vision*, October 2009.
- [14] J.-F. Lalonde, *Understanding and Recreating Visual Appearance Under Natural Illumination*. PhD thesis, Carnegie Mellon University, January 2011.
- [15] P. S. Strauss, "A realistic lighting model for computer animators," *IEEE Comput. Graph. Appl.*, vol. 10, pp. 56–64, November 1990.
- [16] M. Slater, M. Usoh, and Y. Chrysanthou, "The influence of dynamic shadows on presence in immersive virtual environments," in *Selected papers of the Eurographics workshops on Virtual environments '95*, (London, UK), pp. 8–21, Springer-Verlag, 1995.
- [17] P. Debevec, "Image-based lighting," in *ACM SIGGRAPH 2005 Courses, SIGGRAPH '05*, (New York, NY, USA), ACM, 2005.
- [18] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting (The Morgan Kaufmann Series in Computer Graphics)*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2005.
- [19] J. Stumpfel, A. Jones, A. Wenger, C. Tchou, T. Hawkins, and P. Debevec, "Direct hdr capture of the sun and sky," in *ACM SIGGRAPH 2006 Courses, SIGGRAPH '06*, (New York, NY, USA), ACM, 2006.
- [20] J.-M. Frahm, K. Koeser, D. Grest, and R. Koch, "Markerless augmented reality with light source estimation for direct illumination," in *Conference on Visual Media Production CVMP*, 2005.
- [21] A. Panagopoulos, D. Samaras, and N. Paragios, "Robust shadow and illumination estimation using a mixture model," in *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 651–658, IEEE, 2009.
- [22] A. Panagopoulos, C. Wang, D. Samaras, and N. Paragios, "Illumination estimation and cast shadow detection through a higher-order graphical model," in *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.

Author Biography

Ibrahim Arief received his Bachelor of Engineering on Informatics from Institut Teknologi Bandung in 2008, and his Master of Science from the Erasmus Mundus Color in Informatics and Media Technology in 2011 under full scholarship from the European Union Erasmus Mundus program. Academically, his passion is in computer vision and augmented reality. Professionally, he architects and develops cloud-based multi-agent distributed systems.