

Practical Solutions for Calibration of Optical See-Through Devices

Y. Genc[†] M. Tuceryan[‡] N. Navab[†]

[†]Siemens Corporate Research
Imaging and Visualization Department
Princeton, NJ 08540, USA

[‡]Dept. of Computer and Information Sci.
Indiana U. Purdue U. (IUPUI)
Indianapolis, IN 46202, USA

Abstract

One of the most crucial tasks in a see-through augmented reality (AR) system is to register the virtual objects with the real world through a transparent display. The importance stems not only from the fact that registration requires careful calibration of the system but also from the necessity that any calibration procedure should take into account the users.

Tuceryan et al. in [26] proposed a general method for calibrating any see-through AR system. This method is called the single point active alignment method (SPAAM) based on dynamic alignment of a set of displayed points with a single scene point at a time. Our experiences with this method on several optical see-through AR systems showed that although a powerful tool, users find alignment of many points overwhelming. This paper introduces several improvements to simplify the calibration process and increase the success rate. After identifying the causes where calibration parameters change from user to user, we propose ways of preventing most of the change by adopting particular configurations for tracker sensor and the see-through display. Then, a simple user interface is implemented where existing calibrations can be used. Furthermore, we have introduced a simpler model for the calibration that requires less number of user inputs, typically four, to calibrate the system.

1 Introduction

Augmented reality (AR) is a technology in which a user's view of the real world is enhanced with additional information generated from a com-

puter model. The enhancements may include labels, 3D rendered models, or shading and illumination changes. Augmented reality allows a user to work with and examine the physical world, while receiving additional information about the objects in it.

In a typical augmented reality system, a user's view of a real scene is augmented with superimposed graphics. The graphics are generated from geometric models of both virtual and real objects in the environment. In order for the graphics and the video to align properly, the pose and optical properties of the real (human eye) and virtual cameras must be the same. The position and orientation of the real and virtual objects in some world coordinate system must also be known. This is accomplished by tracking the location of the real objects and using this information to update the corresponding transformations within the virtual world. Once these capabilities have been brought together, real objects and computer-generated graphics may be blended together, thus augmenting a dynamic scene with information stored and processed on a computer.

In order for augmented reality to be effective the real and virtual objects must be accurately positioned relative to each other and properties of the components of the system must be correctly specified. This implies that certain measurements or calibrations need to be made. Calibration is one of the most crucial tasks in a see-through augmented reality system. The importance stems not only from the fact that registration requires careful calibration of the system but also from the necessity that any calibration procedure should take

into account the users.

Camera calibration has been studied extensively in the computer vision community [25, 17, 30, 8]. In augmented reality, Azuma and Bishop [1], and Janin et al. [12] describe techniques for calibrating a see-through head-mounted display. Interactive calibration schemes which require the simultaneous alignment of multi-point configurations in order to perform the camera calibration were also explored by McGarrity and Tuceryan [19]. Recently, Kato and Billingham describe an interactive camera calibration method [14] that uses multiple points on a grid. Gottschalk and Hughes [6] present a method for auto-calibrating tracking equipments used in augmented and virtual reality applications.

Some researchers have studied the calibration issues relevant to head mounted displays (HMDs) [2, 3, 1, 10, 11, 13]. Others have focused on monitor based approaches [27, 7, 9, 20, 22, 29]. Both approaches can be suitable depending on the demands of the particular application.

Various tracking modalities, beside the magnetic trackers have been used by numerous researchers. Among those, vision based trackers using fiducials have been implemented [15, 21, 23]. Some researchers have also tried to improve the robustness and accuracy of these trackers using hybrid methods [24].

Evaluation of calibration or registration accuracy of an optical see-through system is a difficult task. Since there is no way to access the augmented image as perceived by the user, traditional image-based measurement methods cannot be used. Several approaches have been suggested. The first is to have the user calibrate the HMD and then report the qualitative accuracy of the alignment of a computer model to its real-world counterpart [4, 12, 28], i.e., “acceptable,” or “unacceptable.” The second method uses a camera in place of the human eye and conducts image-based measurements to assess the accuracy of the alignment [10, 19]. Another approach is to estimate the registration error through error propagation [16]. Finally, the online system described by McGarrity et al. [18] measures the registration accuracy while the users are using the AR system.

User’s actual perception and interaction with the real world is measured by asking the user to indicate the projection of a set of perceived objects on a measurement device. This allows an analysis of the structure of the system error in two or three dimensions.

Tuceryan *et al.* in [28] proposed a method for calibrating any see-through device for an augmented reality application. This method is called the single point active alignment method (SPAAM) that is based on dynamic alignment of a set of displayed points with a single scene point. Our experiences with this method on several optical see-through AR systems showed that although a powerful tool, users find alignment of many points overwhelming.

Any calibration scheme for an optical see-through system should be user friendly and robust. User friendliness is needed since calibration is an online process and needs to be done almost every time the HMD is put on. The particular method needs to be robust to user errors during calibration as well as reduce its requirements for the online calibration step. The method also should be robust against the tracker errors even though their causes cannot easily be determined.

We, in this paper, propose several improvements over the existing calibration techniques. The first of these improvements is based on the usage of the existing calibrations. In this, the user is asked to choose among any number of existing calibrations while his/her view is augmented with some virtual objects which the user is supposed to associate with their real counterparts in the scene. Once the user has chosen one calibration that works approximately, he/she can do several more SPAAM alignments to correct it further using the visual feedback provided.

The second improvement concerns the modification of the models for the virtual camera. The SPAAM calibration method assumes a model that can change between each user and between each use of the same user. An alternative model is established to identify invariant parameters of the model for different users etc. This way a two-stage calibration process is designed such that the first stage involves in off-line estimation of the fixed

parameters and the second stage involves an on-line task for the particular user. The advantage of this modified method is that one can be as careful as he wants to be in the off-line stage with a lot of measurement but can be relaxed in the second stage. For example, it could be possible to perform the second stage of the calibration with much less alignments than required by the regular SPAAM method, typically four.

This paper is organized as follows. Sections 2 and 3 explain the proposed two-stage calibration process. After presenting the results of our experiments with the proposed calibration methods in Section 4, we conclude in Section 5 with some discussions.

2 Use of Existing Calibration

We propose to use the existing calibrations that have been performed on the system by the current user in the following way:

- Let the user scroll through a set of existing calibrations for the current system to choose one that registers the virtual object to the world in the best way.
- Let the user perform as few alignments as necessary for regular SPAAM (or SPAAM2 as explained in the next section) on top of the calibration chosen in the first step.
- Add the successful calibration into the list of calibrations to be used later.

In order to achieve the first step, we need to have some assumptions regarding the augmented-reality system. First of all, we need to assume that the tracker sensor (e.g., the tracker camera), should be fixed with respect to the virtual display. For I-Glasses type of HMD, enforcing this is easy since the displays are rigid. The only thing we need to know is which eye the users use. We can classify the existing calibrations according to the left or right eye and let the user choose between one of them. For MicroVision type of the HMDs, this can be achieved when the tracker sensor is placed on top of the moving display component. The second assumption is not so strong and suggests that the user put the goggles on in

an expected manner so that one of the earlier calibrations can approximately fit.

Once we know the current display position with respect to the tracker sensor, we take a subset of the existing calibration fitting the display position and let the user choose amongst these that overlays some virtual objects on some real world objects. To facilitate this process further, we have built a calibration grid where the user can place it anywhere. This grid will have an integrated marker and a set of other grid like structures for the user to see where the overlay take place in order to assess the quality of the calibration/registration (a user using a prototype of such a calibration device is shown in Figure 1).




Figure 1: An example calibration grid with attached marker to help the user to pick a set of calibration parameters from existing ones and do further SPAAM alignments.

Once the user chooses an acceptable calibration from the existing one, he/she can perform additional alignments to correct the calibration further. The resulting calibration can be time-stamped and saved in a file for future use as an existing calibration.

3 A Modified Calibration Model

The “Two Stage” Single Point Active Alignment (SPAAM2) calibration of see-through head-mounted displays (HMD’s) has based on the following:

1. An off-line calibration of the HMD is performed using SPAAM. This will include the estimation of all 11 parameters of the projection matrix. A subset of these parame-

ters is fixed and does not change based on the HMD moving on one’s head or change of person. For example, the rigid transformation between the tracker mark that is rigidly attached and the display will not change by moving the goggles.

2. The HMD is moved on one’s head or when a different user wears it, the parameters that are affected can be re-estimated using a simpler procedure and by collecting a smaller number of data points.




Figure 2: Tracker and virtual camera coordinate systems for illustration of the proposed simplified model for SPAAM2 calibration procedure.

We have the following observations regarding the calibration process (see Figure 2):

1. What we want is the transformation R_E, T_E . The first step gives us R_C, T_C and the camera intrinsics.
2. Ideally, the camera and the eye should coincide. I.e., R_{CE}, T_{CE} should be identity and the camera intrinsics should be the same for the camera and the eye.
3. Moving the goggles on the same person would change R_{CE}, T_{CE} but would not affect the camera intrinsics.
4. Moving the goggles to a new person would potentially change R_{CE}, T_{CE} , and the camera intrinsics.

Method: With the above observations we propose the following method for calibration: Approximate the changed transformation in the eye

by a warp. This can be justified because the most likely changes due to moving HMD on a user’s head would be a shift in the 2D rendered display. Any change in focal length would result in a scaling of the 2D rendered image. Therefore, one can model the camera O_C to eye O_E transform as a simple linear scale and shift of the image in 2D and estimate the small number of parameters needed to accomplish this transformation. Let $P_C = (x_C, y_C)^T$ be a point in camera coordinates and let $P_E = (x_E, y_E)^T$ be a point in the eye coordinates. We can model the transformation as:

$$\begin{aligned} x_E &= \alpha_u x_C + u_0 \\ y_E &= \alpha_v y_C + v_0 \end{aligned}$$

or, in matrix form

$$P_E = \begin{pmatrix} x_E \\ y_E \\ 1 \end{pmatrix} = \begin{pmatrix} \alpha_u & 0 & u_0 \\ 0 & \alpha_v & v_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_C \\ y_C \\ 1 \end{pmatrix} = K P_C.$$

α_u, α_v, u_0 and v_0 are the parameters representing a stretching and translation in the image as shown in Figure 3.




Figure 3: Stretching effect of the model discussed in the text.

This means that we only have to estimate four parameters when the HMD is moved to a new position or a new person. One-point correspondence results in two equations, therefore, obtaining two point correspondences would give us sufficient data to estimate the necessary parameters to correct the calibration.

Method of collecting the correspondences:

The pre-calibration using SPAAM gives us the projection matrix Π for the display, allowing us to generate the point P_C from the coordinates of the single calibration point P_W in the world. Using the same method of collecting data as in SPAAM,

we collect two points, $\mathbf{P}_{E,1}$ and $\mathbf{P}_{E,2}$, that the user aligns with the calibration point in the world. Using the pre-calibrated projection matrix Π and the tracker data when the points were aligned, we compute the images of the point, $\mathbf{P}_{C,1}$ and $\mathbf{P}_{C,2}$. Now we have the point correspondences that give us the following equation:

$$\begin{pmatrix} x_{C,1} & 1 & 0 & 0 \\ 0 & 0 & y_{C,1} & 1 \\ x_{C,2} & 1 & 0 & 0 \\ 0 & 0 & y_{C,2} & 1 \end{pmatrix} \begin{pmatrix} \alpha_u \\ u_0 \\ \alpha_v \\ v_0 \end{pmatrix} = \begin{pmatrix} x_{E,1} \\ y_{E,1} \\ x_{E,2} \\ y_{E,2} \end{pmatrix}$$

From this equation, we can now solve for the parameters of the K matrix. The new projection matrix then is:

$$\Pi' = K\Pi.$$

Alternative model: Alternatively, we can model the correction as a general 3x3 upper triangular matrix. This would have five parameters to estimate.

$$K = \begin{pmatrix} \alpha_u & s & u_0 \\ 0 & \alpha_v & v_0 \\ 0 & 0 & 1 \end{pmatrix}$$

We can estimate these five parameters using a similar data collection method as above. Only now, we would have to collect at least three points that are not collinear. This models the possible skew as well as the scaling and translation as seen in Figure 4.




Figure 4: The skew and scale effect of the model discussed in the text.

Other models can be introduced as well. For example, we can assume that the translation in the space is dominant with respect to the translation in the image, i.e.,

$$C = \begin{pmatrix} \alpha_u & s & 0 & t_1 \\ 0 & \alpha_v & 1 & t_1 \\ 0 & 0 & 1 & t_3 \end{pmatrix}.$$

The new projection matrix is obtained

$$\begin{pmatrix} \Pi'_{3 \times 4} & 0 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} = C \begin{pmatrix} \Pi_{3 \times 4} & 0 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

4 Experiments

We have tested the methods proposed above on our augmented reality system shown in Figure 5. As in any optical see-through system, assessing the accuracy of the registration, hence calibration is not a trivial task. This is due to the fact that we do not have access to the image formed on the retina of the user.



Figure 5: The augmented reality system used in our experiments.

We have chosen two ways to demonstrate our results. First is a qualitative assessment of the registration accuracy after calibration. We conducted tests with several users and asked them to evaluate qualitatively how good is the registration is. The second is a quantitative measurement of the calibration and registration accuracy. For this we gather an extensive set of data during calibration. In particular, we asked the users to dynamically align a set of 2D image points (crosses) with a 3D scene point the same as the data alignment step of SPAAM calibration. The gathered data is grouped into calibration and test groups for gathering error statistics as explained below.

We had four people test the calibration methods. For this, a fifth person calibrated the system once for each eye. Then, the other four have used the same calibration directly. These people

reported miss-registration on the order of 1-3cm at about 60 cm distance. Then they are asked to input four more alignments. After these, using the second stage of the calibration, the reported accuracy improved to about 1cm.

A second set of experiments is conducted to present some quantitative results. For this, two people calibrated the system. The first user is asked to align 30 points while the HMD is fixed. We label these alignments as set S . Another 20 alignments are done by this user after putting the HMD off and on, labeled set S_1 . Finally, the second person is asked to perform 30 alignments, labeled set S_2 .

Using these alignments, we have tested the accuracy of the calibration as follows. We used the alignments in S to estimate an initial SPAAM projection matrix Π which represents the fixed parameters. For each user ($i = 1, 2$), we used part of S_i (labeled S_{i1}) for calibrating the system for the changed parameters using n points (from 3 to 12). This yields an updated projection matrix. We then estimated the re-projection errors on the rest of the alignments in S_i (labeled S_{i2}). We repeated this process for 100 times by picking n number of alignments from S_i at random. Figure 6 shows the mean and standard deviation of the re-projection errors for S_1 , i.e., the first user. Results of the same experiment for the second user are displayed in Figure 7.

In these experiments, we have six ways of estimating the re-projection errors for different projection matrix estimations representing different models and data. These are (assuming that n alignments are picked at random from S_i):

- Π : Re-projection error on S_{i2} using Π .
- Π' : Re-projection error on S_{i2} using the projection matrix obtained from the alignments S and S_{i1} together.
- $\alpha_u, \alpha_v, s, u_0, v_0$: Re-projection error on S_{i2} using the updated projection matrix obtained from the alignments S_{i1} and Π . This uses the five parameters $\alpha_u, \alpha_v, s, u_0$, and v_0 as described in Section 3.
- $\alpha_u, \alpha_v, u_0, v_0$: Re-projection error on S_{i2} us-

ing the updated projection matrix obtained from the alignments S_{i1} and Π . This uses the four parameters α_u, α_v, u_0 , and v_0 as described in Section 3.

- $\alpha_u, \alpha_v, s, t_1, t_2, t_3$: Re-projection error on S_{i2} using the updated projection matrix obtained from the alignments S_{i1} and Π . This uses the six parameters $\alpha_u, \alpha_v, s, t_1, t_2$, and t_3 as described in Section 3.
- $\alpha_u, \alpha_v, t_1, t_2, t_3$: Re-projection error on S_{i2} using the updated projection matrix obtained from the alignments S_{i1} and Π . This uses the five parameters $\alpha_u, \alpha_v, t_1, t_2$, and t_3 as described in Section 3.




Figure 6: Statistics of the re-projection errors for the calibration done by the first user performing 50 alignments (the last 20 are done after putting the HMD off and on). See the text for details.

As it can be seen from these results, the original SPAAM calibration is not very accurate after the virtual camera is changed (because of a change of user, or due to the replacement of the HMD on the same user). In fact, while it yields an average error of 6 pixels when the HMD stayed fixed, a mean error of 20 pixels is observed in the re-projections when the HMD is moved. With four or more points, the modified model consistently yields better results. The four different models with 5, 4, 6 and 5 parameters yield similar results.

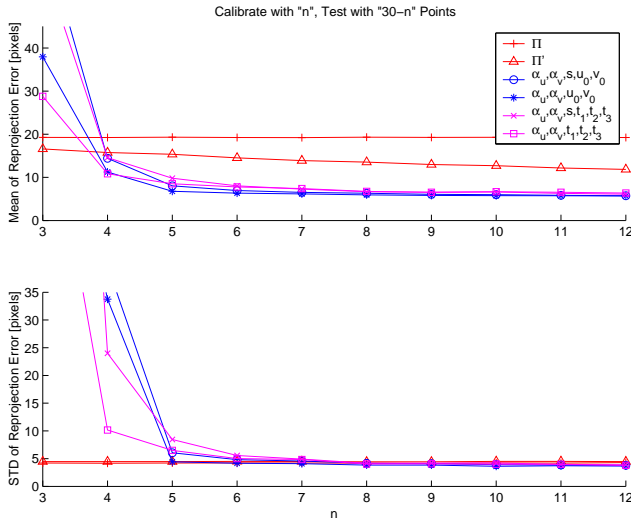


Figure 7: Statistics of the re-projection errors for the second user’s calibration where the 30 alignments from the first user are used along with the 30 alignments from the second person. See the text for details.

The experiments suggests that the less the number of parameters in the model is, the more accurate the calibration gets.

5 Conclusions and Discussions

We have presented practical solutions to the problem of calibrating an optical see-through augmented reality system. First, after identifying the causes where calibration parameters changes from user to user, and eliminating most of the changes, we proposed a simple user interface where existing calibrations can be used. Second, a simpler model for representing the virtual camera is proposed which only considers the changing parameters from user to user.

As the SPAAM calibration allows dynamic alignment, the choices of the scene point or points and of the 2D virtual points are flexible. Our implementation currently considers a stationary 3D scene and a single point for the alignment. This helps us in two important ways. First, the user interface requirements are minimal when only a single scene point is used. In the case of multiple scene points, the system and the user should share the knowledge of what scene point is currently being used. A second advantage is that it requires

a single measurement in the workspace which can later be used for augmentation. Furthermore, by choosing a point in the workspace of the intended application, the user is more likely to get better registration during the system operation.

Acknowledgments We would like to thank Siemens A&D for supporting our research efforts in augmented reality. We would also like to thank to J. Traub and X. Zhang for helping in implementing and testing some of the ideas presented in this paper.

References

- [1] R. Azuma and G. Bishop. Improving static and dynamic registration in an optical see-through display. *Computer Graphics*, pages 194–204, July 1994.
- [2] M. Bajura, H. Fuchs, and Ohbuchi. Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *Computer Graphics*, pages 203–210, July 1992.
- [3] T. Caudell and D. Mizell. Augmented reality: An application of heads-up display technology to manual manufacturing processes. In *Proceedings of the Hawaii International Conference on System Sciences*, pages 659–669, 1992.
- [4] S. Feiner, B. MacIntyre, and D. Seligmann. Knowledge-based augmented reality. *Communications of the ACM*, 36(2):53–62, 1993.
- [5] Y. Genc, E. McGarrity, M. Tuceryan, and N. Navab. An optical see-through augmented reality system for maintenance and training. In *International Symposium for Augmented Reality*, New York, NY, 2001.
- [6] S. Gottschalk and J. Hughes. Autocalibration for virtual environments tracking hardware. *Computer Graphics*, pages 65–72, August 1993.
- [7] W.E.L. Grimson, G. J. Ettinger, S.J. White, P.L. Gleason, T. Lozano-Perez, III W. M. Wells, and R. Kikinis. Evaluating and validating an automated registration system

- for enhanced reality visualization in surgery. In *Proc. of the IEEE Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'95)*, pages 3–12, 1995.
- [8] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge, UK, 2000.
- [9] C. J. Henri, A.C.F. Colchester, J. Zhao, D.J. Hawkes, D.L.G. Hill, and R. L. Evans. Registration of 3D surface data for intra-operative guidance and visualization in frameless stereotactic neurosurgery. In *Proc. of the IEEE Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'95)*, pages 47–58, 1995.
- [10] R. Holloway. *An Analysis of Registration Errors in a See-Through Head-Mounted Display System for Craniofacial Surgery Planning*. PhD thesis, University of North Carolina at Chapel Hill, 1994.
- [11] R. L. Holloway. Registration error analysis for augmented reality. *Presence: Teleoperators and Virtual Environments*, 6(4):413–432, August 1997.
- [12] A. Janin, D. Mizell, and T. Caudell. Calibration of head-mounted displays for augmented reality applications. In *Proc. of the Virtual Reality Annual International Symposium (VRAIS'93)*, pages 246–255, 1993.
- [13] A. R. Kancherla, J. P. Rolland, D. L. Wright, and G. Burdea. A novel virtual reality tool for teaching dynamic 3D anatomy. In *Proc. of the IEEE Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'95)*, pages 163–169, 1995.
- [14] H. Kato and M. Billinghurst. Marker tracking and hmd calibration for a video-based augmented reality conferencing system. In *Proc. IEEE International Workshop on Augmented Reality*, San Francisco, CA, USA, October 1999.
- [15] D. Koller, G. Klinker, E. Rose, D. Breen, R. Whitaker, and M. Tuceryan. Real-time vision-based camera tracking for augmented reality applications. In *Proceedings of the Symposium on Virtual Reality Software and Technology (VRST-97)*, Lausanne, Switzerland, September 1997.
- [16] B. MacIntyre, E. M. Coelho, and S. Julier. Estimating and adapting to registration error in augmented reality systems. *Proceedings of IEEE Virtual Reality*, 2002.
- [17] S. J. Maybank and O. D. Faugeras. A theory of self calibration of a moving camera. *International Journal of Computer Vision*, 8(2):123–151, 1992.
- [18] E. McGarrity, Y. Genc, M. Tuceryan, and N. Navab. A new system for online evaluation of optical see-through augmentation. In *Proc. IEEE International Symposium for Augmented Reality*, New York, NY, USA, October 2001.
- [19] E. McGarrity and M. Tuceryan. A method for calibrating see-through head-mounted displays for ar. In *2nd International Workshop on Augmented Reality (IWAR '99)*, pages 75–84, San Francisco, CA, October 1999.
- [20] J. P. Mellor. Real-time camera calibration for enhanced reality visualizations. In *Proc. of the IEEE Conference on Computer Vision Virtual Reality and Robotics in Medicine (CVRMed'95)*, pages 471–475, 1995.
- [21] U. Neumann and Y. Cho. A selftracking augmented reality system. In *Proceedings of the ACM Symposium on Virtual Reality and Applications*, pages 109–115, July 1996.
- [22] O. Peria, L. Chevalier, A. François-Joubert, J. P. Caravel, S. Dalsoglio, S. Lavallee, and P. Cinquin. Using a 3D position sensor for registration of spect and us images of the kidney. In *Proc. of the IEEE Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'95)*, pages 23–29, 1995.
- [23] F. Sauer, F. Wenzel, S. Vogt, Y. Tao, Y. Genc, and A. Bani-Hashemi. Augmented

- workspace: Designing an AR testbed. In *International Symposium for Augmented Reality*, pages 165–174, Munich, Germany, October 2000.
- [24] A. State, G. Hirota, D.T. Chen, B. Garrett, and M. Livingston. Superior augmented reality registration by integrating landmark tracking and magnetic tracking. In *SIGGRAPH '96*, pages 429–438, New Orleans, LA, August 1996.
- [25] R.Y. Tsai. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras. *IEEE Journal of Robotics and Automation*, RA-3(4):323–344, 1987.
- [26] M. Tuceryan, Y. Genc, and N. Navab. Single point active alignment method (SPAAM) for optical see-through HMD calibration for augmented reality. *Presence: Teleoperators and Virtual Environments*, 11(3):?, August 2002. To appear.
- [27] M. Tuceryan, D. Greer, R. Whitaker, D. Breen, C. Crampton, E. Rose, and K. Ahlers. Calibration requirements and procedures for a monitor-based augmented reality system. *IEEE Transactions on Visualization and Computer Graphics*, 1(3):255–273, September 1995.
- [28] M. Tuceryan and N. Navab. Single point active alignment method (spaam) for optical see-through hmd calibration for ar. In *International Symposium for Augmented Reality*, pages 149–158, Munich, Germany, October 2000.
- [29] M. Uenohara and T. Kanade. Vision-based object registration for real-time image overlay. In *Proc. of the IEEE Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'95)*, pages 13–22, 1995.
- [30] J. Weng, P. Cohen, and M. Herniou. Camera calibration with distortion models and accuracy evaluation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 14(10):965–980, 1992.