

Efficient Verification of Holograms Using Mobile Augmented Reality

Andreas Daniel Hartl, *Member, IEEE*, Clemens Arth, *Member, IEEE*, Jens Grubert, *Member, IEEE*, and Dieter Schmalstieg, *Senior Member, IEEE*

Abstract—Paper documents such as passports, visas and banknotes are frequently checked by inspection of security elements. In particular, optically variable devices such as holograms are important, but difficult to inspect. Augmented Reality can provide all relevant information on standard mobile devices. However, hologram verification on mobiles still takes long and provides lower accuracy than inspection by human individuals using appropriate reference information. We aim to address these drawbacks by automatic matching combined with a special parametrization of an efficient goal-oriented user interface which supports constrained navigation. We first evaluate a series of similarity measures for matching hologram patches to provide a sound basis for automatic decisions. Then a re-parametrized user interface is proposed based on observations of typical user behavior during document capture. These measures help to reduce capture time to approximately 15 s with better decisions regarding the evaluated samples than what can be achieved by untrained users.

Index Terms—Document inspection, holograms, augmented reality, user interfaces, mobile devices.

1 INTRODUCTION

THE purpose of document inspection is to reason about the validity of a document by examination of security elements. Documents of interest are machine-readable travel documents like passports, identification cards and visas, but also checks, vouchers and banknotes. According to a press release by the European Central Bank [8], over half a million counterfeit Euro banknotes were withdrawn in the second half of 2014. This example shows that everyone should be concerned about document counterfeiting and that the public must be educated on how to prevent fraud in everyday cash payment.

A large variety of security features can be incorporated into the substrate of a security document, printed on top, or added as a separate element, sometimes as an additional foil over the entire document. A subset of those features are called optically variable devices (OVD). They have distinct visual properties, which change considerably when varying the viewing angle or the position of incident light sources. Often referred to as holograms by the public, their unauthorized reproduction, but also their verification are challenging tasks. Van Renesse [28] states that, according to the International Hologram Manufacturers Association (IHMA), there have never been any accurate copies of well-designed authentication holograms. Fake documents usually employ substitutes instead of counterfeit holograms.

While trained individuals can identify the majority of fake documents and holograms within a few seconds (according to a domain expert we have consulted), lay people inspect holograms on security documents just by looking for changes in appearance or the pure presence of rainbow

colors. This has no particular value regarding security [28]. First level inspection of holograms is currently based on printed guides, which are often issued by public authorities. This requires an individual to find all relevant patterns listed in a manual by looking at the hologram. However, these manuals often lack an indication on the viewing direction and do not specify requirements on lighting conditions. Also, in real-world situations, manuals are not always at hand, so users fall back to solely looking for appearance changes. In general, there is a tendency to reject questionable documents instead of learning how to inspect them properly [28]. Therefore, the development of tools and algorithms for the automatic detection and verification of holograms by laypeople deserves our interest.

Augmented Reality (AR) can help with inspecting holograms. However, such an approach poses unique challenges regarding image capture, matching and, in particular, user guidance, resulting in high temporal and cognitive effort [11]. Using carefully designed user interfaces along with automatic recording and matching of hologram patches, the efficiency of the process can be improved, but temporal effort is still long and the overall accuracy is limited compared to manual decisions [10]. This impedes any real-world application of mobile hologram verification.

As the main contribution of this work, we provide a thorough evaluation of matching behavior in order to select a suitable similarity measure for robust matching of hologram patches under typical operating conditions. By analysis of typical user behavior during the examination of documents and holograms, a novel parametrization of an efficient user interface for hologram verification is found. The new design is evaluated in a user study and is found to reduce the temporal effort of the process to around 15 s, while using a larger number of views. The mobile prototype delivers perfect verification performance, even surpassing that of human individuals for the evaluated samples.

- A. Hartl, C. Arth and D. Schmalstieg are with the Institute for Computer Graphics and Vision, Graz University of Technology, Inffeldgasse 16, 8010 Graz, Austria.
- J. Grubert is with the Institute for Embedded Systems, Universität Passau, Innstraße 43, 94032 Passau, Germany.



Fig. 1. Mobile prototype for hologram verification (specimen document): Hybrid user interface for capturing image data parametrized to indicate a constrained navigation space corresponding to typical user-behavior (left). Verification results are presented to the user in a summary (right).

2 RELATED WORK

Security holograms show different patterns, depending on the viewing angle and sources of illumination, exploiting various physical phenomena such as diffraction or interference. Rainbow holograms can be viewed by using white light, letting the object appear in all spectral colors. Depending on the number of layers, 2D and 3D images or motion sequence (stereogram, kinegram) can be shown. A more natural appearance can be achieved with true color holograms. Related work on OVD can be divided into approaches suitable for capturing, reconstruction and inspection. Holograms require considerable interaction, even when manual inspection is carried out. Consequently, it is important to guide the user throughout the inspection process and to give appropriate feedback.

2.1 Hologram Capture, Reconstruction and Inspection

Capturing holograms is largely related to capturing a spatially varying bidirectional reflectance distribution function (SVBRDF). This 6D function characterizes the amount of radiance that is reflected at each surface point according to the viewing and lighting directions. Ren et al. describe a portable solution to SVBRDF measurement of flat surfaces using a mobile device, a BRDF chart and a linear light source [22]. Being based on an approach by Dong et al. [6], they locally reconstruct purely specular components which allows for arbitrary per-point variation of diffuse and specular parts. Jachnik et al. [16] conduct real-time surface light-field capture from a single handheld camera with fixed exposure, shutter and gain. They require a static planar scene and illumination and split diffuse and specular components, finally estimating an environment map. They rely on a guidance component in the form of a colored hemisphere, which indicates whether a pixel has already been seen from a particular viewing direction. Being designed for mobile verification, the proposed approach does not require capturing the entire representation, but only a subset of relevant patches. Although the BRDF is not explicitly modeled, sharp changes in appearance as well as the necessary detail in the spatial domain can be preserved. Hartl et al. [9] detect holograms on arbitrary documents by analysis of a registered stack of document images obtained during interaction with a mobile AR setup. This approach does not require prior knowledge about document location

or content and is suitable for real-time operation on off-the-shelf smartphones. They also describe a gaming scenario as an alternative method for gathering image data relevant to hologram detection.

The reconstruction of 3D information from holograms is usually connected to digital holography, where an image sensor is used for recording interference patterns, instead of a photo-platter. Buraga-Lefebvre et al. [3] analyze the diffraction pattern on a hologram (in-line holography) using the wavelet transform in order to reconstruct the location of small particles in 3D. Their setup requires a laser source, a movable hologram, a relay lens and a camera. They state that diffraction can be treated as a convolution between the amplitude distribution in the object plane and a family of Wavelet functions. In contrast to previous approaches, no focusing on individual particles is required, improving overall accuracy. Amplitude reconstructions of holograms are shown by Pitk aho to be suitable for gaining a depth image using stereo reconstruction [20].

Pramila et al. [21] segment the watermark of a dual-layer hologram. Recording is done using a camera and a uniform light source, facing towards a tilt-able plane containing the hologram. They note that the result is very sensitive to the angle of the plane. Holographic patterns are identified from a printed page by Janucki et al. [17]. They create a Wavelet approximation of the intensity distribution of the hologram and use a Wiener filter to eliminate the influence of non-uniform background. This setup is also suitable for quality estimation of a holographic device. Automatic inspection systems for holograms can use sets of patterns illuminated with multiple IR LEDs on a hemisphere [18], [19]. Images are captured with a CCD camera at controlled illumination angle, and correlation-based matching is carried out in the frequency domain. They extend the system with a correction of rotation angles and evaluate it with two Korean banknotes. Soukup et al. [25] sample the BRDF of a diffractive optically variable image device (DOVID) using photometric stereo and light-field-based methods. For this purpose, they propose a tailored feature descriptor which is robust against several expected sources of inaccuracy, but still specific enough for the given task. They demonstrate their approach on the practical task of automated discrimination between genuine and counterfeited DOVID on banknotes.

Hartl et al. have demonstrated the feasibility of capturing and verifying holograms in a mobile AR setting

using an off-the-shelf mobile device [11]. By using the built-in flashlight of the device as a dominant light-source, the appearance of reference patches can be reproduced in a mobile context. Still, this approach requires manual matching of recorded patches by the user. Compared with a reference manual, this approach takes a lot of time and involves heavy physical and cognitive load for the operator. The user is guided towards the required poses using a complex alignment-based user interface, where most parts are augmented onto the target. Proper alignment requires pointing at the base of the reference viewing ray, looking into its direction (iron sights) and adjusting the distance using two appropriately scaled circles at the base and top of the ray. Finally, the orientation around the viewing ray must be matched (virtual horizon).

2.2 User Guidance and Visualization

Since holograms differ in appearance depending on the viewing direction and the presence of light sources in the environment, it is reasonable to guide the user through the inspection process and to give appropriate feedback. Thus, the presentation of information with respect to the real world and guidance for the user deserve special attention. There is a variety of work on guiding the user within a small workspace.

User guidance can be approached by visualization of the view alignment error concerning a given reference pose. Examples are surgical scenarios, where colored augmented coordinate systems are used for easier navigation of the end effector [5]. Pyramidal frustums can also serve as a means of guidance for navigation. This can be seen as a geometric representation of the camera at the time of capture [24]. This approach is used for real-time visual guidance for accurate alignment of an ultrasound probe by Sun et al. [27]. After tracking artificial skin features for probe localization, visual guidance for 6 DoF alignment is provided via an augmented virtual pyramid. Such a pyramidal representation is also related to the Omnidirectional Funnel [2], which is useful for calling attention. Bae et al. [1] use visual guidance for re-photography. They analyze the camera image to determine if a sufficiently similar image was captured. Three visualizations are presented for alignment. First, a 2D arrow indicates the required direction of movement w.r.t. a top-down camera viewpoint. Second, this information is also indicated concerning a back-front camera viewpoint. Finally, they visualize edges for adjustment and feedback of the current camera orientation. Heger et al. [14] perform user-interactive registration of bone with A-mode ultrasound. The pointer is mechanically tracked and a 2D-indicator is used to provide visual feedback about the deviation from the surface normal during alignment of the transducer to the local bone surface.

Alternatively, guidance can be achieved by visualization of a constrained navigation space. Shingu et al. [23] create AR visualizations for re-photography tasks. They use a sphere as a pointing indicator along with a half-transparent cone having its apex at the sphere as an indicator of viewing direction. Once the viewpoint is inside the cone, it is not visible anymore. The sphere changes its color when it is fully visible. This corresponds to a valid recording position.




Fig. 2. Overview of our mobile hologram verification pipeline. For manual verification, images are registered and the extracted hologram patches are rectified. In case of semi-automatic verification, additional processing is required (dotted rectangles).

Sukan et al. [26] propose a wider range of look-from and look-at volumes for guiding the user to a constrained set of viewing positions and orientations, not counting roll (*ParaFrustum*). This can be realized as an in-situ visualization or via non-augmented gauges. In the in-situ variant, the transparency of volumes is modulated depending on the distance and orientation of the current pose. In addition, the general representation of the look-at volume is also changed. Although constrained navigation for inspection tasks is similar, the mobile capture of holograms requires the user not only to enter, but to explore such space in order to get suitable image data.

In this paper we pick up the idea of an efficient navigation approach in small workspaces [10]. In contrast to previous work, we propose a more natural parametrization by analysis of typical user behavior during document inspection. This defines a constrained navigation space which is tailored towards automatic exploration by the user during interaction with the document. In addition, we use a different similarity measure for matching hologram patches, which is shown to be more robust under typical operating conditions.

3 METHOD OVERVIEW

Hologram verification can be seen as a subtask within a document verification process. In the following, we describe a setup for capturing reference data from holograms along with a matching approach that can be used for automatic verification at runtime. Such information can be used in a mobile application for interactive document verification, which performs classification, tracking and augmentation of relevant information (see Figure 2). This allows to select the correct reference information for hologram verification and to make sure that the element is observed from the correct viewpoint. An implementation of this setup serves as a basic building block for goal-oriented mobile hologram verification with appropriate user guidance and visualization.

3.1 Preprocessing

Capturing Reference Data With moderate ambient light, the appearance of a hologram is largely dominated by using the LED flashlight of mobile devices. This essentially means that the workspace consists of a hemisphere centered at the hologram on the document. We use an industrial robot (Mitsubishi MELFA) for capturing all relevant appearances of a view-dependent element. This allows reliable sampling of holograms and eliminates undesired human influence. We spatially sample a hemispherical space using the robot and remotely control the device. We capture the current

video image and the corresponding pose for each position on the hemisphere.

We assume the hologram to be planar and project its bounding box into the image using the recorded pose. We estimate an image transformation with respect to the hologram region on the undistorted template and subsequently warp the sub-image containing the hologram. For increased accuracy, we perform an additional registration step using the template of the document before extraction and rectification of the corresponding patch. The result is a set of registered image patches that represent all observable appearances of the current hologram.

View Selection For successful verification, a series of representative views must be selected using reference information available from the manufacturer or by systematic recording of the hologram and thorough analysis of the captured image data. The choice of reference poses obviously depends on the hologram (e.g., number of transitions) and is constrained by the particular setup being used. We exclude all data recorded $< 5^\circ$ and $> 55^\circ$ away from the orthogonal view in order to avoid artifacts caused by oblique views and tracking failure. From the perspective of security, it seems reasonable to select very different patches having small distances in space. For reasons of usability, a small amount of stable views seems preferable. As a lower bound, at least two visually different views recorded from sufficiently different viewing direction are required. It is reasonable to also consider the typical behavior of users when inspecting documents during view selection (see Section 4.2).

3.2 Runtime Processing

Classification and Tracking The first step in document verification, the identification of the document type, can be achieved by manual selection or by computing the class of the document using a current image. We use visual search running entirely on the mobile device. This avoids the transfer of sensitive information across networks and reduces latency [13]. Afterward, associated reference data relevant to the verification process can be loaded. Tracking works in real-time directly on the mobile device, using natural features obtained from an exemplary template selected during the previous process [29]. In order to allow the verification of slightly bent documents, tracking poses are smoothed in a small ring buffer. Available reference information can be represented by an initial augmentation, providing instant feedback on the presence and location of relevant security features for manual verification (see Figure 3). For increased robustness, the document region can be detected before the actual classification step [12].

Matching Verification of selected reference data demands a suitable similarity measure, which may be used on-line for verification. In an initial iteration we employed normalized cross correlation (NCC) for matching, which was shown to perform reasonably within a feasibility study on hologram verification [11], but did not perform well in a more general setting [10]. In the following, an analysis regarding the performance of different similarity measures is provided in order to find a better basis for automatic matching of patch data obtained from holograms.

We evaluate a series of similarity measures on holograms recorded with the Samsung Galaxy S5 smartphone (see Fig-



Fig. 3. Mobile document verification system tracking a sample instance of the data page used in passport. The position of security features is augmented directly onto the document. Detailed information about an element can be triggered by pointing the camera at it.




Fig. 4. Holograms used in our study. Top Row: Original elements. Bottom Row: Substitutes.

ure 4). In each case, an original, a copy and a substitute were recorded under typical office conditions using the built-in flashlight as a dominant light-source. The settings for the recordings included an office room with light switched off, fluorescent light and the hallway in front of it, which has more daylight influence.

We evaluated several similarity measures such as Sum of Absolute Differences (SAD) and NCC, which are often used for stereo matching [15]. Due to the requirement of correspondence regarding human perception, Structural Similarity Index (SSIM) [30] and Edge-based Structural Similarity (ESSIM) [4] are also included. Additionally, we evaluate SSIM with color patches by reporting the minimum value over all channels (CSSIM). We also perform linear scaling of matching scores using coefficients obtained by analysis of all recorded patch data of a hologram and optionally employ shape matching [7] in order to get rid of false positives (see Figure 2). Certain holograms (e.g., rainbow) show a large variety of colors, which leads to noisy measurements. From our experience, Median pre-filtering of patches with a 3×3 kernel can be used to improve robustness in this case.

Evaluation is carried out as a binary classification task on each reference view of every hologram. The task is to assign the correct class to each recorded patch from an original, copied or substitute hologram based on pair-wise matching. The required matching thresholds are selected automatically based on the difference in scores between original and fake patches for each reference viewing direction. The recording position is considered by matching patches only if the pose




Fig. 5. Left: Performance of various similarity metrics in recognizing fake and original patches when using data from all slices of the orientation map. Right: Relative difference of scores between fakes and originals. Note that the data was taken under various office conditions (no artificial light, fluorescent light, slight daylight (aisle)).

lies within a certain orientation threshold concerning the reference pose. Neglecting this relationship is not desirable, since elements having the same appearance at different viewing positions could not be differentiated anymore. This would weaken the security of the proposed approach.

There are notable differences in patch recognition rate regarding the type of hologram, but also the associated metric (see Figure 5). SSIM-based metrics in general give better results than NCC and SAD. Overall, SSIM is stable, giving patch recognition rates of over 90% for holograms 3 and 4 and over 85% for the remaining ones. Hologram 1 is obviously most difficult to recognize regarding its originality (low SAD and NCC scores). We speculate that this is due to the large amount of rainbow colors present on its patches. As these results do not correspond to the matching performance evaluated within related work (see [10]), further investigation is required.

For robust matching, the margin of the classifier should be as large as possible. Mapping this to the current task, the relative difference of matching scores between originals and fakes should also be large. Normalized relative scoring distances depict considerable differences between the evaluated similarity measures, but also between different holograms (see Figure 5). SAD and NCC only span a very small range compared to SSIM-based measures. So, it is much more difficult to set a reasonable matching threshold for them than it is for SSIM. Based on these insights, it seems more promising to use SSIM for matching hologram patches instead of NCC. This is further backed up by results obtained from performing hologram verification using majority voting on individual patch matching results (see Figure 6). In this case, only SSIM allows to correctly recognize originals and fakes under typical office conditions.

We further investigated the patch matching performance for originals and fakes under optimal office conditions, when using different devices at runtime (see Figure 6). We observed stable verification performance for two off-the-shelf devices (Samsung Galaxy S6 Edge, LG Optimus 4X HD). However, the Xperia Z2 smartphone failed. Further investigation revealed a relatively weak LED light-source coupled with very different sensing characteristics. This device is not able to reproduce different appearances of hologram patches and, thus, cannot be used for hologram verification.

Viewpoint and flashlight act as triggers for different




Fig. 6. Left: Performance in recognizing fake and original holograms by majority voting on patch matching results under various office conditions. Right: Performance in recognizing fake and original patches when using various devices under optimal office conditions.

appearances. Consequently, an initial probing of the current lighting conditions before the actual verification task is required. When the exposure can be fixed on the mobile device, this can be achieved by activation of the flashlight and thresholding the relative amount of saturated pixels in order to reason about the dominance of the built-in light source. From our experience, without using the flashlight as a dominant light source, results are not repeatable.

4 INTERFACE DESIGN

In order to get reasonable input data for verification, the user should be supported throughout the image capture process. An obvious approach is to guide the user to align the mobile device with exactly those view points which are associated with the selected reference data. Alternatively, a portion of space can be visualized for sampling by the user, which requires coverage of a larger region instead of given positions. Combining both approaches leads to a hybrid variant, which uses a comparatively small region for sampling relevant data. The hybrid user interface was evaluated to be overall the fastest one, taking around 40 s for sampling a single hologram [10]. In addition, final decisions by the human operators took approximately another 20 s. However, the success rate of the system was only around 73%, which corresponds to a rate of 90% when treating neutral decisions as correct. This is still a long time span and, together with the non-optimal success rate, such an approach is probably not feasible for a quick check in a real-world situation. It must be noted that users communicated a preference for a constrained navigation approach despite not being the fastest one.

We believe that a more user-friendly parametrization of the hybrid interface together with the updated similarity metric for automatic matching can make hologram verification more accessible to laypeople using off-the-shelf mobile devices. In the following the hybrid user interface will be revisited and then re-parametrized based upon information gained from observing users during document capture.

4.1 Concept and Graphical Representation

The initial step guides the user to point at the hologram as required by the recording setup. We provide guidance using an animated rubber band, which shows a moving arrow, once outside a given radius from the element (see Figure 7). Then, the capture distance needs to be adjusted as a starting




Fig. 7. Geometry of the proposed navigation approach for sampling a hologram. The user is guided to point at the element and a cursor is controlled by the 2D orientation on an augmented pie, divided into slices and tracks (constrained navigation). Alternatively, interesting subspaces can be indicated as circular regions instead of labeling relevant pie slices (hybrid approach).

point for an auto-focus operation, so that the assumption about the flashlight being the dominant light-source holds. For this purpose, we scale the entire widget and require the user to adjust the distance, so that the outer ring of the widget stays within the given distance bounds.

Orientation Map In favor of efficiently treating both originals and fakes, the user should be guided towards different viewing directions or ranges. We propose a 2D orientation map (projection of the conic space) [14] for this task. It is divided into slices that are aligned on one or more tracks. The current position on the map is visualized by a cursor, and the current slice is also highlighted. The cursor position is corrected by the target orientation, so that the movement direction always corresponds to the orientation of the device (see Figures 7, 8). Depending on the selection of reference views, several slices need to be sampled. It is generally not sufficient to just capture a single shot inside each slice. We record several shots per slice that differ at least by a given angle threshold. The exact amount is automatically calculated, taking into account the area of the slice. Consequently, the user can move freely inside the pie slices during the process (constrained navigation). Small arrows around the cursor serve as movement indicators. Whenever the user remains static inside a non-completed slice, flashing arrows remind to move on. The upper arc defined by a (sub-) slice is used as a completion indicator, which switches from red to green with increasing slice coverage.

Circular Regions The location of reference views cannot be mapped straightforward to pie slices. It may be necessary to associate several pie slices with a single reference view, increasing the amount of slices to be checked. Since the number is generally much lower than the total number of pie slices, we use small regions on the augmented map around reference locations, which also serve as local completion indicators (see Figure 8).

If several circular regions with larger distance between them regarding orientation are selected, the hybrid interface moves towards an alignment approach. This is not desirable for the task at hand [10]. Using several circular regions close to each other, the hybrid interface can mimic properties of a constrained navigation space. In the current implementa-



Fig. 8. AR UIs with guidance for interesting subspaces. Either pie-slices (left) or circular regions (right) can be indicated for sampling. In the first case, this corresponds to a constrained navigation space, while in the second case final alignment with the circular regions is required (hybrid).

tion, multiple regions can be affected during image capture by the current pose. This is indicated by visual overlap of circular regions.

During the capture process, automatic matching with reference information is carried out and the results are presented afterward in a summary. This allows an additional inspection of the recorded images by the user (see Figure 1).

4.2 User-Oriented Parametrization

Typical User Behavior With the goal of further reducing temporal effort, the selection of reference views to be checked becomes increasingly important. A selection focusing solely on differences in appearance (as used in a printed or digital manual) could be disadvantageous for mobile applications, since a large range in orientation needs to be spanned by the user. In contrast, a reasonable spatial positioning of reference views could reduce task completion time.

We conducted an experiment in order to gain more insights about the typical behavior of laypersons when recording documents on mobile devices for the purpose of inspection. In the first part of the experiment, participants were asked to record a self-made sample of an ID-document using the Samsung Galaxy S5 smartphone for as long as they deemed appropriate. During recording, the document was tracked, and the pose and video information was logged onto the device including the corresponding timestamps. Users were asked to look specifically at the hologram through the mobile device within two trials. In the first case, users were asked to record the hologram with the document in hand, while in the second case, the document was placed on a table. In order to avoid learning effects, we balanced the order of trials among the participants.

During both trials of hologram inspection, users looked at the hologram and started to tilt the document or the device. On average, users were sampling the hologram in these scenarios for around 33 s ($M=33.35$, $SD=13.59$). Changes in orientation in general took place roughly along the vertical and horizontal axes. However, there is a notable difference in behavior, depending on whether the document is in hand or on the table (see Figure 9). While, in the first case, mainly vertical movements are made into both directions, in the second case, orientation changes take place in the lower direction and to the side (less distinct). The latter seems reasonable, since, otherwise, the user would move the screen of the device away from the field of vision. It must be noted that in the first case, the majority of users tried




Fig. 9. Orientation changes during hologram inspection were considered with the document held in hand (left) and with the document kept on a table (right). This corresponds to tilting the document roughly in the vertical and also in the horizontal direction. In the first case, users did not move the document exactly in the vertical direction. This can also be seen by the visualized Eigenvectors (red and blue lines).




Fig. 10. Alternative layout of reference views (orientation) for hologram verification on movement paths. Reference views should be placed on a vertical path, making the device face towards the user in case the document is lying on a surface (dark-blue circles).

to fix the device in one hand and only tilted the document. From the visualization of the corresponding Eigenvectors it is evident that users did not move the document exactly in the vertical direction, when holding it in their hand, but also rotated it slightly.

Alternative View Selection With the insights gained in the previous experiment, it seems reasonable that the layout of reference views conforms with typical movements of users when examining holograms. In order to allow the inspection of a hologram regardless of whether the document is lying on a desk or held in hand, reference views should be placed in the lower vertical direction of the orientation space (see Figure 10). Due to the observed movement along a path, it also seems reasonable to use a sequence of patterns for verification instead of single spots. In this case, more data is available for matching, which could lead to more robust decisions by the system.

In the following the alternative distribution of reference views is evaluated in a user study regarding accuracy and task completion time. In this case, more reference information is used for matching, while still requiring only small movements by the operator, resulting in low temporal effort and high verification performance.

5 EXPERIMENTAL EVALUATION

We integrated SSIM for matching into the mobile prototype for hologram verification and selected an alternative layout of reference views for the samples used in the patch matching experiment according to typical user behavior (see Figure 1). This prototype was then used in a study with the goal to evaluate the accuracy of decisions by the modified




Fig. 11. Temporal effort (left) and accuracy of user decisions (right) with the updated prototype. Holograms can be assessed by the system in approx. 15 s with all decisions being correct. Users additionally need approx. 20 s for assessment, being correct in 85.94% of all cases.

and re-parametrized system as well as the temporal effort concerning image capture and decisions.

5.1 Study Design and Procedure

Participants were informed about the study purpose and length, followed by a short investigation of demographic data. Then, a training phase was started in order to make the participant familiar with the checking procedure using a fake and an original document. Afterward, four pairs of documents (original, substitute - see Figure 4) had to be checked using the proposed approach. We rotated the sequence of these documents with each participant. During this process, relevant data such as timestamps for various actions, matching scores and system/user decisions on validity were recorded. After each hologram, the users were questioned about their own decision on validity. We consider this a realistic scenario in which the user remains responsible for the final assessment, regardless of the tool being used. After all runs, they were asked to rate the process as a whole and to give comments regarding their experience. For reasons of privacy, only specimen documents were used in our study.

Altogether, 24 users participated in the study (2 female, age $M = 29.54$, $SD = 5.54$). All but one user reported to own a smartphone for at least one year. In general, they described their affinity to technology as high to very high. Half of the participants reported to never have examined a hologram.

5.2 Findings and Discussion

Holograms can be assessed by the system immediately after image capture, which takes approximately 15 seconds ($M = 14.97$, $SD = 8.59$). A subsequent decision by the user takes another 20 s ($M = 20.07$, $SD = 15.17$, see Figure 11). One-way within subjects ANOVA revealed no significant effect of hologram on capture time, but on decision time ($F(7,184) = 2.46$, $p = 0.0196$). Multiple pairwise post-hoc comparisons with Bonferroni correction for hologram revealed that the decision time for hologram 1-o ($M = 27$, $SD = 19.99$) was significantly different from hologram 1-f ($M = 12.83$, $SD = 14.07$). The system was able to assess the hologram correctly in all cases. Users were unsure about the validity of the hologram in 10.94% of all cases and succeeded to give a correct decision in 85.94% of all cases (see Table 1 for details on individual holograms).

Hologram	1-o	1-f	2-o	2-f	3-o	3-f	4-o	4-f
Correct [%]	75.0	100.0	83.3	50.0	91.7	95.8	91.7	100.0
Neutral [%]	25.0	0.0	8.3	37.5	8.3	0.0	8.3	0.0
Wrong [%]	0.0	0.0	8.4	12.5	0.0	4.2	0.0	0.0

TABLE 1

User decisions for hologram inspection using originals and substitutes.

Three users pointed out that they enjoyed using the application (e.g., 'liked it', 'good to use'). Around half of the users mentioned that it was not completely intuitive to use the application (e.g., 'complicated', 'needs practice', 'steep learning curve'). One user suggested to use textual hints or a virtual example. Another user suggested to use the wireframe of a 3D object for alignment or to augment a half-dome on top of the element. Two users mentioned issues with deciding on the validity of a hologram (e.g., 'not clear, when patches are different', 'different colors are irritating').

The modified selection of reference views leads to a reasonable checking time of 15 s when using the system. Due to the fact that three reference views were used instead of two, this is an encouraging result, which confirms that the actual selection is critical to the efficiency of the process.

While participants took another 20 s for coming up with their own decisions based on visual inspection of the recorded data, this is not relevant in our case, since the system always decided correctly. The significant difference in user decision time between Hologram 1-o and hologram 1-f (substitute) is very likely due to a larger visual difference for this pair regarding the original and the substitute. The lower accuracy achieved by the users (85.94%) gives room for speculation that laypeople cannot intuitively assess the evaluated holograms for themselves. Users in particular had issues assessing hologram 2-f correctly, which is a rotated version of the original element. They were also rather unsure about the validity of Hologram 1-o, where the patterns are subject to a larger amount of color noise. Consequently, it seems reasonable to not bother laypeople with the summary of recorded images, except in ambiguous cases.

Several participants pointed out, that the interface was complicated to use. This is due to the complexity of the task, which requires simultaneous monitoring of several parameters and rather fine-grained navigation. This could be improved by using textual instructions or animations throughout the process.

It must be noted that the aforementioned selection of reference views, although natural for the user and beneficial regarding efficiency, may not be possible for arbitrary security elements. The reason is that the complete set of reference patterns does not necessarily become visible when recording with a flash-enabled mobile device and following the suggested path for orientation change (i.e., tilting downwards). Consequently, there is a need for specially designed security elements, which allow the aforementioned selection of viewing directions. This can be considered a realistic demand, since there are already elements on the market which approximately feature this property.

6 CONCLUSION

Mobile AR systems running on off-the-shelf hardware can serve as tools for the verification of holograms by laypeople.

However, previous approaches suffered from high temporal effort and limited accuracy. In order to tackle these issues, we conducted an experiment on matching hologram patches in order to find a suitable similarity measure and modified the spatial distribution of reference views in order to mimic the typical behavior of users observed during document verification. The latter leads to a more user-friendly parametrization, defining a constrained navigation space within the original approach.

A subsequent user study using original and substitute holograms turned out that hologram capture can be done in approx. 15 s, where an automatic decision by the system follows immediately. Consequently, the distribution of reference views is critical for the efficiency of the process. Contrary to decisions on validity made by the users, the system proved to be correct in all cases. From the results obtained in the original study and the evaluation of the improved prototype, it is evident that security elements should be designed with mobile verification by human operators in mind. The results obtained suggest that this would allow a very efficient check of security elements using off-the-shelf mobile devices, while no major changes in the basic production process are required.

Ideally, the type of device should be the same for capturing reference information and on-line verification. Our experiments revealed that for several devices, verification is still possible under optimal office conditions. While further invariance could be handled by using a machine-learning based approach for the comparison of patches, a reasonably large amount of training data is currently not available. Consequently, it lies in the responsibility of the actual implementation to detect the type of device and to retrieve the corresponding data for optimal matching performance.

ACKNOWLEDGMENTS

This work is co-funded by Bundesdruckerei GmbH. Thanks to Christian Reinbacher who provided the scripts and tools for using the robot in sampling the hologram appearance.

REFERENCES

- [1] S. Bae, A. Agarwala, and F. Durand. Computational rephotography. *ACM Trans. Graph.*, 29(3):24:1–24:15, July 2010.
- [2] F. Biocca, A. Tang, C. Owen, and F. Xiao. Attention funnel: Omnidirectional 3d cursor for mobile augmented reality platforms. In *SIGCHI*, pages 1115–1122, 2006.
- [3] C. Buraga-Lefebvre, S. Cotmellec, D. Lebrun, and C. Özkul. Application of wavelet transform to hologram analysis: three-dimensional location of particles. *Optics and Lasers in Engineering*, 33(6):409–421, 2000.
- [4] G.-H. Chen, C.-L. Yang, L.-M. Po, and S.-L. Xie. Edge-based structural similarity for image quality assessment. In *IEEE Int. Conf. on Acoustics, Speech and Signal Proc. (ICASSP)*, volume 2, 2006.
- [5] K. Chintamani, A. Cao, R. Ellis, and A. Pandya. Improved telemanipulator navigation during display-control misalignments using augmented reality cues. *Systems, Man and Cybernetics, Part A: Systems and Humans*, 40(1):29–39, 2010.
- [6] Y. Dong, J. Wang, X. Tong, J. Snyder, Y. Lan, M. Ben-Ezra, and B. Guo. Manifold bootstrapping for svbrdf capture. In *ACM SIGGRAPH 2010 papers*, pages 98:1–98:10, 2010.
- [7] M.-P. Dubuisson and A. Jain. A modified hausdorff distance for object matching. In *IAPR Int. Conf on Pattern Recognition and Image Processing*, volume 1, pages 566–568, Oct 1994.
- [8] E.C.B. Biannual inform. on euro banknote counterfeiting, Jan. 2013.

- [9] A. Hartl, C. Arth, and D. Schmalstieg. Ar-based hologram detection on security documents using a mobile phone. In *Bebis, George et al., editor, Advances in Visual Computing*, volume 8888 of *LNCS*, pages 335–346. Springer International Publishing, 2014.
- [10] A. Hartl, J. Grubert, C. Reinbacher, C. Arth, and D. Schmalstieg. Mobile user interfaces for efficient verification of holograms. In *Proceedings of the IEEE Virtual Reality Annual International Symposium (VR)*, 2015.
- [11] A. Hartl, J. Grubert, D. Schmalstieg, and G. Reitmayr. Mobile interactive hologram verification. In *Proceedings of IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 75–82, 2013.
- [12] A. Hartl and G. Reitmayr. Rectangular target extraction for mobile augmented reality applications. In *Proceedings of International Conference on Pattern Recognition (ICPR)*, pages 81–84, 2012.
- [13] A. Hartl, D. Schmalstieg, and G. Reitmayr. Client-side mobile visual search. In *Proceedings of the International Conference on Computer Vision Theory and Applications (VISAPP)*, pages 125–132, 2014.
- [14] S. Heger, F. Portheine, J. A. K. Ohnsorge, E. Schkommodau, and K. Radermacher. User-interactive registration of bone with a-mode ultrasound. *Engineering in Medicine and Biology Magazine, IEEE*, 24(2):85–95, 2005.
- [15] H. Hirschmuller and D. Scharstein. Evaluation of stereo matching costs on images with radiometric differences. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 31(9):1582–1599, Sept 2009.
- [16] J. Jachnik, R. A. Newcombe, and A. J. Davison. Real-time surface light-field capture for augmentation of planar specular surfaces. In *Proceedings of IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 91–97, 2012.
- [17] J. Janucki and J. Owsik. A wiener filter based correlation method intended to evaluate effectiveness of holographic security devices. *Optics Communications*, 218(4-6):221–228, 2003.
- [18] H.-J. Kwon and T.-H. Park. An automatic inspection system for hologram with multiple patterns. In *SICE*, pages 2663–2666, 2007.
- [19] T.-H. Park and H.-J. Kwon. Vision inspection system for holograms with mixed patterns. In *IEEE Conference on Automation Science and Engineering (CASE)*, pages 563–567, 2010.
- [20] T. Pitkaho and T. J. Naughton. Stereo vision based approach for extracting features from digital holograms. In W. Osten and M. Kujawinska, editors, *Fringe*, pages 1–6. Springer Berlin Heidelberg, 2009.
- [21] A. Pramila, A. Keskinarkaus, E. Rahtu, and T. Seppnen. Watermark recovery from a dual layer hologram with a digital camera. In A. Heyden and F. Kahl, editors, *Proceedings of the Scandinavian Conference on Image Analysis (SCIA)*, pages 146–155. Springer, 2011.
- [22] P. Ren, J. Wang, J. Snyder, X. Tong, and B. Guo. Pocket reflectometry. In *Proceedings of the International Conference on Computer graphics and interactive techniques (SIGGRAPH)*, pages 45:1–45:10, New York, USA, 2011. ACM.
- [23] J. Shingu, E. Rieffel, D. Kimber, J. Vaughan, P. Qvarfordt, and K. Tuite. Camera pose navigation using augmented reality. In *Proceedings of IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 271–272, 2010.
- [24] N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: Exploring photo collections in 3d. *ACM Trans. Graph.*, 25(3):835–846, July 2006.
- [25] D. Soukup, S. Štolc, and R. Huber-Mörk. Analysis of optically variable devices using a photometric light-field approach. In *Proc. SPIE 9409, Media Watermarking, Security, and Forensics*, 2015.
- [26] M. Sukan, C. Elvezio, O. Oda, S. Feiner, and B. Tversky. Parafustrum: Visualization techniques for guiding a user to a constrained set of viewing positions and orientations. In *UIST*, pages 331–340, New York, USA, 2014. ACM.
- [27] S.-Y. Sun, M. W. Gilbertson, and B. W. Anthony. Computer-guided ultrasound probe realignment by optical tracking. In *IEEE Int. Symposium on Biomedical Imaging (ISBI)*, pages 21–24, 2013.
- [28] R. L. van Renesse. *Optical Document Security*. Artech House, third edition, 2005.
- [29] 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 (TVCG)*, 16(3):355–368, 2010.
- [30] H. R. S. Zhou Wang, A.C. Bovik and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 13(4):600–612, April 2004.



Andreas Daniel Hartl is a postdoctoral researcher at the Institute for Computer Graphics and Vision, Graz University of Technology, from which he received the Bakkalaureus, Dipl.-Ing. (with distinction) and Dr. techn. (with distinction) in 2007, 2010 and 2015 respectively. He is author and co-author of several peer-reviewed publications at ISMAR, VR, ICPR, ISVC, VISAPP and CVPRW including relevant patents. He is also a reviewer for ISMAR, ACE and MUM. Andreas gave live demos at ISMAR and VR, demonstrating recent research results on off-the-shelf mobile devices. Initially focusing mainly on object detection, recognition and reconstruction, his current interests as a researcher include the inspection of security documents by means of Mobile Augmented Reality.



Clemens Arth is senior research scientist at the VRVis and at Graz University of Technology, and he is deputy director of the Christian Doppler Laboratory on Handheld Augmented Reality. His main research interests are Computer Vision algorithms, Augmented Reality technology and Machine Learning. The current focus of his work is on accurate global localization of mobile devices for AR applications. In the past, he worked for several companies, such as Frequentis or Zydacron GmbH. He received Dipl.-Ing. (2004) and Dr. techn. (2008) degrees from Graz University of Technology. He was an organizer of the AR Supermodels workshop at IEEE ISMAR in 2010, and of the CDL Workshop on Tracking Technology for AR in 2014. He is author and co-author of numerous peer-reviewed publications at ISMAR, VISAPP, CVPR, ACCV, ICPR and VR. He is also reviewer for several journals and international conferences, such as RTIP, TVCG, VR, C&G, ACE and MVA.



Jens Grubert is a postdoctoral researcher at the Chair of Embedded Systems, University of Passau, Germany. Previously, he held positions at Graz University of Technology and University of Applied Sciences Salzburg, Austria and Fraunhofer Institute for Factory Operation and Automation IFF, Germany. He received his Dr. techn (2015) with highest distinction at Graz University of Technology, his Dipl.-Ing. (2009) and Bakkalaureus (2008) with highest distinction at Otto-von-Guericke University Magdeburg, Germany. He is author of more than 30 peer reviewed publications and patents and published a book about AR development for Android. His current research interests include interaction with body proximate display ecologies, mobile augmented reality, around-device interaction, multi-display environments and cross-media interaction.



Dieter Schmalstieg is full professor and head of the Institute for Computer Graphics and Vision at Graz University of Technology, Austria. His current research interests are augmented reality, virtual reality, computer graphics, visualization and human-computer interaction. He received Dipl.-Ing. (1993), Dr. techn. (1997) and Habilitation (2001) degrees from Vienna University of Technology. He is author and co-author of over 300 peer-reviewed scientific publications, with over ten best paper awards and nominations. His organizational roles include associate editor in chief of IEEE Transactions on Visualization and Computer Graphics, member of the editorial advisory board of computers & graphics and of the Springer Virtual Reality journal, member of the steering committee of the IEEE International Symposium on Mixed and Augmented Reality, chair of the EUROGRAPHICS working group on Virtual Environments (1999-2010), key researcher of the K-Plus Competence Center for Virtual Reality and Visualization in Vienna and key researcher of the Know-Center in Graz. In 2002, he received the START career award presented by the Austrian Science Fund. In 2012, he received the IEEE Virtual Reality technical achievement award for seminal contributions to the field of Augmented Reality. He was elected as a senior member of IEEE, as a member of the Austrian Academy of Sciences and as a member of the Academia Europaea. Since 2008, he is also director of the Christian Doppler Laboratory for Handheld Augmented Reality.