

Spatial Gaze Markers: Supporting Effective Task Switching in Augmented Reality

Mathias N. Lystbæk
Aarhus University
Denmark
mathiasl@cs.au.dk

Ken Pfeuffer
Aarhus University
Denmark
ken@cs.au.dk

Tobias Langlotz
University of Otago
Dunedin, New Zealand
tobias.langlotz@otago.ac.nz

Jens Emil Grønbæk
The University of Melbourne
Australia
Aarhus University
Denmark
jensemil@cs.au.dk

Hans Gellersen
Lancaster University
United Kingdom
Aarhus University
Denmark
h.gellersen@lancaster.ac.uk

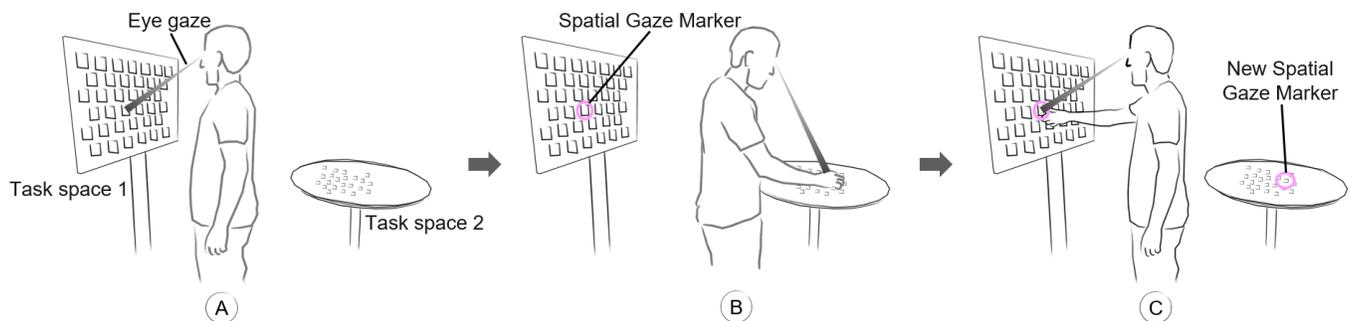


Figure 1: We propose Spatial Gaze Markers as implicit spatial reminders in augmented reality. The system tracks user attention (A) and detects when the user shifts their attention from a current task space (e.g., a fuse board) and places a visual marker in the last gaze position (B). When the user returns their attention, the marker guides them to where they had left off (C). Spatial Gaze Markers are application-independent and can support task switching over different timescales.

ABSTRACT

Task switching can occur frequently in daily routines with physical activity. In this paper, we introduce Spatial Gaze Markers, an augmented reality tool to support users in immediately returning to the last point of interest after an attention shift. The tool is task-agnostic, using only eye-tracking information to infer distinct points of visual attention and to mark the corresponding area in the physical environment. We present a user study that evaluates the effectiveness of Spatial Gaze Markers in simulated physical repair and inspection tasks against a no-marker baseline. The results give insights into how Spatial Gaze Markers affect user performance, task load, and experience of users with varying levels of task type and distractions. Our work is relevant to assist physical workers with simple AR techniques and render task switching faster with less effort.

CCS CONCEPTS

• **Human-centered computing** → **Mixed / augmented reality.**

KEYWORDS

augmented reality, eye-tracking, gaze interaction, task switching, attention switching

ACM Reference Format:

Mathias N. Lystbæk, Ken Pfeuffer, Tobias Langlotz, Jens Emil Grønbæk, and Hans Gellersen. 2024. Spatial Gaze Markers: Supporting Effective Task Switching in Augmented Reality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3613904.3642811>

1 INTRODUCTION

Spatial memory plays a crucial role in human cognition and our ability to switch between tasks [32]. When we remember where objects of interest are located, we are able to turn to them efficiently without visual search. However, a user's capacity to remember spatial locations can be interrupted in many ways – many similar-looking objects in a crowded environment, the complexity of multi-tasking in a given task, or encountering distractions along the way can all shake up our memory of objects in space.

As a novel approach, we propose *Spatial Gaze Markers* (SGM) as implicitly generated spatial reminders in augmented reality (AR). SGM is conceived as a system that operates in the background of the user's activity without any knowledge of the tasks in which they engage. The system leverages eye-tracking to detect attention shifts when a prior gaze fixation is no longer in the user's field of view (FOV) and generates a visual marker in its place (Figure 1A). This goes unnoticed by the user until they return their attention to the task space, where the marker will guide them back to exactly where they had last attended (Figure 1B). Once a marker has been noticed, it will be removed. As a result, SGM is subtle but effective in supporting task resumption.

Prior work has shown gaze cues to be effective for task switching between screens [25]. SGM extends the concept to leaving cues in the real world, based on information about the physical environment captured through a depth camera of the AR head-worn display. The system is designed to be easily deployed and does not rely on any knowledge of user activity or objects they interact with. Attention shifts are inferred from relative eye and head movements without task knowledge, and markers are placed based on gaze and depth of the AR scene. SGM lends itself to be task-agnostic, for support of any situation in which tasks are resumed after an attention switch – from casual activity left behind and returned to later to intensive multi-tasking.

For evaluation of SGM, we developed a real-world simulation of repair and inspection tasks. The tasks require the user to inspect a workspace, switch to a tool space for any simulated fault discovered, and then return to the workspace. This simulates a situation in which attention needs to be shifted between different spaces and where the task is demanding as it contains a larger number of objects that look the same until inspected closely. Based on these tasks, we conducted an experiment with 20 participants, comparing SGM against a baseline of no-marker. We found SGM reducing task time and task load when users return to the workspace, which shows that low- and high-demand tasks can benefit from the use of simple spatial reminders.

In summary, our contributions include (1) SGM as a technique to assist workers via AR by placing visual markers in the real world as reminders and to lower visual search demands; (2) evaluation tasks allowing investigating user performance with physical work, with controlled factors of visual search complexity and task distractions, useful to study task-switching in physical work assisted by AR; (3) insights on the effects of SGM on task distractions and spatial memory build up over time, and when it is beneficial regarding speed, error, and perceived usability.

2 RELATED WORK

How we organise, coordinate, and switch tasks in daily life is related to human action and thought [30]. Task transitions and multi-tasking involve complex mental processes, and studies have shown effects with regard to task and location factors [24], historical information [46], and interference phenomena [26]. Task switches occur frequently throughout daily activity but are prone to fail users when users encounter distractions [26]. The importance of spatial memory for interaction across tasks is long recognised in HCI [37]. Cockburn and McKenzie found that task performance

degrades in 3D when compared to 2D tasks [6, 7]. 3D tasks are more demanding than 2D tasks because of the additional dimension, vast space, and the highly situation-dependent setting of objects and spaces in the physical environment that can affect spatial memory. In this paper, we examine SGM as a generic approach to provide users with “spatial reminders” of where they had left a task.

Task-switching and task resumption have been supported, for instance, with visual links across applications [45] and replay of interactions [29], however, our work is specifically inspired by the use of gaze cues as an implicit support mechanism [38]. Kern et al. introduced Gazemarks to support attention switches between multiple screens and demonstrated reduced search time on one display after being presented with a disruption on the other [25]. EyeBookmark supports readers in recovering from interruptions by highlighting the last known gaze position in the text [23], and SmoothGaze integrates similar support for gaze-assisted task resumption [5]. Gaze cues have also been found effective for navigation in digital maps to counter spatial context lost during zooming [15]. In this work, we build on the principal idea of Gazemarks but extend it from screen-based interaction to interaction in the real world, mediated by AR.

Many of the real-world tasks that AR aims to support require users to shift attention between different spaces in their environment. Among the most widely studied scenarios is order picking, a prototypical task for visual guidance in AR [10, 21, 39, 40]. The task requires attention shifts between a central work area and other spaces from which parts are collected [11, 13, 17]. Numerous works have explored different visual cues to guide users to the parts they need to locate [35, 39, 40]. Other work has considered saliency modulation as an alternative to graphical cues [1, 42, 43]. These approaches have in common that they rely on detailed knowledge of tasks, objects, and their position in the environment. SGM presents a principally different approach as it is more flexible for various use cases. As such, it can support any situation in which a task is being resumed, but limited to generic provision of a gaze marker in the position last attended before a shift from the task occurred.

The tasks we designed for evaluation of our approach resemble order-picking in that they require alternation between a central workspace and an area from which to pick tools. However, our tasks are distinct as they simulate inspection and repair scenarios in which users need to check parts in the workspace and find matching tools. In contrast to order picking, this places particular demand on spatial memory, with a main workspace that contains similar objects while the intermediate task induces cognitive demand to match objects.

A wide range of other work relates to ours in using gaze to support interaction in AR [33]. Gaze is, for instance, adopted to provide hands-free control [28], assist manual input [31, 44], infer attention [41], and adapt interfaces [14]. Gaze has also been explored for collaboration tasks in AR, where gaze is visualised to provide collaborators with an indication of where others are attending [18, 22]. Our work also uses gaze visualisation, however, as implicit cues to oneself, sampled when significant shifts in attention are detected.

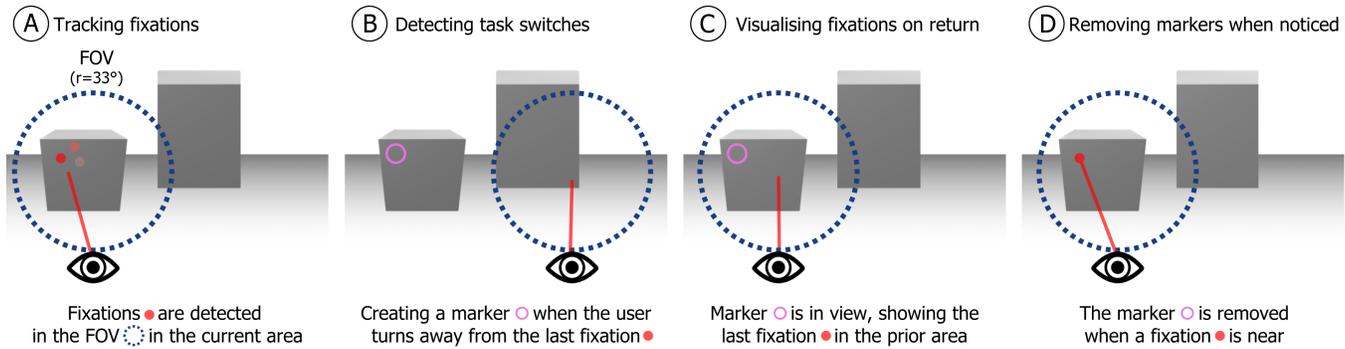


Figure 2: Spatial Gaze Markers are an AR tool that aids return to an earlier task implicitly based on eye-tracking and transient gaze cues without requiring any task or activity knowledge. Fixations are detected in the user’s FOV (A) and are used to detect the user turning away when the last fixation falls out of the FOV and spawns a visual marker in the last fixated position (B). When the user returns their attention to the task space, they are guided by the marker to where they had last attended (C). The system detects when the marker has been noticed, upon which it is removed (D).

3 SPATIAL GAZE MARKERS

The concept of Spatial Gaze Markers is straightforward: whenever attention is removed from a task, a virtual marker is left in its place and remains in place to aid the return to the task. As our approach is task-agnostic, it has no notion of users’ actual tasks and places markers when large gaze shifts occur, that take the user attention away from a space in which they previously focused. The system has no notion of whether the user intends to return but a marker is left in place, to be of potential use.

The system requirements for SGM are an AR head-mounted display (HMD), eye-tracking, and 3D environment tracking. These requirements are, for instance, met by the Microsoft HoloLens 2, which we used for the implementation of our concept. The HoloLens 2 provides optical see-through AR with a $43^\circ \times 29^\circ$ FOV, up to 90Hz eye-tracking, and a real-time tracked 3D environment mesh. The system has four components:

- (1) Fixation detection: the system continuously tracks eye movement to segment gaze fixations.
- (2) Attention shift detection: large shifts in gaze fixation are detected as attention shifts.
- (3) Marker visualisation: the last gaze fixation before the shift is visualised as a marker.
- (4) Marker removal: the marker is removed when it has been noticed by the user.

Fixation detection. The HoloLens 2 provides eye-tracking but does not segment fixations. We use a real-time version of the I-DT algorithm (Algorithm 1) to capture fixations within the FOV (Figure 2A), taking newly captured gaze rays and returning 3D fixation points based on a fixation angle and fixation duration. We set the fixation angle to 1.5° to account for eye-tracking inaccuracy and fixation duration to 250ms to account for natural fixation time and tracking latency [36].

Attention shift detection. We detect when users turn from a task area by checking whether the most recent fixation falls outside of a specified angle from the centre of the HMDs FOV, thereby signifying a larger head movement (Figure 2B). In our prototype,

we set the angle to 33° (Algorithm 2). Given the limited FOV of the HoloLens 2 HMD, this implies an attention shift is only detected when the last point of focus is no longer in the display’s FOV

Marker visualisation. When the user returns their attention to a previous task area, they will be presented with a visualisation of where they last fixated (Figure 2B-C). The gaze marker is presented as a transparent circle outline with a static size of 1.5cm radius. The outline is rendered transparent (12.5%), magenta (FF00FF), and a border thickness of 4mm. Although we considered other visualisations in pilot tests, such as crosshairs and pins, we found that the current design is sufficient to provide a good balance between visual saliency and little obstruction of the scene, is suitable for the purpose of our proof-of-concept implementation, and builds upon findings of prior work on AR guidance cues [39].

Marker removal. We remove markers when a new fixation is detected within 1.5° visual degrees (Figure 2D). This is suitable for a task model where the marker has fulfilled its purpose, i.e., guiding the user back, and then having no need anymore.

4 EVALUATION TASKS

The following describes the design and procedures of two tasks that we developed for empirical evaluation of SGM. We envision that these tasks have the potential for wider use for evaluating other techniques and systems in contexts of AR guidance in physical environments.

4.1 Rationale on Experimental Design Decisions

The challenge is to balance controlled study design (internal validity) and realistic insights (external validity). Focusing on a three-fold task operation – (1) locating the work area and target, (2) turning to find the required tool in the tools area, and (3) returning to the work area for tool use – we emphasized the return phase due to its potential impact on efficiency and task load.

Another crucial challenge was creating balanced task difficulty, requiring tasks to be sufficiently difficult yet still achievable and repeatable. We focused on remembering object locations in a visually consistent grid to introduce visual search difficulty. We constructed

Algorithm 1: Real-time I-DT

```

1 Function GetFixationPoint(NewGP):
2   WindowSize  $\leftarrow \lceil \frac{\text{FixationDuration}}{\text{Time.deltaTime}} \rceil$ 
3   if  $|GPQueue| \leq \text{WindowSize}$  then
4     push NewGP to GPQueue
5     return null
6   end
7   while  $|GPQueue| > \text{WindowSize}$  do
8     pop from GPQueue
9   end
10
11  TmpCtrd  $\leftarrow$  GetCentroid(GPQueue)
12  Dispersion  $\leftarrow$  maxNewGP VisualAngle(NewGP, TmpCtrd)
13  if Dispersion > FixationAngle then
14    clear GPQueue
15    return null
16  end
17  return GetCentroid(GPQueue)
18 End Function
19
20 Function VisualAngle(p1, p2):
21  d1  $\leftarrow$  direction from eyes to p1
22  d2  $\leftarrow$  direction from eyes to p2
23  return Vector3.Angle(d1, d2)
24 End Function

```

Algorithm 2: Spatial Gaze Markers

```

1 while True do
2   HeadDir  $\leftarrow$  head direction
3   PFDir  $\leftarrow$  direction from eyes to PreviousFP
4   if Vector3.Angle(HeadDir, PreviousFP) > 33° then
5     SGM  $\leftarrow$  spawn new marker at PreviousFP
6     add SGM to SGMList
7     continue
8   end
9
10  GazePoint  $\leftarrow$  GetGazeRayCollision(GazeRay, Collider)
11  FP  $\leftarrow$  GetFixationPoint(GazePoint)
12  if FP is null then
13    continue
14  end
15
16  foreach SGM in SGMList do
17    if VisualAngle(GazePoint, FP) > 1.5° then
18      Destroy(SGM)
19    end
20  end
21
22  PreviousFP  $\leftarrow$  FP
23
24 end

```

an inspection/repair-like setup using Post-its to mark points in the work area where work could be done, while the tools area contains the tools required to complete the task.

Making the work area visually complex required some iteration; with too few points, the task is trivial, and with too many, the task becomes too hard and takes too long. We designed the grid with alternating, offset rows with gaps to not overwhelm participants (Figure 3). The grid uniformity is likely the main factor making the task more visually complex, while the offset would make it harder to remember locations.

4.2 Task Design

We designed two simulated physical AR tasks to emulate repair and inspection with switches between a work area and a tools area. Both tasks use the same work and tools areas, with slight variations in procedure. The work area is a 75" TV screen with a grid of 105 Post-its, each covering a rendered shape. The screen controls the physical elements of the task and is separate from any guidance. The screen controls the physical elements of the task and is separate from any guidance. Users flip up Post-its to reveal shapes, categorised as "fine" (85 black shapes) or "faulty" (20 red shapes). The covered shapes are randomly arranged by utilising seeded randomness on the participant ID, condition number, and task type. There are four different shape types, each with a randomly selected two-character string, written in the centre, from a set of 10 options ($4 \cdot 10 = 40$ unique "tools"). The tools area is a small table (108cm tall and 90cm diameter), positioned 190cm in front of the work area, with the participant in between. One of each "fine" shape was printed out, placed in the tools area, and manually shuffled.

Following the task procedures, outlined in Figure 3, participants must: (1) look at the work area (either with target highlighting or manual searching for the next target), (2) find the faulty shape and remember what and where it is, (3) locate and pick up the matching "fine" shape in the tools area, (4) return to the work area and flip up the correct Post-it, and (5) place the matching "fine" shape on the "faulty" shape, thereby marking the once "faulty" shape location as "fine". Ideally, participants immediately remember and inspect the correct Post-it, otherwise, we note a Recheck and let the participant check other Post-its until either the correct Post-it is checked or they recheck more than 5 Post-its, marking the trial as an error.

Errors trigger additional "faulty" shape repairs, ensuring at least 8 successful trials.

4.2.1 Task 1: Repair. In *Repair*, the next Post-it to check is highlighted with an AR border to emulate a broken part and for controlling the study and disappears 2 seconds after initially being seen. While highlighting may seem unrealistic and could narrow the external validity, one could see this task as resembling an "immediate repair" scenario, where something stands out and must be fixed quickly. The *Repair* task is designed to emulate tasks with limited spatial memory, as such we limited targets to only appear within two Post-its from the edges, limiting reliance on edge references. Fixing parts that "stand out" may limit the development of spatial memory, leading to longer time spent searching for the previous location.

In addition, this task incorporates the factor of task distractions, allowing to study its impact on user performance under different AR

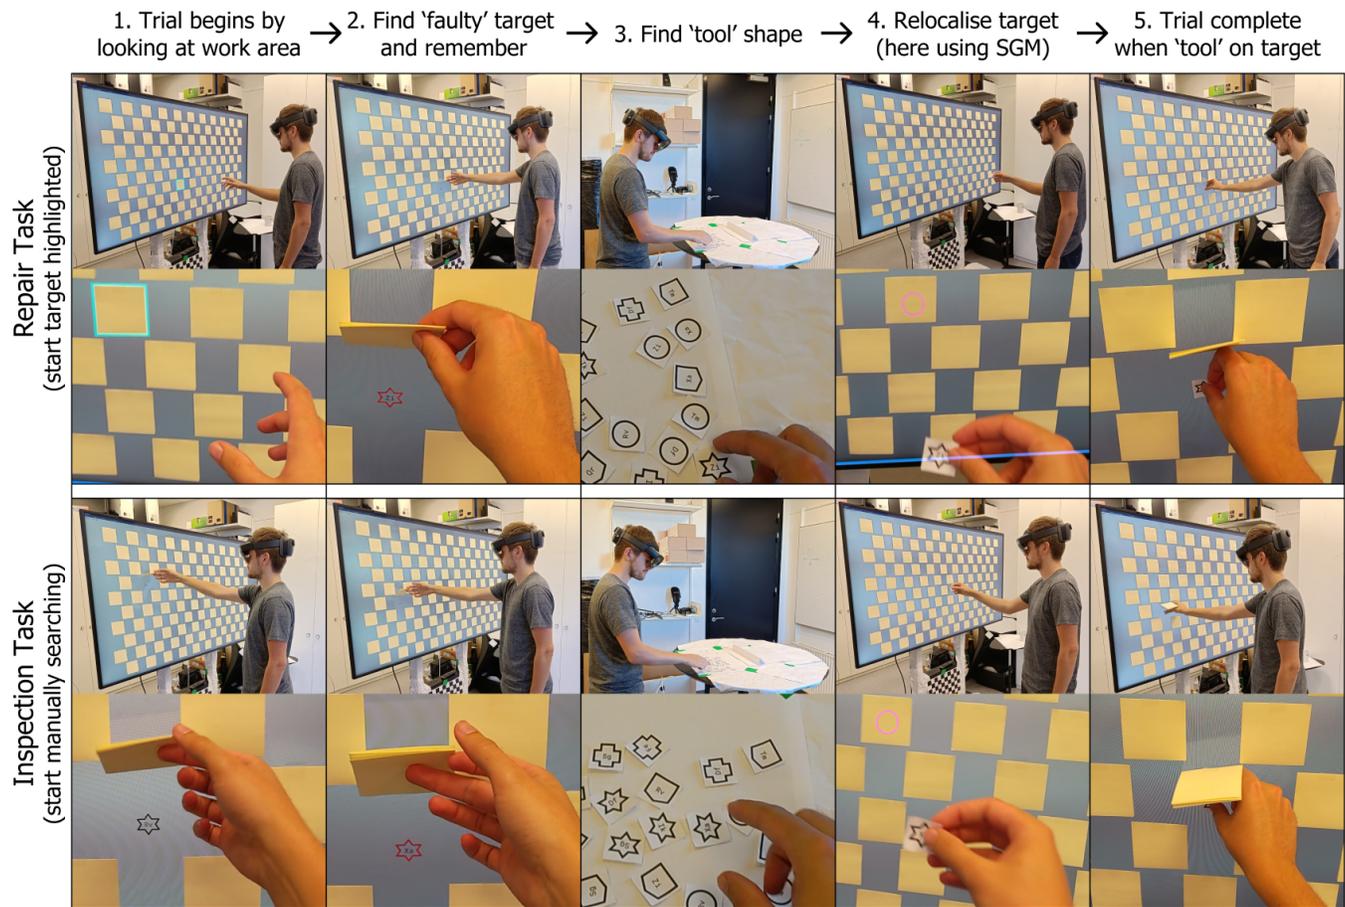


Figure 3: Illustration of the task procedures (leaving out the distraction part of the *Repair* task). Note that visuals have been slightly enhanced for the figure.

guidance concepts. It remains an open question whether SGM will prove beneficial in scenarios involving frequent task interruptions. Mathematical questions are employed as distractions, following a precedent in previous studies that effectively controlled cognitive load during experiments [9]. In our task sequence, distraction can be introduced after the participant has found the tool, and just before they turn around to the work area. While facing the tools area, participants subtract 17 from a random four-digit number three times and complete the steps out loud, ensuring that the calculations were performed. This distraction, in theory, impairs spatial memory due to mental capacity being occupied by the mathematics task.

4.2.2 Task 2: Inspection. In *Inspection*, Post-its are not highlighted, instead, participants manually examine Post-its to find “faulty” shapes, after which the procedure is the same. Inspection is a more flexible, continuous process where participants can start anywhere in the work area. One can think of it as a search where the location of issues is unknown. Participants move from one Post-it to the next until a “faulty” shape is found, potentially building spatial memory. For example, starting from the left and progressing systematically

can help participants recall inspected areas, aiding in locating the correct position in the work area.

5 USER STUDY

We designed a user study aiming to evaluate the usability and performance of the SGM concept in two abstract inspection-like tasks. We compare to a no-marker condition and involve task variations to assess how task distractions can further affect user performance.

5.1 Study Design

The user study is split into two tasks, both designed as within-subject. The first task, *Repair*, has two independent variables, namely *Technique* and *Distraction*, counterbalanced using a Latin Square [3] in two steps, first on *Technique* and then on *Distraction*. We compare having SGM versus a baseline of No Marker and having the user complete the task with No distraction versus being Distracted by a secondary task before finishing (Section 4.2.1). This results in four conditions for the *Repair* task: (1) No Marker+No Distraction (NM+ND), (2) SGM+No Distraction (SGM+ND), (3) No Marker+Distraction (NM+D), (4) SGM+Distraction (SGM+D). The second task of the study, *Inspection*, has one independent

variable, namely *Technique* (NM vs. SGM), as in the *Repair* task, counterbalanced using a Latin Square [3]. In sum, $2 \text{ Techniques} \times 2 \text{ Distractions} \times 8 \text{ trials} + 2 \text{ Techniques} \times 8 \text{ trials} = 48$ successful trials per participant. We had 20 participants complete our study, totalling 960 successful trials.

5.2 Apparatus and Implementation

The study is implemented with Unity2021 for the Microsoft HoloLens 2 using the Mixed Reality Toolkit v2.8.2. The eye-tracking refresh rate of the HoloLens 2 was extended from 30 Hz to 90 Hz using the Extended Eye Tracking API¹. Another instance of the study ran on a Windows 10 laptop outputting video to the big TV screen (3180 × 2160), also implemented in Unity2021, controlling the physical part of the study, namely which shapes to show behind the Post-its and whether the shapes are “fine” or “faulty”. Both instances of the study use the same seeded randomness; both the HoloLens 2 variant and the screen variant of the studies have the exact same shapes. Although, the shapes on the HoloLens 2 are not rendered, only the target highlighting in *Repair*. Additionally, to capture where participants look inside the two areas, the study requires a pre-setup phase on the HoloLens 2 where the experimenter marks the work and tools areas. These planes are only necessary for capturing the participants’ gaze for logging and conducting the study and are by no means integral to the function of SGM, which can work on any 3D mesh.

5.3 Procedure

Participants received a study briefing, completed consent and demographics forms, and watched a video of the *Repair* task with SGM, ensuring that participants understood the study elements before starting. Participants then wore the HoloLens 2 and underwent eye-tracking calibration.

Starting with *Repair*, participants had three training trials following the counterbalancing and were instructed to be as fast and as accurate as possible. Subsequently, participants completed 8 successful study trials for the current *Repair* condition, followed by a post-condition questionnaire consisting of a NASA-TLX [20] with one additional Eye Demand question. After each condition, participants waited until the screen was cleared of stuck-on shapes before proceeding to the next condition according to the counterbalancing. Additional training sessions were completed when new elements of the study design were reached.

Ending with *Inspection*, participants had three training trials, followed by 8 study trials according to the counterbalancing, and completed the same post-condition questionnaire, just as for *Repair*. Then, after completing both *Inspection* conditions, the participants completed a final post-study questionnaire to gauge how the participants felt about the system and the setup. The study lasted on average around 60 minutes.

5.4 Evaluation Metrics

For dependent variables, we include the following measures. *Relocalisation Time* measures how much time it takes for the participant to find the correct Post-it after returning to the work area, counting

from the first time gaze hits the work area after returning from the tools area until the correct tool is placed under the target Post-it. *Relocalisation Time* is thereby not directly affected by the additional distraction task in *Repair*, but only indirectly affected by how the distraction task impacts the participants’ memory. *Rechecks* count how many Post-its the participant had to check before finding the correct Post-it, manually noted by the study conductor. An *Error Rate* was calculated as the fraction of error trials out of total trials, with trials marked as an error if the participant checked more than 5 Post-its or picked an incorrect printed-out shape. To get insight into the participants’ task load we used the *NASA-TLX* in a 7-point Likert scale variant² [8, 16, 34] conducted as Raw TLX [4, 19] answered immediately after each condition, and *User Feedback* on SGM was gathered in a post-study questionnaire after all six conditions were completed.

5.5 Participants

We recruited 20 participants (13 male, 6 female, and 1 other³) from the local university, consisting mainly of Computer Science researchers and Master’s students. Participants’ age ranged from 22 to 37 ($M = 26.86$, $SD = 3.73$). On a scale between 1 (low) and 5 (high), participants rated themselves as having average experience with VR/AR ($M = 3.48$, $SD = 1.29$) and eye-gaze ($M = 2.38$, $SD = 1.32$).

5.6 Results

After outlier removal ($\approx 3.8\%$), the data contained non-normally distributed data (as reported by Shapiro-Wilk tests). Therefore, we ran a series of Friedman tests with posthoc Bonferroni corrected Wilcoxon signed-rank tests on both logged objective data and subjective survey data. Statistical significance is shown in graphs as * for $p < .05$, ** for $p < .01$, and *** for $p < .001$. Note that significance tests were not carried out between *Repair* and *Inspection*.

This section focuses on Relocalisation Time, Rechecks, Error Rate, NASA-TLX, and User Feedback. Full analysis results on logged data can be found in supplementary materials, including results on additional measures that did not add to the discussion, such as Task Completion Time (i.e. the time from the start of the trial until the participant completes the trial). However, as reference points for Relocalisation Time, the average Task Completion Time for each condition was, in the *Repair* task, 22.84s for NM+ND, 16.63s for SGM+ND, 50.29s for NM+D, and 38.07s for SGM+D. In the *Inspection* task, NM took 19.59s and SGM took 18.30s.

5.6.1 Relocalisation Time (Figure 4a). We found significant differences in Relocalisation Time in the *Repair* task ($\chi^2(3) = 53.1$, $p < 0.001$), both SGM+ND ($M = 3.01$, $SD = 0.425$) and SGM+D ($M = 3.39$, $SD = 0.683$) conditions were faster than both NM+ND ($M = 5.14$, $SD = 1.508$) and NM+D ($M = 7.13$, $SD = 2.33$) conditions (all $p < 0.001$), NM+ND was faster than NM+D ($p < 0.001$), and SGM+ND was faster than SGM+D ($p = 0.042$). In the *Inspection* task, no significance was indicated ($\chi^2(1) = 3.2$, $p = 0.074$).

5.6.2 Rechecks (Figure 4b). In terms of Rechecks, we found significant differences in the *Repair* task ($\chi^2(3) = 31.194$, $p < 0.001$); both SGM+ND ($M = 0.1$, $SD = 0.447$) and SGM+D ($M = 0.65$, $SD =$

¹<https://learn.microsoft.com/en-us/windows/mixed-reality/develop/unity/extended-eye-tracking-unity> - last accessed 17th of November 2023.

²7 points (1-7) neatly divide the 21 original points (0-20) and is easier for participants.

³Anonymized.

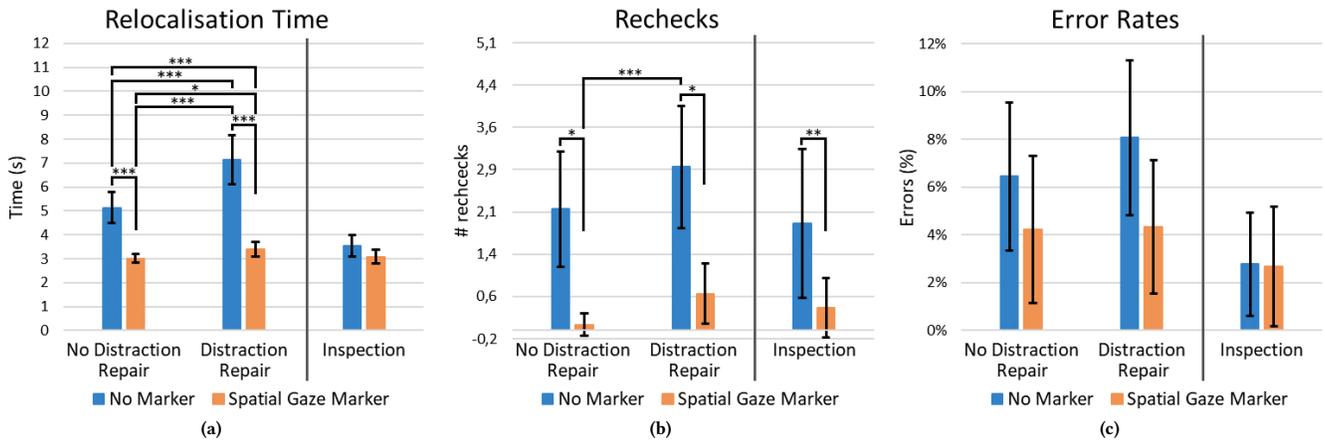


Figure 4: Results on Relocalisation Time (a), Rechecks (b) and Error Rate (c). No comparisons were done between *Repair* and *Inspection*.

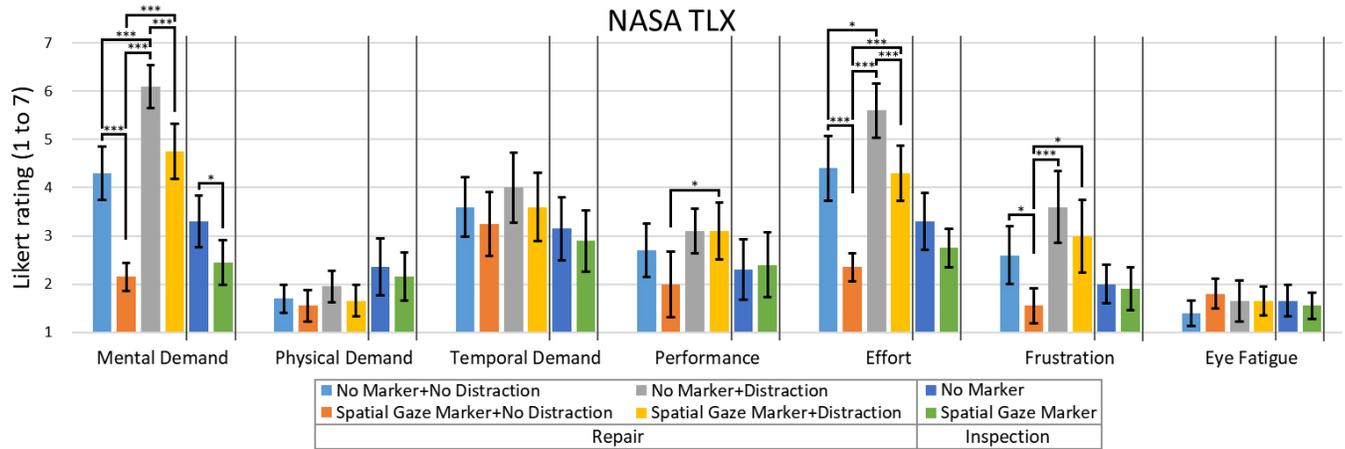


Figure 5: Results on NASA-TLX. No comparisons were done between *Repair* and *Inspection*.

1.226) conditions exhibited significantly fewer rechecks than the equivalent NM+ND ($M = 2.15, SD = 2.323$) and NM+D ($M = 2.9, SD = 2.47$) conditions within the same level of Distraction (all $p \leq 0.018$) and SGM+ND exhibited fewer rechecks than NM+D ($p < 0.001$). We also found significant difference in the *Inspection* task ($\chi^2(1) = 5.333, p = 0.021$) with SGM ($M = 0.4, SD = 1.188$) exhibiting significantly fewer rechecks than NM ($M = 1.9, SD = 3.007$).

5.6.3 *Error rate (Figure 4c)*. No significant differences were indicated in Error Rate in the *Repair* task ($\chi^2(3) = 4.244, p = 0.236$) or in the *Inspection* task ($\chi^2(1) = 0.667, p = 0.414$).

5.6.4 *NASA-TLX (Figure 5)*. As for the NASA-TLX scores, we found significant differences in Mental Demand in the *Repair* task ($\chi^2(3) = 52.175, p < 0.001$), SGM+ND ($M = 2.15, SD = 0.67$) was rated significantly less mentally demanding than all three other conditions (all $p < 0.001$), SGM+D ($M = 4.75, SD = 1.29$) was less mentally demanding than NM+D ($M = 6.1, SD = 1.02$)

($p < 0.001$), and NM+ND ($M = 4.3, SD = 1.26$) was less mentally demanding than NM+D ($p < 0.001$). Significant difference was also found in Mental Demand in the *Inspection* task ($\chi^2(3) = 6.250, p < 0.012$) with SGM ($M = 2.45, SD = 1.05$) being less mentally demanding than NM ($M = 3.3, SD = 1.22$). Regarding Performance in the *Repair* task ($\chi^2(3) = 12.677, p = 0.03$), SGM+ND ($M = 2.0, SD = 1.56$) was rated better (lower rating) than SGM+D ($M = 3.1, SD = 1.33$) ($p = 0.018$). In terms of Effort in the *Repair* task ($\chi^2(3) = 43.235, p < 0.001$), SGM+ND ($M = 2.35, SD = 0.67$) was rated significantly less effortful than all three other conditions (SGM+D ($M = 4.3, SD = 1.3$), NM+ND ($M = 4.4, SD = 1.54$), NM+D ($M = 5.6, SD = 1.27$)) (all $p < 0.001$), NM+ND required less effort than NM+D ($p = 0.012$), and SGM+D required less effort than NM+D ($p = 0.03$). Finally, we found significant differences in Frustration in the *Repair* task ($\chi^2(3) = 28.938, p < 0.001$), SGM+ND ($M = 1.55, SD = 0.83$) was rated less frustrating than all three other conditions (SGM+D ($M = 3.0, SD = 1.72$),

NM+ND ($M = 2.6, SD = 1.35$), and NM+D ($M = 3.6, SD = 1.7$) (all $p \leq 0.018$).

5.6.5 User Feedback. 13 participants mentioned that they appreciate that SGM reduced task load, making tasks easier to manage. 12 participants mentioned that SGM helped them remember target locations. 3 participants mentioned feeling that SGM improved task efficiency, especially in Distraction conditions. However, our participants also noted several challenges. 5 participants mentioned that SGM would occasionally provide incorrect guidance. However, these incorrect placements were infrequent and most often caused an error trial and are thereby captured by the Error Rate. 2 participants mentioned SGM disappeared too quickly, causing them to lose track of it. 4 participants mentioned that SGM was less useful in the *Inspection* task. Finally, 1 participant mentioned that the limited FOV of the HoloLens 2 required them to move their head more.

6 DISCUSSION

Our goal with SGM was to develop a task-agnostic AR tool that would guide users back on track, especially in cases where spatial reminders have the potential to render physical tasks easier to perform. The following discusses the main findings on objective and subjective findings and relates these findings in the context of related work. Furthermore, we briefly discuss envisioned use cases and the quick deployment through the task-agnostic nature of SGM.

6.1 Main Study Insights

SGM effectively guided participants across all factors in the *Repair* task, reducing Relocalisation Time by 41.39% and 52.44%, and the number of Rechecks by 95.35% and 77.59%, for No Distraction and Distraction respectively. Participants noted that the markers were occasionally positioned incorrectly, either through not being aware that they looked there last or possibly by tracking inaccuracies. This effect could potentially be mitigated through a smart visual indicator before the user switches, to communicate the return point even before one turns away. Notably, the factor of distraction had an effect on the Relocalisation Time for both our SGM and the No Marker baseline. Here the main insight is that distractions affect the baseline condition more pronounced than with using SGM. Furthermore, in *Repair*, SGM effectively lowered Mental Demand, Effort, and Frustration (Figure 5), indicating no penalty for the additional marker visualisation in the scene.

In *Inspection*, users of SGM needed 78.95% fewer Rechecks (Figure 4b) and exhibited lower Mental Demand. As the task allowed the participants to build up spatial memory over time, the effect was lower but still significant across multiple factors. It is interesting to think about how SGM may transfer to the many ways that our spatial memory operates in multi-tasking activities across time. In principle, this study inspected the immediate gaze fixation as a reminder. With different spatial memory demands across tasks, future gaze markers may dynamically utilise historical eye-gaze and other context information to improve the intent recognition of visually returning to a point in space.

The original Gazemarks, with a screen-based car navigation design, improved Relocalisation Time by 68.7% [25]. In our case,

the limited FOV of the HoloLens 2 constrained the utility of SGM – participants already had to be looking in the general direction of where they thought that the marker should be. Additionally, while our task-switch detection was chosen based on the limited FOV of the HoloLens 2, it may have been too small of a threshold. This was mostly evident in the *Inspection* task, where participants were doing a lot of head movement (e.g., going left to right, then back to left, like a typewriter), placing the previous fixation off to the side, possibly placing a marker in a less relevant location. Nonetheless, the subtle design of the marker makes it no problem to leave further markers in distinct areas. In case this is not preferred, the parameter could be adjusted to fit the task, as one option to optimise for users.

Our simulation of a real-world task for evaluation provides a new perspective on related experiments in the prior art of AR support for order picking [39, 40], that showed how AR reminders can aid the user in the visual search tasks. The use of eye-tracking enables both a mechanism to where and when to leave a cue to the user. Our evaluation task captures a dynamic scenario where users shift between different states of visual memory allocation. In there, we show that even with a minimalistic design, such as SGM, people's physical work can be positively supported in a simple and continuous way. This is novel in contrast to prior work, as we firstly demonstrate that gaze can provide context for an effective return point, and demonstrate interaction benefits for markers across task variations. At its core, our eyes play a fundamental role in our actions during spatial tasks – it remains a highly interesting question of how an active presentation of our gaze trajectory over time can improve spatial memory.

6.2 Application Examples

Beyond the scenario we studied in this paper, SGM is principally applicable to many other scenarios. For example, a LEGO-based assembly task as demonstrated in Figure 6. We deployed SGM for a concrete scenario involving multi-tasking, users can spontaneously use it when needed, with no overhead of a lengthy setup or prepared mark-ups of areas. This kind of assembly work simulation is often found in prior work [2, 12, 27], engaging the user in a continuous task of finding the right LEGO piece from a tool area to the work area. For each LEGO piece, a user shifts their attention and a marker is left as a way for the user to immediately return to the relevant location. In principle, such a multi-tasking scenario between two or more task areas extends to various tasks, as further illustrated in Figure 7. This ranges from professional activities such as maintenance or repair tasks (a-c) to casual activities such as cooking or other tasks with task interruptions (d-e). For example, when cooking, switching between the recipe book and stove would allow the user to get back to the last item in the recipe. SGM can, in principle, be useful for many use cases that we are unaware of, considering how complex human spatial memory and visual search are. There could be many cases we do not expect at this stage, that become only obvious after use.

6.3 Limitations & Future Work

Our study has several limitations. In designing our study, we aimed for a balance between internal and external validity. While our tasks are more externally valid than a virtual counterpart, the tasks

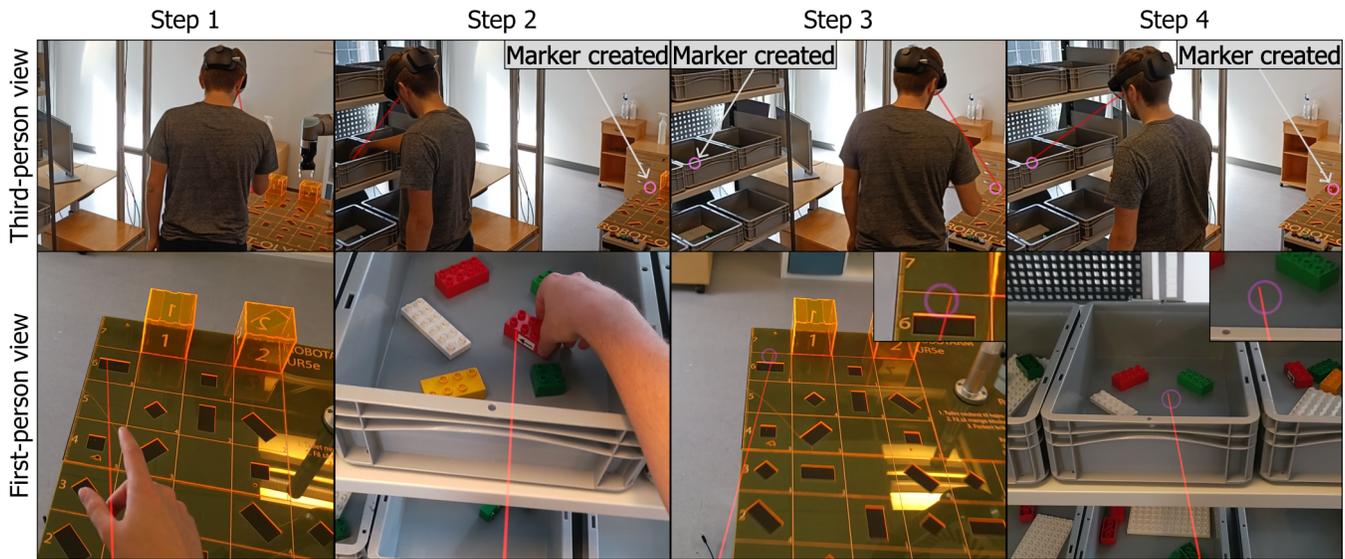


Figure 6: We integrated the system into a LEGO assembly task, demonstrating its adaptability for use in another scenario without modification. Here the user shifts between work (Step 1, 3) and tool area (Step 2, 4) – with markers laid in both areas’ last fixation points.

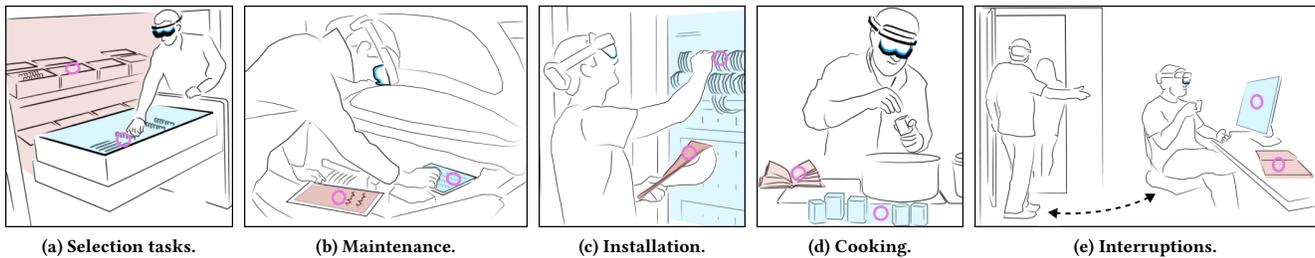


Figure 7: SGM has great potential to aid when switching between spaces, highlighted in blue and red. We envision use cases for SGM to include complex professional contexts with task or focus switching (a-c), and casual scenarios (d-e).

remain abstract and relatively niche. Future work could investigate the use of SGM in a different study design, such as a longitudinal study in everyday use. Another limitation of the task design is how we decide that the Post-it should hide the visual information, which is rather uncommon in reality. It could be interesting to see how well SGM perform in tasks where the visual information remains. Our study task focused on areas with few spatial references for remembering and precluded for instance easy locations such as at the borders. While this represented the scope of this work, it would be interesting to expand to exploration of designs for other visual memory demands. We also see improvements for the study apparatus, to reduce manual organisation and data logging. Furthermore, our findings are limited by the relatively low number of samples we gathered and validation with more people and more diverse backgrounds will provide further insights into the applicability. While our marker design was sufficient for our study, a greater design exploration may find other visualisations more beneficial. It could be interesting to investigate more explicit activation (more explicit

than turning one’s head, e.g., through a gesture, (AR) button press, or blinking).

7 CONCLUSION

This paper investigated SGM, a task-agnostic, automatic, AR marker placement concept, building upon a prior concept from Kern et al. [25]. We described a general set of conceptual and technical requirements for effectively detecting when users turn their attention away from an area, through the users’ fixations, instantiating a Spatial Gaze Marker at the last fixation, then visualising the marker upon return and removing the marker after use. Through a two-task user study, we showed that SGM effectively guide users to more efficiently complete tasks, even with lower task load placed upon the user, although, more effectively when spatial memory is limited, as in our *Repair* task. Overall, this work demonstrates the usefulness of simple eye-tracking systems for aiding users in their tasks without requiring the user to set up their workspace or be in

a specific environment. SGM could become highly used in everyday tasks and industrial environments.

ACKNOWLEDGMENTS

This work has received funding by the Innovation Fund Denmark (IFD grant no. 9090-00002B) for the MADE FAST project – part of the Manufacturing Academy of Denmark (MADE), the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant no. 101021229 GEM-INI), by the Independent Research Fund Denmark (DFF) (grant ID. 10.46540/3104-00008B), by a Marsden Fund Council from Government funding administered by the Royal Society of NZ (UOO1834), and the Pioneer Centre for AI, DNRF grant number P1.

REFERENCES

- Reynold Bailey, Ann McNamara, Nisha Sudarsanam, and Cindy Grimm. 2009. Subtle gaze direction. *ACM Transactions on Graphics (TOG)* 28 (9 2009), 1–14. Issue 4. <https://doi.org/10.1145/1559755.1559757>
- Jonas Blattgerste, Benjamin Streng, Patrick Renner, Thies Pfeiffer, and Kai Essig. 2017. Comparing Conventional and Augmented Reality Instructions for Manual Assembly Tasks. In *Proceedings of the 10th International Conference on Pervasive Technologies Related to Assistive Environments (Island of Rhodes, Greece) (PETRA '17)*. Association for Computing Machinery, New York, NY, USA, 75–82. <https://doi.org/10.1145/3056540.3056547>
- James V. Bradley. 1958. Complete Counterbalancing of Immediate Sequential Effects in a Latin Square Design. *J. Amer. Statist. Assoc.* 53 (1958), 525–528.
- Ernesto A. Bustamante and Randall D. Spain. 2008. Measurement Invariance of the Nasa TLX. <http://dx.doi.org/10.1177/154193120805201946> 3 (9 2008), 1522–1526. <https://doi.org/10.1177/154193120805201946>
- Shiwei Cheng, Jing Fan, and Anind K Dey. 2018. Smooth gaze: a framework for recovering tasks across devices using eye tracking. *Personal and Ubiquitous Computing* 22 (2018), 489–501.
- Andy Cockburn and Bruce McKenzie. 2002. Evaluating the Effectiveness of Spatial Memory in 2D and 3D Physical and Virtual Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Minneapolis, Minnesota, USA) (CHI '02)*. Association for Computing Machinery, New York, NY, USA, 203–210. <https://doi.org/10.1145/503376.503413>
- A. Cockburn and B. McKenzie. 2004. Evaluating Spatial Memory in Two and Three Dimensions. *Int. J. Hum.-Comput. Stud.* 61, 3 (sep 2004), 359–373. <https://doi.org/10.1016/j.ijhcs.2004.01.005>
- Valerio De Luca, Laura Corchia, Carola Gatto, Giovanna Ilenia Paladini, and Lucio Tommaso De Paolis. 2022. An Augmented Reality Application for the Frescoes of the Basilica of Saint Catherine in Galatina. In *The Future of Heritage Science and Technologies: ICT and Digital Heritage*, Rocco Furferi, Lapo Governi, Yary Volpe, Kate Seymour, Anna Pelagotti, and Francesco Gherardini (Eds.). Springer International Publishing, Cham, 112–125.
- Andrew T. Duchowski, Krzysztof Krejtz, Izabela Krejtz, Cezary Biele, Anna Niedzielska, Peter Kiefer, Martin Raubal, and Ioannis Giannopoulos. 2018. The Index of Pupillary Activity: Measuring Cognitive Load Vis-à-Vis Task Difficulty with Pupil Oscillation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173856>
- Ralf Elbert and Tessa Sarnow. 2019. Augmented Reality in Order Picking—Boon and Bane of Information (Over-) Availability. In *Intelligent Human Systems Integration 2019*, Waldemar Karwowski and Tareq Ahram (Eds.). Springer International Publishing, Cham, 400–406.
- Wei Fang and Zewu An. 2020. A scalable wearable AR system for manual order picking based on warehouse floor-related navigation. *The International Journal of Advanced Manufacturing Technology* 109 (2020), 2023–2037.
- Markus Funk, Thomas Kosch, Scott W. Greenwald, and Albrecht Schmidt. 2015. A Benchmark for Interactive Augmented Reality Instructions for Assembly Tasks. In *Proceedings of the 14th International Conference on Mobile and Ubiquitous Multimedia (Linz, Austria) (MUM '15)*. Association for Computing Machinery, New York, NY, USA, 253–257. <https://doi.org/10.1145/2836041.2836067>
- Markus Funk, Alireza Sahami Shirazi, Sven Mayer, Lars Lischke, and Albrecht Schmidt. 2015. Pick from Here! An Interactive Mobile Cart Using in-Situ Projection for Order Picking. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (Osaka, Japan) (UbiComp '15)*. Association for Computing Machinery, New York, NY, USA, 601–609. <https://doi.org/10.1145/2750858.2804268>
- Christoph Gebhardt, Brian Hecox, Bas van Opheusden, Daniel Wigdor, James Hillis, Otmar Hilliges, and Hrvoje Benko. 2019. Learning Cooperative Personalized Policies from Gaze Data. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology (New Orleans, LA, USA) (UIST '19)*. Association for Computing Machinery, New York, NY, USA, 197–208. <https://doi.org/10.1145/3332165.3347933>
- Ioannis Giannopoulos, Peter Kiefer, and Martin Raubal. 2012. GeoGazemarks: Providing Gaze History for the Orientation on Small Display Maps. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction (Santa Monica, California, USA) (ICMI '12)*. Association for Computing Machinery, New York, NY, USA, 165–172. <https://doi.org/10.1145/2388676.2388711>
- Jeffrey Goderie, Rustam Alashraf, Pieter Jockin, Lu Liu, Xin Liu, Marina A. Cidota, and Stephan G. Lukosch. 2017. [POSTER] ChiroChroma: An Augmented Reality Game for the Assessment of Hand Motor Functionality. In *2017 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*. Institute for Electrical and Electronics Engineers (IEEE), New York City, United States, 115–120. <https://doi.org/10.1109/ISMAR-Adjunct.2017.44>
- Anhong Guo, Shashank Raghu, Xuwen Xie, Saad Ismail, Xiaohui Luo, Joseph Simoneau, Scott Gilliland, Hannes Baumann, Caleb Southern, and Thad Starner. 2014. A Comparison of Order Picking Assisted by Head-up Display (HUD), Cart-Mounted Display (CMD), Light, and Paper Pick List. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers (Seattle, Washington) (ISWC '14)*. Association for Computing Machinery, New York, NY, USA, 71–78. <https://doi.org/10.1145/2634317.2634321>
- Kunal Gupta, Gun A. Lee, and Mark Billinghurst. 2016. Do You See What I See? The Effect of Gaze Tracking on Task Space Remote Collaboration. *IEEE Transactions on Visualization and Computer Graphics* 22, 11 (2016), 2413–2422. <https://doi.org/10.1109/TVCG.2016.2593778>
- Sandra G. Hart. 2006. Nasa-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50 (10 2006), 904–908. <https://doi.org/10.1177/154193120605000909>
- Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Advances in Psychology* 52 (1 1988), 139–183. Issue C. [https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9)
- Steven J. Henderson and Steven Feiner. 2009. Evaluating the benefits of augmented reality for task localization in maintenance of an armored personnel carrier turret. In *2009 8th IEEE International Symposium on Mixed and Augmented Reality*. IEEE, Orlando, FL, USA, 135–144. <https://doi.org/10.1109/ISMAR.2009.5336486>
- Allison Jing, Kieran May, Brandon Matthews, Gun Lee, and Mark Billinghurst. 2022. The Impact of Sharing Gaze Behaviours in Collaborative Mixed Reality. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 463 (nov 2022), 27 pages. <https://doi.org/10.1145/3555564>
- Jaemin Jo, Bohyoung Kim, and Jinwook Seo. 2015. EyeBookmark: Assisting Recovery from Interruption during Reading. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (Seoul, Republic of Korea) (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 2963–2966. <https://doi.org/10.1145/2702123.2702340>
- James F Juola, Juan Botella, and Antonio Palacios. 2004. Task-and location-switching effects on visual attention. *Perception & Psychophysics* 66 (2004), 1303–1317.
- Dagmar Kern, Paul Marshall, and Albrecht Schmidt. 2010. Gazemarks: Gaze-Based Visual Placeholders to Ease Attention Switching. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Atlanta, Georgia, USA) (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 2093–2102. <https://doi.org/10.1145/1753326.1753646>
- Iring Koch and Andrea Kiesel. 2022. *Task Switching: Cognitive Control in Sequential Multitasking*. Springer International Publishing, Cham, 85–143. https://doi.org/10.1007/978-3-031-04760-2_3
- Thomas Kosch, Markus Funk, Albrecht Schmidt, and Lewis L. Chuang. 2018. Identifying Cognitive Assistance with Mobile Electroencephalography: A Case Study with In-Situ Projections for Manual Assembly. *Proc. ACM Hum.-Comput. Interact.* 2, EICS, Article 11 (jun 2018), 20 pages. <https://doi.org/10.1145/3229093>
- Mikko Kytö, Barrett Ens, Thammathip Piumsomboon, Gun A. Lee, and Mark Billinghurst. 2018. Pinpointing: Precise Head- and Eye-Based Target Selection for Augmented Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3173574.3173655>
- Luis A. Leiva. 2011. MouseHints: Easing Task Switching in Parallel Browsing. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems (Vancouver, BC, Canada) (CHI EA '11)*. Association for Computing Machinery, New York, NY, USA, 1957–1962. <https://doi.org/10.1145/1979742.1979861>
- Gordon D Logan. 1985. Executive control of thought and action. *Acta psychologica* 60, 2-3 (1985), 193–210.
- Mathias N. Lystbæk, Peter Rosenberg, Ken Pfeuffer, Jens Emil Grønþæk, and Hans Gellersen. 2022. Gaze-Hand Alignment: Combining Eye Gaze and Mid-Air

- Pointing for Interacting with Menus in Augmented Reality. *Proc. ACM Hum.-Comput. Interact.* 6, ETRA, Article 145 (may 2022), 18 pages. <https://doi.org/10.1145/3530886>
- [32] Stephen Monsell. 2003. Task switching. *Trends in cognitive sciences* 7, 3 (2003), 134–140.
- [33] Alexander Plopski, Teresa Hirzle, Nahal Norouzi, Long Qian, Gerd Bruder, and Tobias Langlotz. 2022. The Eye in Extended Reality: A Survey on Gaze Interaction and Eye Tracking in Head-Worn Extended Reality. *ACM Comput. Surv.* 55, 3, Article 53 (mar 2022), 39 pages. <https://doi.org/10.1145/3491207>
- [34] Jarkko Polvi, Takafumi Taketomi, Atsunori Moteki, Toshiyuki Yoshitake, Toshiyuki Fukuoka, Goshiro Yamamoto, Christian Sandor, and Hirokazu Kato. 2018. Handheld Guides in Inspection Tasks: Augmented Reality versus Picture. *IEEE Transactions on Visualization and Computer Graphics* 24 (7 2018), 2118–2128. Issue 7. <https://doi.org/10.1109/TVCG.2017.2709746>
- [35] Patrick Renner and Thies Pfeiffer. 2017. Evaluation of Attention Guiding Techniques for Augmented Reality-Based Assistance in Picking and Assembly Tasks. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces Companion* (Limassol, Cyprus) (*IUI '17 Companion*). Association for Computing Machinery, New York, NY, USA, 89–92. <https://doi.org/10.1145/3030024.3040987>
- [36] Dario D. Salvucci and Joseph H. Goldberg. 2000. Identifying Fixations and Saccades in Eye-Tracking Protocols. In *Proceedings of the 2000 Symposium on Eye Tracking Research & Applications* (Palm Beach Gardens, Florida, USA) (*ETRA '00*). Association for Computing Machinery, New York, NY, USA, 71–78. <https://doi.org/10.1145/355017.355028>
- [37] Joey Scarr, Andy Cockburn, Carl Gutwin, et al. 2013. Supporting and exploiting spatial memory in user interfaces. *Foundations and Trends® in Human-Computer Interaction* 6, 1 (2013), 1–84.
- [38] Christina Schneegass and Fiona Draxler. 2021. Designing Task Resumption Cues for Interruptions in Mobile Learning Scenarios. In *Technology-Augmented Perception and Cognition*, Tilman Dingler and Evangelos Niforatos (Eds.). Springer International Publishing, Cham, 125–181. https://doi.org/10.1007/978-3-030-30457-7_5
- [39] Bjorn Schwerdtfeger and Gudrun Klinker. 2008. Supporting order picking with Augmented Reality. In *2008 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*. IEEE, Cambridge, UK, 91–94. <https://doi.org/10.1109/ISMAR.2008.4637331>
- [40] Bjorn Schwerdtfeger, Rupert Reif, Willibald A. Gunthner, Gudrun Klinker, Daniel Hamacher, Lutz Schega, Irina Bockelmann, Fabian Doil, and Johannes Tumlner. 2009. Pick-by-Vision: A first stress test. In *2009 8th IEEE International Symposium on Mixed and Augmented Reality*. IEEE, Orlando, FL, USA, 115–124. <https://doi.org/10.1109/ISMAR.2009.5336484>
- [41] Ludwig Sidenmark, Christopher Clarke, Joshua Newn, Mathias N. Lystbæk, Ken Pfeuffer, and Hans Gellersen. 2023. Vergence Matching: Inferring Attention to Objects in 3D Environments for Gaze-Assisted Selection. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (, Hamburg, Germany,) (*CHI '23*). Association for Computing Machinery, New York, NY, USA, Article 257, 15 pages. <https://doi.org/10.1145/3544548.3580685>
- [42] Jonathan Sutton, Tobias Langlotz, Alexander Plopski, Stefanie Zollmann, Yuta Itoh, and Holger Regenbrecht. 2022. Look over There! Investigating Saliency Modulation for Visual Guidance with Augmented Reality Glasses. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (*UIST '22*). Association for Computing Machinery, New York, NY, USA, Article 81, 15 pages. <https://doi.org/10.1145/3526113.3545633>
- [43] Eduardo E. Veas, Erick Mendez, Steven K. Feiner, and Dieter Schmalstieg. 2011. Directing Attention and Influencing Memory with Visual Saliency Modulation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (*CHI '11*). Association for Computing Machinery, New York, NY, USA, 1471–1480. <https://doi.org/10.1145/1978942.1979158>
- [44] Uta Wagner, Mathias N. Lystbæk, Pavel Manakhov, Jens Emil Sloth Grønbaek, Ken Pfeuffer, and Hans Gellersen. 2023. A Fitts' Law Study of Gaze-Hand Alignment for Selection in 3D User Interfaces. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (, Hamburg, Germany,) (*CHI '23*). Association for Computing Machinery, New York, NY, USA, Article 252, 15 pages. <https://doi.org/10.1145/3544548.3581423>
- [45] Manuela Waldner, Werner Puff, Alexander Lex, Marc Streit, and Dieter Schmalstieg. 2010. Visual Links across Applications. In *Proceedings of Graphics Interface 2010*. Canadian Information Processing Society, CAN, 129–136.
- [46] Glenn Wylie and Alan Allport. 2000. Task switching and the measurement of “switch costs”. *Psychological research* 63 (2000), 212–233.