Effects of Predictive Navigation Assistance on Location Learning in Visual Workspaces

Morgan Beattie
morgan.beattie@usask.ca
University of Saskatchewan
Saskatoon, Canada

Carl Gutwin
gutwin@cs.usask.ca
University of Saskatchewan
Saskatoon, Canada

Andy Cockburn
andrew.cockburn@canterbury.ac.nz
University of Canterbury
Christchurch, New Zealand

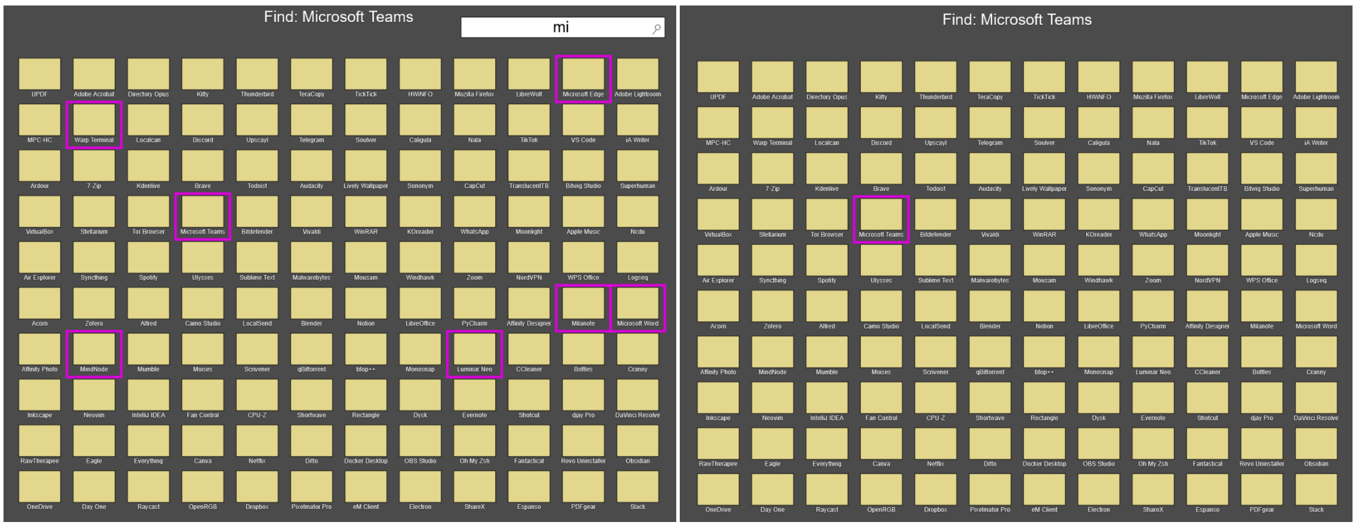

Figure 1: Two navigation aids for locating an item in an interface. *Search Assistance* (left) is a user-directed aid that highlights every item predicted from a text query. *Automatic Assistance* (right) is a simulated fully-automatic aid that highlights the predicted item with perfect accuracy.

Abstract

Complex visual information spaces require a considerable amount of time and effort to visually search for items. To help inexperienced users locate items faster, some systems predict and visually emphasize the item the user is looking for. However, while designers of these navigation aids rely on the *rehearsal principle* to help users acquire spatial knowledge of item locations, in situations where the prediction is unavailable or incorrect, it is unclear whether performance will suffer due to users being overly reliant on the assistance. To investigate the effect of predictive navigation assistance on location learning, we ran a study to compare three navigation techniques: Automatic Assistance highlighted the target item with perfect accuracy, Search Assistance allowed users to filter items by name, and No Assistance provided no help to find items. Results showed that while Search Assistance had comparable memory performance to the unassisted baseline condition, Automatic Assistance had substantially reduced location learning. Our work contributes new understanding of the spatial-memory effects of prediction and rehearsal, and provides designers with new opportunities to support location learning in visual workspaces.

Keywords

Location learning, prediction, navigation assistance, visual workspaces, rehearsal interfaces.

ACM Reference Format:
Morgan Beattie, Carl Gutwin, and Andy Cockburn. 2026. Effects of Predictive Navigation Assistance on Location Learning in Visual Workspaces. In *Proceedings of Graphics Interface 2026 (GI '26)*. ACM, New York, NY, USA, 11 pages. https://doi.org/XXXXXXX.XXXXXXX

GI '26, Waterloo, ON, Canada

ACM ISBN 978-1-4503-XXXX-X/2026/06
https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Acquiring spatial knowledge of a visual information space (such as a network visualization, a catalogue of items, a map, or a GUI) is essential for users to make selections efficiently. In complex visual spaces, inexperienced users must rely on costly visual search to find items, which can feel slow and frustrating. For information spaces that are spatially stable, users can build *spatial knowledge* of the workspace through repeated selections, allowing experienced users to quickly retrieve item locations from memory [9].

To alleviate the cost of visual search and to help beginners find items, some systems predict the item that the user needs next, and visually emphasize that item in the interface. Predictive models can use a variety of inputs, including explicit user actions like search text or filters, implicit information like the user's past behaviour, as well as non-user information such as task models and constraints. Statistical methods exist for predicting user actions, including Markov [61] and Bayesian models [27] – and with the introduction of large language models, prediction of user actions is becoming better (e.g., [37, 52]). However, because current prediction systems are not perfectly accurate, and because users prefer to be shown the steps rather than simply have the system complete the task [32], it is important that predictive systems also help the user learn the locations of items in the interface.

Researchers have already considered how predictive interfaces should be designed to aid location learning: in particular, Fitchett et al. state a critical principle in the design of predictive feedback that is required for users to learn locations: the feedback should guide the user's actual navigation task in the information space, rather than providing a separate representation that does not involve the actual interface (e.g., a separate results list from a 'find' functionality) [16]. This allows users to *rehearse* the actions needed to navigate to an item, contributing to incidental learning – this is the *rehearsal principle* [33]. Several previous studies have demonstrated the value of the rehearsal-based approach (e.g., Marking Menus [33], Octopocus for gesture learning [5], and FastTap for command location learning [21]); however, other studies have suggested that rehearsal alone might not adequately support the transition to memory-based interaction [20].

While systems that predict and visually emphasize the user's next item can reduce the need for visual search, prediction techniques also raise problems when the prediction is either unavailable or incorrect. Previous work has indicated that predicting user actions is difficult [15], and even with the improvements that come with LLM-based prediction models, it remains difficult to always predict accurately what item the user is going to need next (e.g., the output of Shaikh et al.'s Next Action Prediction system aligned with user actions only 17.1% of the time [52]). Further, the prediction may be completely unavailable to the user in scenarios such as planning a sequence of actions before interacting with the system, using a different view (e.g., an in-application file dialog versus the main file explorer), helping another person who doesn't have the prediction turned on, or switching to a different device where the prediction is not set up or the device doesn't have the computational power to run the AI model. In these situations, if the user has become over-reliant on the prediction and therefore has not learned the item locations, then they will be forced to return to laborious visual search when the prediction is either wrong or not available.

This problem leaves us with two open questions about predictive navigation aids: first, how well a prediction technique works to reduce search time, and second, how well a technique supports location learning and subsequent performance when the prediction is not available.

To answer these questions, we carried out a study to test the effect of two types of predictive navigation aids on location learning in a GUI containing text-labelled items. The first navigation technique, *Search Assistance*, used user-directed prediction that allowed participants to search the labels with a text query (all items matching the query were dynamically highlighted with a pink box). The second navigation technique, *Automatic Assistance*, simulated a navigation aid capable of automatically predicting and highlighting the next item with perfect accuracy. Both navigation aids were compared to an unassisted baseline condition which received no predictive guidance (conditions were assigned between-participants). Participants completed three experiment phases: first, an *assisted use phase*, in which participants completed 8 blocks of selections for 6 targets with the assigned assistance technique present; second, a *memory test* in an unlabelled version of the interface to assess absolute spatial learning of the target locations; and finally, an *unassisted use phase*, where participants completed 8 blocks of selections for the same 6 targets with no prediction assistance (in order to assess performance when prediction was not available).

Our results show major differences between the two types of predictive assistance. While both Search and Automatic assistance provided substantial performance benefits during assisted use compared to the unassisted baseline, only Automatic Assistance resulted in significantly reduced location learning on the memory test and significantly reduced performance in the unassisted use phase. Further, subjective responses showed that both navigation aids improved user experience during assisted use compared to the unassisted baseline, with Automatic Assistance having the best perceived performance but worst perceived location learning (consistent with performance data).

This work provides two main contributions. We provide novel empirical results about how different prediction approaches affect navigation-assisted location learning in a visual workspace – in particular, we show that the use of a perfectly-accurate automatic prediction harmed learning, but the use of search-directed assistance did not, even though it substantially improved selection performance when turned on. Our results suggest that designers of predictive navigation aids can use search assistance to boost the performance and user experience of novices, while still facilitating location learning. Overall, our work provides researchers and designers with valuable new information about the application of predictive assistance in complex visual spaces.

2 Related Work

We examine three areas of related work: location learning and rehearsal in visual workspaces, in-application prediction and copilots, and predictive navigation assistance.

2.1 Location Learning and Rehearsal in Visual Workspaces

Gaining spatial knowledge of a visual information space is an essential part of building expertise [49]. To make selections, novice users of complex interfaces must perform slow and frustrating visual search over all available items, whereas expert users can rely on memory to find target items quickly and easily [9]. Even when visual search can be bypassed by various interaction techniques (for instance, text search), learning the layout is valuable to gain an understanding of what items are available for selection, and where they might be located in the interface (a recent study of visual search in GUIs found that users tend to begin in the top-left corner of the interface, which commonly contains information about command locations) [46].

Spatial stability is a necessary characteristic of visual workspaces in order to build spatial knowledge (since, for instance, re-flowing causes item locations to change) [22, 47]. When users repeatedly select an item in a spatially-stable visual workspace, they begin to encode the item's location in memory and experience performance improvements characterized by considerable early gains that eventually reach a stable asymptote (i.e., they transition from visual search to recall of the item's location). This predictable pattern of performance improvement with repetition is termed the *power law of practice* [44].

Leveraging the performance improvements that occur as a function of practice, rehearsal-based interfaces aim to support a smooth transition from a slow, visually-guided interaction to an expert memory-based mode (i.e., rehearsal interfaces remove the "performance dip" that occurs between the two interaction modes [48]). This 'rehearsal principle', which states that the novice mode of selection should require the same actions as the expert mode (such that the expert actions are *rehearsed* by novices), was fundamental in the design of Kurtenbach's marking menus [33–35]. Other rehearsal-based interfaces have been designed to test the benefit of rehearsal for learning interfaces. Octopocus was developed to help users learn gesture sets by providing dynamic feedforward guidance of the possible gesture space given the current input; once a gesture is learned, experts can execute commands without referring to the guide [5]. Schramm et al. showed that spatial memory could be leveraged in rehearsal-based interfaces with hidden toolbars, where location-based shortcuts (e.g., swiping across the screen bezel) could replace traditional visually-guided selections [51]. Finally, the FastTap interface allows novices to make visually-guided selections by pressing and holding a menu invocation button with the thumb and making a selection with a finger, and experts can accelerate these selections using thumb-and-finger chorded input [21].

A potential risk of rehearsal interfaces is that users will become over-reliant on the visual guidance provided in the novice mode, fail to learn the interface, and never be able to switch to the expert memory-based strategy – a phenomenon termed the *guidance hypothesis* [50]. Gutwin et al. tested the rehearsal principle in a longitudinal study with two FastTap interfaces, and found that participants were not guaranteed to switch to memory-only interaction [20]. However, there is evidence that spatial learning can occur incidentally [3, 8, 23, 58], indicating that rehearsal of selecting item locations could be beneficial even when visual guidance is present. Further, previous work on the guidance hypothesis in games showed no significant disadvantage of guided navigation on spatial learning of 3D environments [29, 30].

2.2 In-Application Prediction and Copilots

Research into predicting user actions has occurred for decades (see [12, 24, 36, 41, 60] for examples of early work on predicting sequences of user actions, automatic filling of repetitive forms, goal recognition, text prediction, and predicting commands). Many traditional prediction algorithms, such as *Most Recently Used* and *Most Frequently Used* (e.g., the algorithms used for Gajos' adaptive GUIs [18]), are based on simple patterns of user behaviour – an effective approach due to the highly repetitive nature of many forms of interaction [12, 55]. Based on this principle, Fitchett et al. developed AccessRank, a revisitation-based prediction algorithm which utilizes recency, frequency, temporal clustering, and time of day [15]. Some systems also consider actions performed by other community members, such as the approach taken by Matejka et al. in the CommunityCommands command recommendation system [42] or in the use of data mining for recommending websites [19].

As prediction of user actions continues to improve, the next possibility is full automation of tasks that are traditionally driven by the user. Early intelligent agents required a significant investment from the user – for instance, Maes and Kozierok's intelligent agents learned how to assist the user by imitating user actions, receiving feedback on incorrect actions, and practising user-guided hypothetical examples [39]. However, researchers saw the potential value of these fully-automated agents for decreasing mental load for users, or to help automate repetitive tasks [38]. At the same time, others suggested that full automation would reduce user control, and instead pushed for *mixed initiative interaction* such that "each agent can contribute to the task what it does best" [2].

Today, full automation of some user-driven tasks is becoming increasingly possible with large language model-based assistants, or *copilots* (examples of popular copilots include GitHub Copilot for code prediction [17], and Microsoft 365 Copilot for content generation and productivity across the Microsoft Office suite [54]). As these recent intelligent agents continue to gain popularity, recent work has focused on the role of user control in copilot use: Khurana et al. found that users scored a semi-automated copilot better than a fully-automated copilot for user control, software learnability, and software utility, indicating that participants preferred to learn by doing the task with the guidance of the copilot (rather than having the task done for them) [32].

2.3 Predictive Navigation Assistance

Navigation assistance can utilize a variety of approaches to accelerate selections in complex visual workspaces. These aids usually include some form of visual emphasis of a predicted item – instead of having to perform slow visual search [9], the user can have their attention drawn to locations of the most-likely next items with a "pop-out" effect (according to Treisman and Gelade's feature integration theory [56], the strength of this effect is determined by the features of the item and the surrounding interface – such as colour, shape, size, and movement). For instance, hierarchical file navigation is a slow yet common task; Fitchett et al. designed and examined multiple visual navigation aids (shown in Figure 2) for assisting file retrieval [16]. Predictive navigation aids can be fully automatic (such as Fitchett's *Icon Highlights*, which emphasized the locations of predicted items in the current folder, and *Hover Menus*, which provided shortcuts to predicted items within sub-folders), or may be user-directed (like Fitchett's *Search Directed Navigation*, which allowed users to locate items based on a text query).

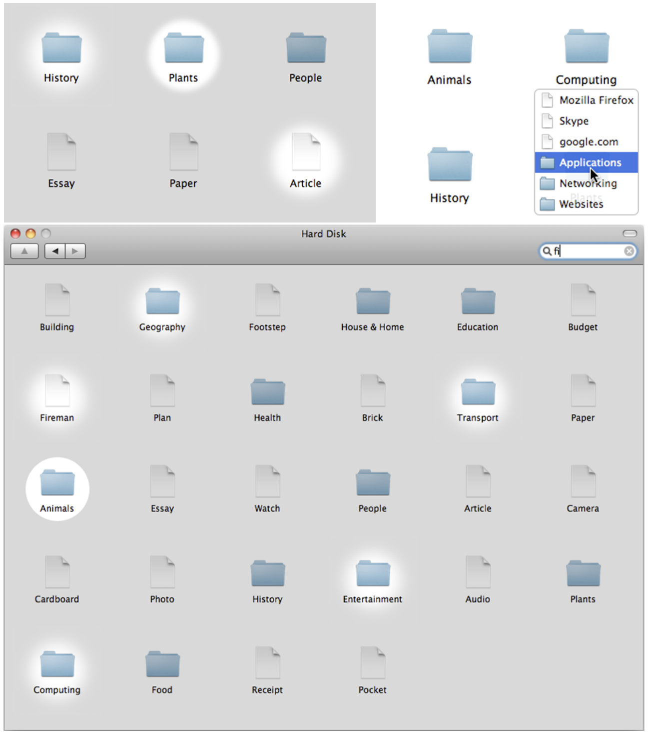

Figure 2: Navigation aids by Fitchett et al. [16]. Automatic *Icon Highlights* are shown top-left, *Hover Menus* are shown top-right, and *Search Directed Navigation* is shown in the bottom panel.

Automatic navigation assistance can include visualizations of a user's history of actions or prediction of a user's future actions based on previous actions. Visual indicators of a user's history facilitate revisitation by showing where previous interaction has occurred. For instance, Hill et al. developed attribute-mapped scroll bars with Edit Wear (i.e., markers of "authorship history") and Read Wear (i.e., markers of "readership history") to guide navigation to particular document locations [25]. Landmarks, whether provided intentionally (e.g., "artificial landmarks" in two-dimensional user interfaces [57] and annotated document scrollbars [43]) or unintentionally (including the "accidental landmarks" that occur as a user continues to interact with and mark up an interface [40]) can also assist in navigation by acting as spatial anchors of the overall layout. Other automatic predictive aids, like Fitchett's icon highlights [16], use more advanced algorithms to guess what actions the user might take next, as discussed in the previous section.

User-directed navigation aids are intentionally triggered by the user, and predict the most-likely next items based on specific input. For example, systems may include interaction techniques that allow the user to mark important items to return to later – like brushing, which allows users to keep track of certain data points in a visualization by manipulating their appearance (often, colour) [6] and bookmarking, which allows users to save items by creating an explicit mark or by defining a shortcut [1]. Other techniques may allow the user to input characteristics of the item to help guide the prediction, including dynamic query, which provides slider and button filters to reduce the result set [53], and search, which allows users to find items based on a text query [16] (note that these techniques may result in multiple items being identified). The main disadvantage of user-directed navigation aids is the potential high mental demand – e.g., users must realize that they will need to revisit brushed or bookmarked items, need to manage brushed sets or bookmark lists, need to know the characteristics of items to define filter settings for dynamic query, and are required to recall text for search queries.

3 Study Methods

To assess the effect of predictive assistance on location learning in a complex visual workspace, we compared three techniques for finding targets in a spatially-stable visualization that was similar to a file explorer application with icons and text labels.

3.1 Navigation Techniques

We tested three mechanisms for workspace navigation that served as conditions in the experiment:

Search Assistance allowed participants to search for items with a text query. All items with labels containing the query substring were highlighted with a pink box (the search set was dynamically updated as characters were typed). An example of the Search Assistance condition, including the search box, is shown in the left panel of Figure 1.

Automatic Assistance automatically highlighted the target item with a pink box. The goal was to simulate a copilot assistant which can identify the next-most-likely item with perfect accuracy. An example of the Automatic Assistance condition is shown in the right panel of Figure 1.

No Assistance provided no help to find the target item (i.e., participants had to perform unassisted visual search to find items).

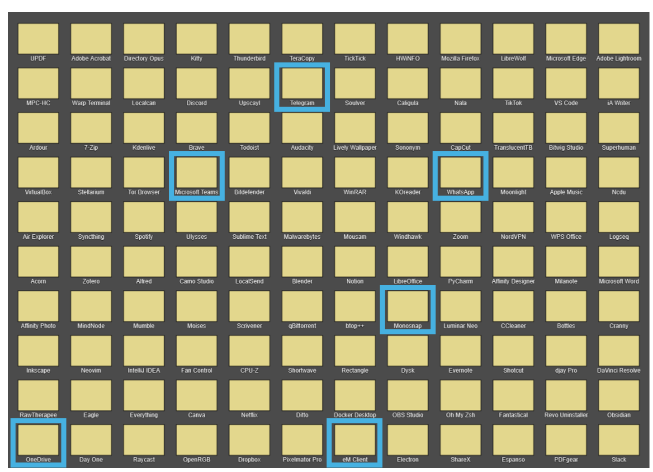

Figure 3: Interface and Targets used in the Study. Target items are indicated with blue box highlights.

3.2 Apparatus and Procedure

The study used a custom web application built with P5JS to administer experiment tasks and log participant actions. In the visual workspace, 120 items were randomly arranged in a 12×10 grid and were each displayed with a yellow rectangle icon and a text label. Six targets were selected such that they were spatially spread throughout the workspace, and required a varying number of prefix characters (in the Search Assistance condition) to filter the search set down to one item (the interface and selected targets are shown in Figure 3). The study system, including the consent form, questionnaires, and task application was deployed on the web using an open-source experiment management system [28]. Participants completed the study on a desktop computer, and performed the experiment tasks in a full-screen environment.

Participants first completed informed consent and demographics forms, and then were shown pages containing task instructions and instructions for their assigned between-participants experiment condition (Search Assistance, Automatic Assistance, or No Assistance). They then completed the three main experiment phases. First, to simulate real use of the in-application navigation aids, participants completed the *assisted use phase*, which involved eight blocks of randomly-ordered selections for the six targets. Each condition was provided navigation assistance as described in Section 3.1, except for the unassisted baseline condition. After completing the assisted use phase, participants filled out an experience questionnaire which contained the NASA-TLX survey (questions only, no ranking), with additional questions that asked about perceived speed, accuracy, learning, and effort put into learning in the task.

To assess absolute spatial knowledge, participants then completed a *memory test* which required a selection for the location of each target from an unlabelled grid of rectangle icons (i.e., without the application names). Following each selection, participants rated their confidence in their response on a 7-point scale. The memory test contained three attention checks to discourage random inputs.

Finally, participants completed the *unassisted use phase*, which was identical to the first task except that no navigation assistance was provided. The goal was to assess performance in the case that the assistance was no longer available.

3.3 Participants

Participants were recruited from Prolific, a platform that facilitates recruitment and payment for crowd-sourced online studies. Participants were required to reside in the US, UK, or Canada, be fluent in English, and have a task approval rate of at least 90% (an indication of participant quality). Participants were paid £3.00 for the 15-minute session.

A total of 90 participants were recruited (46 men, 38 women, 1 non-binary, mean age 39.7), with 30 being assigned to each condition. All participants were familiar with mouse-and-windows systems, and reported a mean of 42 hours of use per week.

3.4 Study Design

Separate analyses were conducted for each experiment phase: for the assisted use phase, we were interested in changes to performance during use of each navigation technique; in the memory test, we measured absolute spatial knowledge of the interface; and in the unassisted use phase, we assessed performance in the absence of the aid.

Mixed-factorial and one-way comparisons were made across the main between-participants factor, *Navigation Technique* (No Assistance, Search Assistance, and Automatic Assistance), in each of the three study phases. In the assisted use phase, we used a 3×8 mixed-factorial ANOVA over Navigation Technique and secondary within-participants factor Block. Memory test data were analyzed by one-way ANOVA on factor Navigation Technique. Finally, we used a 3×8 (Navigation Technique $\times$ Block) ANOVA in the unassisted use phase.

The main dependent variables were trial completion time in assisted use and unassisted use phases, and accuracy on the memory test. Secondary measures included error rate, number of characters typed for the Search Assistance technique, confidence ratings on the memory test, and subjective questionnaire responses.

Each participant completed 102 trials (6 targets $\times$ 8 blocks in the assisted use phase; 6 targets on the memory test; 6 targets $\times$ 8 blocks in the unassisted use phase), for a total of 9,180 total trials observed over 90 participants.

4 Results

Analyses are organized by study phase, and then by dependent variable (completion time, errors, and subjective responses for the assisted use phase, accuracy and confidence for the memory test, and completion time and errors for the unassisted use phase).

In the assisted and unassisted use phases, trials were removed if completion times were greater than three standard deviations from the mean for each Navigation Technique $\times$ Block group (98 of 4,320 assisted use trials, and 90 of 4,320 unassisted use trials – 2.2% overall). Data for one participant was removed from the memory test for failing more than one attention check.

Since trial completion times had unequal variances across conditions (Levene's test: $p<.001$), data was log-transformed for ANOVA and follow-up tests in the assisted and unassisted phases. Holm-Bonferroni correction was used for follow-up t-tests. Subjective data were analyzed by applying the Aligned-Rank Transform [59] and performing a one-way ANOVA on factor Navigation Technique. All charts show non-transformed data with error bars representing 95% confidence intervals.

4.1 Assisted Use Phase

Figure 4 shows the trial completion times in the assisted use phase. As can be seen in the figure, the No Assistance condition had high completion times which gradually improved over the eight blocks of selections. The Search Assistance and Automatic Assistance conditions had relatively stable performance, with Automatic Assistance being consistently fastest overall. A 3×8 (Navigation Technique $\times$ Block) ANOVA was used to compare trial completion times for the three conditions; results are summarized in Table 1. There were significant main effects of Navigation Technique and Block, with an interaction effect between the two factors. A follow-up test revealed that Automatic Assistance (mean 1.18s) was significantly faster than Search Assistance (mean 3.52s), which was significantly faster than No Assistance (mean 8.41s), with $p<.001$. The interaction between Navigation Technique and Block is primarily explained by

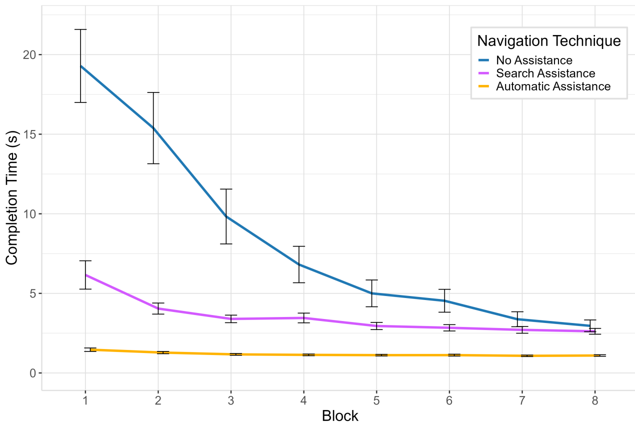

Figure 4: Trial completion time by block for each navigation technique in the assisted use phase.

Factor	DF (n,d)	F	p	η^2
Navigation Technique	2, 87	218.75	**<.001**	0.76
Block	7, 609	191.78	**<.001**	0.45
Navigation Tech. × Block	14, 609	50.06	**<.001**	0.30

Table 1: Factorial Analysis of Trial Completion Times in the Assisted Use Phase.

the substantial performance advantage of training with navigation assistance (i.e., while No Assistance had a steep learning curve, Automatic Assistance had constant, low completion times, and Search Assistance fell in-between these curves).

Though completion time was the primary measure, we also analyzed differences in errors between the three Navigation Technique conditions. Error rates were low in the assisted use phase (mean 2.6%), and a one-way ANOVA revealed no significant differences in errors made between the three conditions ($F_{2,87} = 2.16$, p=.12, $\eta^2 = 0.05$).

To better understand how participants used the search assistance, we monitored the characters typed in each trial. Figure 5 displays the mean number of characters typed in each block; the figure shows that participants reduced their use of the search assistance as the study progressed. The dashed line in Figure 5 shows the proportion of trials in each block where no characters were typed. Consistent with the decreasing number of characters typed during the task, the proportion of memory-only trials rose gradually; together, these results indicate that participants were beginning to rely on spatial memory recall instead of needing the navigation aid.

Immediately following the assisted use phase, participants completed a questionnaire that included NASA-TLX questions and custom questions that asked about speed, accuracy, learning, and effort put into learning. Post-phase questionnaire responses are shown in Figure 6, and corresponding factorial analyses are given in Table 2. One-way ANOVA on factor Navigation Technique revealed that participants found No Assistance was more mentally demanding, effortful, frustrating, and slower than Automatic Assistance and Search Assistance. On the other hand, Automatic Assistance was

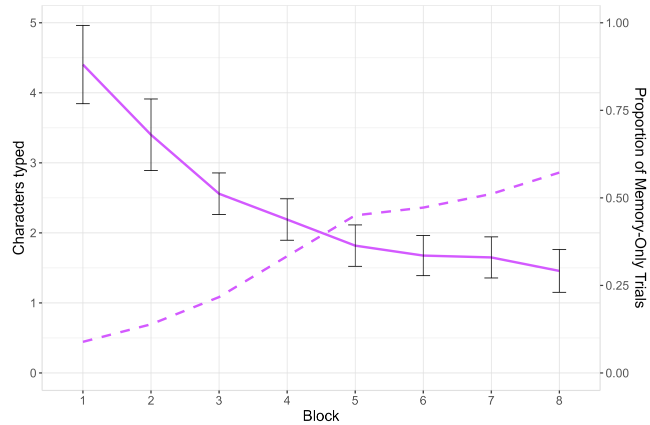

Figure 5: A summary of characters typed in the Search Assistance condition during the assisted use phase. The solid line shows the mean number of characters typed in each block. The dashed line shows the proportion of trials at each block that were completed with zero characters (i.e., without using the search assistance, indicating that participants used memory rather than the aid).

less mentally demanding than Search Assistance and No Assistance, with a higher perceived level of performance, but participants put less effort into learning and felt worse about their learning. These findings indicate that while navigation assistance was valuable for user experience, the fully-automatic predictive aid may have hindered participants' ability to learn the item locations.

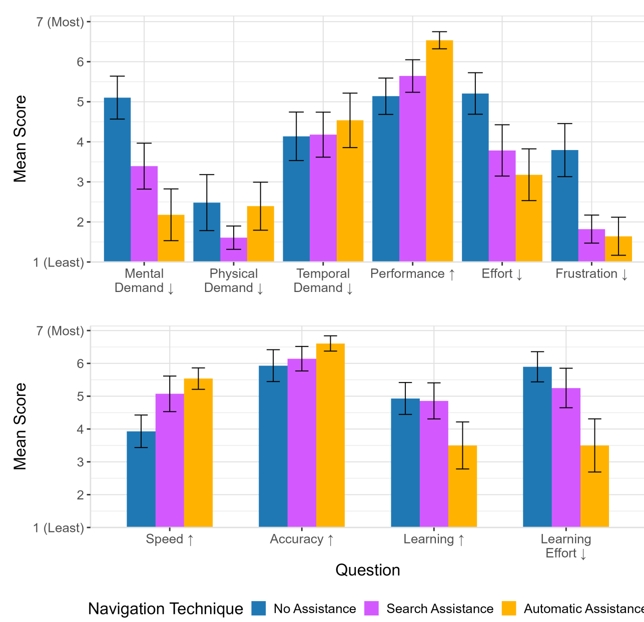

Figure 6: Mean Subjective Questionnaire Responses for NASA-TLX questions (top) and custom questions (bottom) following Assisted Use.

Question	DF (n,d)	F	p	η^2
Mental Demand ↓	2, 82	25.75	**<.001**	0.39
AA (2.18) < SA (3.39), **p<.05** AA (2.18), SA (3.39) < NA (5.10), **p<.001**				
Physical Demand ↓	2, 82	1.38	.26	0.03
Temporal Demand ↓	2, 82	0.75	.47	0.02
Performance ↑	2, 82	14.88	**<.001**	0.27
SA (5.64) < AA (6.54), **p<.01** NA (5.14) < AA (6.54), **p<.001**				
Effort ↓	2, 82	12.38	**<.001**	0.23
SA (3.79) < NA (5.21), **p<.01** AA (3.18) < NA (5.21), **p<.001**				
Frustration ↓	2, 82	21.03	**<.001**	0.34
AA (1.64), SA (1.82) < NA (3.79), **p<.001**				
Speed ↑	2, 82	12.29	**<.001**	0.23
NA (3.93) < SA (5.07), **p<.01** NA (3.93) < AA (5.54), **p<.001**				
Accuracy ↑	2, 82	2.57	.083	0.06
Learning ↑	2, 82	6.12	**<.01**	0.13
AA (3.50) < SA (4.86), **p<.05** AA (3.50) < NA (4.93), **p<.01**				
Learning Effort ↓	2, 82	11.93	**<.001**	0.23
AA (3.50) < SA (5.25), **p<.01** AA (3.50) < NA (5.90), **p<.001**				

Table 2: Factorial Analysis of ART-Transformed Survey Responses. Arrows indicate whether higher or lower values are better. Where significant differences exist, pairwise contrasts with group means are given below each factor (AA = Automatic Assistance, SA = Search Assistance, NA = No Assistance).

4.2 Memory Test

Accuracy in the memory test was measured as Manhattan distance between the target and the selection; results are shown in Figure 7. While No Assistance and Search Assistance resulted in comparable memory test performance, Automatic Assistance had the least accuracy (i.e., highest mean Manhattan distance). A one-way ANOVA found a significant effect of Navigation Technique on memory test accuracy ($F_{2,86}$ = 31.92, p<.001, η^2 = 0.43), with follow-up tests indicating that Automatic Assistance (mean Manhattan Distance 4.53) was less accurate than No Assistance (mean 1.21) and Search Assistance (mean 1.69), both p<.001.

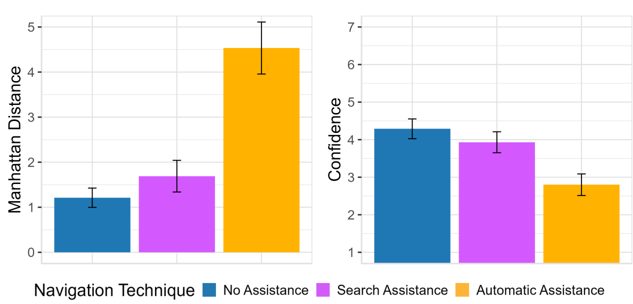

Figure 7: Accuracy and Confidence Ratings in the Memory Test. Accuracy (left) was measured as Manhattan Distance between the target and the selection (lower is better), and confidence scores (right) were self-reported by participants on a scale of 1-7 immediately following each selection (higher is better).

Confidence scores (shown in the right panel of Figure 7) generally mirrored the accuracy results, with Automatic Assistance resulting in the lowest confidence scores of the three Navigation Technique conditions.

Overall, the memory test results indicate that while training with Search Assistance resulted in comparable performance in recalling item locations to the unassisted baseline, use of Automatic Assistance significantly reduced spatial learning.

4.3 Unassisted Use Phase

After completing the memory test, participants completed eight blocks of selections with no assistance in order to assess performance in the absence of the navigation aid. Trial completion times are shown in Figure 8. While performance was comparable in the conditions that used Search Assistance and No Assistance, participants that used Automatic Assistance took about seven seconds longer to find item locations in the first block. Performance for the Automatic Assistance condition only converged with the other conditions after all eight blocks of repetitions.

A 3 × 8 (Navigation Technique × Block) ANOVA found a significant main effect of Navigation Technique and Block on trial

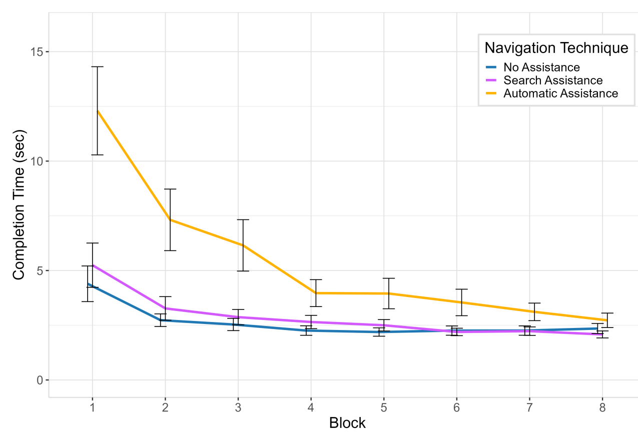

Figure 8: Trial completion time by block for each navigation technique in the unassisted use phase.

Factor	DF (n,d)	F	p	η^2
Navigation Technique	2, 87	20.26	**<.001**	0.22
Block	7, 609	81.61	**<.001**	0.27
Navigation Tech. × Block	14, 609	9.84	**<.001**	0.08

Table 3: Factorial Analysis of Trial Completion Times in the Unassisted Use Phase.

completion time, with an interaction effect (see Table 3). The interaction is primarily explained by the high starting point of trial completion times in the Automatic Assistance condition. Follow-up analyses indicated that No Assistance (mean 2.62s) and Search Assistance (mean 2.88s) were significantly faster than Automatic Assistance (mean 5.38s), with p<.001; however, there were no significant differences between No Assistance and Search Assistance.

Errors remained low in the unassisted use phase (mean 3.0%), with a one-way ANOVA identifying no significant differences in errors between the three Navigation Technique conditions ($F_{2,87} = 2.65$, p=.077, $\eta^2 = 0.06$).

5 Discussion

This work provides new results that contribute to our understanding of how predictive navigation assistance affects location learning:

- Both Automatic Assistance and Search Assistance made visual search substantially more efficient than the unguided baseline during assisted use.
- The number of characters typed decreased throughout use of Search Assistance, indicating that participants were switching to a memory-based strategy.
- Subjective responses indicated that prediction substantially improved user experience, with Automatic Assistance being perceived as best for performance and effort level but worst for learning during the assisted use phase.
- Automatic Assistance performed substantially worse on both the memory test and the unassisted trials, suggesting that perfect prediction hindered spatial learning.
- In contrast, we found no differences between Search Assistance and the No-Assist baseline on either the memory test or the unassisted trials, indicating that participants did learn item locations with this condition.

Overall, our results indicate that user-directed predictive navigation aids such as Search Assistance can successfully improve performance for inexperienced users, while still supporting a transition to memory-based location retrieval. In the following sections, we provide potential explanations for the results, generalize the findings to other useful contexts, and discuss limitations of the study which open avenues for future work.

5.1 Explanations for Main Results

5.1.1 Performance of No Assistance. In studies of location learning, no assistance is considered as the standard for comparison with other conditions (e.g., [11, 14]). In our experiment, the No Assistance navigation technique performed as expected: when participants had no experience with the items, visual search was slow, difficult, and frustrating; but the high level of effort required with the No Assistance resulted in robust spatial learning – on the memory test, selections were only one item away from the true location, and performance in the unassisted phase reached a stable asymptote after only one block.

5.1.2 Performance of Automatic Assistance. In the Automatic Assistance condition, we saw a marked trade-off in performance across the phases of the study: while there was a large improvement in performance during the assisted-use phase, there was a substantial reduction in performance both on the memory test and in the unassisted-use phase. Although the Automatic Assistance technique follows Fitchett's principle of rehearsal for navigation aids, our results suggest that the *type* of rehearsal may be more important than previously considered. Automatic Assistance provides rehearsal of the physical actions of moving the cursor to the target and clicking on it – but this "mechanical rehearsal" did not appear to be sufficient for participants to encode and reinforce the target's location in memory. Previous research has suggested that locations can be learned incidentally (i.e., without intention or effort) [23], but although participants in the Automatic condition did learn to some degree (that is, their answers on the memory test were substantially better than random), their learning was not nearly as strong as in the Search Assistance or No Assistance conditions, suggesting that other kinds of actions beyond mechanical rehearsal may be required as well.

Our study results suggest that the perceptual cues provided by Automatic Assistance may have allowed participants to simply "follow the highlight" without committing the location to memory. This aligns with previous research suggesting that users will adopt the lowest-cost strategy when finding information [11, 14], and that mechanical rehearsal alone does not lead to the adoption of memory-based selection [20]. Further, our results suggest that mechanical aspects of rehearsal may not trigger the cognitive processes that lead to complete learning of locations. Previous research on spatial learning suggests that only some elements of location memory are encoded automatically: in particular, Postma and colleagues suggest that while the experience of being in a location may be remembered automatically, learning the mapping between an object and its location may benefit from intentionality and effort [45].

The performance of Automatic Assistance raises questions about what would be required for users to successfully learn locations in this condition. In particular, the likelihood of less-than-perfect prediction could mean that a real-world version of this condition would demand a higher level of user engagement in the navigation, since people would need to perform an additional mental check to confirm whether the emphasized item is correct. Similarly, if the automatic predictor emphasized more than one item in the workspace (e.g., all items with scores above a threshold, similar to what is seen in Fitchett's techniques [16]), the additional mental effort required to decide which is the correct target could reduce the user's ability to simply "follow the highlight," potentially improving the encoding of locations. In future work, we plan to test the effects of different levels of accuracy in the automatic predictor, as well as emphasis of a set of items rather than a single item.

5.1.3 Performance of Search Assistance. Unlike Automatic Assistance, Search Assistance did not show a strong trade-off across the different phases of the study. Search Assistance substantially improved completion time when it was available (almost as much as Automatic Assistance), but this improvement was *not* coupled with a reduction in performance on either the memory test or the unassisted-use trials, compared to No Assistance. This suggests that despite Search Assistance making the visual-search task easier during assisted use, the technique still allowed participants to learn the target locations. The disparity in location learning between

Search Assistance and Automatic Assistance raises the question of what features of search-directed navigation aids may have led to the strong performance differences.

Unlike the Automatic condition, Search Assistance provides variable amounts of feedback depending on the user's actions. First, if the user does not type any search text, then Search Assistance is identical to No Assistance, and would require that the user carry out a full visual search. Second, if the user types one or two characters, then several potential targets are highlighted, meaning that the user must carry out a partial visual search (rather than simply moving to a single highlighted target). Third, the need to type search text in order to get location feedback meant that there was always a period of time between the appearance of the cue and the highlighting of the target; this time period may have given users the opportunity to think about the target's location, and thus engage in an additional kind of rehearsal than what was carried out with Automatic Assistance. That is, participants may have engaged in mental rehearsal in addition to mechanical rehearsal, simply by thinking about the target's location, which previous research has shown to be useful for reinforcing memory traces [4, 26]).

5.1.4 What do our results say about rehearsal? Our work helps to refine the understanding of the HCI principle of rehearsal. While Automatic Assistance only facilitated mechanical rehearsal for selecting items, and resulted in reduced location learning, Search Assistance provided an opportunity for both mental and mechanical rehearsal, which may have allowed participants to better reinforce spatial memories. Although previous techniques that employ the rehearsal principle (e.g., by Kurtenbach [33] and Fitchett [16]) may have implicitly involved both mechanical and mental rehearsal, the distinction between these types has not been clearly identified. Our results show that rehearsing an expert method's actions alone may not be enough, and that the opportunity for mental rehearsal of retrieving item locations may be an important part of the transition to expert memory-based interaction.

5.2 Generalizing the Findings and Lessons for Designers

There are several reasons we believe our results will generalize to the real world. First, the design of the GUI used in the study was based on a standard file explorer visualization which is familiar for most users, and the navigation aids used a simple visual emphasis (comparable to the highlighting used in systems like the Mac OS X Spotlight Search [7]). Second, we tested location learning in two ways: absolute spatial learning on the memory test, and the unassisted-use trials which tested how performance would be affected if the navigation assistance became unavailable – a scenario which is a realistic possibility for real-world users. Third, participants were crowd-sourced from a population with a wide range of ages and backgrounds, and the experiment was conducted remotely, meaning that participants completed the study on a wide variety of hardware, in their own familiar environments, and away from experimenter supervision.

Although this work indicates several directions for follow-up research, designers of navigation assistance techniques can already begin to use our findings to improve user experience for finding and selecting items while still supporting location learning. First, complex but spatially stable visual workspaces should include user-directed navigation aids such as search-directed prediction to show item locations. Second, if a prediction mechanism is highly accurate, automatic navigation aids may be problematic due to the possibility that users will become over-reliant on the perceptual highlighting, reducing location learning.

5.3 Limitations and Future Work

Although we analyzed the navigation techniques over three experiment phases and used a large number of participants, there are limitations to our study which provide directions for future work:

- *Prediction accuracy.* Predictions in Automatic Assistance were 100% accurate, which hindered location learning in our study. However, it is possible that less-accurate prediction could lead to different results – for example, if the prediction was 0% accurate, then users would ignore the navigation aid altogether and revert to visual search, which would lead to better location learning. Future studies will explore how different levels of prediction accuracy affect learning.
- *Visual Complexity.* Visual spaces can have a varying degrees of complexity, and may rely on images, icons, structural landmarks, or layout to present information to the user. While the simple grid visualization used in our study was similar to a standard file explorer and contained a large number of items, future studies are needed to test the effects of navigation aids in a variety of real-world visual workspaces.
- *Non-textual data.* Our Search Assistance condition relies on the availability of text-based filtering, but some workspaces may not have text labels attached to each item (or, the user may not use these as a primary lookup key). Further studies are needed to explore other user-directed navigation aids such as slider-based filters (although earlier studies suggest that these GUI controls can work well with spatially-stable visualizations [22]).
- *Long-term learning.* Our study used a blocked drill-and-practice approach for repeated selections, but real systems involve different frequencies for selecting items, and over various periods of time. Since we only assessed spatial knowledge following a fixed number of repetitions, it is pertinent to study the effects of navigation assistance on location learning over greater periods of time. Further, spaced practice over multiple learning sessions is known to impact acquisition and retention of skills [13]; therefore, we will conduct future studies to understand how navigation aids impact location learning over sequences of selections that more accurately reflect real interaction.
- *User strategies.* In our study, participants in the Search Assistance condition were instructed to select items without using the aid once they could remember the locations, but it is unknown whether participants would have discovered and switched to this memory-based strategy on their own (previous work has indicated that users do not always successfully transition to expert modes of interaction [10]). Therefore, we will conduct a future study to understand natural user behaviours that arise when using navigation aids, and assess differences between strategies.

- *Linking items and hierarchies.* Our study used a simple grid-based interface, but navigation aids can also be useful for navigating through hierarchies (like Fitchett's folder highlights [16]) or linking series of items (such as the procedural learning in Kelleher's stencil-based tutorials [31]). Future work will examine the robustness of predictive navigation aids for location learning over more complex and deeper visual information spaces.

6 Conclusion

Spatial knowledge is essential for users to make efficient selections in complex spatially-stable visual workspaces, but novices must rely on costly visual search to locate items. While some systems use predictive navigation aids to show beginners item locations, evidence suggests that users may never transition to memory-based interaction, which could mean performance detriments when the prediction is wrong or unavailable. We compared three navigation techniques in a study – Automatic Assistance always highlighted the target item, Search Assistance showed the subset of items matching a text query, and No Assistance provided no help to find target items. The techniques were assessed over three study phases: first, participants completed an assisted use phase with their assigned navigation technique; second, we administered a memory test to assess spatial learning; and third, participants completed an unassisted use phase to assess performance when assistance is unavailable. While both Automatic and Search assistance enhanced performance when the predictions were available, results showed a severe disadvantage for location learning with Automatic Assistance, but Search Assistance was as effective as the unassisted baseline. Overall, our results provide new understanding about how different forms of prediction can impact spatial learning, and suggest that designers should consider using Search Assistance to boost the performance of beginners while still supporting location learning in visual workspaces.

References

[1] David Abrams, Ron Baecker, and Mark Chignell. 1998. Information archiving with bookmarks: personal Web space construction and organization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Los Angeles, California, USA) *(CHI '98)*. ACM Press/Addison-Wesley Publishing Co., USA, 41–48. doi:10.1145/274644.274651

[2] J.E. Allen, C.I. Guinn, and E. Horvtz. 1999. Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications* 14, 5 (1999), 14–23. doi:10.1109/5254.796083

[3] Jackie Andrade and Peter Meudell. 1993. Short report: Is spatial information encoded automatically in memory? *The Quarterly Journal of Experimental Psychology Section A* 46, 2 (1993), 365–375. doi:10.1080/14640749308401051

[4] Edward Awh, John Jonides, and Patricia A Reuter-Lorenz. 1998. Rehearsal in spatial working memory. *Journal of Experimental Psychology: Human Perception and Performance* 24, 3 (1998), 780. doi:10.1037/0096-1523.24.3.780

[5] Olivier Bau and Wendy E. Mackay. 2008. OctoPocus: a dynamic guide for learning gesture-based command sets. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (Monterey, CA, USA) *(UIST '08)*. Association for Computing Machinery, New York, NY, USA, 37–46. doi:10.1145/1449715.1449724

[6] Richard A. Becker and William S. Cleveland. 1987. Brushing Scatterplots. *Technometrics* 29, 2 (1987), 127–142. doi:10.1080/00401706.1987.10488204

[7] Christopher Breen. 2005. *Searching in the Spotlight.* Retrieved May 20, 2026 from https://www.macworld.com/article/175478/tigerspotlight.html

[8] Elizabeth R Chrastil and William H Warren. 2012. Active and passive contributions to spatial learning. *Psychonomic bulletin & review* 19, 1 (2012), 1–23. doi:10.3758/s13423-011-0182-x

[9] Andy Cockburn, Carl Gutwin, and Saul Greenberg. 2007. A predictive model of menu performance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '07)*. Association for Computing Machinery, New York, NY, USA, 627–636. doi:10.1145/1240624.1240723

[10] Andy Cockburn, Carl Gutwin, Joey Scarr, and Sylvain Malacria. 2014. Supporting Novice to Expert Transitions in User Interfaces. *ACM Comput. Surv.* 47, 2, Article 31 (Nov. 2014), 36 pages. doi:10.1145/2659796

[11] Andy Cockburn, Per Ola Kristensson, Jason Alexander, and Shumin Zhai. 2007. Hard lessons: effort-inducing interfaces benefit spatial learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '07)*. Association for Computing Machinery, New York, NY, USA, 1571–1580. doi:10.1145/1240624.1240863

[12] Brian D Davison and Haym Hirsh. 1998. Predicting sequences of user actions. In *Notes of the AAAI/ICML 1998 workshop on predicting the future: AI approaches to time-series analysis*. 5–12.

[13] John J Donovan and David J Radosevich. 1999. A meta-analytic review of the distribution of practice effect: Now you see it, now you don't. *Journal of applied psychology* 84, 5 (1999), 795. doi:10.1037/0021-9010.84.5.795

[14] Brian D. Ehret. 2002. Learning where to look: location learning in graphical user interfaces. 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, 211–218. doi:10.1145/503376.503414

[15] Stephen Fitchett and Andy Cockburn. 2012. AccessRank: predicting what users will do next. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) *(CHI '12)*. Association for Computing Machinery, New York, NY, USA, 2239–2242. doi:10.1145/2207676.2208380

[16] Stephen Fitchett, Andy Cockburn, and Carl Gutwin. 2013. Improving navigation-based file retrieval. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 2329–2338. doi:10.1145/2470654.2481323

[17] Nat Friedman. 2021. Introducing GitHub Copilot: your AI pair programmer. Retrieved May 20, 2026 from https://github.blog/news-insights/product-news/introducing-github-copilot-ai-pair-programmer/

[18] Krzysztof Z. Gajos, Mary Czerwinski, Desney S. Tan, and Daniel S. Weld. 2006. Exploring the design space for adaptive graphical user interfaces. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (Venezia, Italy) *(AVI '06)*. Association for Computing Machinery, New York, NY, USA, 201–208. doi:10.1145/1133265.1133306

[19] Mathias Géry and Hatem Haddad. 2003. Evaluation of web usage mining approaches for user's next request prediction. In *Proceedings of the 5th ACM International Workshop on Web Information and Data Management* (New Orleans, Louisiana, USA) *(WIDM '03)*. Association for Computing Machinery, New York, NY, USA, 74–81. doi:10.1145/956699.956716

[20] Carl Gutwin, Andy Cockburn, and Benjamin Lafreniere. 2015. Testing the rehearsal hypothesis with two FastTap interfaces. In *Proceedings of the 41st Graphics Interface Conference* (Halifax, Nova Scotia, Canada) *(GI '15)*. Canadian Information Processing Society, CAN, 223–231.

[21] Carl Gutwin, Andy Cockburn, Joey Scarr, Sylvain Malacria, and Scott C. Olson. 2014. Faster command selection on tablets with FastTap. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 2617–2626. doi:10.1145/2556288.2557136

[22] Carl Gutwin, Michael van der Kamp, Jeremy Storring, Andy Cockburn, and Cody Phillips. 2020. Testing the Limits of the Spatial Approach: Comparing Retrieval and Revisitation Performance of Spatial and Paged Data Organizations for Large Item Sets. In *Proceedings of Graphics Interface 2020* (University of Toronto) *(GI 2020)*. Canadian Human-Computer Communications Society, 215 – 224. doi:10.20380/GI2020.22

[23] Lynn Hasher and Rose T Zacks. 1979. Automatic and effortful processes in memory. *Journal of experimental psychology: General* 108, 3 (1979), 356. doi:10.1037/0096-3445.108.3.356

[24] L.A. Hermens and J.C. Shlimmer. 1994. A machine-learning apprentice for the completion of repetitive forms. *IEEE Expert* 9, 1 (1994), 28–33. doi:10.1109/64.295135

[25] William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. 1992. Edit wear and read wear. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Monterey, California, USA) *(CHI '92)*. Association for Computing Machinery, New York, NY, USA, 3–9. doi:10.1145/142750.142751

[26] L. Himmer, M. Schönauer, D. P. J. Heib, M. Schabus, and S. Gais. 2019. Rehearsal initiates systems memory consolidation, sleep makes it last. *Science Advances* 5, 4 (2019), eaav1695. arXiv:https://www.science.org/doi/pdf/10.1126/sciadv.aav1695 doi:10.1126/sciadv.aav1695

[27] Eric J Horvitz, John S Breese, David Heckerman, David Hovel, and Koos Rommelse. 2013. The Lumiere project: Bayesian user modeling for inferring the goals and needs of software users. *arXiv preprint arXiv:1301.7385* (2013). doi:10.48550/arXiv.1301.7385

[28] Colby Johanson. 2024. *colbyj/bride-of-frankensystem.* doi:10.5281/zenodo.11176739

[29] Colby Johanson, Carl Gutwin, and Regan Mandryk. 2023. Trails, rails, and over-reliance: How navigation assistance affects route-finding and spatial learning

in virtual environments. *International Journal of Human-Computer Studies* 178 (2023), 103097. doi:10.1016/j.ijhcs.2023.103097

[30] Colby Johanson, Carl Gutwin, and Regan L. Mandryk. 2017. The Effects of Navigation Assistance on Spatial Learning and Performance in a 3D Game. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play* (Amsterdam, The Netherlands) *(CHI PLAY '17)*. Association for Computing Machinery, New York, NY, USA, 341–353. doi:10.1145/3116595.3116602

[31] Caitlin Kelleher and Randy Pausch. 2005. Stencils-based tutorials: design and evaluation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, Oregon, USA) *(CHI '05)*. Association for Computing Machinery, New York, NY, USA, 541–550. doi:10.1145/1054972.1055047

[32] Anjali Khurana, Xiaotian Su, April Yi Wang, and Parmit K Chilana. 2025. Do It For Me vs. Do It With Me: Investigating User Perceptions of Different Paradigms of Automation in Copilots for Feature-Rich Software. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 880, 18 pages. doi:10.1145/3706598.3713431

[33] Gordon Kurtenbach. 1993. *The design and evaluation of marking menus*. Ph. D. Dissertation. University of Toronto.

[34] Gordon Kurtenbach and William Buxton. 1993. The limits of expert performance using hierarchic marking menus. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems* (Amsterdam, The Netherlands) *(CHI '93)*. Association for Computing Machinery, New York, NY, USA, 482–487. doi:10.1145/169059.169426

[35] Gordon Kurtenbach and William Buxton. 1994. User learning and performance with marking menus. In *Conference Companion on Human Factors in Computing Systems* (Boston, Massachusetts, USA) *(CHI '94)*. Association for Computing Machinery, New York, NY, USA, 218. doi:10.1145/259963.260376

[36] Neal Lesh. 1997. Adaptive goal recognition. In *Proceedings of the Fifteenth International Joint Conference on Artifical Intelligence - Volume 2* (Nagoya, Japan) *(IJCAI'97)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1208–1214.

[37] Zhengxi Lu, Yuxiang Chai, Yaxuan Guo, Xi Yin, Liang Liu, Hao Wang, Han Xiao, Shuai Ren, Pengxiang Zhao, Guangyi Liu, and et al. 2026. UI-R1: Enhancing Efficient Action Prediction of GUI Agents by Reinforcement Learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 40, 21 (Mar. 2026), 17608–17616. doi:10.1609/aaai.v40i21.38816

[38] Pattie Maes. 1994. Agents that reduce work and information overload. *Commun. ACM* 37, 7 (July 1994), 30–40. doi:10.1145/176789.176792

[39] Pattie Maes and Robyn Kozierok. 1993. Learning interface agents. In *Proceedings of the Eleventh National Conference on Artificial Intelligence* (Washington, D.C.) *(AAAI'93)*. AAAI Press, 459–464.

[40] Aristides Mairena, Md. Sami Uddin, and Carl Gutwin. 2022. Accidental Landmarks: How Showing (and Removing) Emphasis in a 2D Visualization Affected Retrieval and Revisitation. In *Proceedings of Graphics Interface 2022* (Montréal, Quebec) *(GI 2022)*. Canadian Information Processing Society, 254 – 270. doi:10.20380/GI2022.25

[41] Toshiyuki Masui and Ken Nakayama. 1994. Repeat and predict—two keys to efficient text editing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) *(CHI '94)*. Association for Computing Machinery, New York, NY, USA, 118–130. doi:10.1145/191666.191722

[42] Justin Matejka, Wei Li, Tovi Grossman, and George Fitzmaurice. 2009. CommunityCommands: command recommendations for software applications. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology* (Victoria, BC, Canada) *(UIST '09)*. Association for Computing Machinery, New York, NY, USA, 193–202. doi:10.1145/1622176.1622214

[43] Ehsan Sotoodeh Mollashahi, Md. Sami Uddin, and Carl Gutwin. 2018. Improving revisitation in long documents with two-level artificial-landmark scrollbars. In *Proceedings of the 2018 International Conference on Advanced Visual Interfaces* (Castiglione della Pescaia, Grosseto, Italy) *(AVI '18)*. Association for Computing Machinery, New York, NY, USA, Article 30, 9 pages. doi:10.1145/3206505.3206554

[44] Allen Newell and Paul S Rosenbloom. 2013. Mechanisms of skill acquisition and the law of practice. In *Cognitive skills and their acquisition*. Psychology Press, 1–55.

[45] Albert Postma, Roy P.C. Kessels, and Marieke van Asselen. 2008. How the brain remembers and forgets where things are: The neurocognition of object–location memory. *Neuroscience & Biobehavioral Reviews* 32, 8 (2008), 1339–1345. doi:10.1016/j.neubiorev.2008.05.001 Special section: The European Workshop in Imagery and Cognition: Neurocognition and Visual Imagery.

[46] Aini Putkonen, Yue Jiang, Jingchun Zeng, Olli Tammilehto, Jussi P.P. Jokinen, and Antti Oulasvirta. 2025. Understanding visual search in graphical user interfaces. *International Journal of Human-Computer Studies* 199 (2025), 103483. doi:10.1016/j.ijhcs.2025.103483

[47] Joey Scarr, Andy Cockburn, Carl Gutwin, and Sylvain Malacria. 2013. Testing the robustness and performance of spatially consistent interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 3139–3148. doi:10.1145/2470654.2466430

[48] Joey Scarr, Andy Cockburn, Carl Gutwin, and Philip Quinn. 2011. Dips and ceilings: understanding and supporting transitions to expertise in user interfaces. 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, 2741–2750. doi:10.1145/1978942.1979348

[49] Joseph Laurence Scarr. 2014. *Understanding and Exploiting Spatial Memory in the Design of Efficient Command Selection Interfaces*. Ph. D. Dissertation. University of Canterbury. Computer Science and Software Engineering. doi:10.26021/3553

[50] Richard A Schmidt, Douglas E Young, Stephan Swinnen, and Diane C Shapiro. 1989. Summary knowledge of results for skill acquisition: support for the guidance hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15, 2 (1989), 352. doi:10.1037/0278-7393.15.2.352

[51] Katherine Schramm, Carl Gutwin, and Andy Cockburn. 2016. Supporting Transitions to Expertise in Hidden Toolbars. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 4687–4698. doi:10.1145/2858036.2858412

[52] Omar Shaikh, Valentin Teutschbein, Kanishk Gandhi, Yikun Chi, Nick Haber, Thomas Robinson, Nilam Ram, Byron Reeves, Sherry Yang, Michael S Bernstein, et al. 2026. Learning Next Action Predictors from Human-Computer Interaction. *arXiv preprint arXiv:2603.05923* (2026). doi:10.48550/arXiv.2603.05923

[53] B. Shneiderman. 1994. Dynamic queries for visual information seeking. *IEEE Software* 11, 6 (1994), 70–77. doi:10.1109/52.329404

[54] Jared Spataro. 2023. Introducing Microsoft 365 Copilot – your copilot for work. Retrieved May 20, 2026 from https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/

[55] Linda Tauscher and Saul Greenberg. 1997. How people revisit web pages: Empirical findings and implications for the design of history systems. *International Journal of Human-Computer Studies* 47, 1 (1997), 97–137. doi:10.1006/ijhc.1997.0125

[56] Anne M. Treisman and Garry Gelade. 1980. A feature-integration theory of attention. *Cognitive Psychology* 12, 1 (1980), 97–136. doi:10.1016/0010-0285(80)90005-5

[57] Md. Sami Uddin, Carl Gutwin, and Andy Cockburn. 2017. The Effects of Artificial Landmarks on Learning and Performance in Spatial-Memory Interfaces. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 3843–3855. doi:10.1145/3025453.3025497

[58] Marieke Van Asselen, Eva Fritschy, and Albert Postma. 2006. The influence of intentional and incidental learning on acquiring spatial knowledge during navigation. *Psychological Research* 70, 2 (2006), 151–156. doi:10.1007/s00426-004-0199-0

[59] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. 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, 143–146. doi:10.1145/1978942.1978963

[60] K. Yoshida. 1994. User command prediction by graph-based induction. In *Proceedings Sixth International Conference on Tools with Artificial Intelligence. TAI 94*. 732–735. doi:10.1109/TAI.1994.346414

[61] Ingrid Zukerman and David W Albrecht. 2001. Predictive statistical models for user modeling. *User Modeling and User-Adapted Interaction* 11, 1 (2001), 5–18. doi:10.1023/A:1011175525451