Perceptual Grouping: Selection Assistance for Digital Sketching

David Lindlbauer[1], Michael Haller[1], Mark Hancock[2], Stacey D. Scott[3], Wolfgang Stuerzlinger[4]

[1]Media Interaction Lab, University of Applied Sciences Upper Austria, Hagenberg Austria;
[2]Management Sciences, [3]Systems Design Engineering, University of Waterloo, Waterloo, Canada;
[4]York University, Toronto, Canada

{david.lindlbauer|michael.haller}@fh-hagenberg.at;{mark.hancock|stacey.scott}@uwaterloo.ca;
wolfgang@cse.yorku.ca

ABSTRACT

Modifying a digital sketch may require multiple selections before a particular editing tool can be applied. Especially on large interactive surfaces, such interactions can be fatiguing. Accordingly, we propose a method, called Suggero, to facilitate the selection process of digital ink. Suggero identifies groups of perceptually related drawing objects. These "perceptual groups" are used to suggest possible extensions in response to a person's initial selection. Two studies were conducted. First, a background study investigated participant's expectations of such a selection assistance tool. Then, an empirical study compared the effectiveness of Suggero with an existing manual technique. The results revealed that Suggero required fewer pen interactions and less pen movement, suggesting that Suggero minimizes fatigue during digital sketching.

Author Keywords

perceptual grouping, digital sketching, sketch analysis, selection assistance, interactive whiteboard

ACM Classification Keywords

H5.2 User Interfaces: Graphical User Interfaces; I.2.10 Vision and Scene Understanding: Perceptual Reasoning

General Terms

Human Factors; Design; Measurement.

INTRODUCTION

Digital sketching environments offer many possibilities for creating and modifying content. The ability to freely modify complex sketches at any point enables more fluid work compared to paper-based sketching. However, performing such modifications can still be time-consuming and cumbersome with existing digital sketching tools.

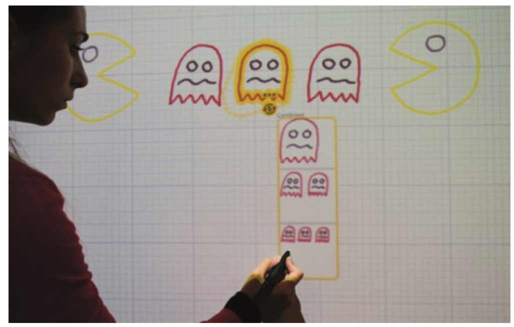

Figure 1: Suggero in action on an interactive whiteboard.

Current sketching applications require precise and potentially repetitive selection of drawing objects before a desired modification can be made. This problem is exacerbated on large interactive surfaces [22], where arm fatigue may play a role. To help minimize the interaction required, previous research investigated techniques to infer perceptual structures from drawings [14,29,30], to analyze selection gestures for perceptual information [6], and to provide suggestions based on one's selection [10]. In contrast to these algorithms, humans are very skilled at visually identifying object groups. We visually perceive a set of drawing objects as grouped, based on different perceptual features such as proximity or similarity [8,25,26,33,34,35,36]. In sketches, such *perceptual groups* are frequently the target of modifications, such as moving, rotating or recoloring. Insights from perception research have been used in computer science research, including sketch recognition algorithms (e.g., [5,18]), interactive beautification (e.g., [13]) and perceptual organization [14,23,28,29,30,32].

In this paper, we explore the ability of a selection assistance tool to leverage perceptual grouping principles to identify and suggest potential selection options during digital sketching. We draw on insights from Gestalt Theory [8,36] and Feature Integration Theory [33,34,35] to identify perceptual groups in real-time during digital sketching. These groups are then used to suggest potential extensions when a person begins a manual selection (Figure 1).

Our main contribution is a new selection method, Suggero, for digital sketching that uses a new three-step grouping approach. We also provide empirical evidence that Suggero requires less physical effort for selection of perceptual groups, which is especially beneficial for reducing fatigue on large interactive vertical surfaces.

RELATED WORK

Perceptual Psychology

One of the best-known theories on human perception and perceptual grouping is the Gestalt Theory, which identifies important factors for perceptual groups of visual objects, such as proximity, common region, and similarity [8,26,36]. Feature Integration Theory, by Treisman and Gelade, describes how similarity cues, such as color or shape, are pre-attentively processed [33,34,35]. Additionally, they introduced *distractors* and *object boundaries,* structures that are visually salient for humans because of their properties, such as rotation or similarity. Regardless of the theoretical explanation, humans are quite skilled at perceiving visual structures. Such structures visually "pop-out" of images and sketches, a phenomenon used in this work and many others.

Perceptual Organization

Results from human perception research have been used in a variety of fields, particularly to infer underlying structures and groupings. Thórisson [32] used object proximity and similarity to find perceptual groups and discusses the usage of such groups for interaction. Igarashi et al. [14] used proximity and regularity to find structures in card stacks. Similarly, Shipman et al. [30] identified visual structures in their pen-based whiteboard system and then used gestures to interact with them. Additionally, their system permitted people to create borders and to use them to delimit groups. Rome's work [28] built on Feature Integration Theory to deal with similarity, but used also proximity. Saund et al. [29] extracted line art and blobs from input images and identified groups, based on different kinds of paths, including closure. Nan et al. [23] used several Gestalt cues to simplify architectural drawings. They developed a model for Gestalt principles and showed ways to combine them. Our work contributes to this literature by similarly using these principles for selecting such perceptual groups and we present empirical evidence of the fatigue reducing benefits.

Perceptual Selection

The research closest to ours uses both perceptual grouping and intelligent gesture interpretation to facilitate selection. Dehmeshki and Stuerzlinger [6,7] focused on the selection of perceptual groups, based on proximity, regularity, and path continuity. Then they analyzed gestures to detect the best-matched grouping for selection. Xu et al. [37] used proximity, shape similarity, and common region to identify perceptually salient groups in sketches consisting of non-overlapping closed polygons. They applied their Lazy Selection tool to groups that best matched the path, area, and speed of the selection gesture. Lazy Selection requires all intended objects to be touched by the selection stroke, which can be fatiguing, especially on large surfaces. In our work, users need only select a single object and then interact mainly with presented suggestions (which appear close by). Our grouping algorithm also uses different analysis methods to capture additional potential perceptual groups.

SUGGERO: PERCEPTUAL-BASED OBJECT GROUPING

Large interactive walls are a great medium for sketching, giving users the chance to sketch ideas, draw graphs, brainstorm, and annotate content. Here, we refer to the strokes users draw during sketching as *objects*. For editing, selecting desired objects in sketches is more complex as current systems do not provide good tools to interact with objects that are perceptually grouped by humans. Especially on large interactive walls, selecting groups of objects in a cluttered or wall-spanning sketch or objects that overlap can be cumbersome. The increased workload can exacerbate the well-known problem of fatigue on vertical interactive displays. To address this, we present Suggero, a technique that automatically suggests groups based on the perceptual theories discussed above. We use the perceptual cues of *proximity, endpoint connectivity, parallelism,* and *similarity of shape, color and thickness* as features for identifying all perceptually related objects. The identified perceptual groups are presented to users as *selection suggestions*, based on an initial selection of one or more objects.

We first present the underlying algorithms used for identifying these groups. Suggero uses a three-step grouping approach, consisting of *Pre-processing*, *Feature Extraction*, and *Dynamic Grouping*. We then discuss how the identified perceptual groups are then used to assist selection.

Pre-processing

We implemented Suggero in an existing sketching application on a large interactive wall, where users can draw freely. Suggero extends the Lasso and Harpoon tools in that system. Since Suggero is targeted at hand-drawn sketches, input strokes are collections of 2D points (polylines). Each stroke is re-sampled immediately to ensure reasonably uniform sample density. The coordinates, color, and thickness of strokes are stored for later processing. Additionally, a 2D bounding circle is calculated for subsequent object proximity calculations.

Feature Extraction

In Suggero, sketches are analyzed for different perceptual features that represent the perceptual relation(s) between drawing objects. For each feature, the pairwise perceptual relation between objects is expressed in an *affinity value*. We define affinity values as normalized values ranging from 0 to 1, where 0 means no relation between objects and 1 stands for highly related objects (e.g. objects of identical color have a value of 1 for the color affinity). These affinity values are used later for Dynamic Grouping.

Feature Choice in Suggero

A combination of features from perception research [25,26,33,34,35,36] is used during Feature Extraction. The primary ones are *proximity* and *similarity*, as Gestalt Theory identifies them as key features; they have also been successfully applied in previous work [6,14,32]. In addition to proximity, *endpoint connectivity* is used as a visually important feature. Also, Feature Integration Theory identifies *similarity of shape and color* as critical features for human perception [33,34,35]. Given the application context, *similarity of thickness* is also analyzed, since it is a strong visual feature in a sketch. Finally, *parallelism* is included as it is also a strong visual feature [5].

Proximity

Two different proximity measures are used: *global* and *local proximity*. Both measures are used as separate features, thus making Suggero able to detect enclosing structures as well as perpendicular objects (e.g. lines). *Global proximity* refers to the distance between two objects, which also considers the length and shape of each object. The distance between two strokes can vary, as different ends may have different distances. Thus, measuring the distance for only a single point (e.g. center point or central moment) or an endpoint is not sufficient. For a more robust measure, we compute the average of the distance of 10 equally spaced point pairs along the two strokes (Figure 2, *left*). *Local proximity* refers to the situation where an object is contained within another one, which is perceived as spatially close [26]. In Suggero, we compute a *local proximity* measure from the distance between the centers of the bounding circles of both objects (Figure 2, *right*).

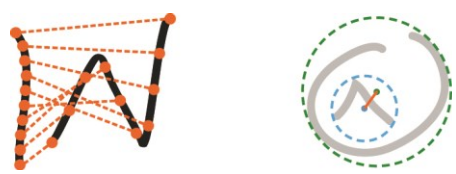

Figure 2: Global proximity (left) and local proximity (right).

Endpoint Connectivity

Connected strokes can be an important feature for visual grouping. Two line segments can be perceived as connected when they intersect or when their endpoints connect. Yet, endpoints of objects do not need to have a real connection for the whole shape to be perceived as closed. To detect and automatically close such gaps, we use *tolerance zones* [31], originally propsed to merge objects in sketch recognition. We use this method to compute affinity values based on endpoint connectedness. For this, we modified the original algorithm by computing the size of the tolerance zones as the average distance of an endpoint to all other points of an object combined with the average distance to all other tolerance zones. This modified version is faster and can handle strokes that are non-equidistant sampled.

Parallelism

Parallel structures of sketched objects are often non-accidental and are easily recognized by humans [5]. Parallel lines are perceived as such because they are at the same angle and no intersections. The difference in angle is a good measure for the *degree of parallelism*. No-difference means that the lines are parallel and a difference of 90° means that the lines are perpendicular. The same concept can be applied to non-straight strokes and other shapes. The similarity between shapes is also important if objects are perceived as parallel (Figure 3). An object can also be parallel to a sub-segment of another (Figure 3, *right*). In Suggero, the pairwise degree of parallelism is computed for all objects by combining the difference in rotation and similarity [17].

Figure 3: Parallelism is perceived due to common orientation (left) and similarity (middle), also sub-segments (right).

Similarity

The perceptual relationship between objects strongly depends on their similarity in shape, color, and other properties (e.g. stroke thickness). In Suggero, pairwise shape similarity is computed between all objects using Fourier shape descriptors [4]. Color similarity is computed in the CIELab space, as it optimally represents human perception of color differences. Finally, the similarity for stroke thickness is computed from the pairwise difference of the average thickness. The resulting affinity values are normalized to the largest stroke thickness in the sketch.

The output of the Feature Extraction phase is a collection of pairwise affinity values for all objects based on all the features. The affinity values for each feature are then normalized (with respect to all objects in a sketch) and used as input to the Dynamic Grouping phase.

Dynamic Grouping

The groups we perceive in sketches may change with every object that is added, removed, or modified. In the Dynamic Grouping phase the output from the Feature Extraction phase is processed and perceptual groups are identified.

Similarity-based grouping

Suggero groups all objects based on their perceptual relationship in a hierarchical manner, based on insights into how human perception works [15]. Dynamic Grouping begins by identifying the object pair with the strongest affinity and grouping them. In the next step, the next closest object pair is found and assigned to a new group and so on. Once created, each group is treated like a regular object; that is, pairwise relations are calculated between the remaining (un-grouped) objects and existing groups. If the closest pair is an object and a group, the object is added to the group. This process is continued until all objects are grouped accordingly (Figure 4). All identified groups are stored for later ranking. This grouping method is an implementation of *Hierarchical Agglomerative Clustering* (*HAC*) [20], a greedy, bottom-up grouping approach.

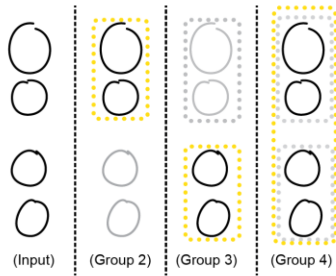

Figure 4: Illustration of the dynamic grouping

The *HAC* technique meets the requirements of grouping the objects based on their relative similarity rather than using fixed thresholds or other rigid methods. Additionally, we also calculate a group-quality value (*confidence value)* for every group as discussed below.

Input Preparation - Feature Combination

In Suggero, affinity values from multiple features are used. Thus, the input for *HAC* has to be pre-processed. The outputs from the feature extraction phase are multiple matrices with affinity values (one per feature). Two different strategies are used to combine them for feature combination. The first strategy combines all matrices into a single matrix via a weighted sum, which is then used by HAC. The weights for this are empirically determined and were initially set by hand and later tuned based on the results of the preliminary study. Since different features, such as shape similarity and parallelism, are combined in this step, it may be hard for users to tell which features formed a particular group. Although this strategy produces complex groups, combining all features into a single matrix often matches the grouping that users intend to see. The second strategy creates an affinity matrix for each feature and processes each separately with *HAC*. The resulting groups are more specialized, since they cover only a single feature. The computed groups from both strategies are then collected, ranked, and merged to remove duplicated groups. By combining both strategies, we cover a greater number of scenarios, making Suggero more flexible, while still being robust.

Group Confidence Value

In order to show users only the best groupings, they have to be ranked by quality (Figure 5). In Suggero, the confidence value of a group is calculated by averaging the pairwise affinity values of a group's objects combined with a penalty function for group size (*confidence value* = *average affinity values* × *group size* × γ). In our implementation, γ is set to 0.9 to favor smaller groups over large ones. Since users are presented with smaller groups first, they can grow the selection easily as needed.

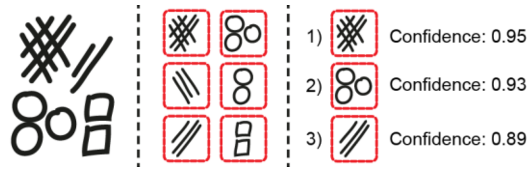

Figure 5: Input sketch for dynamic grouping (left), identified groups (middle) and ranked groups (right).

Using Suggero to Assist Selection

Suggero assists selection by presenting a set of suggested extensions after an initial manual selection occurs. We describe here how these suggestions are determined and visualized, and how the initial manual selection occurs.

Suggestions

When an object is manually selected, we search the precomputed perceptual groups for all those containing this object. Groups are ranked based on their confidence value and then presented as *suggestions*. Users can also select multiple objects; Suggero then searches for groups containing all those objects. In general, selecting more objects gives Suggero more information and results in more exact suggestions. We chose to provide users with multiple suggestions, a common approach to resolve ambiguities [19].

Visualization: Linear Menu

A linear menu is used to display the suggestions in decreasing order of confidence (best on top). Each suggestion shows a perceptual group of objects. All groups are scaled to the same size (70 × 70 pixels) and are shown in their original color to facilitate identification.

Since Suggero is designed for large interactive surfaces, the menu is placed adjacent to the initial selection to minimize user movement. If a user taps a suggestion, Suggero then selects all objects contained in said suggestion (Figure 6). Afterwards, the suggestions shown in the linear menu are updated based on the revised selection. This allows users to quickly grow their selection to very large groups.

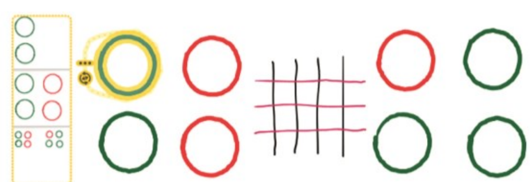

Figure 6: When the green circle is manually selected, Suggero then provides suggestions.

Initial Selection: Harpoon

Users can tap an object to select/deselect it. Beyond this, Suggero uses the manual selection technique Harpoon [16] to facilitate initial selection(s). Harpoon enables the selection of on-screen objects by "crossing", i.e., drawing through, them with a pen. Each time the Harpoon tool crosses an object its selection state is toggled. Harpoon is speed-dependent: the faster the stylus moves, the bigger the selection area and the more strokes are selected. Harpoon selection is faster than tapping or lassoing for more than one object [16]. It is suitable for both small, specific object selection in cluttered sketches as well as large-scale selection. Still, selections in cluttered areas or selections of overlapping objects can be cumbersome, even if objects form a perceptually salient group easily identifiable at a glance.

Performance and Computational Load

Suggero is constantly analyzing user input, therefore it is running in a separate background thread to avoid blocking the user interface. Additionally, calculated properties such as bounding spheres get preserved to increase performance and operations like stroke comparison are running in multiple parallel threads. As a result of these optimizations, Suggero runs in real-time, needing about 30-100ms for feature analysis and clustering, depending on the number of strokes in a sketch (up to a few hundred strokes). This is sufficiently fast, since users also have to switch tools between drawing and selection mode.

STUDYING ASSISTED SELECTION

We first conducted a preliminary study to elicit people's expectations about perceptual grouping and to gather participant-generated drawings for the second study. We also used results from the preliminary study to fine-tune parameters of the Suggero algorithm. We then performed an empirical laboratory-based user study to evaluate performance and to discover how people use the Suggero technique.

PRELIMINARY STUDY

In the preliminary study, users drew sketches on an interactive display without Suggero. We then asked them to manually select objects to provide training data for Suggero.

Participants

Ten participants (4 female), between 20 and 39 years old (*Mdn*=25.5), were recruited from a local university. All had experience with pen- or touch-based devices, and two had experience with interactive whiteboards.

Apparatus

This study was conducted on a 70-inch interactive whiteboard with a Hitachi CP-AW251N 1280×800 pixel ultra short-throw projector (~8.3 pixels/cm). Input was provided with an Anoto digital pen (ADP-301). Participants could freely draw, change stroke thickness, color or erase strokes, and select items with Harpoon in the sketching application.

Procedure

In a 10-minute training period, the whiteboard and sketching application was explained and participants drew a training sketch. Tasks included replicating four template drawings (black and white) using at least 4 different colors and stroke thicknesses to generate variety in the collected sketches. After each drawing was completed, the participant performed five trials in which they provided sample selections. In each trial, a candidate drawing objects was randomly chosen and the participant was asked to identify four groups of objects that included the candidate, in descending order of relevance. Each session lasted approximately 60 minutes, and each participant performed a total of 80 trials (4 drawings × 5 candidates × 4 groups).

Results & Discussion

Participants produced a variety of drawings based on the four provided templates (Figure 7), two of which were also used for the second study.

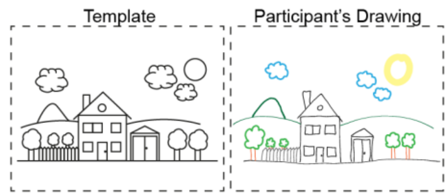

Figure 7: One of the four provided templates (left) and the participant's drawing (right).

We observed that participants often had difficultly identifying a third and fourth group based on a single candidate, indicating that, given the right suggestion set, 2-3 options may be sufficient. We also used the suggestions elicited from participants to manually fine-tune parameters of the grouping algorithm. In the future, we intend to use machine learning to automate the process of parameter tuning.

EXPERIMENT 2: PERFORMANCE AND USAGE

It may seem clear that for the selection of large numbers of objects some form of assistance may prove useful. Yet, the question whether providing suggestions for selections will hinder or help, remains open. Specifically, the cognitive load necessary to identify the appropriate group and to task-switch between selecting strokes and identifying that group may outweigh any performance gain achieved by reducing a large number of selections to a single tap. Moreover, on large interactive walls, objects targeted by user selections may be spread over a large area, which may lead to increased fatigue. Thus, determining both the cognitive as well as physical workload (in terms of movement time and distance) is important. Our second study was designed to examine this tradeoff and permitted us to observe how people use Suggero. This study consisted of two parts: in the first part, we compared Suggero to Harpoon; in the second, we observed participants while modifying a realistic drawing with Suggero for additional insights.

Participants and Apparatus

We recruited 18 paid participants (8 female) from a local university (22 – 40 years, *Mdn*=26.5). Participants (2 left-handed) controlled the stylus with their dominant hand. Five participants had experience with interactive drawing applications. Seven had experience working with interactive whiteboards. The apparatus was the same as before, except that algorithm parameters had been tuned. Drawings and target selections were displayed on an adjacent wall to the right of participants, in the same scale as their sketch.

Experimental Design
We used a 2 (technique) × 3 (complexity) repeated-measures design.

Factor 1 - Technique: We compared Suggero with the Harpoon manual selection technique. We chose Harpoon as a baseline since the initial, and all subsequent, selections in Suggero can be performed with it. Harpoon outperforms Tapping and Lassoing [16] and is a state-of-the-art technique, especially on large surfaces.

Factor 2 - Complexity: Suggero requires cognition for the initial selection, but improves performance by suggesting potential completions. The tradeoff between them revolves around complexity. Specifically, we suspected that there would be a "sweet spot", where selection was sufficiently complicated that manual selection would be too tedious, but simple enough that Suggero would still be capable of providing useful suggestions for more performance. Thus, in addition to varying whether or not we provided suggestions, we also adjusted the level of complexity. We first describe the theoretical foundation for group complexity in the study. Using different complexities for target selections is a common approach for selection research [6,10,16,21].

Two important aspects of complexity are the *visual complexity* of the content and *selection complexity* of the target selection. The *visual complexity* of a sketch can be increased by increasing the number of sides (or turns) [2,3] of an object, its "figural goodness" [16, p. 398], or by increasing the quantity of objects with different properties like shape or color [24]. Another way is to scatter objects. Sketches are perceived to be most complex if no object equals another and no visually salient subsets are present. *Selection complexity* can be modeled as a combination of Fitts' law [9], the Hick-Hyman law [11,12], and the Steering law [1], and depends on the selection technique used.

Visual and selection complexity were combined to create three levels of complexity: *simple*, *challenging*, and *arbitrary*, with increasing number of objects. We did not examine selections with high visual complexity but low selection complexity since we believe this is a rare scenario. *Simple* sketches had low visual and low selection complexity for both techniques. Here the target drawing consisted of 42 objects (Figure 8, *left)*. *Challenging* sketches were slightly more visually complex. They had low selection complexity for our Suggero technique, but a high selection complexity for manual selection (83 objects, Figure 8, *middle*). This challenging condition represents circumstances when many, perceptually related objects have been drawn, but the configuration might interfere with later selection. *Arbitrary* sketches had both high visual and high selection complexity (132 objects, Figure 8, *right*) for both selection techniques. These sketches were generated by arbitrarily choosing objects, thus minimizing perceptual relationships.

Our expectation was that the *challenging* sketches would be the "sweet spot" where Suggero would outperform manual selection, but that manual selection would outperform Suggero for *simple* selections. We expected the benefit of Suggero to no longer hold for the *arbitrary* condition.

We used an *abstract pattern* (Figure 8) to control the different factors of *visual* and *selection complexity*. Participants were asked to replicate the patterns themselves to avoid a potential bias and to be able to analyze the performance of Suggero selections in sketches that were drawn differently compared to our templates.

Procedure
Participants were briefly introduced to the setup and the purpose of the experiment, followed by a 15 minute training session. During the training, participants were guided through the process of creating a sketch using a practice drawing and performed a minimum of 5 selections with both techniques. Participants were trained with both selection techniques and were quickly able to produce good or optimal selection of targets, even with Harpoon. For each level of complexity, participants were asked to select objects on a sketch provided by the experimenter. Participants began each trial by tapping the start button, then the target selection was shown at the participant's right, and once they had performed the selection, ended the trial by tapping the end button. Participants were instructed to perform selections as quickly and accurately as possible. With Suggero, participants were additionally asked to *select as few objects as possible with manual selection*, in order to encourage them to use the provided suggestions. After completing all trials for each technique, participants filled out a questionnaire. The order of levels of complexity was counterbalanced using a Latin square. Each complexity corresponded to an abstract pattern and target selections specific to that pattern. Selections were determined in advance (the same way for each technique and all participants).

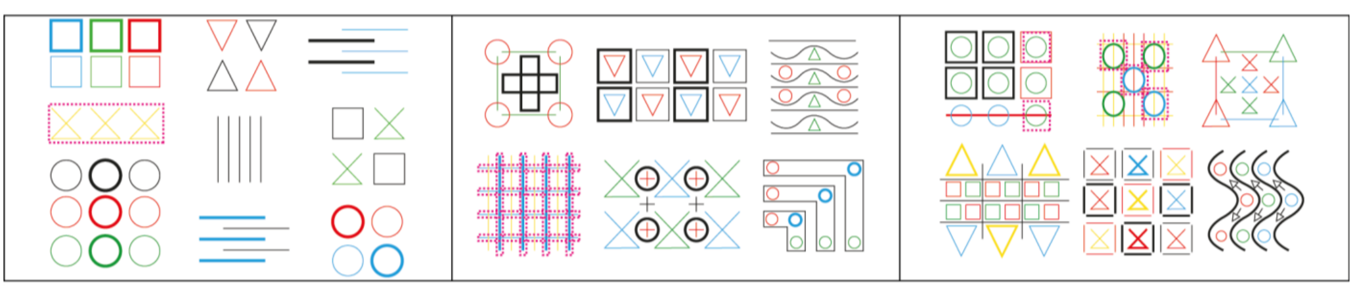

Figure 8: Simple (left, 42 objects), Challenging (middle, 83 objects) and Arbitrary (right, 132 objects). Target selections were marked with dotted, red lines. Figure shows one (of twenty) target selections for each complexity.

Each participant performed a total of 120 trials (2 techniques × 3 complexities × 20 selections). In the next phase of the study, an incomplete sketch, with 30% of the strokes manually removed from the initial drawing (a participants' drawings from the preliminary study), was provided and participants were asked to complete the drawing as desired for 5 minutes. Participants were then asked to perform selections with Suggero and to interact with the selected objects (moving, rotating) in a 5-minute speak-aloud session. The entire session lasted about 90 minutes. Every stylus movement on the interactive whiteboard and every action in the software was logged. All sessions were audio and video recorded.

RESULTS & DISCUSSION

Performance Results

We performed a 3 (*Complexity*) × 2 (*Technique*) repeated measures ANOVA (α=.05) on four dependent measures: task completion time, movement time, interaction count, and movement distance. The Greenhouse-Geisser correction was used when Mauchly's test of sphericity was violated (influencing *df*, *F*, and *p* values). Bonferroni adjustments were used for post-hoc analyses.

Trial completion time

Trial completion time was defined as the time between tapping the start and end buttons. There was a main effect of complexity ($F_{2,34}$=168.4, p<.001). Post-hoc pairwise comparisons revealed that all three levels of complexity were significantly different (p<.05). Participants were fastest for *simple* (M=4.6 s, SE=0.3 s), followed by *challenging* (M=8.8 s, SE=0.6 s) and *arbitrary* (M=19.1 s, SE=1.0 s). We found a main effect of *technique* ($F_{1,17}$=88.3, p<.001) with Harpoon (M=7.9 s, SE=0.6 s) being faster than Suggero (M=14.0 s, SE=0.7 s). There was also an interaction between complexity and technique ($F_{2,34}$=29.5, p<.001). Post-hoc tests revealed that for each level of complexity, all pairwise differences between techniques were significant (p<.05); however the difference between the two techniques was larger for *arbitrary* (Suggero: M=25.4 s, SE=1.6 s; Harpoon: M=12.8 s, SE=1.0 s) than for *simple* (Suggero: M=6.2 s, SE=0.4 s; Harpoon: M=3.6 s, SE=0.3 s) and *challenging* (Suggero: M=10.3 s, SE=0.7 s; Harpoon: M=7.2 s, SE=0.6 s). We suspect that Harpoon was faster due to the experimenter's instruction to minimize the number of manually selected strokes when using Suggero, as participants were observed sometimes to spend time on determining a strategy to perform a target selection. This led to participants being faster for simple and challenging sketches with the Harpoon technique (which omitted these instructions). To better understand the components of action required to perform selections, we broke our dependent measure down into: real pen movement time, movement distance, and interaction count. Measuring time from the first pen down event would exclude any cognitive time in the measurements, which would bias in Suggero's favor.

Movement Time

Movement time was calculated as the total amount of time per trial that the stylus was touching the surface. Results showed a main effect of complexity ($F_{2,34}$=10.7, p<.001) with increasing time between *simple* (M=0.6 s, SE=0.1 s), *challenging* (M=0.7 s, SE=0.1 s) and arbitrary (M=1.2 s, SE=0.2 s), and all pairwise differences were significant (p<.05). Additionally, we found a main effect for technique ($F_{1,17}$ =7.7, p<.05) with Suggero (M=0.7 s, SE=0.1 s) requiring significantly less movement time than Harpoon (M=1.0 s, SE=0.1 s). Pairwise post-hoc tests showed that Participants spent less time with Suggero than with Harpoon for the *simple* (p<.001) and *challenging* (p<.05) condition (Figure 9). For *arbitrary,* the difference was not significant (p=.280). While Suggero had a longer trial completion time, a closer look revealed that Suggero required less movement time for *simple* and *challenging*. Movement time is an important factor in performance and fatigue.

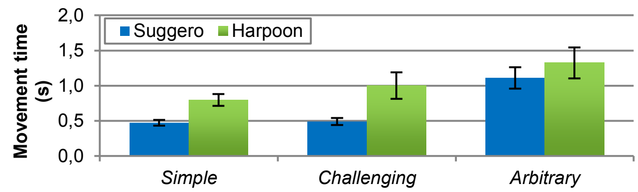

Figure 9: Movement time by Complexity and Technique

Movement Distance

Movement distance was defined as the distance participants moved the stylus on the interactive whiteboard in pixels. There was a main effect of complexity ($F_{2,34}$=199.4, p<.001). Post-hoc pairwise comparisons revealed that the difference between *simple* (M=153.1 px, SE=9.8 px) and *challenging* (M=94.5 px, SE=13.0 px) was significant (p<.01) as well as the difference between *challenging* and *arbitrary* (M=134.2 px, SE=13.7 px, p<.05). The difference between *simple* and *arbitrary* was not significant (p=.579). We also found a main effect of technique ($F_{1,17}$=56.3, p<.001), with participants moving the stylus significantly less with Suggero (M=72.5 px, SE=5.7 px) than with Harpoon (M =180.7 px, SE=15.6 px). There was also an interaction between *complexity* and *technique* ($F_{1.4,23.6}$=18.9, p<.001). Pairwise post-hoc tests revealed that Suggero required less movement for all three complexity conditions (p<.05, Figure 10). As with movement time, Suggero requires significantly less pen movement in terms of distance. Since selections of more objects normally requires more movement, avoiding it is important, especially for distant objects like in the *challenging* condition.

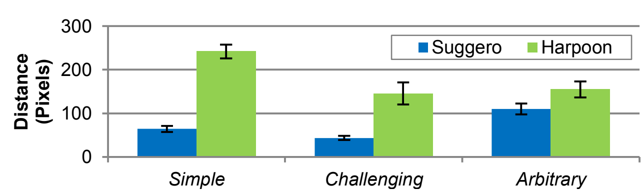

Figure 10: Stylus movement distance by Complexity and Technique measures in pixels.

Interaction Count

Interaction count was defined as the number of times participants performed a stroke (touch, optional move, and lift of the pen). Interactions with the Suggero's suggestions were also included in this count. Results showed a main effect of complexity ($F_{2,34}$=155.9, p<.001) with increasing interactions per complexity (*simple*: M=1.9, SE=0.1; *challenging*: M=4.1, SE=0.2; *arbitrary*: M=6.3, SE=0.26), which were all pairwise significantly different (p<.05). We found a main effect of *technique* ($F_{1,17}$=15.1, p<.001) with Suggero (M=4.4, SE=0.2) needing more interactions than Harpoon (M=3.8, SE=0.2). There was also an interaction between complexity and technique ($F_{1.1,19.4}$=12.8, p<.01). Pairwise post hoc tests revealed that Harpoon needed significantly fewer interactions for *simple* (p<.001) and *arbitrary* (p<.05). For *challenging,* Suggero needed significantly fewer interactions (p<.05).

This interaction can be seen in Figure 11, which also indicates through shading when Suggero interactions were with suggestions. Harpoon required fewer interactions for *simple* and *arbitrary*, while Suggero requires less interaction for *challenging*. Taking a deeper look at the kind of interaction reveals that a large part of the interactions for *simple* and *challenging* were with suggestions. These are essentially just tapping actions and thus require neither much time nor effort. Although the design of Suggero makes selection easier by using exactly these kinds of actions, the necessary cognitive effort still increases the overall selection time.

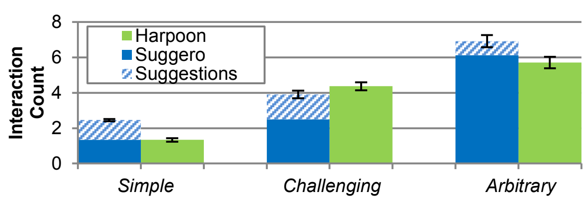

Figure 11: Interaction count by Complexity and Technique. Shaded areas indicate interactions with suggestions.

Suggero Accuracy

With Suggero, it is possible to manually select a few strokes, then choose a suggestion in Suggero to expand the selection, then to again select additional strokes manually, to use suggestions again, and so on. To better understand people's strategies when using Suggero, a more detailed analysis of all Suggero trials was conducted. Trials in which no Suggero suggestions were used were classified into a *no suggestions* category, with remaining trials classified as *high accuracy* (1-3 interactions), *medium accuracy* (4-5 interactions) and *low accuracy* (6+ interactions). Among the *high accuracy* trials, trials with 1 manual selection + 1 suggestion were further classified as *perfect accuracy*.

Participants used Suggero in 1080 trials (360 trials per complexity). Figure 12 shows the breakdown of these categories. The reduction of movement time and distance and the low number of interactions for *simple* and *challenging* is due to the high accuracy of Suggero in these conditions.

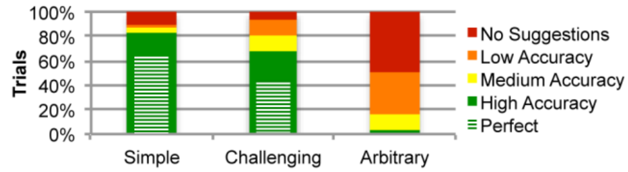

Figure 12: Accuracy of Suggero by Complexity.

Being able to provide users with likely correct results is the main goal of Suggero and any perceptual grouping system. This advantage is not present in the *arbitrary* condition, which was expected since the objects in the target selections were not perceptually related.

Observations & Participant Feedback

A series of Wilcoxon Signed-Rank tests were used to compare ratings between techniques on the post-condition questionnaire. For *arbitrary,* Harpoon (*Mdn=6*) was ranked significantly better than Suggero (*Mdn*=3, *z*=3.5*, p*<.001). There was no significant difference in ratings for *simple* (Suggero *Mdn=5; Harpoon Mdn*=6) and *challenging* (Suggero: *Mdn*=6; Harpoon: *Mdn*=5), as seen in Figure 13. Participant feedback was generally consistent with the ratings. Participants understood and were able to use Suggero well for *simple* and *challenging* complexity. For *arbitrary*, participants reported difficulties performing selections, with some being frustrated because they could not achieve the desired selections with Suggero. Yet, Suggero was found to be helpful for composing complex selections. Two participants reported having difficulties identifying the correspondence of the suggestions with the sketch, because the suggestions were scaled and showed no context. Both mentioned that displaying surrounding objects would be helpful. One participant mentioned that more than three suggestions would be helpful (6 to 8), since he was able to identify at least that many possible groupings for a certain object. In the second phase of the study, participants could express their feedback verbally to the experimenter while drawing and performing selections. Participants tended to select semantically related objects (e.g. the car, house or person they added to the provided sketch). Also, participants experimented with their understanding of Suggero by selecting objects sharing the same properties like color or shape. No participant tried arbitrary selections of perceptually unrelated objects. When participants were unable to perform their intended selection with Suggero, they often reported comments such as "*that was too complex for it*" or "*I will try it [Suggero] for something simpler*". This further indicates that participants were aware of Suggero functionalities, advantages, and limitations.

Results Summary

Although Suggero required higher trial completion times, it required less movement time and distance. This has promising implications for avoiding fatigue, particularly important on large wall displays. Detailed analysis showed that these benefits arise from requiring fewer interactions to perform target selections. An important takeaway is that Suggero

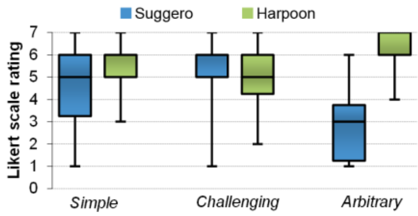

Figure 13: Participants rating on the question if they would use Suggero or Harpoon per *Complexity* (0: never, 7: always).

has the potential to accelerate certain types of group selection; in particular, perceptual groups that are not arranged in compact blocks or that are enclosed or overlapped by other objects, as discussed below. Such selections are more difficult for manual selection tools like Harpoon because of the complexity of selecting them. Study results showed that participants were able to predict the behavior and success rate of Suggero. This finding is consistent with people's behavior in other applications. For example in Adobe Photoshop, selection tools such as the Magic Wand, enable the selection of pixels based on tone and color. People seem to understand they address a very specific use case. Similarly, we believe that rather than having it "always on", it may be better to add Suggero as an additional tool to applications, which users can choose to use explicitly. Omitting incorrect suggestions is as important as providing good suggestions, as our observations from the *arbitrary* condition highlighted. Providing suggestions instead of trying to automatically select groups is important to avoid distraction or confusion and may make selection easy and effortless.

ADVANTAGES AND ERROR HANDLING

Here we outline several use cases where Suggero provides particular benefits for selection on interactive walls.

Covered Objects

Selection of covered objects (Figure 14) can be tedious with existing selection tools. If the selection consists of perceptually related objects, Suggero can greatly assist due to its analysis and suggestions to resolve ambiguity. Editing sketches during group discussions on digital walls may particularly benefit from selection of covered objects.

Figure 14: The underlying structure is overlapped by a large number of objects and can easily be selected with Suggero.

Hierarchical Groups

Sketches are used in many situations. People frequently use hierarchical structures in sketches to add visual complexity. Suggero supports selection of all parts of this hierarchy by exposing the structures in its suggestions (Figure 15).

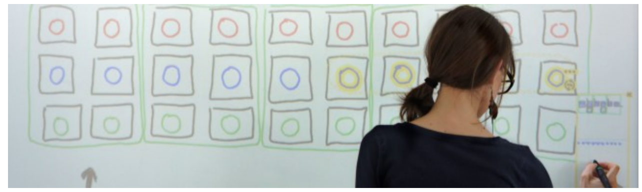

Figure 15: Selecting the blue circles can easily be achieved by selecting one circle and navigation through the suggestions.

Large Structures

Suggero offers advantages for selecting many objects contained in a perceptual group. Especially in a cluttered sketch or a compound of perceptual groups our technique performs better than other ones. E.g., imagine tapping every single object contained in the target selection in Figure 16.

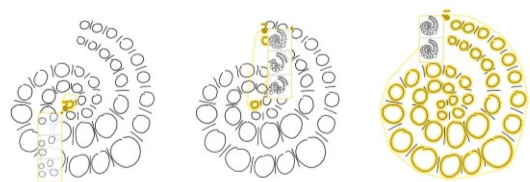

Figure 16: The initial selection (left) results in a local selection of related objects. By adding another selection (middle) the intended selection (right) is achieved with 3 actions overall. Note that only circles are selected on the right.

Error Handling

Although Suggero shows the desired result in its menu in most cases, errors may occur due to incorrect classification or because the desired selection contains perceptually unrelated objects. Thus it is essential to enable users to easily fix incorrectly selected groups. This is why we combined Suggero with Harpoon. Users can refine and correct a selection easily with Harpoon, since it does not require explicit mode switching to toggle selection. After adjusting the selection, Suggero updates its suggestions accordingly.

CONCLUSION AND FUTURE WORK

We presented Suggero, a new perceptual grouping tool, that assists with perceptually related selections in hand-drawn digital sketches by analyzing the content and suggesting possible completions to a given selection. A preliminary study gave valuable insights for the design. A second study found that Suggero decreased selection effort, interactions and stylus movement. These factors decrease fatigue – a well-known problem on large, vertical displays.

In the future, we plan to add more perceptual features to Suggero, and explore automatic parameter tuning for the weights in Dynamic Grouping. Also, we intend to optimize the suggestions Suggero shows to the user by providing, for example, more context within the suggestions. Whenever a complex suggestion cannot be depicted well, we plan to investigate showing a simplified version in the menu. Finally, we will explore automatic mode switching to disable Suggero when the suggestions are not beneficial or the sketch is too complex to infer valid perceptual groups.

ACKNOWLEDGEMENTS

This research was funded by the European Union's Seventh Framework Programme FP7/2007- 2011 (grant n° 318552 (Idea Garden)), by the Austrian Research Promotion Agency (FFG) (BRIDGE program, grant n° 2579917 (Active Office)), by the GRAND NCE, and by the Canada-EU LEIF Exchange Program. We also thank the reviewers for their comments.

REFERENCES

1. Accot, J. and Zhai, S. Beyond Fitts' law: models for trajectory-based HCI tasks. *Proc. CHI'97*, 295–302.
2. Attneave, F. and Arnoult, M.D. The quantitative study of shape and pattern perception. *Psychol Bull 53*, 6 (1956), 452 – 471.
3. Attneave, F. Physical determinants of the judged complexity of Shapes. *Q.J.Exp.Psychol. 53*, 4 (1957), 221–227.
4. Burger, W. and Burge, M.J. *Principles of Digital Image Processing*. Springer-Verlag, 2013.
5. Cates, S.J. Combining representations for improved sketch recognition. *PhD Thesis*, (2009).
6. Dehmeshki, H. and Stuerzlinger, W. Design and Evaluation of a Perceptual-Based Object Group Selection Technique. *Proc. BCS HCI'10*, 365–373.
7. Dehmeshki, H. and Stuerzlinger, W. GPSel: A Gestural Perceptual-Based Path Selection Technique. *Smart Graphics*, (2009), 243–252.
8. Ellis, W.D. *A source book of Gestalt psychology*. London, 1938.
9. Fitts, P.M. The information capacity of the human motor system in controlling the amplitude of movement. *J Exp Psychol [Gen] 121*, 3 (1992), 262–269.
10. Grossman, T., Baudisch, P., and Hinckley, K. Handle Flags: Efficient and Flexible Selections for Inking Applications. *Proc. GI'09*, 167–174.
11. Hick, W.E. On the rate of gain of information. *Q. J. Exp. Psychol. 4*, 1 (1952), 11–26.
12. Hyman, R. Stimulus information as a determinant of reaction time. *J Exp Psychol 45*, 3 (1953), 188–96.
13. Igarashi, T., Matsuoka, S., Kawachiya, S., and Tanaka, H. Interactive beautification: a technique for rapid geometric design. *Proc. Siggraph'07*.
14. Igarashi, T., Matsuoka, S., and Masui, T. Adaptive recognition of implicit structures in human-organized layouts. *Proc. VL'95*, 258–266.
15. Kahneman, D. and Henik, A. Effects of visual grouping on immediate recall and selective attention. *Attention Perform VI*, (1975), 307–332.
16. Leitner, J. and Haller, M. Harpoon selection: efficient selections for ungrouped content on large pen-based surfaces. *Proc. UIST'11*, 593–602.
17. Lindlbauer, D. Perceptual Grouping of Digital Sketches. *Master's Thesis*, (2012).
18. Lowe, D.G. and Binford, T.O. Perceptual organization as a basis for visual recognition. *Proc. AAAI'83*, (1983), 255–260.
19. Mankoff, J., Hudson, S.E., and Abowd, G.D. Interaction techniques for ambiguity resolution in recognition-based interfaces. *Proc. UIST'00*, 11–20.
20. Manning, C.D., Raghavan, P., and Schütze, H. An Introduction to Information Retrieval. (2009).
21. Mizobuchi, S. and Yasumura, M. Tapping vs. Circling Selections on Pen-based Devices : Evidence for Different Performance-Shaping Factors. *Proc. CHI'04*, 607–614.
22. Moran, T.P., Chiu, P., and van Melle, W. Pen-based interaction techniques for organizing material on an electronic whiteboard. *Proc. UIST'97*, 45–54.
23. Nan, L., Sharf, A., Xie, K., et al. Conjoining Gestalt Rules for Abstraction of Architectural Drawings. *ACM TOG 30*, 6 (2011), 1.
24. Oliva, A., Mack, M.L., Shrestha, M., and Peeper, A. Identifying the perceptual dimensions of visual complexity of scenes. *Proc. CogSci'04*, 1041–1046.
25. Palmer, S.E. and Rock, I. Rethinking perceptual organization: The role of uniform connectedness. *Psychon B Rev 1*, 1 (1994), 29–55.
26. Palmer, S.E. Common region: A new principle of perceptual grouping. *Cognitive Psychol 3*, 24 (1992), 436–441.
27. Palmer, S.E. *Vision science: Photons to phenomenology*. Bradford Books, 1999.
28. Rome, E. Simulating perceptual clustering by gestalt principles. *Proc. OEAGM/AAPR'01*, (2001), 191–198.
29. Saund, E., Mahoney, J. V, Fleet, D., Larner, D., and Lank, E. Perceptual organization as a foundation for intelligent sketch editing. *AAAI Spring Symposium on Sketch Editing*, (2002), 118–125.
30. Shipman, F.M.I., Marshall, C.C., and Moran, T.P. Finding and Using Implicit Structure in Human-Organized Spatial Layouts of Information. *Proc. CHI'95*, 346–353.
31. Shpitalni, M. and Lipson, H. Classification of Sketch Strokes and Corner Detection Using Conic Sections and Adaptive Clustering. *J Mech Design 119*, 2 (1996), 131–135.
32. Thórisson, K.R. Simulated perceptual grouping: An application to human-computer interaction. *Proc. CogSci'94*, 876–881.
33. Treisman, A. and Gelade, G. A Feature-Integration Theory of Attention. *Cognitive Psychol 12*, (1980), 97–136.
34. Treisman, A. Perceptual Grouping and Attention in Visual Search for Features and for Objects. *J Exp Psychol [Hum Percept] 8*, 2 (1982), 194–214.
35. Treisman, A. Features and objects in visual processing. *SciAm 255*, 5 (1986), 114–125.
36. Wertheimer, M. Untersuchungen zur Lehre von der Gestalt. II. *Psychol Research 4*, 1 (1923), 301–350.
37. Xu, P., Fu, H., Au, O.K.-C., and Tai, C.-L. Lazy Selection: A Scribble-based Tool for Smart Shape Elements Selection. *ACM TOG 31*, 6 (2012), 1.