NEAT: A Set of Flexible Tools and Gestures for Layout Tasks on Interactive Displays

Mathias Frisch, Ricardo Langner, Raimund Dachselt
User Interface & Software Engineering Group
Otto-von-Guericke University Magdeburg, Germany
{mfrisch, rlangner, dachselt}@ovgu.de

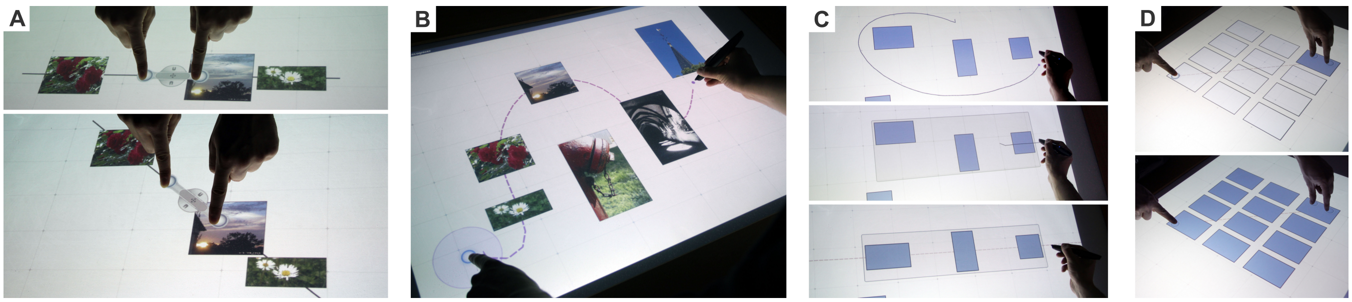

Figure 1: NEAT extends multi-touch alignment guides and provides a set of pen and touch gestures for creating layouts without invoking tools: (a) Objects bound with their centers to a guide keep their orientation during rotation. (b) Guides of arbitrary shape can be sketched with the pen. Crossed objects are bound automatically. (c) *Align-by-Crossing*: Grouped objects can be aligned by crossing one of them with the pen. (d) A grid of cloned objects can be created by holding the original object with two fingers and dragging the clones from it.

ABSTRACT

Creating accurate layouts of graphical objects is an important activity in many graphics applications, such as design tools, presentation software or diagram editors. In this paper, we are contributing **N**atural and **E**ffective L**a**yout **T**echniques (NEAT). The system provides a consistent set of multi-touch tools and gestures for aligning and distributing graphical objects on interactive surfaces. NEAT explicitly considers expert requirements and supports a rich and consistent set of layout functions. Amongst others, it minimizes visual distraction by layout tools, combines separate steps of interaction to compound ones and allows effective interaction by combining multi-touch and pen input. Furthermore, NEAT provides a set of bimanual gestures for achieving layout tasks in a quick and effective way without explicitly invoking any tools. From initial expert user feedback we derive several principles for layout tools on interactive displays.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical User Interfaces.

General terms: Design, Human Factors

Keywords: Multi-touch, pen input, bimanual interaction, tabletops, expert mode, gestures, alignment, constraints

ITS'11, November 13-16, 2011, Kobe, Japan.

INTRODUCTION

Creating accurate layouts is an essential activity in many applications, such as graphics design tools. Thereby, *alignment* is one of the most important design principles and supports the fundamental design concept of visual unity [24]. It ensures that no graphical object is placed arbitrarily and that objects are visually connected with other objects. Another important principle is *proximity*. It means that related graphical objects should be placed close together [26]. Examples for the application of these principles can be found in every well-designed and *neat* piece of graphical work. Similar rules can also be found in other application domains. For example, in graph visualization, nodes with a high degree should be centered, or child nodes should be placed below their parent node in a symmetric way [22]. Applying these rules during editing leads to more comprehensible graph structures. Experts who consider these design principles invest much effort to position pictures, text boxes and other graphical objects to create pleasant layouts. Thus, they make wide use of layout functions offered by respective graphical applications such as graphics design tools [1], diagram editors or presentation software [19, 2]. Since these tools are based on the desktop metaphor, invoking the functions is typically achieved by selecting commands from menus or toolbars, invoking hotkeys or by entering values with either mouse or keyboard. However, this often involves many steps of interaction and can be very time-consuming and cumbersome.

Modern interactive surfaces such as tabletops or tablets are very promising for alleviating these problems of graphics applications. First tools have already been ported and adopted to devices like the iPad [2]. Multi-touch interaction or even the combination of touch and pen allow a more direct and natural way of interaction. These modalities can be applied to create objects by sketching with the pen and to create them in a more structured way by touch [9]. Furthermore, tasks can be performed simultaneously by both hands, and mode switches are possible without invoking menus [14].

With respect to the applications mentioned before, corresponding systems for interactive displays have to be carefully designed. In particular, special requirements from domain experts have to be considered, as they need more functionality than basic transformation of graphical objects.

In this paper, we are contributing a system for **N**atural and **E**ffective **La**yout **T**echniques (NEAT). It offers multi-touch tools and gestures for effectively aligning and distributing graphical objects on interactive displays, and it was designed to meet experts' needs. In particular, the NEAT system considers aspects such as minimizing the visual distraction of layout tools. Furthermore, it provides interaction techniques which seamlessly combine separate steps of interaction. This makes it possible to achieve layout tasks in a more effective and fluent way. We also considered the combination of multi-touch and pen input which allows natural interaction by sketching and quick mode switches. Parts of NEAT are based on the multi-touch alignment guides developed in our previous work [10]. For them, we realized advanced techniques such as the combination of guides and free-form guides which allow the arrangement of objects along arbitrary curves (see Figure 1 a,b). In addition to that, NEAT also contributes a novel set of effective gestures for layouting graphical objects (see Figure 1 c,d). This approach is an alternative to the usage of distinct tools such as alignment guides. It was explicitly designed as an effective way of interaction for experts.

It is important to mention that NEAT is not a collection of isolated techniques, but a consistent set of tools and gestures. It offers layout functionality which can be seamlessly applied in different application scenarios. In particular, we integrated them in two prototypes: an application for manipulating pictures and a node-link diagram editor.

The remaining paper is structured as follows: After presenting related work in this area, we briefly reflect on the functionalities and principles of the multi-touch alignment guides [10]. Then, we present the extensions and additional concepts for the guides. After that, the novel, purely gestural approach for creating layouts is presented. Finally, we present initial expert user feedback and reflect and summarize basic principles our NEAT system is based on. We believe that these principles can serve as general design principles for layout tools on interactive displays.

RELATED WORK

Alignment in mouse-based applications

In existing mouse-based graphics tools alignment tasks are mainly achieved in an indirect way. For example, commands have to be invoked from menus to create proper alignments. For setting the spacing between graphical objects, concrete numerical values have to be entered. Beyond that, techniques such as *snap-dragging* [5] are commonly applied in modern tools. They allow precise positioning in a more direct way. Users can drag objects and the system provides automatic assistance by snapping objects to background grids, to bounds of other objects or to previously created guides. These objects then serve as constraints. Wybrow et al. [28] presented multi-way constraints for node-link diagrams. This means that a graphical object is not only constrained by a single rule (such as right alignment), but by several rules. Wybrow et al. conducted two usability studies. They showed that multi-way constraints are more beneficial than one-way constraints in many situations. Furthermore, some techniques have been developed to extend *snap-dragging*. *Snap-and-Go* [3] provides a solution for positinging objects close to snapping points. This is achieved by using additional motor space. Another advanced snapping technique is *HyperSnapping* [18]. With this approach, the cell size of the background grid is for example changed dynamically during dragging.

Alignment in digital sketching tools

Several pen-operated digital sketching applications offer automatic assistance for working on interactive displays. One example is the automatic alignment of hand-drawn strokes. For that, Igarashi et al. [15] propose interactive beautification for digital sketching. Their system automatically recognizes geometric relations between strokes, such as parallelism, connections of strokes, and symmetry. Beyond that, Fung et al. [11] presented a template-based technique for guiding the cursor movement along paths. With this approach it is possible to draw visual patterns like parallel lines or spirals in an easy way. The drawing application *Lineogrammer* [32] also implements beautification techniques such as snapping sketched strokes to existing objects. Beyond that, it integrates a ruler tool. Besides alignment and distribution, it supports mirroring of objects or cloning objects along the ruler. It is based on the *Alignment Stick*, a ruler tool for aligning graphical objects [21]. The *Alignment Stick* was designed for bimanual input with mouse and trackball. The techniques presented in this work realize similar functionality, but instead they are designed for bimanual interaction by multi-touch and pen.

Alignment on mulit-touch enabled displays

Graphical objects can be translated and rotated in various ways by means of multi-touch interaction. Different techniques for that are discussed by Hancock et al. [13]. This work also considers snapping of objects for precise alignment, e.g. to the edge of a tabletop display. As multi-touch techniques usually allow the unconstraint translation, scaling, and rotation in a single interaction, it is difficult to perform one of these actions separately. For example, scaling an object often results in slight unwanted rotation. Nacenta et al. [20] addressed this problem and compared five different techniques. Their results showed that the distinction of gestures by thresholding and matching gestures against models can improve precise transformations.

At CHI 2011 two techniques for aligning graphical objects on interactive displays were presented: Rock & Rails [25] and Grids & Guides [10]. In contrast to the techniques mentioned before, these approaches do not only consider the manipulation of single objects. They focus on creating precise spatial relationships between objects such as alignment or the arrangement along paths. Rock & Rails applies hand shapes to scale graphical objects in a non-uniform way or to isolate rotation from scaling. The authors summarize their approach with: *fingertips manipulate and hand shapes constrain*. Furthermore, proxy objects are used to overcome the problem of occlusion. Grids & Guides presents two specific tools for supporting layout tasks on interactive displays. *Interactive grids* allow adjusting their cell sizes by direct touch input and can be seamlessly changed to radial grids. *Multi-touch alignment guides* support the distribution and alignment of graphical objects along lines and circles. In this paper, we present advanced techniques for the multi-touch alignment guides. In addition to that, we propose a set of pen and touch gestures to achieve layout tasks without invoking any tools.

Simultaneous use of pen and touch

All systems which realize simultaneous pen and touch input consider the findings of Guiard [12]. He investigated how the hands collaborate to accomplish tasks. In particular, the non-dominant hand frames the action of the dominant hand, and the preferred hand performs more precise actions. The combination of pen input and single touch was investigated by Yee [29]. He proposes for example panning the canvas with the finger while drawing with the pen or operating with a file browser in a bimanual way.

The usage of digital pens and multi-touch on tabletops has been studied by Brandl et al. [7]. The authors suggest general design principles and present interaction techniques for a graphics application. They found that the combination of pen and touch is superior concerning speed, accuracy and user preference. Frisch et al. [9] proposed a set of multi-touch and pen gestures for diagram editing. It is based on a user elicitation study, and the pen can be used for sketching diagram elements. In addition, bimanual multi-touch input can be applied for functionalities such as copying nodes. Further systems which consider touch and pen input were presented by Zeleznik et al. [31] and Hinckley et al. [14]. The first one is a basic algebra system, and the second one is a digital drafting table application. Hinckley et al. explicitly differentiate between both modalities and introduce an interaction principle which they summarize with *pen writes, touch manipulates and pen + touch yields new tools*.

BACKGROUND AND DESIGN GOALS

Parts of NEAT are based on the first iteration of our multi-touch alignment guides [10]. In this section, we briefly describe these tools and discuss their shortcomings. Subsequently, we present five design goals for the novel NEAT system. They are derived from initial user feedback, including two professional architects [10].

A multi-touch alignment guide consists of a line and an attached handle (Figure 1a). The line can have an arbitrary shape (the first prototype only realized straight lines and circles). Objects can be bound to it, and the line serves as a constraint for them. The handle can be used for repositioning the guide and for switching modes. Interacting with multi-touch alignment guides follows the principle *Create–Bind–Arrange*. These three tasks can be achieved with three distinct actions.

Create (A1): Guides can be created by sketching the desired shape, holding two fingers on the background or laying down a pen on the surface.

Bind (A2): After a guide was created, graphical objects can be bound to it by snap-dragging. Besides that, two further techniques were realized. *Flick & Snap* allows the binding of objects by flicking them towards the respective guide. With *Collide & Snap* objects can be collected by dragging a guide across the surface. For that, the guide becomes sticky after the respective mode was activated by a button at the handle.

Arrange (A3): When objects are bound to a guide, they form a group. Therefore, by interacting with the attached handle, all bound objects can be translated and rotated in a single step. It is also possible to manipulate them directly, whereby the guide serves as a constraint for them. For example, bound objects can be dragged along the guide to adjust their spacing or across the guide to align them properly.

Besides these three steps, two general principles have also been mentioned in [10]: Some of the actions are based on physical metaphors, such as flicking. Furthermore, actions can be performed sequentially by single touch but also within a single flow by multi-touch input.

In this work, we considerably improve and extend the concept of multi-touch alignment guides. Based on the results of our initial user evaluation (see [10]), we analyzed them deeper to determine problems and drawbacks. We did this by considering the requirements of expert users such as efficient workflows and precise interaction. This was the first part of the iterative design process we applied during the development of the NEAT-system. As a result of this analysis we set up the following five design goals for the NEAT system presented in this paper:

Combining separate actions (D1): As mentioned above, interacting with the guides takes place in three distinct steps (A1–A3). Typically, this is not very efficient. Several of these actions should be combined to a single one, as suggested in [8]. Thus, tasks can be achieved in a more fluent way of interaction.

Making layout tools more flexible (D2): The multi-touch alignment guides are limited to straight lines and circles. However, in many applications graphical objects need to be arranged more freely along arbitrary paths and to be combined for more powerful layouts. Layout tools should support this.

Minimizing visual distraction (D3): Alignment guides can clutter the workspace if complex graphical layouts are created. As a result, it becomes difficult for users to focus on the task at hand. Therefore, the visual representation of the applied layout tools and the feedback given by the system should be as unobtrusive as possible.

Supporting expert users (D4): To accomplish recurring tasks, experts are used to apply shortcuts instead of invoking explicit tools. The multi-touch alignment guides do not support this. Therefore, we aimed at developing an alternative way for creating layouts which is designed for expert users. However, this additional expert mode and the alignment guides should complement each other in a seamless way.

Awareness of work artifacts (D5): Many domain experts are used to work with physical tools such as pens, rulers and sten-

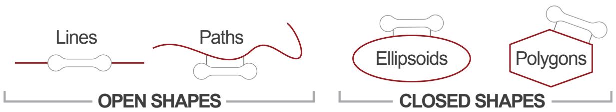

Figure 2: Two basic types of guides: open shapes (e.g., lines, paths) and closed shapes (e.g., ellipsoids, polygons).

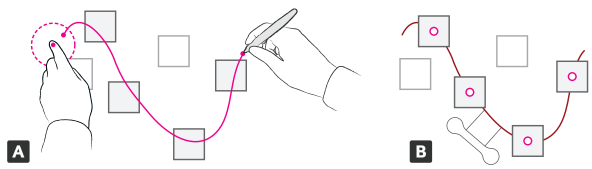

Figure 3: *Cross-and-snap* technique enables users to (a) create a guide and (b) bind crossed objects to it simultaneously.

cils. Thereby, the pen is certainly the most important tool, as it allows sketching in a natural way and the precise selection of small graphical objects. Thus, pens should be considered for interacting with graphics applications on interactive displays. Since the combination of touch and pen can be applied for fast mode switches [14], the system should be able to explicitly distinguish between both modalities.

During the iterative development of the NEAT system, we considered all the mentioned design goals (D1-D5). The goal was to provide consistent layout techniques which can be applied in various tools, such as graphics design software or diagram editors. Whereas the multi-touch alignment guides [10] served as a starting point in our design process, we considerably extended the original concept by building software prototypes and continuously refining them. For that, expert reviews involving four interaction design experts of our department were repeatedly applied similar to the pluralistic usability walkthrough method [4]. Special care has been taken to design techniques which are not only effective in themselves, but also combinable. The extended alignment guide concepts are described in the following section. The subsequent section then presents novel gestural layout techniques for experts, where tools entirely vanish in favor of an efficient and fluent interaction embedded in an expert's workflow. Finally, we discuss initial user feedback.

GUIDES: FLEXIBLE LAYOUT TOOLS

It is an essential part of the *Grids & Guides* system [10] to bind multiple objects to a certain geometric shape. It serves as a proxy for a visual connection. Combined with a handle, this shape can be used as an interactive tool to arrange objects. In the following subsections we describe several major extensions of the original concept.

Guides based on arbitrarily curved shapes

In addition to the straight line and circular guides of our first prototype [10], the NEAT system provides free-form guides which can be of arbitrary shape. We thereby distinguish between guides consisting of an *open shape* (e.g., lines, curves, and paths) and guides featuring a *closed shape* (e.g., ellipsoids, rectangles, or polygons), as can be seen in Figure 2. For creating guides of arbitrary shapes and manipulating them afterwards we introduce several novel interaction techniques. They mainly address the design goals of combining distinct actions (D1) and making guides more flexible (D2).

Cross-and-Snap. Common tools and techniques depend on performing separate actions successively. For creating multi-touch alignment guides and simultaneously binding objects to it, we propose a technique called *Cross-and-Snap* which applies a combination of touch and pen input (see design goal D5). To make use of *Cross-and-Snap*, the user holds a finger on the background and creates a guide by sketching the desired shape from this position (see Figure 1b and Figure 3a). Considering real world sketching activities, this is achieved by using the pen. The combination of touch and pen is necessary to distinguish creating a guide from common sketching. The NEAT system detects all objects crossed by the stroke. These objects will be bound to the shape immediately. Thereby, the objects are translated so that they are centered on the guide (Figure 3b). To revoke the binding as well as to restore the previous position of an object users can move the pen or finger back along the stroke. That allows corrections of the current selection and the sketched stroke on the fly. In addition to creating guides consisting of an open shape, NEAT provides *Cross-and-Snap* for closed shapes as well. For that, the sketched stroke needs to be finished close to its starting point. By applying *Cross-and-Snap* the two separate activities A1 (creating a guide) and A2 (binding objects) are combined to a single action (see design goal D1).

Modifying the Shape. In order to make *Cross-and-Snap* more effective and flexible, NEAT supports different mechanisms to modify the shape of a guide, which are based on physical metaphors. Users are able to modify the shape of a guide by *shortening (a)*, *straightening (b)* or *bending (c)*. *(a)* A guide can be shortened by sketching a stroke across its shape with the pen. This results in cutting off the segment facing away from the handle (for *open shapes*). If a *closed shape* is crossed with the pen, it turns to an *open shape*. Furthermore, a guide can be shortened by dragging a shape's end point along the shape. Objects passed through will be unbound.

(b) Dragging one or both end points apart continuously converts an arbitrary curved guide into a line guide (Figure 4a). Thereby, bound objects stick to the changing shape. This physical metaphor is inspired by pulling a cord apart. If dragging apart the end points of a line guide is continued beyond the line's original size, this results in extending the guide like stretching a rubber band. Especially in combination with a previously performed *Cross-and-Snap* technique (mentioned above), the shape of a free-form guide as well as the alignment of bound objects can be easily straightened.

(c) A straight line guide can also be continuously transformed into a guide with a curved shape. To achieve that, the end points can be rotated by simply touching them simultaneously and rotating the fingers (Figure 4b). For that, the orientation of the touch points (provided by input devices or software algorithms [17, 23]) can be applied. For this technique, bending a leaf spring served as underlying metaphor.

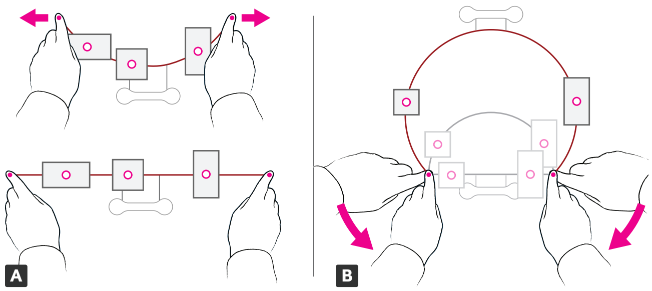

Figure 4: Modifying the shape of a guide: (a) Pull apart the endpoints of the guide's shape to continuously convert a curved guide into a line-guide. (b) Rotating the endpoints bends the shape of a guide.

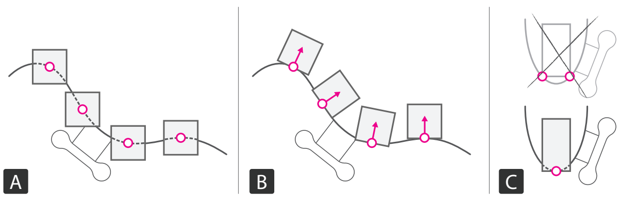

Figure 5: Different orientation of bound objects based on their alignment: (a) Centered objects keep their original alignment. (b) Sidewise aligned objects inherit the orientation of the guide (perpendicular to the tangent at the docking point). (c) Only one single docking point for each object.

Behavior of bound objects

The layout manipulation functions in current graphics and presentation software [1, 2, 19] tend to utilize two different strategies: *transform all* (a group of objects is transformed as a single object) and *transform each* (each object of the group is transformed separately). In NEAT, *transform all* is achieved by selecting the respective objects (e.g., by encircling) and rotating or scaling this selection. For *transform each*, objects have to be bound to a guide. Scaling or rotating one of the bound objects transforms all the others in the same way.

Beyond that, NEAT realizes an additional interactive layout mechanism which cannot be found in other tools. If objects are bound with their centers to a guide (Figure 5a), they will keep their original orientation even if the guide is translated or rotated (Figure 1a). Opposed to this, objects which are aligned sidewise (to the "left" or to the "right") will change their orientation when the guide is rotated (Figure 5b). Objects are rotated so that their orientation matches the normal vector of the corresponding docking point. Note that there is only one single docking point for each bound object (Figure 5c).

Moreover, NEAT applies further physical metaphors to objects which are bound to a guide. We suggest that these objects do not overlap. NEAT defines a bound object as a rigid body and uses simple collision detections to prevent overlapping objects [27]. Beyond that, users can unbind objects by using a bimanual gesture involving a magnetic metaphor. Holding the guide using one hand and simultaneously dragging away a bound object using the other hand results in unbinding it.

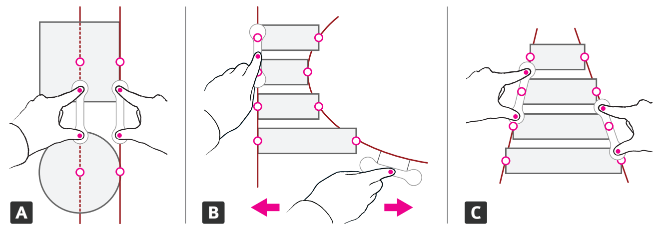

Figure 6: Manipulative combination of guides: (a) Uniform scaling if objects are centered on the guide. (b) Combination of guides of different shapes.(c) Non-uniform scaling of bound objects by rotating guides.

Combination of Guides

By default, guides are not combined and can be dragged across the surface without affecting each other. A combination of guides can be used to accomplish more complex tasks. In [10] we sketched a first mechanism to resize bound objects using two guides. However, this is only one example for combining guides. We believe that combining guides in multiple ways can enable users to effectively manipulate the positioning of multiple objects and their shapes. In the following subsections we introduce multiple major extensions of the original idea. In general, we distinguish between two different scenarios for the combination of guides: *manipulative combination* and *constraining combination*.

Manipulative Combination. A *manipulative combination* of guides can be used to interactively change the size, orientation or shape of multiple objects. To enable the *manipulative combination*, the *Collide & Snap* mode [10] of a guide needs to be activated in advance. Moving this guide across objects which are already bound to another guide results in binding these objects to both guides (Figure 6). Dragging the guides can result in manipulating the arrangement, size, and orientation of bound objects. If objects are bound with their centers to one of the guides (the left one in Figure 6a), uniform manipulations, e.g. scaling, will be applied. The guide containing the centered objects can be understood as an axis of symmetry. This is similar to the mirror mode for symmetric editing on the ruler tool introduced by Zeleznik et al. [32]. Guides of different shapes can be combined as well (Figure 6b) to the outline of a group of objects (e.g., to match with other graphical elements of a document). Figure 6c shows the non-uniform scaling of graphical objects by rotating the guides.

Constraining Combination In contrast to the manipulative type mentioned above, the *constraining combination* does not require any additional activation of modes. *Constraining combination* is intended for the alignment of objects using multiple guides and does not affect the size of bound objects, but their positioning. This can be achieved by dragging a guide across another one, which is held with a second finger. As a result, bound objects are moved according to the constraints of both guides. An example for *constraining combination* is illustrated in (Figure 7). As both guides are touched and moved across each other, object 2 is bound to both guides. When the guides are moved, object 2 sticks to both of them and will always be aligned to its left with object 1 and to its top with objects 3 and 4. This technique realizes multi-way constraints as discussed and evaluated by Wybrow et al. [28].

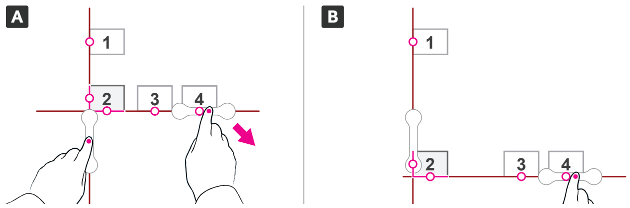

Figure 7: Constraining combination: (a) Object 2 is bound to both guides. (b) While dragging the right guide to the bottom right, object 2 sticks to both guides.

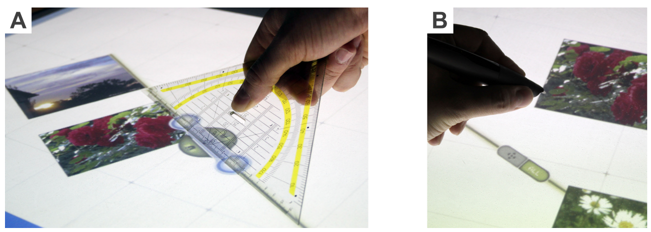

Figure 8: (a) Example for tangible interaction: a guide is created by putting a triangle on the surface. (b) Minimized version of the handle.

Tangible guides

Guides cannot only be created by sketching. In [10] we proposed to put down a pen onto the surface in order to create a line guide temporally. In fact, professionals use different physical tools for different tasks, such as rulers, stencils, or pens. We propose to take these work artifacts and their utilization into account while designing interaction techniques and interfaces (see design goal D5)(Figure 8a). By simply putting down a work artifact onto the table, we can create a guide of its very shape or contour. After binding objects to the guide created by the physical object, it can be positioned by moving the physical object itself or by using the attached virtual handle. This is similar to setting a constraint by putting a hand onto the surface [25]. To a certain extent, multiple "physical guides" can be combined freely with virtual guides.

Minimizing visual distraction

According to our design goal D3, layout tools should be as unobtrusive as possible. To meet this goal, we further adapted the handle of the guide and developed a smaller version (Figure 8b). It consists of only two areas: one for translating the guide by dragging, and a second area for buttons. If the translation area is tapped, the handle will become bigger in an animated way and areas for rotating will become visible. Furthermore, guides are blended out completely if the bound objects are not touched for more than five seconds. This prevents cluttering the workspace with alignment lines. Gestural layout techniques, which entirely avoid explicit tools and thus visual distraction, are presented in the next section.

GESTURES: LAYOUTS WITHOUT EXPLICIT TOOLS

In the previous section we presented a variety of advanced techniques for multi-touch alignment guides. As visual tools they provide affordances for available actions, but tend to clutter the workspace if applied extensively. This contradicts design goal D3 and also prevents fast interaction for trained users (see design goal D4).

In this section, we therefore contribute several novel expert gestures for aligning, distributing and cloning graphical objects. They are applied without invoking tools or using handles and buttons. Along with the multi-touch alignment guides, they are also tightly integrated into the NEAT system. We did not intend to design these gestural techniques as a walk-up-and-use functionality, but for experienced expert users. Thus, bimanual gestures of a more abstract nature are applied which have to be learned and trained. As one basic principle we explicitly distinguish whether graphical objects are touched with one or two fingers. Single fingers can be used for example to move or rotate a group of objects. Touching an object with two fingers indicates a mode switch and activates constraint manipulation. The manipulation itself is then performed with the second (dominant) hand. Beyond that, we explicitly distinguish between pen and touch input (see design goal D5). Thereby, the pen serves for sketch-based input and precise interaction with small graphical objects. Special care has been taken to match design goal D1 in that the techniques allow to efficiently perform compound tasks with a single fluid interaction.

Aligning multiple objects

For a quick alignment of graphical objects we propose a technique which we call *align-by-crossing*. The basic idea of *align-by-crossing* is that in a group of objects, one particular object is crossed with the pen. This object then serves as a reference for the alignment of all other objects of that group. The direction of the crossing and the location where the object is crossed determine the way of alignment. The pen is used to distinguish crossing from dragging by touch. Figure 9a illustrates an example for this technique: a group is created by encircling the respective objects. After that, one of the objects is crossed with the pen vertically along its right border. This object now serves as a reference. As a result, all other objects of the group are moved in an animated way so that they align with the right border of the crossed object.

Of course, it is also possible to cross an object vertically at its left or at its center. As a result, other objects are aligned to its left or center, respectively. In the same way, objects can be aligned to the top, center or bottom by crossing horizontally (Figure 1c). To realize this technique, objects are virtually partitioned in three horizontal areas (top, center and bottom) as well as in three vertical areas (left, center and right) (Figure 9b). The alignment takes places depending on the area which was hit by the crossing stroke. For more complex objects (e.g., with concave parts or compound objects) the partition is realized according to their bounding rectangle.

Whereas the *align-by-crossing* technique follows the principle of first selecting the objects and then the desired action to be applied, we also support first choosing the action and then the objects (Figure 10). For that, a single reference object is chosen for which the crossing will be applied. We propose that the object must be held with two fingers while it is crossed as described above. In that way, *align-by-crossing* can be distinguished from common sketching or other functions such as cutting an object with the pen [14]. Again, the crossing specifies the way of alignment (the right border in 10a). After that, tapping other objects while still holding the reference

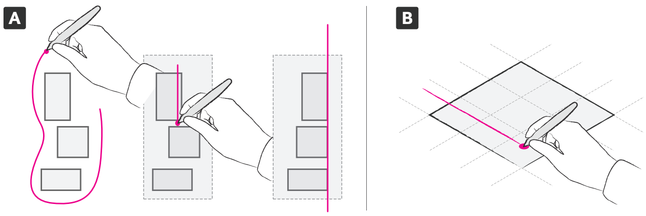

Figure 9: *Align-by-crossing* on a group of objects: (a) Objects are grouped by encircling, one of them is crossed with the pen, as a result all other objects are aligned to the right border of the crossed one. (b) For realizing *align-by-crossing* a graphical object is divided into three vertical and three horizontal areas.

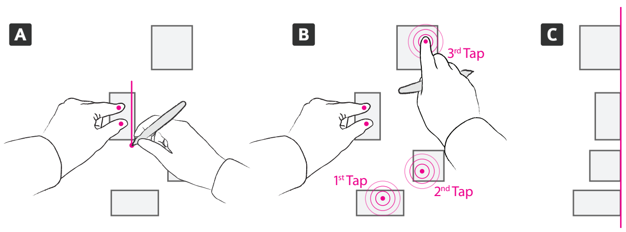

Figure 10: *Align-by-crossing* on separate objects: (a) The reference object is touched with two fingers and crossed with the pen. (b) Then, all other objects are tapped. (c) They are aligned according to the reference object. A feedback line indicates the alignment.

object (Figure 10b) results in aligning the tapped objects according to the specified alignment (Figure 10c). To make this more comprehensible, the objects are moved to their target position in an animated way. If the reference object is released, *align-by-crossing* is disabled again and tapping results in selecting the respective objects, which is the default mode.

For pure gestural interaction it is especially important to provide sufficient feedback. In this case, a line indicates the alignment axis and is blended in as soon as the crossing was finished. The feedback line is not an alignment guide and does not constrain the aligned objects. Furthermore, it disappears if none of the aligned objects are touched for more than three seconds. However, if one of the objects is touched again, the feedback line is blended in again to indicate that these objects were aligned before. *Align-by-crossing* is beneficial for quickly aligning selected objects without further manipulating their positions and layout. This approach is similar to invoking a menu in existing tools and selecting the respective alignment function. As a result, the selected objects are aligned but not further constrained.

Transforming Objects

Manipulating grouped objects with fingers transforms the whole group (*transform all*). For example, rotating a group results in rotating all objects around the center of the group. However, in many situations users want to perform the same transformation on each object (*transform each*), such as rotating all grouped objects around their centers. In NEAT, this cannot only be achieved by binding the respective objects to guides, but also by applying gestures. In particular, linear transformations such as translation, rotation and scaling are supported. To distinguish between *transform all* (the default) and *transform each*, one of the grouped objects has to be touched with two fingers simultaneously. This again implies a mode switch. The actual gesture (such as pinching for scaling or rotating two fingers for rotation) is then performed on another object of the group. As a result, the respective transformation takes place according to the objects' centers.

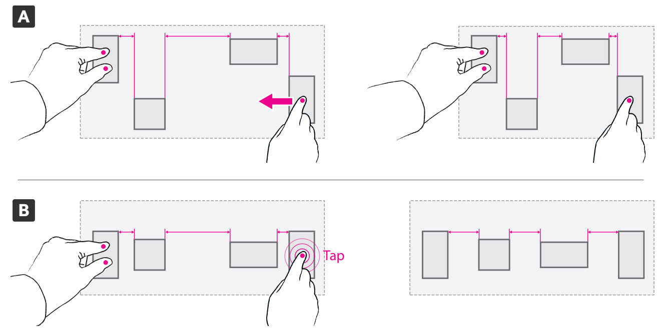

Figure 11: Adjusting the spacing of grouped objects by touching one object with two fingers. (a) *Dragging* another object adjusts the spacing of all objects in between. (b) *Tapping* another object results in distributing them evenly.

Holding one object with two fingers and translating another one by dragging results in adjusting the spacing between grouped objects (Figure 11a). As a result, the distances between all objects are adjusted accordingly. Thereby, the distance to their closest neighbors is shown in pixels in the same way as interacting with the multi-touch alignment guides. Dragging horizontally adjusts the horizontal spacing and dragging vertically adjusts the vertical spacing. It is also possible to evenly distribute grouped objects (Figure 11b). For that, it is necessary again to hold an object with two fingers. Tapping another object of the group equally distributes all objects in between. The distribution occurs in an animated way to make it more comprehensible. According to our design goal D1, all these activities can be performed within a single step of interaction. For example, objects of a group can all be scaled by the same amount by holding two fingers and performing a pinch gesture simultaneously. Then, their spacing can be adjusted by dragging without releasing the fingers.

Cloning Objects

Aligning several objects of the same shape and size is essential in many graphics applications. Expert users usually achieve this by pressing modifier keys and dragging a clone from the original object. In previous work [9], we proposed cloning elements on interactive displays by holding the original object with one finger and dragging the copy from it with a second finger (similar to Hinckley et al.[14]). This is also possible in our prototype. In addition to that, several clones of an object can be created with only one step of interaction. This is similar to the techniques for mouse-based interfaces presented by Zaman et al. [30].

Again, we use two fingers to activate this mode: Holding an object with two fingers and starting a drag gesture from this object creates a copy of the held object. When this clone is

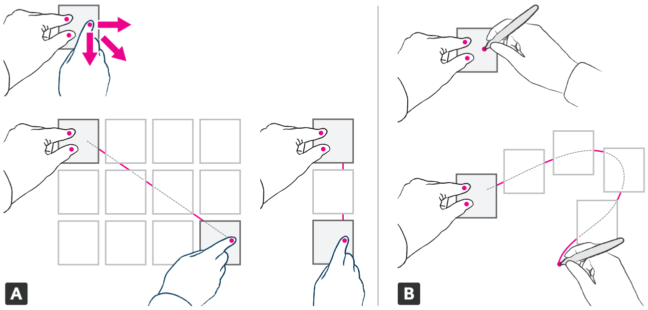

Figure 12: (a) Cloning an object by touching it with two fingers and dragging with a third finger. Straight dragging results in vertically and horizontally aligned clones. Diagonal dragging results in a grid. (b) Cloning objects along a path by sketching with the pen.

dragged further, as many clones as possible will be created between the original and the dragged one (Figure 12a). Creating clones towards the bottom or top results in vertically aligned clones. Creating clones towards the left or right results in aligning the clones along a horizontal path. As a third option, dragging diagonally creates a grid of clones. In all cases a feedback line appears like the one for *align-by-crossing*.

During dragging, the clones are shown as previews. Previews are semi-transparent and have the same size and shape as the original object. If the dragging stops for more than one second, the previews are changed to real objects. When the user starts dragging again without lifting the finger, no further clones are created. Instead, the spacing between the created clones can be adjusted dynamically. Again, this realizes design goal D1 of combining several actions.

In addition to arranging cloned objects along straight lines, it is also possible to create clones along free paths. This is done by holding an object with two fingers and performing the drag gesture with the pen (Figure 12b). As a result, as many clones as possible are created along the sketched path.

Further Considerations

Transition to guides As mentioned before, objects arranged by gestures are temporally constraint. This means, the constraint takes effect only while the gesture is performed. After aligning or cloning objects, just a temporary feedback line is shown, but the objects are not constrained any more. However, tapping this line with one finger changes it to a multi-touch alignment guide. The associated objects are then automatically bound to this guide. In that way, it is easy to seamlessly switch from the gestural approach to the tool centric approach which offers persistent constraints and more functionality.

Working with small objects In some situations graphical objects can be too small to touch them with two fingers simultaneously. To solve this problem we propose to simply place the fingers in a way that the respective object is located between the fingers. After that, the gestures for aligning or cloning can be performed as described before.

These two features are always available in the NEAT system, if applicable. This underlines again our goal to provide consistent principles throughout all different gestural techniques.

USER FEEDBACK AND DISCUSSION

NEAT has been implemented in Java and was designed as a flexible software component which can be integrated into other graphics applications. Currently we successfully integrated it in two programs: an application for manipulating photographs and an editor for creating node-link diagrams. Both programs run on any TUIO-enabled[1] multi-touch device and on SMART Tables[1]. For pen input we are utilizing digital pens based on the Anoto[1] technology. In that way, it is possible to distinguish between multi-touch and pen input.

Expert Feedback Up to now we did not run an extensive evaluation of NEAT. However, besides the continuous evaluation of interaction design experts, the system was tested by three domain experts who are no HCI professionals. One of them is an architect (she already gave feedback for the first version of the prototype [10]). The other two are bioinformaticians. They are involved in the development of a diagram-editing tool[2] for biologists which can be used for example to create biological pathways. Both have long-term experience concerning the needs and requirements of biologists for such applications. All domain experts confirmed that creating proper layouts is an essential task in their domain and that they apply respective functions regularly in their desktop tools. Each test took place in a separate session and lasted about 45 minutes. We showed and explained the whole functionality of NEAT to the three expert users. After that, they could try NEAT themselves and we asked them to perform particular tasks (e.g., align the present nodes of a diagram). After that, we conducted a concluding interview with each test user and asked them about their opinions concerning NEAT.

Discussion of Results

Feedback concerning guides All users quickly understood how the guides worked and overall they liked their features. They stated that providing feedback such as distance indicators between bound objects and rotating guides in fixed angles is important for many tasks. Furthermore, all test users mentioned that making appearance less distracting is a crucial feature. Free-form guides including *Cross-and-Snap* (one of our main extensions) was especially appreciated by the architect but not by the bioinformaticians. The reason for that is certainly that for diagram layouting predefined geometric shapes such as straight lines and circles are more important than arbitrary curves. The bioinformaticians asked for templates which can be reused. Thus, for example creating copies of guides with predefined anchor points can be beneficial. Furthermore, one of them commented that "when I drag this guide other groups of objects should move along". This indicates the need for multi-way constraints (e.g., by combining guides) which were not yet implemented in our prototype.

Feedback concerning gestures The layout gestures were not only considered as "pretty cool", but also as powerful and productive. All domain experts stated that they liked the gestures because of their ability to create layouts without further constraints in a quick way. Persistent constraints such as provided by the multi-touch alignment guides "are not always necessary". Using two fingers to activate the layout

[1] www.tuio.org, www.smarttech.com, www.anoto.com
[2] VANTED: http://vanted.ipk-gatersleben.de/

gestures seemed a reasonable approach for our domain experts. However, all of them confirmed that they were not able to spontaneously guess the gestures. We expected this, as our goal was to design gestures explicitly for experts and not as a walk-up-and-use feature. Thus, users need to learn and train the techniques. Moreover, we observed that the distinction of touch and pen was not always clear to our users. In many cases the system did not provide enough feedback. As stated by Hinckley et al. [14], adding feedback for this kind of multimodal interface is difficult because many gestures (e.g., *align-by-crossing* or distributing objects by tapping) are of a discrete nature. Thus, their effect is not shown until the pen or finger is lifted.

The expert users liked the idea of switching from the unconstraint gestural approach to the guides by tapping on the feedback line. However, we also observed that in some situations the users expected persistent constraints for the gestural techniques as well. As gestures allow just temporal constrained interaction, performing several gestures consecutively can lead to unintended results. For example, when aligning objects (e.g., by *align-by-crossing*) and then distributing them, our test users expected that the previous alignment of the grouped objects is still considered for the next gesture. In our current implementation this is not the case. Each gesture works separately without considering the outcome of a previous gesture. Generally speaking, the current version of NEAT offers two approaches: *guides* as persistent constraints and *gestures* to set non-persistent and temporal constraints. Both approaches coexist in a consistent way within a single system. In future versions we plan to bring both closer together. During gestural interaction at least a few constraints should be considered to make work more productive. On the other hand, while objects are bound to a guide, it should also be possible to manipulate them with abstract gestures (e.g., they could be distributed by shaking and by holding and tapping).

Derived Design Principles and Recommendations

One basic principle for designing tools like NEAT should be to support casual users and experts likewise. To achieve that, in our system we seamlessly combine distinct layout tools providing affordances (the multi-touch alignment guides) and a set of abstract gestures designed for expert users. For our system, we considered techniques which are already known from mouse-based layout tools, such as snap-dragging or showing previews of graphical objects. However, from the initial expert user feedback we derived more promising design principles which are especially tailored to touch- and pen-enabled displays. These principles will be summarized in the following subsections and will be studied deeper in future evaluations.

Physical metaphors The multi-touch alignment guides realize gestures and techniques which are based on physical metaphors (e.g., flicking of objects to bind them to a guide or making a guide sticky to collect objects). The early user feedback indicates that these metaphors are easier to memorize than more abstract techniques and contribute to a natural user experience. However, a system like NEAT should not react to user input like a real physics simulation (as the one in [27]), as this makes the outcome often unpredictable. It should realize naïve physics [16], but the actual reaction should happen in a metaphorical and constrained way to support productivity.

Abstract gestures for complex tasks The pure gestural approach does not allow for switching modes by buttons as in the multi-touch alignment guides. Thus, alternatives have to be found to trigger more complex or mode-dependent actions. In addition, gestures need to be distinguishable from the perspective of both the user and the system. Our proposal is to use more abstract gestures for these complex tasks. Within NEAT we followed the continuous principle of touching objects with two fingers to support major or alternative actions. Furthermore, the touched objects serve as a references (e.g., for alignment). In contrast to the techniques presented by Wigdor et al. [25], we do not apply distinct proxy objects which have to be created in an additional step and can clutter the workspace.

Seamless phrasing With our techniques separate actions (see A1-A3) are combined, as suggested by Buxton [8]. For example, for aligning objects only two steps are necessary (grouping and performing *align-by-crossing*). Beyond that, cloning objects along a path and even changing their spacing can be achieved within a single step. Especially bimanual gestures allow for a more fluent style of interaction. However, they require a higher level of skills and experience.

Pen vs. touch interaction Many of the presented techniques explicitly distinguish between pen and touch input. Thereby, the pen serves for creating objects by sketching and precise interaction (e.g., cutting guides or crossing objects to specify alignments). The combination of both modalities serves for quick mode switches [14] such as creating free-form guides by holding down a finger and sketching with the pen. As mentioned above, this approach also allows to combine separate tasks [8] (e.g., cloning objects along a sketched path or automatically binding them to a sketched guide). Besides treating the pen just as a tool for drawing and writing, we consider it in a more general way - as a tangible. An example for this is laying down the pen on the surface for creating a guide. Beyond that, other objects can be used as well (Figure 8a). We recommend to only apply objects which are used by designers in their daily work, such as physical rulers, triangles or erasers. In this way, the themes of *environment awareness & skills* [16] are considered.

CONCLUSION & FUTURE WORK

We presented NEAT – **N**atural and **E**ffective L**a**yout **T**echniques – a consistent set of tools and gestures for supporting layout tasks on interactive displays. An important design goal was to support novices as well as expert users such as graphics designers. We see NEAT as a first step towards the usage of tools for professionals on interactive displays. For that, it is not enough to just provide a good user experience, but to design effective and productive techniques. For NEAT we considerably extended the multi-touch alignment guides [10]. The novel guides are more flexible and allow the arrangement of graphical objects along arbitrary curves. Novices are supported by interacting in sequential steps and by the application of physical metaphors. These separate steps of interaction can be effectively combined to a single one which is beneficial for experts. Additionally, NEAT provides a set of gestures. It allows the creation of layouts without invoking tools and is explicitly designed for experts. It applies abstract gestures to achieve complex tasks such as aligning objects or cloning objects along paths. For all of our techniques we considered the simultaneous use of multi-touch and pen input. Thereby, pens are applied for sketching and precise input. The combination of both modalities serves for quick mode switches. Furthermore, a seamless transition from the gesture to the tool-centric approach is possible. The presented techniques are integrated in the NEAT system without ambiguities and are successfully implemented in two applications. For future work, we plan to run studies with expert users to verify our design recommendations. Furthermore, we will compare the differences between novices and experts while they are using our NEAT system for given layout tasks.

REFERENCES

1. Adobe Creative Suite, http://www.adobe.com/.
2. Apple Inc. Keynote, http://www.apple.com/.
3. Baudisch, P., Cutrell, E., Hinckley, K., and Eversole, A. Snap-and-go: helping users align objects with-out the modality of traditional snapping. In *Proc. of CHI '05*, ACM, 2005, pp. 301–310.
4. Bias, R. The pluralistic usability walkthrough: coordinated empathies. JSAI Workshops, 1994.
5. Bier, E. A. and and Stone, M. C. Snap-dragging. In *Proc. of SIGGRAPH '86*, ACM, 1986, pp. 233–240.
6. Bragdon, A., Zeleznik, R., Williamson, B., Miller, T., LaViola, Jr., and Joseph J. GestureBar: improving the approachability of gesture-based interfaces. In *Proc. of CHI '09*, ACM, 2008, pp. 2269–2278.
7. Brandl, P., Forlines, C., Wigdor, D., Haller, M., and Shen, C. Combining and measuring the benefits of bimanual pen and direct-touch interaction on horizontal interfaces. In *Proc. of AVI '08*, ACM, 2008, pp. 154–161.
8. Buxton, W. Chunking and phrasing and the design of human-computer dialogues. In *Human-computer interaction*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995, pp. 494–499.
9. Frisch, M., Heydekorn, J., and Dachselt, R. Diagram editing on interactive displays using multi-touch and pen gestures. In *Proc. of Diagrams '10*, Springer, 182–196.
10. Frisch, M., Kleinau, S., Langner R., and Dachselt, R. Grids & guides: multi-touch layout and alignment tools. In *Proc. of CHI '11*, ACM, 2011, pp. 1615–1618.
11. Fung, R., Lank, E., Terry, M., and Latulipe, C. Kinematic templates: end-user tools for content-relative cursor manipulations. In *Proc of UIST '08*, ACM, 2008, pp. 47–56.
12. Guiard, Y. Asymmetric division of labor in human skilled bimanual action: the kinematic chain as a model. Journal of *Motor Behavior*, Vol. 19/4., Dec. 1987, pp. 486–517.
13. Hancock, M., Vernier, F. D., Wigdor, D., Carpendale, S., and Shen, C. Rotation and translation mechanisms for tabletop interaction. In *Proc. of TableTop '06*, IEEE Computer Society, 2006, pp. 79–86.
14. Hinckley, K., Yatani, K., Pahud, M., Coddington, N., Rodenhouse, J., Wilson, A., Benko, H., and Buxton, B. Pen + touch = new tools. In *Proc. of UIST '10*, ACM, 2010, pp. 27–36.
15. Igarashi, T., Matsuoka, S., Kawachiya, S., and Tanaka, H. Interactive beautification: a technique for rapid geometric design. In *Proc. of UIST '97*, ACM, 1997, pp. 105–114.
16. Jacob, R., Girouard, A., Hirshfield, L., Horn, M., Shaer, O., Solovey, E. and Zigelbaum, J. Reality-based interaction: a framework for post-WIMP interfaces. In *Proc. of CHI '08*. ACM, 2008, pp. 201–210.
17. Malik, S. and Laszlo, J. Visual touchpad: a two-handed gestural input device. In *Proc. of ICMI 04*, ACM, 2004, pp. 289–296.
18. Masui, T. HyperSnapping. In *Proc. of HCC '01*, IEEE Computer Society, 2001, pp. 188–194.
19. Microsoft Office (PowerPoint & Visio), http://office.microsoft.com/.
20. Nacenta, M. A., Baudisch, P., Benko, H., and Wilson, A. Separability of spatial manipulations in multi-touch interfaces. In *Proc. of GI '09*, Canadian Information Processing Society, 2009, pp. 175–182.
21. Raisamo, R. and Räihä, K.-J. A new direct manipulation technique for aligning objects in drawing programs. In *Proc. of UIST '96*, ACM, 1996, pp. 157–164.
22. Sugiyama, K. Graph Drawing and Applications. In *World Scientific Series on Software Engineering and Knowledge Engineering, Vol. 11*, ISBN 978-981-02-4879-6.
23. Wang, F., Cao, X., Ren, X., and Irani, P. Detecting and leveraging finger orientation for interaction with direct-touch surfaces. In *Proc. of UIST 09*. ACM, pp. 23–32.
24. White, A. W. The elements of graphic design: space, unity, page architecture, and type. Allworth Press, NY, USA, 2002.
25. Wigdor, D., Benko, H., Pella, J., Lombardo, J., and Williams, S. Rock & rails: extending multi-touch interactions with shape gestures to enable precise spatial manipulations. In *Proc. of CHI '11*, ACM, 2011, pp. 1581–1590.
26. Williams, R. The Non-Designer's Design Book. Third Edition, Peachpit Press, 2008
27. Wilson, A. D., Izadi, S., Hilliges, O., Garcia-Mendoza, A., and Kirk, D. Bringing physics to the surface. In *Proc. of UIST '08*, ACM, 2008, pp. 67–76.
28. Wybrow, M., Marriott, K., Mciver, L., and Stuckey, P. Comparing usability of one-way and multi-way constraints for diagram editing. In *ACM Trans. Comput.-Hum. Interact. 14, 4, Article 19*, ACM, 2008.
29. Yee, K. Two-handed interaction on a tablet display. In *Ext. Abs. CHI 04* ACM, 2004, pp. 1493–149.
30. Zaman, L. and Stuerzlinger, W. A new interface for cloning objects in drawing systems. In *Proc. of GI '10*, Canadian Inform. Processing Society, 2010, pp. 27–34.
31. Zeleznik, R., Bragdon, A., Adeputra, F., and Ko, H.-S. Hands-on math: a page-based multi-touch and pen desktop for technical work and problem solving. In Proc. of UIST '10, ACM, 2010, pp. 17–26.
32. Zeleznik, R., Bragdon, A., Liu, C.-C., and Forsberg, A. Lineogrammer: creating diagrams by drawing. In Proc. of UIST '08, ACM, 2008, pp. 161–170.