Around-Device Interaction for Exploring Large Information Spaces on Mobile Devices

By

Mohammad Khalad Hasan

A Thesis Submitted to the Faculty of Graduate Studies of

The University of Manitoba

in partial fulfillment of the requirements of the degree of

Doctor of Philosophy

Department of Computer Science

University of Manitoba

Winnipeg

Abstract

The standard approach for browsing information on mobile devices includes touchscreen gestures such as pinch and flick. These gestures often require minute operations such as repetitive panning to browse contact lists on a mobile device. Using these gestures to explore large information spaces to facilitate decision-making tasks often involves considerable effort and the user has to deal with screen occlusion and fat-finger situations. However, the void space around mobile devices is much larger than the small touch screen. Researchers have demonstrated that such in-air space can be used as an alternative to touch input for fundamental operations, such as answering and rejecting phone calls. While such prior work has laid the foundation for around-device input, a complete mobile application that deploys and benefits from such an input modality had not been investigated prior to this thesis.

In this thesis, we explored how in-air space around a mobile device can be used to structure mobile interfaces to facilitate complex goals such as making a purchase decision with a smartphone. To achieve this goal, we began with investigating various design factors that influence the performance of accessing content that can virtually exist around the device. We then explored users' and spectators' perceptions of using around-device gestures to access on-device information as their readiness of performing such gestures could lead to rapid adoption of this interaction style. Finally, we used these prior findings to design and structure a complete mobile commerce application with around-device space and compared it to traditional touch interfaces. Study results revealed that using an in-air mobile interface can be more efficient than standard touchscreen interactions. Overall, this research took the first successful step in empirically showing the practical value for using the around-device space for exploring large information spaces on mobile devices.

Acknowledgements

First, I would like to express my sincere gratitude to my academic advisor, Dr. Pourang Irani, for his continual guidance and for being a ceaseless motivator. I am fortunate to have had him as an academic advisor. He has put his trust in my skills, has helped me to discover my strengths, and has nurtured me to become self-dependent. This dissertation would not have been possible without his guidance and inspiration to keep moving ahead with my PhD work.

I owe a debt of gratitude to Dr. David Ahlström (Alpen-Adria-Universität Klagenfurt, Austria) for his support in my PhD research. He has spent countless hours with me discussing study designs and has helped in every way possible in regards to data analysis and reporting. For every question related to my research, I always found him to be there with an answer. His assistance has played a key role in finishing my PhD program.

I would like to thank my committee members, Dr. Yang Wang (University of Manitoba, Canada), Dr. Jason Leboe-McGowan (University of Manitoba, Canada), and Dr. Edward Lank (University of Waterloo, Canada) for their invaluable feedback to improve my dissertation.

I would like to express my appreciation to Dr. Pourang Irani, the Government of Manitoba, the Faculty of Graduate Studies, the Faculty of Science, and the Department of Computer Science for providing financial support during my graduate program. I extend my gratitude to them for travel grants, which have enabled me to attend pertinent conferences.

The HCI lab at the University of Manitoba, equipped with the state-of-the-art technologies, has provided me with a uniquely excellent research environment. I am very fortunate to have a team of extraordinary lab members: Dr. Barrett Ens, Dr. William Delamare, Teng Han, Junhyeok Kim, Fereshteh Amini, Gaganpreet Singh, Fouad Alallah and Ali Neshati. I would like to thank them for their inputs, ideas, and suggestions in my research. I would also like to thank John Swystun (Faculty of Management, University of Manitoba) for proofreading my thesis.

I would like to sincerely acknowledge the unconditional support from my family members during my endeavour to successfully finish my graduate program. My parents have always been a tremendous source of inspiration during my PhD program. They left no stone unturned to support my education, have always kept faith in me, and have helped me to believe in my abilities. I also thank my sisters, Mahmuda Siddiqua, Saima Siddiqua, and Ayesha Siddiqua and my brothers-in-law Md Omar Farooq and Dr. K.M. Ashraful Huda for their endless love. I would like to make a special mention of my wife Afrina Rahman and my daughter Anayah Hasan for their unconditional support throughout my graduate study. They sacrificed a lot to ensure that my endeavours have gone smoothly. I am what I am only due to my family members' efforts.

I am also very fortunate to have met some extraordinary people who have made my stay in Winnipeg very pleasant. I would like to mention two names upfront: Dr. Debajyoti Mondal and Farjana Ferdous. They spent countless hours with me drinking coffee, sharing leisure time, and debating on different topics. Additionally, I would like to mention a few names whom I consider to be my extended family in Canada: Dr. Mohammad Al-Amin Sadek, Afzal Hossain Rupak, Kaiser Moon, Kazi Rushdi, Zubair Ahmed, Rubiat Shakil, Abul Khair (and their families), Ferdous Ahmed, Purbasha Mistry and Jawata Afnan. I thank them for standing by me in all circumstances.

Finally, I would like to dedicate this thesis to my wife, Afrina Rahman, for her support in the pursuit of my lifelong dream.

Copyright Notices and Disclaimers

Three chapters of this thesis have already been published in conference proceedings and one chapter has been accepted to publish in a conference. Permissions to reuse the published chapters have been obtained from publishers. Permission details, and the terms and conditions for these are included in Appendix A.

Khalad Hasan, Junhyeok Kim, David Ahlström, and Pourang Irani. AirPanes: Two-Handed Around-Device Interaction for Pane Switching on Smartphones. In *Proceedings of the 31st International Conference on Human Factors in Computing Systems* (CHI 2017) Denver, CO. ACM. (To appear).

Khalad Hasan, Junhyeok Kim, David Ahlström, and Pourang Irani. Thumbs-Up: 3D Spatial Thumb-Reachable Space for One-Handed Thumb Interaction on Smartphones. In *Proceedings of the ACM Symposium on Spatial User Interaction 2016* (SUI 2016), Tokyo, Japan. ACM. 103-106.

David Ahlström, Khalad Hasan, and Pourang Irani. 2014. Are You Comfortable Doing That?: Acceptance Studies of Around-Device Gestures in and for Public Settings. In *Proceedings of the 16th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI 2014), Toronto, Canada. ACM. 193-202.

Khalad Hasan, David Ahlström, and Pourang Irani. 2013. AD-Binning: Leveraging Around Device Space for Storing, Browsing and Retrieving Mobile Device Content. In *Proceedings*

of the 31st International Conference on Human Factors in Computing Systems (CHI 2013), Paris, France. ACM. 899-908.

Contents

Abstract II
Acknowledgements III
Copyright Notices and Disclaimers V
Chapter 1: Introduction 1
Motivation 2
Research Objective 4
Research Questions 4
Research Approach 6
Contribution 10
Chapter 2: Related Work 14
Around-Device Interaction 14
Around-Device Sensing 21
Infrared proximity sensor 22
Vision-Based Sensing 23
Magnetic-Field Based Sensing 26
On-Board Sensor Based Sensing 27
Non-Speech Voice and Acoustic-Based Sensing 27
Electric Field Sensing Approaches 29
Other Sensing Approaches 29
Around-Device Feedback 31
On-Device Visual Feedback 31
In-Air Visual Feedback 33
Non-Visual On-Device Feedback 36
Non-Visual In-Air Feedback 37
AD-Interactions and Social Acceptance 40
Social Acceptability 41
Thumb Input on Mobile Devices 44
Bimanual Interaction 46

Information Browsing on Mobile Devices 47
Chapter 3: Design and Evaluation of the Around-Device Space for Accessing Mobile Device Content.................... 51
AD-Binning Design Framework 52
Design Factors 52
Selection methods....................52
Bin size 53
Visual feedback 53
Space discretization for bins 54
Input range around the device....................54
Ideal binning locations 54
Binning methods....................54
Mode switching....................55
AD-Binning Apparatus and Input Range 55
Apparatus 55
AD-Binning Input Range....................56
Study 1: Design Factor Exploration....................57
Selection Methods 58
Session A – Selection Methods and Bin Size 59
Results 60
Summary – Session A....................63
Session B – Space Discretization & Binning Locations....................63
Task and experimental design.................... 64
Results 65
Summary – Session B 68
Study 2 – Binning & Retrieval 68
Participants, Task and Experimental Design.................... 69
Results 71
Summary 73
Study 3 – Analytic Task 74

Participants, Task and Experimental Design....74
Results....77
Error rate....77
Trial time....77
Preference ratings....78
Summary....79
Discussion....79
Design Considerations....79
Applications....81
Summary....82
Chapter 4: Acceptance Studies of Around-Device Gestures in and for Public Settings....83
Study 4 – Comfortable Region and Distance....85
Results....88
Summary....93
Study 5 – Comfortable Gesture Size and Duration....94
Results....96
Summary....99
Study 6 – Spectators Attitude about AD-gestures....100
Results....101
Discussion....103
AD-Input Design Considerations & Recommendations....103
Adaptations to AD-Interactions....105
Improved Methods for Acceptance Studies....106
Summary....107
Chapter 5: Thumbs-Up: 3D Spatial Thumb-Reachable Space for Around-Device Interaction....109
Study 7: Thumbs-Up Input Range....110
Participants and Apparatus....111
Tasks....112
Results....113

Study 8: Selection and Item Placement Style on Thumbs-Up Space 115
Task.. 117
Results 118
Summary 120
Design Recommendations 121
Summary... 121
Chapter 6: AirPanes: Two-Handed Around-Device Interaction for Pane Switching on Smartphones 123
Background... 124
Study 9: AirPane Layout and Pane Switching...127
Task.. 130
Results 132
Study 10: Using AirPanes in an m-commerce interface.................................... 134
Task 138
Results.. 139
Summary 142
Discussion... 143
Design Recommendations.. 143
Considerations for Integration of AirPanes in Applications 144
Summary... 146
Chapter 7: Conclusion 148
Summary... 149
Limitations... 151
Future Work 155
Around-device sensing 155
Around-device feedback.. 156
Around-device display 157
Final words...157
References 159
Appendix A... 187

List of Tables

Table 1: A comparison of different sensing approaches used in around-device interaction 21

Table 2 Feedback on around-device items can be presented with on-device or in-air with visual and non-visual approaches. 31

Table 3: A comparison of different feedback mechanism for around-device interaction .. 39

List of Figures

Figure 1: To make a hotel reservation, a user's query puts hotel information in AD-Bins around the mobile device. As (a) the finger hovers on an AD-Bin (b) its content is shown on the screen allowing the user to browse and compare alternatives. 6

Figure 2: Around-Device input, in public can create feelings of discomfort, but only on specific gesture parameters. 7

Figure 3: Two-handed around-device interaction in an m-commerce application using AirPanes. a) Previously tagged favorite products are access in-air with the thumb of the hand holding the phone. b) Query result lists and c) filter options are accessed using in-air panes reachable with the other hand. 8

Figure 4: An overview of the thesis.. 13

Figure 5: (a) Vicon apparatus. (b) The AD-Binning space. 56

Figure 6: The AD-Binning Input Range, i.e. the range accessible by the right hand around the device (center) in centimeters 57

Figure 7: (a) Trial start screen. (b) Visual feedback overview. 59

Figure 8: (a) and (b) Mean error rates. Error bars: ± 1 S.E. (c) and (d) Geometric mean trial times. Error bars: 95% CI. 61

Figure 9: Discretization techniques and distance mapping. 64

Figure 10: (a) and (b) Interaction effects, error rate. (c), (d), and (e) Interaction effects, trial time... 65

Figure 11: (a) and (b) Sector utilization. (c), (d), (e) and (f) Geometric mean retrieval times. Horizontal axis in (e) and (f) represents the first, second and third item retrieval. Error bars: 95% CI.. 72

Figure 12: (a) On-screen interface with 10 hotels at low density. (b) AD-Bins with 15 hotels at high density.. 75

Figure 13: (a) numbers of items x interface interaction. (b) numbers of items x density interaction. 77

Figure 14: AD-notes. (a) Creating, arranging and storing physical notes. (b) Browsing notes stored in AD-space. 81

Figure 15: Example task prompts for (a) a gesture in the Above region at the Close distance and for (b) a gesture in the Right region at the Far distance. (c) Gesture distances and regions.. 85

Figure 16: Study 4 questionnaire 87

Figure 17: Acceptance rates for audiences and locations. 89

Figure 18: Acceptance rates for audiences and locations split by introverted and extroverted users........ 90

Figure 19: Comfort-ratings of the 15 gesture positions (segments around each middle circle show individual ratings, best viewed in color)........ 91

Figure 20: Left: task screen. Right: Study 2 questionnaire........ 95

Figure 21: Comfort-ratings for all gesture size-duration combinations (segments around each middle circle show individual ratings, best viewed in color)........ 96

Figure 22: Acceptance rates for audiences (top) and locations (bottom). Rates aggregated across size and duration (left), split by size (middle), and split by duration (right).. 98

Figure 23: Study 6 questionnaire........ 100

Figure 24: Acceptance rates for audiences and locations........ 101

Figure 25: (a) Thumbs-Up interactions occur in the thumb-reachable in-air space around a smartphone. A 3D printed shape of this thumb-reachable space is shown in (b)..... 110

Figure 26: (a) Map navigation and (b) Space filling task........ 112

Figure 27: TouchSpace (a) the left and (b) the right thumb. Thumbs-Up space with (c) the left and (d) the right thumb. In-air thumb-reachable space with (e) the left and (f) the right thumb. The red areas indicate TouchSpace and the blue areas represent Thumbs-Up space........ 114

Figure 28. a) On screen visualization with a colored cursor to indicate the position of the thumb inside the in-air region in relation to the items, arranged in two and b) three layers........ 115

Figure 29. Selection methods. (a) Dwell: requires users to keep their finger over an item for a 600ms to trigger a selection; (b) Tilt: rotating the device clockwise or counter-clockwise direction with a certain angle triggers selection; and (c) Touch: a quick tap on the screen from the mid-air item position invokes the selection........ 116

Figure 30. Number of errors for item positions from left to right in each layer with (a) 2-Layer and (b) 3-Layer layout........ 119

Figure 31. Geometric mean time for item positions from left to right in each layer with (a) 2-Layer and (b) 3-Layer layout........ 119

Figure 32: The 24 most visited m-commerce services (top 15 unique providers marked with a *). a) General UI structure for the surveyed interfaces. b) example responsive website's overview list and c) a detail view in a mobile app........ 126

Figure 33. Five AirPane layouts........ 129

Figure 34. a) A participant moves the 'pinched' in-air fingers to steer the on-screen cursor over the target. b) Releasing the pinch with the cursor over the target selects it. c) Start screen........ 131

Figure 35. Geometric mean trial times, a) layout, b) target size........ 132

Figure 36. AirPanes: a) product details, b) filter options, and c) product detail for a bookmarked item. Classic touch interface: d) a scrollable overview list of search results, e) a view for product details, f) a view with filters, and g) a modal view with filter options. 136

Figure 37. Geometric mean times. a) Trial time, b) filter time, c) results time, and d) favorite time. Error bars: 95 CI. 140

Figure 38. In-air pane organization: a) and b) for one-handed interaction with nested panes, c) for two-handed interaction using in-air thumb space for pane switching. 145

Figure 39. AirPanes can leverage a nearby surface for haptic feedback, e.g., (a) a knee, (b) table, or (c) wall. 146

Chapter 1: Introduction

Jennifer plans to purchase a camera for her husband as a birthday gift. She wants a camera that meets the following criteria: i) the price is within her budget; ii) it has a good user rating; and, iii) it is either from the Canon or Nikon brand. During her coffee break at a café, she opens a mobile commerce application on her smartphone, applies the criteria and starts browsing available camera options displayed in a list. When she finds a prospective camera, she explores it by tapping on the item, which switches to a new view with detailed camera information (i.e., detailed view). After a quick look, she bookmarks it for later comparison. At this point, she switches back to the list view and taps the next prospective item. This action again opens a new detailed view with corresponding camera information. This multi-view user interface structure forces her to go back and forth between the list and detailed view. After her exploration, she decides to check the bookmarked items to make a final purchase decision.

This scenario is common in our everyday life where we use mobile devices for analytic tasks, in this case, to look for a specific set of cameras. It involves browsing, comparing and re-inspecting previously visited items to make a final decision. Current mobile interfaces are not suitable to perform such tasks as switching between different windows (e.g., the list view and detailed view) imposes additional cognitive load and breaks the seamless interaction.

Motivation

Mobile devices are an indispensable part of our everyday life, replacing the ways we interact with information. Instead of using it for communication purposes, such as making phone calls or sending messages, it has now been adopted to replace the functionalities of traditional desktop systems. It is gradually becoming popular for interacting with large information repositories. Estimates suggest that over 50% of smartphone owners browse and research products on their mobiles before making a purchase, 60% of last minute hotel bookings are made on mobile devices, and search on mobiles will generate 27.8 billion more queries than on desktops by 2016 [1]. Additionally, mobile devices are preferred over the traditional desktop by business people for conducting research on commercial products [141].

Though mobile devices are now considered a primary medium to access information, they have several limitations. Mobile devices are typically equipped with a small display where touch is considered natural to interact with the device as it provides direct access to an item of interest. To accommodate large information, mobile app designers most often structure the interface into multiple panes or views [205]. Each such view is either dedicated to a specific functionality or for displaying multiple related information items, in most cases using a scrollable list. Such list views are quick to use if the necessary detail information about each individual item is visually accessible while scrolling the list. However, in many interfaces this is not the case: the user can only access item details by selecting the desired item in the list view, which displays the details in a following full-screen detail view. With a tap on a back-button in the detail view, the user returns to the list view. Switching back and forth between the list view and detail views quickly becomes tedious if the user needs to inspect details of several items. Another common limitation of mobile devices is occlusion where a significant amount of information presented on the small screen is occluded under the finger while interacting with the devices. Additionally, small widgets specially designed for mobile screens are commonly error prone due to the coarse-grained finger input which is commonly known as the 'fat finger' problem. Tasks with high information bandwidth often require that users quickly browse and compare the breadth of available choices before making a decision. With the limitations mentioned above, on mobile interfaces such tasks require many minute operations, such as flicking through screens and opening and closing items of interest, resulting in less efficient information exploration and browsing.

The surrounding void space with mobile devices is considerably larger than the small screen space they are built with. Ideally, mobile devices could be augmented with the virtual interactive plane that extends beyond their physical form factor for accessing information. Instead of switching between multiple views on mobile interfaces, a user can use around-device space for interacting with views or panes that are placed in off-screen space for fast access. Shifting input intensity away from the screen, into a much larger space around the device, provides larger proxy objects for interacting with smaller on-screen items, minimizes the likelihood of having the input hand occlude important screen content, and reduces the need for frequent repetitions of small on-screen manipulations. In this thesis, we propose and study this idea of using ***'Around Device' (AD)*** space for supporting analytic decision-making tasks that require browsing through large items located in that space. To this end, we design, implement, and study the design space available to this input modality.

Research Objective

Around-device interaction has opened new lines of inquiry that resolve some of the challenges with mobile device interactions. With rapid advances in optical sensing and finger tracking technologies [28,150,193], researchers have explored the use of the in-air space around a mobile device for input. Prior studies have demonstrated the use of in-air space for fundamental on-screen interactions, such as selecting on-screen [70] and off-screen items [71], switching modes [179], text-entry [111], and zooming and panning [99]. While such prior work has laid the foundation for around-device in-air input, a complete mobile application that deploys and benefits from such an input modality has yet to be demonstrated.

The main research objective of this thesis aims to step beyond the design and study of in-air alternatives for standard on-screen interactions and to explore how in-air input can enhance user performance in a complete mobile scenario. In particular, we focus on the ability to facilitate a complex goal, such as making a decision through information exploration and interaction. To achieve this goal, we examine and address the following three research questions:

Research Questions

Research Question 1: How can the around-device space be designed for browsing virtual items that are placed in the surrounding area of mobile devices?

Research Question 2: What are users' and spectators' attitudes about using such new interaction metaphors in ecologically valid settings such as in public places?

Research Question 3: How can a mobile application be designed to leverage around-device space for exploring large information quantities?

In this thesis, we start by exploring the first research question and progressively move towards the other two. The first question focuses on identifying and resolving key human-factors issues for using the surrounding space of a mobile device for content browsing. For instance, we examine suitable around-device target size, item selection, and item placement techniques. The second question is interlinked with the first question, where we investigate the factors that influence users' and spectators' attitudes and acceptance of interacting in this void space in different public locations and settings. We investigate how gesture properties (e.g., duration and distance), users' contexts (e.g., private and public places) and audience types (e.g., familiar and non-familiar audience) influence users' and spectators' willingness to perform in-air gestures. Based on this knowledge, we finally explore how to design the around-device space to facilitate the functionalities of mobile commerce applications that require browsing and interacting with large information content before arriving at a decision. These explorations push the boundaries for around-device interactions to facilitate an entire mobile workflow: browsing through large information (e.g., long query results) with in-air mobile interfaces for decision-making tasks that involved frequent switching, browsing, and revisiting of items.

In brief, this thesis aims to gather meaningful findings and design suggestions for using around-device space for supporting analytic tasks on mobile devices where users are required to browse large information quantities on mobile applications before making a final decision.

Research Approach

The research approach taken in this thesis is strongly connected with the research questions that are listed previously:

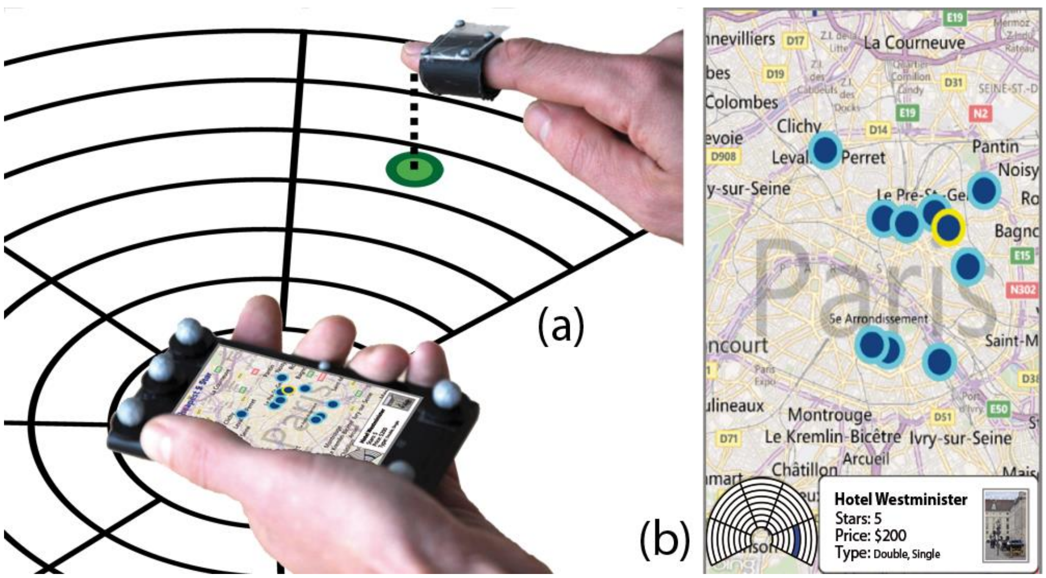

Figure 1: To make a hotel reservation, a user's query puts hotel information in AD-Bins around the mobile device. As (a) the finger hovers on an AD-Bin (b) its content is shown on the screen allowing the user to browse and compare alternatives.

Question 1: To address the first question, we propose and study Around-Device Binning, or AD-Binning, a novel mobile user interface that allows users to off-load mobile content into the space around the device. With AD-Binning, a user can directly access the off-loaded items by moving her hand around the device. We informed our implementation of AD-Binning by exploring various design factors for placing content in off-screen space.

Through two user studies, we identify key human-factor issues, such as (a) What are suitable methods for placing items off-screen?; (b) What selection methods provide efficient retrieval? and, (c) how small can targets be before affecting performance? With this knowledge, we design a novel interface, AD-Binning, where a user can directly access in-air items by moving his hands around the device. In a task requiring content browsing for making a decision, we show that participants were more efficient with AD-Binning than with on-screen exploration.

Figure 2: Around-Device input, in public can create feelings of discomfort, but only on specific gesture parameters.

Question 2: The first research question looks at using the around-device mid-air space for accessing on-screen items. However, little is known about users' attitudes using these innovative interaction styles. Particularly when performed in a public setting, hand movements and finger gestures around the device may attract by-passers' undesired attention or intrude into areas 'owned' by others (e.g., when sitting on a bus), and thus may evoke feelings, such as embarrassment or discomfort. Accordingly, the acceptance and willingness to perform AD-gestures may be limited to certain settings. In the chapter 4, we explore how socially comfortable (we term this as 'comfortable' throughout the remainder of the thesis) users feel when performing AD-gestures in a public place. We also survey users for which locations and in front of whom they would feel comfortable using AD-gestures. With two studies, we examine the influence of fundamental AD-gesture features - the distance from the device, the position relative to the device, gesture size and gesture duration - on users' level of comfort. We further examine whether such perceptions are related to a user's introversion/extroversion personality trait. We then switch to a spectator's point of view and examine peoples' reactions when having observed someone else using AD-gestures. We elicit opinions from people observing others using AD-gestures in public.

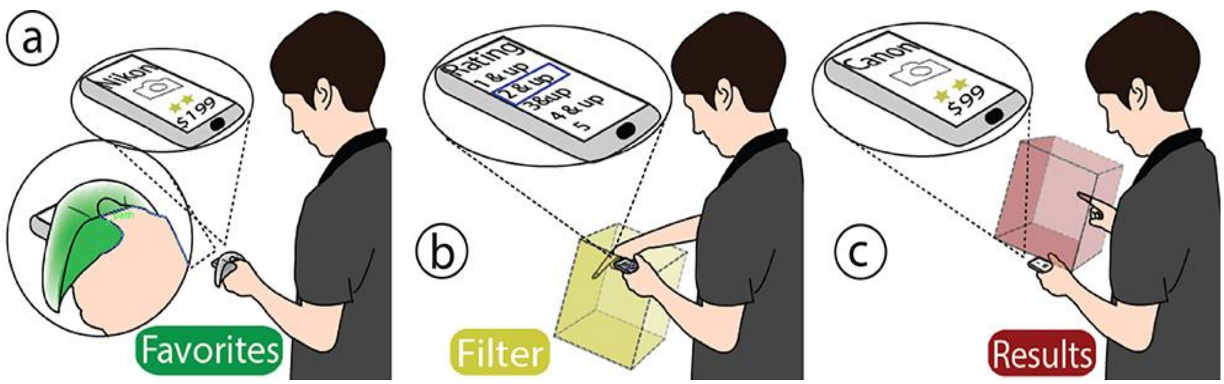

Figure 3: Two-handed around-device interaction in an m-commerce application using AirPanes. a) Previously tagged favorite products are access in-air with the thumb of the hand holding the phone. b) Query result lists and c) filter options are accessed using in-air panes reachable with the other hand.

Question 3: This research question aims to explore how in-air space can be used to enhance user performance in more complete mobile scenarios. In particular, we focus on the ability to facilitate a complex analytic task, such as making a decision through information exploration and interaction. While designing an application with around-device space, we realized that prior research has commonly considered the right, top and left regions surrounding the mobile device for around-device interaction. However, the void space just above the device reachable by the thumb of the hand holding the phone has been heretofore an unexplored interaction space. Therefore, we start with exploring the suitable input range of thumb movements while the smartphone is held in the hand. We refer to this space as Thumbs-Up. With Thumbs-Up, the in-air space could be used to access on-screen content through directly pointing with the thumb, as shown in Figure 3a. After having defined the accessible in-air thumbs-up region, we explore ways to utilize this in-air space in conjunction with other around-device spaces. We envision that it useful for triggering commands and for storing, selecting and browsing information items. Accordingly, we explore various ways to arrange items within the accessible thumb-space and methods to select such in-air items.

Based on this information, we next investigate how to structure interfaces of a mobile application with these in-air spaces. In the context of mobile applications, GUI designers generally organize large information spaces into multiple views or panes. Typically, each pane serves a specific function, such as providing interactive controls to query a large dataset, showing the query results, or showing details of a selected list item. This use of a multi-pane UI structure (analogous to Tabs or Windows in desktop applications) forces users to frequently switch back and forth between views, which quickly becomes tedious for even common tasks such as looking at details of items in a list.

To this end, we propose AirPanes, a novel strategy to structure a mobile interface, using panes located in mid-air around the device. As an exemplary scenario, we pick mobile commerce (also known as m-commerce) applications as these are used to purchase products through mobile devices by millions of users [139] and require browsing and interacting with large information content before arriving at a decision. We demonstrate the benefits of AirPanes in an analytic decision marking task where the user browses products, applies filters, inspects result lists, and bookmarks interesting items before making a final purchase. We design AirPanes to take full advantage of the spaces around the device, i.e., both the spatial region accessible by the thumb of the hand holding the device, as shown in Figure 3a, and the in-air space reachable by fingers of the non-holding hand, as shown in Figure 3b and c. Via a user study, we first optimize the design parameters for organizing around-device interfaces and its components. A further user study confirms that using in-air mobile interface can be more efficient than standard touchscreen interactions, specifically when it concerns analytic decision marking tasks that require exploring a large information space.

Contribution

My Ph.D. thesis makes the following seven contributions:

1. It introduces the concept of Around-Device Binning or AD-Binning, a novel technique that leverages the surrounding void space of a mobile device for storing and browsing content through direct interaction with around-device space. It also offers design guidelines for AD-Binning that other applications can benefit from using AD-space for accessing information on mobile devices.

2. It demonstrates a prototype of AD-Binning for a complex information exploration and decision-making task. With the prototype application, a user can off-load mobile contents around the device and retrieve the information by moving his finger to the AD-space. This capability allows the user to perform analytic tasks such as browsing hotel information from many alternatives to make a reservation with around-device space.
3. The thesis provides insight on aspects of AD-gestures that influence user comfort and acceptance in public settings. Our investigation informs the AD-researchers about critical parameters for designing AD-gestures, such as gesture properties (e.g., duration and distance), users' contexts while using the gestures (e.g., private and public places) and surrounding audience types (e.g., familiar and non-familiar audience). We reveal that people are selective concerning the settings where they would use such AD-gestures. Additionally, we reveal that the gesture properties have an influence on users' comfort levels.
4. This thesis offers key design guidelines and recommendations for creating socially acceptable AD-gestures in both public and private settings. These guidelines are critical for the successful adaptation of AD-interaction on mobile devices as our results indicate that micro-level AD-gesture features are important in designing the gestures.
5. This thesis introduces and defines Thumbs-Up space, the in-air space just above the mobile device reachable via thumb of the hand holding the device. We show that

with Thumbs-Up, the in-air space could be used to access on-screen content on mobile devices.

6. This thesis introduces the concept of AirPanes, a strategy to structure mobile interfaces using content panes located in the air around the device. Additionally, we explore and reveal insights on key AirPanes design factors.
7. Finally, this thesis offers the first mobile prototype application that leverages around-device space. Additionally, we show the benefit of using AirPanes over tradition touch interfaces for an analytic decision-making task that requires browsing through large item sets and frequent switching between interface views.

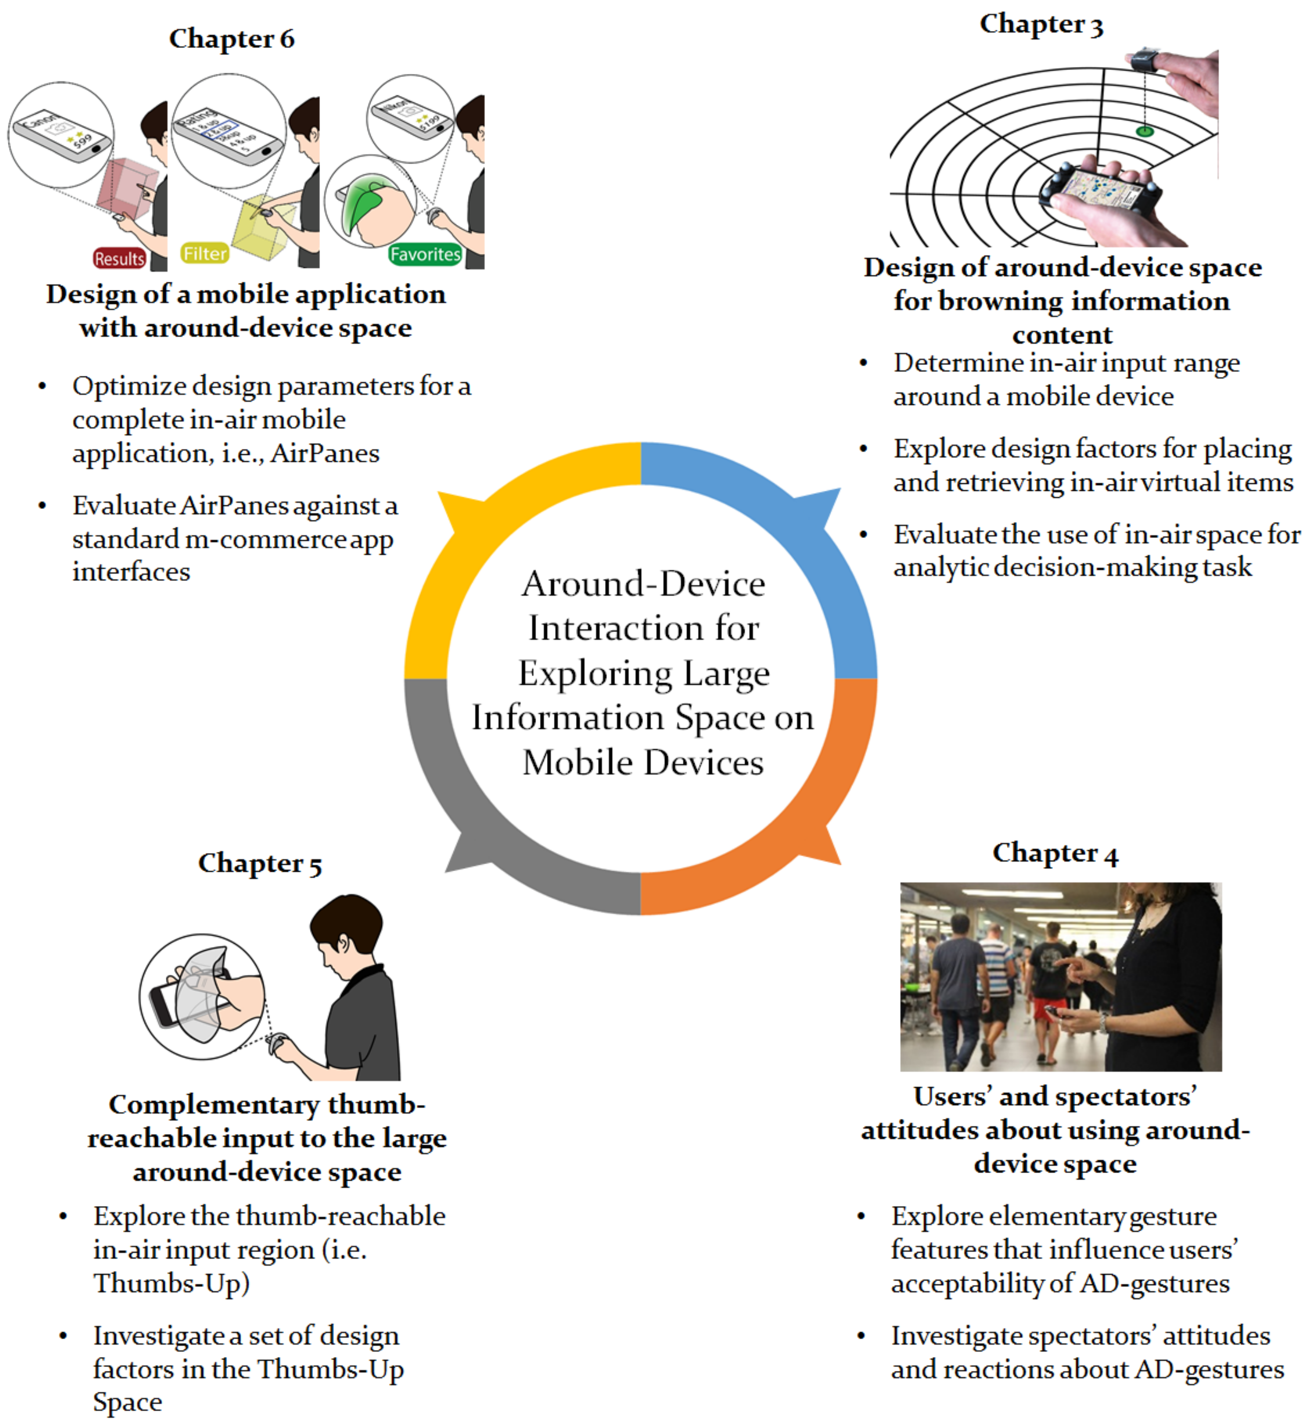

Figure 4: An overview of the thesis.

Chapter 2: Related Work

In this thesis, we examine use of the surrounding space of a mobile device to facilitate complex analytic tasks that require users to browse and interact with large set of information. This work is primarily inspired by previous research on around-device interaction. In this chapter, we start with a comprehensive literature review of in-air interaction with mobile and wearable devices. We find that in-air interaction relies on robust sensing technologies to track hand and finger movement surrounding the device. Thus, we discuss the current progress on around-device sensing techniques. In addition, in-air interaction benefits from feedback mechanisms that guide users to explore off-screen content. We review work done on feedback mechanisms for around-device items. Our investigation of social acceptance of in-air interaction is inspired by recent results on social acceptability studies on body- and mobile-based gestures which we review next. Additionally, in this thesis, we explore thumb-based in-air interaction to support two-handed input on mobile devices. We briefly cover related work in the area of thumb and two-handed input. We then conclude this chapter with a review of recent research on information browsing on mobile devices, a task that we frequently apply in evaluating the efficiency of in-air interactions for information exploration.

Around-Device Interaction

The limited input space of mobile devices has inspired researchers to explore the void space around devices to develop richer interactions with the device. Prior work on around-device interaction has demonstrated that the in-air space surrounding a device can be used as an alternative to touch input [70,115,120,208]. Interactions explored include accessing application features, workspace navigation, game or imaginary object controller or tangible interaction which are discussed below.

Accessing Application Features: Among the most basic interactions supported by AD-space include highlighting and selecting on-screen items or changing discrete parameters within an application (e.g., the volume of a music player). For instance, Harrison and Hudson [70] showed that in-air input can be used for controlling an on-device cursor and making item selections. They demonstrated that in-air selection can also be applied to small devices such as smartwatches to invoke on-screen widgets with high accuracy. Kratz and Rohs [120] also showed that in-air input can be used for selection purposes. They used in-air hand movement to highlight and select a colour from a colour palette. In addition, around-device space has also been explored for interacting with User Interface (UI) widgets, such as interacting with navigation buttons to change the song playing in a music player [115], and to control a slider to change the volume [208].

Researchers have also shown how to leverage AD-space in more complex interactions, such as supporting multi-touch operations, controlling call management fractures, and entering text on devices without using the touchscreen. For instance, SideSight [20], an early project on around-device interactions, showed that the in-air space around a mobile device can be used for multi-touch interactions. The authors demonstrated that around-device space supports activities that require multi-finger operations (e.g., zooming a map with two-finger pinch gesture on the touch screen). A similar idea was later explored in other device platforms: Bi et al. [18] explored how to support multi-touch interactions surrounding a desktop computer where regions around the computer are used as multi-touch surfaces. Researchers also showed the use of in-air hand gestures to call management purpose, such as answering or rejecting phone calls [112,114]. They showed such features can be useful while controlling call management from mid-air in situations where the users' hands are not clean enough to interact with the touchscreen, or the phone is in a pocket. Ketabdar et al. [111] showed that the 3D space around the device can be utilized for text or digit entry purposes. Instead of using an on-device soft keyboard or touchpad, their application allows users to draw digits with a magnet. In addition, researchers also showed that text entry can be done by drawing stroke in this space, which helps reduce the screen occlusion issues [27,111,147].

Around-device space has also been used for supporting more advanced interaction such as triggering shortcuts on a mobile device, task switching on applications, or as a game input controller. Li et al. [129] developed Virtual shelves, a technique that uses a circular hemisphere in front of a user's body to place shortcuts in 7×4 grid. Users can trigger an application on the mobile device by moving it into that a grid. With a study, they revealed that accessing items with Virtual Shelves is faster than the native interface on the mobile device (i.e., touch interface). Similarly, Hsieh et al. [197] implemented a technique, Pile Across Space, that allows users to store virtual piles around a device with flick gestures. They revealed that users can efficiently retrieve such in-air information by leveraging their spatial memory. Switching tasks on mobile devices commonly requires multiple touch and pan actions. Grubert et al. [58] showed that application switching can be avoided by using around-device space. Items can be placed around the device in a specific position; moving the hand to that position triggers the corresponding item. Similarly, Hakoda et al. [67] showed that in-air hover gestures could be useful for switching tabs in a web browser. Recently, Song et al. [179] demonstrated that around-device gestures can enrich the existing interaction dictionary by exploring the use of AD-Space for a number of tasks, such as mode switching, menu selection, navigation, and application management.

Workspace Navigation: Navigating a large workspace with around-device input has been shown to be a promising alternative to standard touch gestures such as flick and pinch [108,152,181]. In a recent study, Spindler et al. [181] showed that using in-air space around a mobile device can significantly improve workspace navigation performance. They revealed that users leverage a larger in-air interaction space with in-air input while browsing a workspace. This space is commonly much smaller with a touchscreen than the in-air space as the touch space is limited by the devices' physical boundary. Similarly, Hasan et al. [73] examined two map navigation techniques that leverage around-device space and compared their performance with the standard pinch-and-flick gestures. Their results showed that navigating from one map location to another with the in-air technique is faster over flick and pinch. Jones et al. [99] also investigated workspace navigation performance with in-air input where they showed that in-air gestures can be as good as traditional touch input for workspace navigation. Hwang et al. [90] developed a prototype to track a user's fingernail around the device. They showed that nail input can be used for browsing webpages. In a recent work, Chen et al. [28] showed that map or document navigation can be performed with in-air hand movement, as users can interact with the workspace without occluding the screen.

Game Controller: In-air space has also been demonstrated as an input space for game controllers. Steins et al. [183] developed a technique called Imaginary device that allows users to control a game with in-air hand gestures and postures such that users exploit their previous experience of using physical gaming input devices (i.e. Joystick or driving wheel) to control games. As examples, they showed that a steering wheel holding posture can be mapped to drive a car or an imaginary joystick movement can be used to fly a plane. With user two studies, Ketabdar et al. showed that their prototype can detect participants' hand posture with a high accuracy (~97%).

Imaginary Object Controller: In-air space has also been shown to be useful for controlling imaginary objects, such as imaginary phones, imaginary interfaces, imaginary gaming controls and imaginary devices. Gustafson et al. [63] initially explored the concept of using in-air space for interacting with imaginary interfaces such as drawing on a virtual canvas with in-air gestures. They examined users' performance using such imaginary interfaces in scenarios where no devices are available but users can use their hand as a reference point to a virtual device. Their results revealed that having a reference point is critical for such imaginary interfaces and users' performance decreases when interactions are further away from that point. In a later study, Ens et al. [38] showed that user performance when accessing virtual interfaces can be improved by using different visual feedback mechanisms such as static and dynamic visual cues.

As an extension of work on imaginary interface, Gustafson et al. [66] explored whether imaginary interfaces could benefit by leveraging the spatial knowledge that user develop using a real-world interface. They developed an Imaginary Phone interface which leverages users' spatial knowledge of mobile widgets (e.g., home screen app icons) to access an imaginary phone that is placed on the user's palm. Via two studies, they showed that users can recall (with an average accuracy of 68%) the position of widgets on their mobile phone. In additional follow-on work, Gustafson et al. [65] also showed that visual sense (i.e., fingers position in the hand) plays an important role in these interfaces as it assists users to interact with the virtual mobile widgets with their spatial knowledge. They revealed that when the visual sense is removed by obscuring the users' view, they depend on the tactile feedback to access widgets on such interfaces.

Tangible Interaction: Around-device space has been demonstrated to be effective for tangible interactions. Avrahami et al. [10] built a system called Portico that uses two cameras above the device display to enable tangible interactions around a portable device (e.g., tablet). The system detects and recognizes different physical objects around the device and users interact with digital elements on the device by manipulating the physical objects around the device. Kane et al. [104] extended this idea by incorporating output capabilities in a self-contained device. They added two micro-projectors to provide an interactive display around the device. The camera and micro-projector work together to provide input (e.g., recognizing objects, gestures) and output capabilities (e.g., projecting information to the around-device space). Hwang et al. [88] developed a prototype that allows tangible interaction with magnetically driven controllers around the device. The controllers can be operated without power or wireless connection and can be used for a wide range of application scenarios. Examples of such scenarios include designing external attachments (e.g., a physical joystick or a physical slide with the controllers) interacting with mobile UI widgets, performing gestures with the controllers for authentication purpose, or attaching a digital pen for supporting drawing application in around-device space.

Though around-device interaction has been investigated for many scenarios, it has received very little attention in the context of using the space as storage. Even research that has explored storing items around the body [129] or creating virtual piles[197], has not explored how to use the around-device space for accessing large information repositories (e.g., browsing long query results) with in-air mobile interfaces. In addition, to our knowledge, around-device interactions have not been evaluated with advanced analytic tasks in which users are required to interact with large information content on mobile devices for exploratory searching tasks that involved frequent switching, browsing, and revisiting of items.

Around-Device Sensing

There has been substantial previous work from industry and academia investigating different sensing mechanisms that allow users to track hand movement in real-time around the device. There are several industry solutions, such as Vicon motion capture system [193], optiTrack motion capture systems [150], the Microsoft Kinect[117] and the Leap motion [127] that allow researchers to track users' hand and finger movements surrounding a device. Several prior projects have used these commercial solutions [39,47,49,65,72,73,168,185]. These systems commonly use multiple infrared cameras to track an object in 3D. Though such commercial solutions provide precise real-time motion capture data, miniaturizing for portability on mobile devices is not possible.

		Augmentation									
		Device		Finger/Body		Range			Accuracy		
	Sensing Solutions	Yes	No	Yes	No	Low	Mid	High	Low	Mid	High
	Commercial		X	X	X			X			X
	Infrared Sensor	X			X	X				X	
	Depth sensing camera	X			X		X			X	
	External Camera	X			X		X			X	
On-board Sensors	Magnetic sensor		X	X			X			X	
	On-Device Camera		X		X	X			X		
	IMU		X		X	X			X		
	Microphone		X		X		X		X		
	Electric Field Sensing		X		X		X			X	

Table 1: A comparison of different sensing approaches used in around-device interaction

Accordingly, a number of research projects have mounted external sensors on a mobile device to detect in-air hand movement in the device's vicinity. Researchers have explored a number of around-device sensing technologies which are discussed below (a quick overview of the techniques is listed in Table 1)

Infrared proximity sensor

Infrared (IR) proximity sensors have been widely used to detect the presence of the finger around the device. These sensors commonly include IR emitters and receivers, and detect the presence of objects by monitoring the IR signal reflected back to the receivers caused by finger/obstacle interruptions. SideSight [20] is one of the earlier projects that employed IR sensors along the edge of a mobile device to detect the presence of a finger around the device. With an array of IR sensors, the system was capable of detecting finger movement approximately 8cm around the device. Kratz and Rohs [120] used IR distance sensors to facilitate around-device sensing capabilities on mobile devices. Their prototype system allowed the mobile device to detect finger and hand activities in a 3D space input it. The system consisted of six IR distance sensors attached along the edge of the device, facing upwards, to detect coarse hand gestures 5cm to 7cm above the device. The authors also showed that the system can recognize seven above-device gestures (e.g., sweep right, sweep left, rotate hand, and move top-down) with an accuracy of 88.6%. IR sensors have also been used to detect the hand and finger movements around small devices such as smartwatches. Nakatsuma et al. [145] attached seven IR sensors at the side of a smartwatch to enable the back-of-palm as an extended input space for the device. In a recent study, Withana et al.

[203] showed that arrangements of such sensors on different devices (e.g., smartphones, smartglasses and smartwatches) could play an important role in saving space, processing power, and energy while achieving high in-air gesture detection accuracy.

Vision-Based Sensing

On mobile and wearable device platforms, vision-based systems have been used to recognize the devices' surrounding activities. These devices are now equipped with high-resolution cameras and advanced processing power that allows vision-based around-device interaction in real-time. Song [179] showed that a smartphone's RGB camera can be used to recognize in-air hand movement behind or in front of the camera without relying on any external sensor. They found that processed on-board camera images can be used to detect around-device hand gestures at a short distance with a high accuracy (average 93%) in real-time. Niikura et al. [147] also used a camera-based approach to track the 3D space above a mobile device for text entry. They showed that images captured with embedded camera can be used to detect a user's fingertip, which they later mapped to enter text on a mobile device. Hwang et al. [90] observed that when a user applies pressure on a surface with a finger, it creates tension and pressure on the blood vessel under the nail of the finger. They used the on-device camera to detect this change to track finger posture (such as an extended or bent finger). In a recent project, Grubert et al. [59] showed that a mobile device's front facing camera can be used to capture the reflected image in the user's sunglasses. This image could further be processed to track the device itself as well as the user's hand around the device. They noted that wearing reflective glasses is very common in some scenarios (e.g., skiing, or biking) and users could leverage reflection on the glasses to interact with mobile devices from around-device space. On wearable devices, an on-board camera has been used to detect around-device activities. Watchme [192] used an embedded camera on a smartwatch to sense users' finger activity around the device. With a cloud-based optical character recognition engine, the system allows user input on a canvas with mid-air finger strokes.

Researchers have also demonstrated a method to enhance a camera's limited field-of-view by adding an external attachment, such as omnidirectional mirrors. Yang et al. proposed [208] Surround-See, a self-contained smartphone equipped with an omnidirectional camera that enables the peripheral vision of its surroundings. The prototype device captures a real-time image of the 360° surrounding view of the device. The image is further processed to detect users' around-device activities. The authors demonstrated that Surround-See is capable of identifying certain activities in the vicinity of the device, such as when a user walks away from it or when a user remotely waves at it to alter its state. Portico [10] used two cameras on foldable arms positioned above a tablet. The cameras were configured to look down at the tablet screen and its surrounding space. With such a configuration, Portico extends the device's sensing capabilities by monitoring a large flied of view. Bonfire [104] also used cameras attached to a mobile device to capture users' activities around the device. While Portico enables extended input capabilities to the device, Bonfire supports extended input as well as output capabilities around the device with two attached laser micro-projectors.

Depth cameras have also been used in a number of projects to support a wider range of around-device sensing capabilities. For instance, Tango [186] used a short-range depth camera on a tablet for gesture and periphery detection. The system allows users to detect the position and orientation of the device in a 3D environment. It also provides details about the surrounding environment, such as distance and size of any scanned objects in its periphery. Kratz et al. [119] used a short-range depth camera to estimate finger pose when it comes in contact with the device's touchscreen. They proposed an algorithm to extract finger pose information, such as finger rotation and finger tilt angles relative to the device's touch screen. Results from two user studies confirmed that their algorithm can reliably estimate finger poses on mobile devices. In another project, Kratz et al. [121] mounted a depth camera on a mobile device to track hand gestures in front of the camera. The authors showed that in-air gestures could be mapped to manipulate 3D objects on the device, a more intuitive mapping of 3D manipulations rather than using a small 2D touch-screen on a mobile device. In similar work, Chen at al. [28] used a depth camera mounted on top of the device to track finger activities above the screen. Finally, Ens et al. [39] used a SoftKinetic depth camera [177] to track finger activities in front of a head-worn display (HWD). The authors showed that camera input can reliably detect fine and course grained gestures (e.g., grasping, pointing and flicking gestures). With a set of applications, they demonstrated that such gestures could be used to interact with virtual content displayed on the HWD.

Magnetic-Field Based Sensing

Magnetically driven input has been shown as a promising alternative to detect around-device activities. This around-device sensing style uses a magnet attached to the user's finger. Based on the magnetic field shift recorded by the magnetometer integrated on the device, it detects the presence and angular location of the finger holding the magnet. This input style has been widely used for around-device interaction as such input does not require external power sources.

Harrison et al. [70] leveraged magnetically-driven input to detect finger activities around a watch. They proposed a technique called Abracadabra which uses a small magnet attached to the user's fingertip in combination with an on-device multi-axis magnetometer. When the finger moves around the device, the magnetometer detects the magnetic field changes, which are further processed to track the finger input around the watch. With a user study, the authors demonstrated that magnetic field changes can be used for item selection on the smartwatch with a high accuracy (above 92%) even when the targets are small. Similarly, a number of projects investigate the use of magnetic sensors in different contexts to interact with mobile devices including detecting gestures around a mobile device for text entry proposes [111,114], sending commands to mobile devices such as zooming a map or selecting an item [112], controlling virtual musical instruments to compose music on mobile devices [109], authentication to access mobile devices with in-air 3D gestures [110], creating tangible interactions around a mobile device [88], or controlling call alerts [113]. In recent work, Hwang et al. [89] showed that a small piece of a magnet can be attached to a pen to create new capabilities for pen-based interaction, such as tracking the pen's orientation and pen spin around mobile devices.

On-Board Sensor Based Sensing

Alongside augmentation of the user or environment, researchers have presented a number of projects showing in-air interaction capabilities with the device's on-board sensors, but in a limited context. Wen et al. [199] proposed Serendipity that senses fine-grained in-air finger gestures with the hand wearing a smartwatch. Serendipity processed onboard sensor data (i.e., accelerometer, and gyroscope) to detect a number of hand and finger gestures, such as pinching, tapping, rubbing, squeezing and waving. The authors showed that onboard sensors can be used to detect the above mentioned gestures with an average accuracy of 87%. Similarly, Chen et al. [27] developed techniques that leverage a phone's front camera, accelerometer, and inertia measurement units (IMU) data, to detect the position of the phone relative to a user's body. Their system utilizes the front-facing camera to extract the distance between the device and user's body, on-board compass to detect horizontal orientation and accelerometer to find vertical orientation. They evaluated the performance of these sensors by attempting to locate 27 around-body positions; the results revealed that the sensors achieved 100% accuracy in locating these positions.

Non-Speech Voice and Acoustic-Based Sensing

Non-speech voice and acoustic sensing capabilities have advanced to the point of detecting sound source locations to create novel around-device interactions. Researchers have demonstrated interaction alternatives with a wide range of non-voice input, such as humming [182], whooshing [161] and blowing [26,45,154] at close proximity to the device. Such non-speech voice input is commonly acquired by an on-device microphone or with an array of a microphones attached to the device. Researchers have also explored the time distance of arrival (TDOA) approach as a promising acoustic localization technique that relies on input from a set of acoustic sensors, i.e., microphones, attached to the device. In TDOA, the time difference at which the acoustic signals get processed allow the system to localize the sound source.

In a recent project, Xiao et al. [204] applied TDOA principles to detect hand tap events around mobile devices. Their system, Toffee, used four piezo sensors at the four corners of a mobile device or a laptop to detect sound source around the device. With a set of microphones, the system can detect sound source location within $1m^2$ within the centre of the device. Harrison and Hudson [69] proposed Scratch Input that leverages acoustic based input for around-device interaction. They constructed a prototype with a modified stethoscope and showed that their system is capable of detecting sound produced by finger scratches with a high accuracy (~90%).

Despite their potential to be used as a promising input modality for mobile and wearable devices, acoustic based sensing has several limitations. For example, the acoustic based input is highly prone to interference from ambient noise. Additionally, such input may not be appropriate in social settings such as when a person is in a conversation with others or in a meeting.

Electric Field Sensing Approaches

GestIC [50] is one of the earlier commercial technologies that deploys electric field sensing to detect near field around-device gestures. Several commercial products provide tracking at a range up to 15cm around the attached device.

Alongside commercial technologies, researchers have also explored electric field sensing technology to enable around-device interaction on mobile and wearable devices. Goc et al. [126] used an electric field to achieve hand and finger localization around a mobile device. Based on GestIC's chip, they built a low-cost, thin, transparent prototype which is used on top of the touch screen to allow the device to sense hand and finger activities surrounding the device. The authors demonstrated that the prototype can track 3D motion gestures and provide precise hand and finger localization around the device. In recent work, Zhou et al. [217] showed that electric field technology can also be embedded in small devices such as smartwatches to enable around-device interaction.

Other Sensing Approaches

In recent years, researchers have investigated unconventional sensing mechanisms, such as using GSM signals or radio waves to track hand and finger movements around the device. Zhao et al. [214] used reflected GSM pulses from hand movements around a mobile device to recognize AD-gestures. They showed that, with a robust algorithm, the system is capable of detecting and recognizing gestures with a high accuracy. Nandakumar et al. [146] proposed fingerIO, an active sonar system to track a user's hand movement to provide fine grained input for around-device interaction. They showed that fingerIO supports 2D finger tracking around the device with an average accuracy of 8mm, even when the device is occluded by objects or is in the user's pocket. Recently, Google lunched Soli [156], a small portable device that use radio waves generated by hand activities received by a small radar. Researchers have shown that in spite of its small form factor, Soli can be reliably used to detect fine-grained finger gestures with high precision [198].

The Samsung Galaxy S4 [166] is the first commercial mobile device that tracks in-air hand movement above the device. It contains two capacitive sensors to support touch and hover state on the device. The hover feature allows users to perform hand gestures above the device to control basic on-device functionalities, such as answering or rejecting a call without touching the screen. However, to activate the sensing, a close proximity to the deivce (approximately 6mm) is required and only limited interaction capabilities with a small number of in-air gestures are supported.

All of these efforts to develop around-device sensing capabilities inspired us to anticipate that future smartphones will come equipped with advanced 3D around-device sensing technologies. Therefore, in this thesis, we assume that self-contained reliable tracking in 3D around the device will become possible in the near future. As a result, we emulate the around-device sensing environment with a Vicon motion tracking system [193] (MX system with eight TSeries cameras) to track participants' hand and finger movements around mobile devices. We choose this sensing solution as it provides precise movement data in real-time around the device.

Around-Device Feedback

Feedback about around-device items is very crucial as it provides guidance to point at the off-screen items. Researchers have proposed a wide range of visual and non-visual feedback mechanisms that can present around-device items either on the device or in mid-air, as shown in Table 2.

	Visual	**Non-Visual**
On-Device	Visual Cue	Contact Tactile Audio
In-Air	Projection Light Emitting Diode (LED) Fog Display In-Air Bubble	Non-Contact Tactile

Table 2 Feedback on around-device items can be presented with on-device or in-air with visual and non-visual approaches.

On-Device Visual Feedback

On-device visual feedback commonly uses screen space to represent off-screen items. Researchers have explored a number of different solutions for presenting off-screen items on the device using the following visual feedback mechanisms:

Overview: is a widely used approach [15,62,71–73,86,94,96] to present around-device feedback. With Overview, a large workspace is shown with a miniature view on the device's screen where the remaining on-screen space is used to provide a detailed view of the workspace. When a user manipulates the detailed view with flick or pinch gestures (to pan or zoom the workspace), corresponding changes are reflected on the overview. Researchers have demonstrated that using Overview to provide around-device feedback to localize relative item and finger positions is essential for efficient navigation in a large workspace [38,40,71,72]. However, overview consumes a portion of screen space and occludes the detailed view presented underneath it. To overcome such limitations, researchers have proposed the use of contextual cues [14,62,213] where abstract shapes are used near the edge of the screen to minimize such issues.

City Lights [213]: is one of the earlier works that used the contextual cue concept. With City Lights, the around-device information is displayed using thick lines. The lines are projected at the edge of the screen based on the orthogonal projection of the around-device items. Additionally, line properties, such as thickness or colors, are used to represent the distance of the off-screen items from the screen. One of the major drawbacks of this technique is the limited feedback regarding the exact position of the around-device items.

EdgeRadar [64]: is another cue based technique that extends the idea of overview and City Light. In this technique, the overview is distributed into four regions and positioned along the four edges of the screen. The around-device objects are commonly represented by small proxies or dots on the four docked sub-overviews. These proxies convey relative distance and direction of the off-screen items. Hossain et al. [86] extended this idea to provide selection capabilities on these off-screen items. They proposed a technique, EdgeSplit, which splits the docked sub-overviews into multiple small regions, such that each region contains one off-screen item. The authors showed that accessing off-screen items with EdgeSplit is faster than the traditional touchscreen input.

Halo [14]: is another well-known contextual cue-based representation, in which a circle is drawn, centered at an around-device item. A small portion of this circle (i.e., arc) is displayed on the edge of the screen. The arc shrinks or grows as the user moves the device from close to or far away from the item. Though this representation provides continuous guidance of off-screen items by changing its size, it can lead to clutter and overlap when dealing with a large number of off-screen items.

Wedge [62]: To overcome this limitation, Gustafson et al. proposed a cue-based representation called Wedge. Instead of an arc, a wedge shape is used on the edge of the screen to present an off-screen item. This around-device feedback mechanism has been demonstrated to be more accurate for finding off-screen items than Halo as it reduces overlap and minimizes clutter related issues. Burigat [19] compared the performance of Overview and Wedge with a large number of around-device items for navigation tasks. They found that users were faster and more accurate in finding around-device content with Overview than Wedge. In another study, Gonçalves et al [52] showed the overview technique provides better information about items that around the device and assists users in navigating such items.

In-Air Visual Feedback

In-air visual feedback commonly takes places around the device where the items are actually located. With the advent of new technologies, researchers have demonstrated that feedback regarding the off-screen items can be provided in the around-device space. Several previous projects used commercial solutions that support an immersive virtual reality or augmented reality environment to display around items around the device. For instance, Visbox cave system [194], which is a projection-based VR display, is commonly used to show around-device items and contents [42,43]. Additionally, commercial head-word displays (e.g., Epson Moverio) were also shown to be useful to provide around-device feedback [56–58,172]. However, due to the high cost, these solutions are not widely accessible to a general audience. Additionally, miniaturizing to make them portable for a mobile device is not yet possible.

Small projection devices have been shown to be a viable alternative to provide around-device feedback such as pico- or laser projectors [104,125] or arrays of LEDs [49,143,157]. Laput et al. [125] developed a smartwatch prototype that used a laser projector to provide visual feedback regarding items on users' skin. With a user study, they found that their system provides reliable output (with an item-recognition accuracy of 98%) around the device. Qin et al. [157] showed that dynamic ambient lighting can be used to provide feedback on items that are placed around the device. They developed a prototype system with 40 RGB LEDs attached to the edge of a mobile device and used a light aura to convey information. Muller et al. [143] investigated the design space of ambient light to display off-screen information. They built a prototype system with 50 LEDs to provide feedback for off-screen items. With two user studies, they explored different design issues such as how best to represent distance and direction of off-screen items with light. Their results revealed that using light feedback reduced user workload and increased the system's usability. Freeman et al. [48] also explored the design space for such interactive output modalities. They built a prototype system where they attached an array of 60 LEDs to the edge of a device. They showed that these LEDs have two major properties, brightness and hue, which can be varied to effectively display around-device information.

Other approaches to provide feedback directly around the device include using fog [41,135,158], mid-air bubbles [144,170], levitated lightweight object and hologram [81] to support projection. Plasencia et al. [135] proposed MisTable where a see-through and interactive surface is created with fog. They used a fog distribution system to ensure a continuous flow of fog, projectors to project content on the fog, and a Kinect to track users' movement to provide interaction capabilities with the content. The authors showed that the fog surface can be used as a personal space between the user and the tabletop to display 2D and 3D content.

Seah et al. [170] showed that fog-filled bubbles can also be used to display information in mid-air. They built an in-air display system, SensaBubble that creates different sized fog-filled bubbles with visual information displayed on them to convey information to users. The system tracks mid-air bubbles with a Microsoft Kinect [117] and projects information on the bubbles with an external projector. With a user study, the authors examined how well users could extract information (such as color, text, and digits) from in-air display. Their results revealed that colors can be detected with high accuracy; however, text and digit recognition decreases with increasing characters or digits. In a recent project, Sahoo et al. [165] showed that mid-air feedback could be provided with small and lightweight objects that can be levitated with an acoustic levitation technique. They used an electrostatic field-based approach to position and rotate the lightweight objects to project light along a direction.

Most of the feedback mechanisms discussed in this section are heavily dependant on external and bulky attachments. Therefore, such solutions are yet not feasible for integration with small portable devices.

Non-Visual On-Device Feedback

Non-visual feedback mechanisms do not occupy screen space to provide feedback on around-device items, thus saving valuable screen space on the small device for other interaction.

While exploring the previous research on non-visual on-device feedback for around-device interaction, we found that researchers used audio and non-contact tactile feedback for providing around-device feedback. One of the major advantages of audio feedback is that it doesn't require users' visual attention and it is also useful when there is no display available to provide visual feedback. Zhao et al. [216] developed earPod, which used audio to provide feedback on items to allow eye-free menu selection. They divided the touchpad dial of an Apple iPad into multiple sectors, similar to a Pie menu [74,215] or marking menu [122,123], and placed the virtual menu items (e.g., color, job, and instrument) on the sectors. When a user moved their finger on the touchpad, feedback on the currently selected item was provided with audio to an external headset. Ashbrook et al. [8] built a wearable smart-ring prototype that used audio to provide feedback on around-device items. The items were placed around the device and users selected items one-by-one by rotating the ring. A similar strategy was used by Imaginary Phone [66] where users selected virtual items that were placed on a user's palm. Kajastila and Lokki [102] also used audio to provide feedback on items that are placed in an in-air circular menu. When users accessed the menu items by moving their hand on the circular menu, the feedback on the currently selected item was provided with visual cues and audio. With a user study, they showed that participants can effectively access menu items with audio feedback and visual cues. However, some of the users were in favor of audio feedback as it does not require any visual attention. Though the audio provides accurate and reliable feedback, such forms of feedback may not be socially acceptable as it might intrude on other people around the device.

Niikura et al. [147] developed a typing interface with in-air space around the device where tactile feedback was provided on the device. They attached a vibration motor on the back of the device to provide vibration feedback. When the user types a letter, it provides feedback with a short vibration. One of the limitations of such feedback mechanisms is limited distance and directional feedback on the around-device items.

Non-Visual In-Air Feedback

Prior work on around-device interaction has shown a number of different approaches to provide non-visual in-air feedback on the around-device items. A number of projects have explored the benefits of providing tactile feedback for around-device interaction. Iwamoto et al. [95] fabricated a prototype with an array of 91 ultrasound transducers that are placed in a hexagonal arrangement to provide tactile feedback in 3D space above a device. Later, they extended this idea to provide feedback with a mid-air tactile display [85] where the system can float an image allowing users to freely interact with it. The authors also developed an interactive system where the hand position was tracked using an IR-based hand positioning system. Items (such as bounding balls, raindrops) were positioned and tactile feedback was provided based on the hand positon. The authors developed a number of prototypes varying the number of aerial sensors and the hand tracking mechanism to provide fine-grained feedback in mid-air [79,80,82–84]. In these projects, the feedback was commonly provided in a single localized point. Carter et al. [23] built a system, UltraHaptics, that provides in-air tactile feedback on multiple points using an array of ultrasonic transducers. With two user studies, the authors showed that users could distinguish tactile feedback even with a small variation in the feedback. Wilson et al. [202] explored users' perceptions of using tactile feedback for providing in-air feedback. They showed that users can successfully localize a point of feedback (within 8.5mm), though localization accuracy is decreased along the longitudinal axis, (i.e., the long axis of the body) when compared to the transverse axis (across the body). In further projects, the authors investigated hardware solutions to support feedback on mobile platforms [4,93] and designed guidelines for developing systems with tactile feedback [149].

Air pressure has been shown as an alternate solution to provide haptic feedback. Sodhi et al. [176] developed a system, AIREAL, that used air vortex to provide tactile feedback in free-air. The system relies on a depth camera to track a user's hand and uses a vortex generator to generate vortices and provide feedback. Later, Gupta et al. [61] explored how the different design properties such as vortex velocity, feedback delay, and aperture size could potentially influence the performance of such systems. Though several researchers showed that haptic feedback through air pressure provides rich user experiences [175,191], the resolution of such approaches is low compared to other solutions (e.g., ultrasonic transducers based solutions).

	Feedback Type		Visibility			Display mode		Tangibility	
	Discrete	Cont.	High	Mid	Low	Direct	Indirect	Tangible	Intangible
Visual Cue		X	X				X		X
Contact tactile	X				X		X	X	
Projection		X	X			X			X
Fog display		X		X		X		X	
In-air Bubble		X		X		X		X	
LED (on edge)		X			X		X		X
Audio	X				X		X		X
Non-contact tactile	X				X	X		X	

Table 3: A comparison of different feedback mechanism for around-device interaction

Table 3 shows a comparison of the feedback options discussed in this section. We categorize them based on properties such as (i) feedback type: that specifies whether the information is presented continuously or discretely, (ii) visibility: which indicates whether the presentation technique has capabilities to provide a clear visual feedback or not, (iii) display mode: which indicates whether information is presented direct in off-screen space where the item is actually located or on the screen displayed as an off-screen item; and (iv) tangibility: which indicates whether the presented information is tangible or not.

In the table, we can see that there are a few choices that can be used to provide continuous feedback on the around-device items and hand movement (e.g., visual cues, projection, for display, in-air bubble and LED). As effective cues are crucial to locate off-screen items, we primarily considered two options that provide feedback with higher visibility of around-device items: visual cues and projection-based approaches. Since projection based systems are difficult to miniaturize, embedding them on a mobile device is not possible. Therefore, we decided to go with visual cue approaches.

In this section, we discussed different visual cue options to provide feedback regarding around-device items. Researchers showed that overviews of the entire workspace have better performance for providing feedback on around-device activities. Therefore, in this thesis, we used overview as the feedback mechanism to present around-device items, hand, and fingers.

AD-Interactions and Social Acceptance

In contrast to the significantly large number of works on AD-interactions, AD-Sensing and AD-feedback mechanisms [20,70,111,114,120], the social acceptance, or users' comfort level, of these interactions has received very little attention. Projects have mostly studied AD-interactions in a lab setting. To our knowledge, only Jones et al. [99] and Kratz et al. [121] have considered how users would feel about using AD-interactions in a public setting. Jones et al. [99], who compared various AD-methods for panning and zooming, also asked their study participants how comfortable they would feel using the AD-methods in public places. They report that although the enlarged interaction space provided by the AD-methods was valued, participants said they were less likely to use these in public settings. Similarly, Kratz et al. [121], who studied AD-techniques for rotating on-screen objects, found that participants were split regarding using AD-technique in public. These findings, together with recent studies [142,162,163,200,201] regarding users' concerns and feelings about performing body or device-based gestures in public settings – e.g., tapping the shoulder or shaking the phone to mute a call – warrant caution and further investigations. Probing user perceptions of novel AD-interactions could allow designers to rule out unwanted styles of input. What is needed is an understanding of how AD-interactions are perceived in ecologically valid settings.

Social Acceptability

While AD-gestures have not been studied from the perspective of social acceptability, other gestural interactions have been explored. Ronkainen et al. [163] introduced the idea of studying social acceptance of gesture input on mobile devices. In an online survey they used short video clips of people performing different device-based gestures (e.g., swinging and slapping the device) for various tasks and in different settings (café, library, while walking). Participants were asked to comment on whether they would use the featured gestures themselves. As participants were not explicitly instructed to consider the social setting in their responses, Ronkainen et al. were surprised to find that roughly half of the participants mentioned context-related and social issues in their rationales when rejecting a gesture.

Rico and Brewster [162] expanded on this finding and studied how social setting influences the acceptance of device-based and body-based gestures (e.g., tapping the nose, squeezing the forearm). Again, participants watched short video clips showing a person using the examined gestures while being alone in a room. From this, participants were asked to state where and in front of whom they would use the gestures themselves. Rico and Brewster found that participants were very selective regarding usage context: both locations (e.g., at home, on a sidewalk, in a café) and audience (e.g., colleagues, family, and strangers) affected participants' willingness to use gestures. In a follow-up, user reactions were elicited from eleven persons having performed some of the gestures, both in a private room and on a busy sidewalk. Most participants commented on how different they felt when performing the gestures in public, i.e., feeling of discomfort and feeling uncomfortable and worry about what others might think. Several participants also reported feeling somewhat less uncomfortable in later stages of the exercise. These findings demonstrate the potential limits of letting participants imagine usage without having had a firsthand experience with strangers watching of their own behaviour.

Williamson et al. [200,201] collected user experiences and insights from real-world usage situations. They asked participants to use a mobile phone application operated though body gestures (nods, rotating and shaking the wrist). The majority of registered usage situations took place in private while at home or in semi-private settings at work or when walking. Only a few participants decided to use the application during public transport. Many worried about what others would think and worked out strategies to appropriate or disguise the gestures in order to avoid attention from potential spectators. Williamson et al. also report that several participants noted a large, and unanticipated, difference between how they felt about 'transitory' and 'sustained' spectatorship (e.g., a person catching a glimpse of the interaction while walking by versus a fellow passenger on a train that watches the interaction for a longer time).

Montero et al. [142] used video clip demonstrations of device-based gestures. They asked participants to indicate on a scale from 'embarrassed' to 'comfortable' how they would feel performing the example gestures in a public place. Their results demonstrate that factors such as gesture category, 'suspenseful' (e.g., writing a letter with a large in-air gesture) or 'magical' (e.g., controlling a light from a mobile phone), have an impact on the acceptability of device-based gestures. Interestingly, their results also show that early and late technology adopters perceive suspenseful gestures, those with a clear action that is easily seen by bystanders but without a noticeable outcome, as being less acceptable. These results could suggest that AD-gestures may not be accepted for public use.

In summary, previous work clearly demonstrates that people are very concerned about what others think and how they will react when observing them perform unusual interactions. Previous studies also show that people are selective regarding which gestures they would feel comfortable with in various social settings. However, AD-gestures – as well as many body-based and device-based gestures – possess several attributes, such as the area of input and duration that can affect comfort and acceptance. No prior study has explored whether and how user attitudes and acceptance vary depending on such features, a task necessary to refine and propose novel AD-interactions.

Thumb Input on Mobile Devices

The thumb is commonly the only available finger to interact with mobile devices in one-handed mobile interactions [78]. Researchers have mostly studied the use of thumb for one-handed mobile usage scenarios to facilitate touch interaction. They point out several limitations of one-handed thumb input on mobile devices which are discussed in this section:

Researchers have demonstrated that the length of the thumb sets a significant limit on what screen areas can be reached [16,105,106,212]. This is often referred to as "reachability". A number of research studies have been conducted to explore alternate interface designs to tackle reachability issues. For instance, Karlson and Bederson [105] proposed ThumbSpace, which used a miniature version of the screen (i.e., a proxy of the screen) and was placed in thumb-reachable screen space. Similarly, Hürst and Merkle [87] presented an interface where users can swipe their thumb on an arc-shape thumb-reachable on-screen band to access items that are not reachable via thumb. To facilitate thumb input for one-handed interaction, researchers have also proposed shifting the cursor by a fixed or variable distances [116,151,155], using a virtual on-screen thumb [124], using fisheye and zoomable interface techniques [107], moving interface components based on device tilt direction [25], using back-of-device input to access information [75,209], or using the back and front surface simultaneously [210].

While using the phone with one hand, information on the screen easily gets occluded by the interacting thumb. This scenario is commonly refer to "Occlusion". Consequently, it triggers inaccuracy while selecting an item that is occluded by the thumb [106]. To address this problem, researchers have explored a number of solutions, such as zooming and shifting the occluded items to a visible on-screen space [9,195], triggering an occluded item with a predefined thumb-gesture [206,211], or using back-of-device input [13,171].

Researchers showed that the size of items displayed on the screen is crucial while accessing items with the thumb for one-handed interaction. Parhi et al. [153] studied the optimal target size for discrete item selections such as selecting a radio button or checkbox and repetitive target selection actions such as typing a word with a soft keyboard on a mobile device. Their results revealed that target size should be higher than 9.2mm for discrete item selection and 9.6mm for repetitive target selection. Researchers have also proposed the use of zoom-based techniques in which the canvas and items underneath the finger zoom to provide a larger selection space [164], showing a cursor with an offset from the finger [155,195], and displaying the occluded content with a callout and a cursor above the finger [195].

While a significant amount of work has explored how to use the thumb for one-handed mobile interaction, there has been much less attention devoted to using the thumb to access thumb-reachable in-air space. Schmieder et al. [169] showed that a mobile phone's front-facing camera can be used to recognize 3D thumb movement above the camera. The authors also showed that thumb movement can be used as a continuous input on the device. They conducted one user study with three gesture types: thumb tilt gestures: tilting the thumb either to the right or left; thumb distance gestures: distance between the thumb and the camera; and circular gestures: circular motion in front of the camera) for performance and user satisfaction. Results revealed that the system can recognize each gesture with a high degree of accuracy. Additionally, participants found circular gestures are more difficult to perform than others.

In my thesis, we are the first to explore the use of thumb-reachable around-device space for interacting with the mobile device. We consider this space as complementary to the large around-device space that can be accessible via the other hand. As this space is closer to the touch screen, it could also be used in conjunction with touch input to design new interaction capabilities.

Bimanual Interaction

Researchers have studied users' ability to control two hands in a number of contexts, such as desktop [12,97,131,136], handheld devices [37,138], pen-based interfaces [74] and wearable devices [118]. Buxton and Myers [21] did an early investigation on bimanual interaction and demonstrated that using two hands significantly improves the performance of a navigation and selection task over uni-manual interaction. Inspired by these results, researchers have shown the advantages of using two-handed input over a single-handed input for symmetric and asymmetric bimanual tasks. Symmetric assignments, in which both hands have a similar role and are used in parallel, has been shown to improve performance [12,24,128]. However, Balakrishnan and Hinckley [12] revealed that due to the lack of visual integration, symmetric role assignments could become asymmetric. Asymmetric bimanual interaction was further defined by Guiard's [60] Kinematic Chain (KC) model where the non-dominant hand provides the frame of reference and the dominant hand regulates the finer level of detail. Both styles of interaction have shown useful results on different interactive paradigms including tablets [11] and free-hand mid-air input [35].

Ideally, users can access two around-device spaces while holding the phone with their non-dominant hand: (i) the large in-air space that can be reachable with their dominant hand and (ii) the small in-air space just above the mobile device accessible via the thumb of the hand holding the device. In this thesis, we leverage both in-air spaces for designing an application for a mobile device. We followed the Guiard's KC model to design our tasks to support bimanual interaction. We used asymmetric bimanual interactions on mobile devices, where we asked participants to hold the mobile phone with the non-dominant hand and access in-air space with the dominant hand. Holding the phone with non-dominant hand creates a reference point and the dominant hand uses this reference point to navigate around the in-air space. Additionally, the dominant hand performs fine-grained operations such as exploring items and panes, whereas the non-dominant hand (using thumb) issues course-grained operations.

Information Browsing on Mobile Devices

Though mobile devices provide access to information on-the-go, accessing information on small devices presents several limitations. For instance, mobile devices are commonly equipped with limited display space which can only show a limited amount of information on the screen. Researchers have demonstrated that the limited display space makes information browsing (e.g., accessing weather forecast, browsing news) slower on mobile devices than on a desktop computer [100,174]. They also reported that users completed fewer tasks on a mobile device than a desktop computer in a given period of time.

To enhance the experience of browsing information on small platforms (i.e., mobile), prior research has recommended design guidelines. Examples of such guidelines include transforming a large web page into multiple small blocks and displaying individual information blocks one at a time on the small screen [205], optimizing item placement strategies for navigation elements to provide better user browsing experience [91], and using pagination to reduce the amount of scrolling [51]. Additionally, to better present the information on mobile devices, researchers have suggested using shorter text [32] or key phrases to summarize longer text passages [101].

Prior work showed that user information browsing patterns (e.g., keyword used, search query length) are influenced by several factors, such as the user's device platform, usage time, and user location. Song et al. [180] showed that user browsing patterns, such as query volume, query category, query length and search time, vary significantly from one platform to another (i.e., desktop vs. mobile). Kamvar et al. [103] conducted a similar study to explore user information browsing patterns across different devices and showed that the pattern also differs across similar platforms such as between the iPhone and an android mobile phone. Sohn et al. [178] revealed that users' location, context, and time of mobile phone usage have an influence on user information seeking patterns. With a two-week diary study, they found that users frequently access information on a mobile phone when they are prompted by a conversation with other people, when they want to find information based on location-based artifacts such as a billboard, and when they are looking for route directions. With a similar study, Church and Smyth [31] found that user location, activity, time to access information, and social interactions play a role while accessing information on mobile devices. Other work [29,76,77] confirmed these factors as influential dynamics concerning user information access patterns on mobile devices.

In 2008, Church and Oliver [30] investigated mobile web usage and mobile search patterns. Their results revealed that people commonly use the mobile web when they are at home or at the office. Results also indicated that people access information via mobile phone to share it with another person, such as a friend, colleague, or family member. Nylander et al. [148] conducted a similar study where they found that people commonly use the internet on a mobile device when they are at home (31% of total usage) and do not have a computer near them (49% of total usage). Additionally, they found that people accessed information on mobile devices to perform a wide range of activities such as reading news and social updates, checking email, or making transactions. In a recent study, Carrascal and Church [22] showed that users intensively use browsing, communication, social networking, and e-commerce applications on their mobile devices. The authors revealed that people spend a significant amount of time on analytic tasks, such as searching for products on shopping and retail websites. In addition, a recent survey jointly conducted by Google [53] and Nielsen [134] showed that mobile users spent, on average, more than 15 hours per week researching products and deals on their mobile phone [54]. Additionally, they found that 93% of the searches on mobile devices commonly lead to a final purchase decision. These results show promise for a new era of mobile commerce (i.e., m-commerce), in which mobile applications would be the primary media to reach customers and facilitate transactions.

Inspired by this finding, we demonstrate the potential of around-device interactions for m-commerce applications. In addition, we investigate in-air space to support analytic decision-making tasks, such as purchasing products, that require users to browse and interact with a large amount of information content before arriving at a decision. This decision-making task commonly consists of many subtasks such as searching product categories, applying filters to narrow searches, inspecting alternative products, adding products to a shortlist, and browsing and comparing shortlisted items before making the final purchase decision. These sub-tasks require users to perform many fundamental mobile operations such as selection, navigation and switching between multiple applications windows.

Chapter 3: Design and Evaluation of the Around-Device Space for Accessing Mobile Device Content

In this chapter, we explore the possibilities of using the around-device space for accessing items that are placed around the device. To achieve this goal, we design a novel mobile user interface called around-device binning or AD-Binning that allows users to store and retrieve virtual contents around the device. To design AD-Binning, we start by finding AD-Binning's input range around a mobile device. Our results reveal that users can comfortably reach up to 40cm from the device when they are holding it close to their body. We next explore a set of design factors that are related to placing and retrieving content in around-device space. Through a series of user studies, each building on the results of the previous, we explore a number of design parameters including suitable around-device target (or bin) size, methods for selecting items in around-device space, and mid-air space discretization methods. We also investigate different in-air virtual content placement (i.e., binning) and retrieval strategies. Our investigations provide us insight on how the various design parameters influence the performance of the AD-Binning technique.

With this knowledge, we conduct a further study showing the potential of using the around-device space for a practical AD-Binning usage scenario for browsing information content. In this study, users participated in an analytic task that involves frequent zooming, panning, and re-visitation of items to make a final decision. We compare our AD-binning technique that leverages the around-device space with traditional touch input (i.e., pinch and flick gestures). Results reveal that AD-binning significantly reduces information browsing time over touch input by avoiding frequent selection and flicking mechanisms needed for touch interaction.

AD-Binning Design Framework

AD-Binning is inspired by earlier work demonstrating that around-device input is valuable for interacting with small form-factor devices [70] and for extending the input vocabulary of mobiles [114,120]. Device manufacturers are considering adopting around-device sensing methods in next generation mobile devices [10]. Unlike most prior work on around-device interaction [20,70,99,114,120,121], we focus on direct interaction with off-screen content. Such an interaction style presumes that the mobile device's interaction plane extends beyond the physical boundaries of the device [38], and users can directly point to retrieve items. Prior work on a class of interactions involving around-body input [33,129] suggests that users can leverage their spatial abilities to efficiently recall items through mid-air pointing. We expect similar benefits for AD-Binning. We examine additional prior work to frame AD-Binning's design factors.

Design Factors

Several key factors influence the design of AD-Binning. We explore these factors in relation to prior work.

Selection methods

AD-Binning allows users to explore content by letting the user move their finger in the space around the device. This facilitates rapid item browsing. However, a selection is required to retrieve an item and put it into focus. Researchers have designed similar mechanisms for triggering a selection when iterating through items using auxiliary input channels, such as pressure and tilt. These include dwelling on an item [74], quickly releasing a button [160], or lifting the finger [160]. AD-Binning facilitates item selection through two general methods: interaction on the device for triggering selection (touch or back-tap) or micro-gestures in mid-air around the device. We investigate the suitability of both of these methods.

Bin size

AD-Binning relies on direct off-screen pointing to place and retrieve items, a task influenced by Fitt's law [130]. AD-Binning provides the advantage that items around the device can take on large sizes to compensate for the small size commonly seen on mobile devices. However, little is known of how small targets can be without affecting performance. We investigate suitable bin sizes to facilitate accurate selection.

Visual feedback

To get rapid and accurate access to around-device items, effective on-screen cues are needed to point at (a) off-screen items [38] and (b) the user's moving finger. Overviews of the entire workspace have shown slightly better performance for direct off-screen pointing [38] than visual cues such as Wedge [62]. The differences between these visual techniques are affected by regions in which off-screen items are placed (performance with Wedge is non-uniform across the viewport). The design of AD-Binning uses an overview to show relative item positions and the user's finger in AD-Space.

Space discretization for bins

Closely tied to the input range and target size is the method for breaking up the around-device space into bins or space discretization. With auxiliary input streams (pressure and tilt) input discretization leads to better control [74,160,173]. Due to the bio-mechanical limits of the arm and difference in control at extreme arm ranges, we examine the effect of applying different discretization methods to the task of placing and retrieving around-device items.

Input range around the device

The bio-mechanical properties of the human arm dictate that on average users can extend their arm to about 60cm [133], limiting how many items can be placed off-screen. Little is known about this range when the arm moves around the device, i.e. the right arm on the left side will have a smaller range. We capture this range in a pilot prior to the studies.

Ideal binning locations

Prior work has shown that pointing at items placed in corners around the device is less effective and accurate than pointing at items on the sides [38]. These results were not obtained by evaluating the entire range for placing items and therefore more knowledge of ideal locations can assist in the design of AD-Binning.

Binning methods

Spatial memory and proprioceptive feedback can assist in retrieving information that is laid out spatially [33,167,197]. Ideally, information can be placed, or binned, using techniques that leverage this capability.

Mode switching

AD-Binning requires mode switching to differentiate around-device input from other accidental gestures in space. Mode switching could be explicit, wherein the user sets the device in bin-mode as needed. Alternatively, advanced sensing mechanisms could distinguish users' fingers in space separately from other items around the device.

AD-Binning Apparatus and Input Range

Apparatus

In this work, we assume that finger tracking in 3D around the device will become possible [10]. We emulate such a system using a Vicon MX system with eight cameras (T-Series) to track participants' hand movements (a). We placed markers on a smartphone (Nokia Lumina 800, size = 48.3×80.59mm, resolution = 480×800 pixels) and on a Velcro loop worn on the right-hand index finger (a). A Windows Presentation Foundation server application transferred tracking data every 10ms from the Vicon over Wi-Fi to the experimental software (Silverlight Windows Phone application) running on the smartphone. Advised by previous work [33] on how performance (task speed and accuracy) drastically suffers as the input space extends from 2D to 3D, the current implementation of AD-Binning only considers the space defined by the plane around the device: all interactions above or below the plane (in the z-direction) are projected on the interaction plane (Figure 5b). Future work will investigate the use of 3D space to layer items.

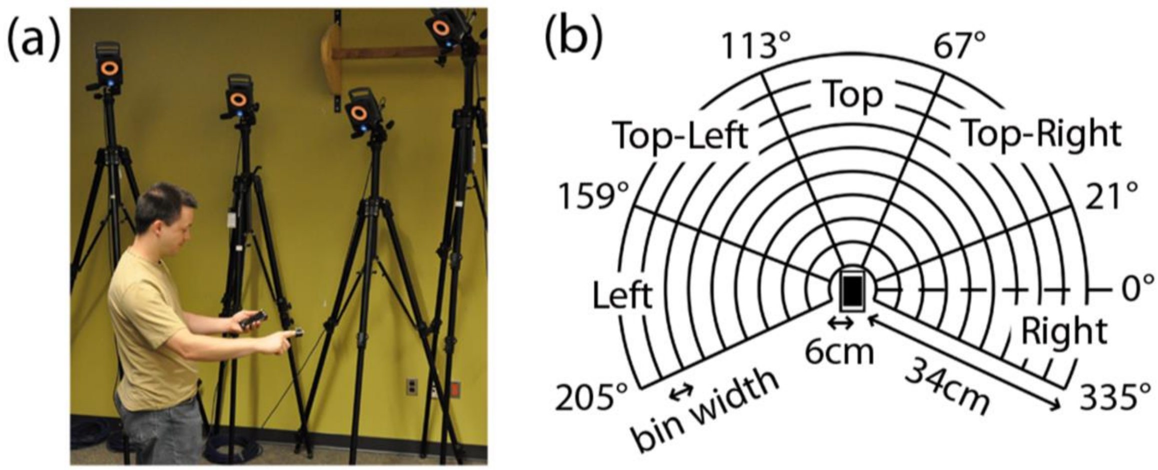

Figure 5: (a) Vicon apparatus. (b) The AD-Binning space.

AD-Binning Input Range

To determine AD-Binning's input range we asked two female and four male adults (these and all subsequent participants were right-handed and stood in the experiments) to hold a smartphone in their left hand and to 'draw' a half-circle around the device with their right hand, going from the left to right, and then back again five times. Participants were asked to perform without reaching their maximum distance. The collected movement data resulted in three design decisions: 1) 40cm was comfortably within reach for all participants and we used this value as the maximum input range (as shown in Figure 6); 2) points within ¾ of a full circle are within reach, and we use this to map bins into a circular layout; 3) we split the circular space into five sectors mapped to cardinal directions (North, North-East or labeled as Top, Top-Right) to facilitate spatial recall of items (Figure 5b). More than five sectors results items that are too small within the inner circle, leading to inefficient selection in these regions. Our exploration was based on these design choices and does not limit the use of other parameter values based on user preference and arm-length.

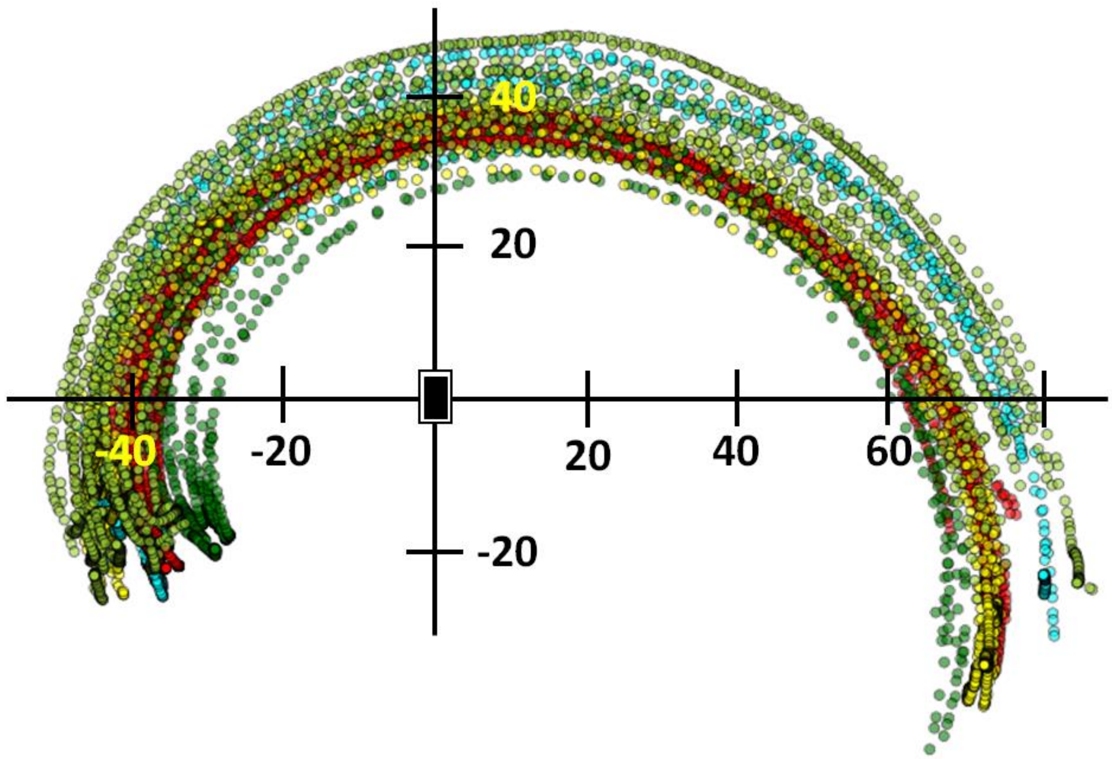

Figure 6: The AD-Binning Input Range, i.e. the range accessible by the right hand around the device (center) in centimeters

Study 1: Design Factor Exploration

In the first study, we identified suitable selection techniques, the minimum bin size for efficient item selection, and appropriate around-device space division methods. We split the study in sessions A and B to reduce study length and the complexity of the analysis, and to focus on a few design parameters at a time.

Selection Methods

Little is known about the specific methods for selecting items in around-device space. We grouped the selection techniques into methods that take place on or off the device. We settled on six candidate methods – two performed on the device with the non-dominant hand, and four using the dominant hand and its pointing finger in the air.

º *Tap* does not restrict the on-screen tapping area (as would be necessary when interactive elements are presented on the screen).

º *BackTap* is based on reading the device's accelerometer data and tapping the back of the device using the index finger of the device hand. After experimenting with various thresholds we found 0.15g (gravitational units) to be suitable for BackTap detection. The obvious advantage of BackTap over Tap is that it eliminates the risk of invoking interactive items on the screen during a selection. Conversely, BackTap cannot be used when the device is placed on a table.

º *Dwell* is often suggested in the literature as an alternative to click, e.g., in eye-gaze input. The dwell time was 600ms.

º *LiftOff* requires an active movement raising the pointing finger. A change in z-position >30mm between two consecutive time cycles triggers a LiftOff. An alternative to LiftOff is to push down, which we did not test as both behave similarly.

º *Pierce* assumes an imaginary horizontal interaction plane defined by the mobile device, which the finger needs to 'pierce' to make a selection.

º *DownUp* uses a down-up motion (30mm down, and up) inside a bin to trigger a selection. The two-stage motion, up and down, allows for a backoff possibility to cancel a started selection, similar to clicking an on-screen button with a mouse. This is the only method with a possibility to reverse in mid-course of the selection.

Session A – Selection Methods and Bin Size

Participants, Task and Experimental Design

Twelve daily computer and touch screen users (3 female) aged 20 to 39 years participated. With short breaks and practice trials, each session lasted approximately 45 minutes.

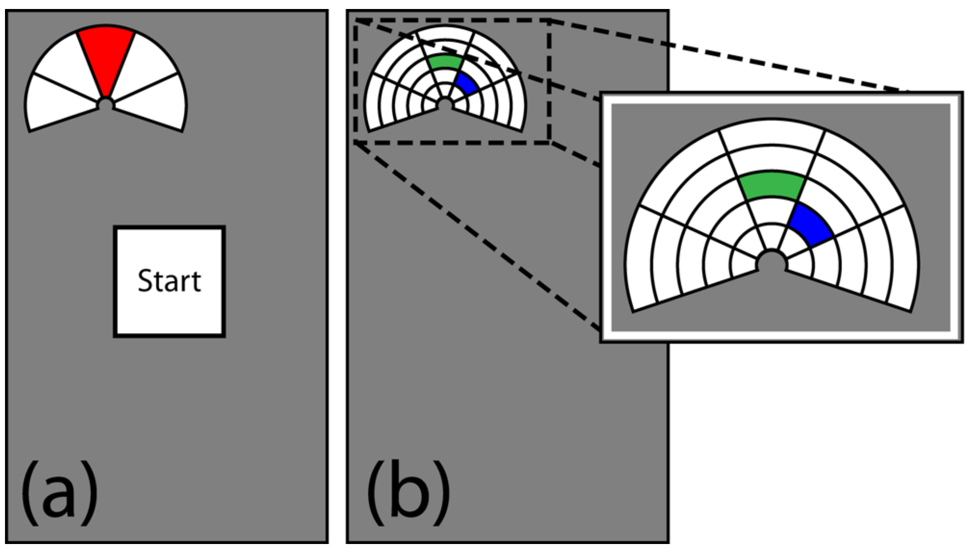

Figure 7: (a) Trial start screen. (b) Visual feedback overview.

A start button and a small overview (1.61cm wide) are displayed on the screen (Figure 7a). A red marker in the overview highlights one of the five sectors to indicate the direction to the next target bin (we chose to indicate direction before trial start to minimize visual search and unaimed arm movements at trial onset). The participant presses the start button with the right-hand index finger to begin the trial (participants were not allowed to time-optimize by first moving their tracked index finger to the estimated target location and then starting the trial with their left thumb). The overview shows the target in green, and a blue cursor in the overview follows the tracked finger (Figure 7b). A correctly performed selection action ends the trial and loads the start screen for the next trial.

Session A used a 6×4 within-subjects design for factors *selection method* (Tap, BackTap, Dwell, LiftOff, Pierce, DownUp) and bin *widths* (68, 38, 26, 20mm). The four bin widths were obtained by dividing each sector into 5, 9, 13 or 17 equally wide bins. Participants performed ten repetitions of each *selection method-width* combination, resulting in a total of 240 trials per participant. Participants completed 20 random practice trials before the test trials.

The presentation order of the six *selection methods* was balanced among participants using an incomplete Latin Square. The order of bin *widths* was randomized for each *selection method*. We kept the distance to the target bin constant (230mm) by only using the middle bin in each bin sector. Two trials from each *method-width* combination were located in each of the five bin sectors (Left, Top-left, Top, Top-right and Right –Figure 7b). The order of target location was randomized. No feedback was given when participants selected a non-target bin or if the intended selection action was not detected. Trials were only terminated after a correct selection occurred in the correct target bin. We asked participants to perform each trial as quickly and accurately as possible. After completing all selection methods participants rated them based on preference.

Results

We used a repeated measures ANOVA and post-hoc pairwise comparisons to analyze trial times. We used Friedman tests with Wilcoxon tests for post-hoc pairwise comparisons to analyze error rates (number of trials with incorrect selections divided by the total number of completed trials). Post-hoc pairwise comparisons were Bonferroni adjusted (α-level = 0.05). The same tests were used in all experiments unless otherwise noted.

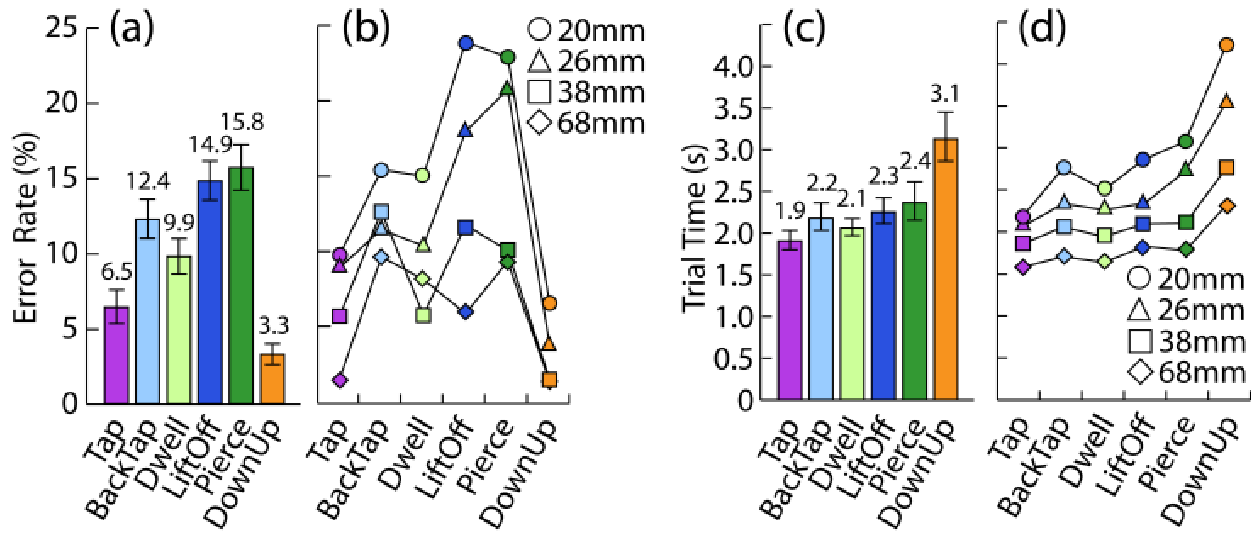

Figure 8: (a) and (b) Mean error rates. Error bars: ± 1 S.E. (c) and (d) Geometric mean trial times. Error bars: 95% CI.

Error rate: The overall error rate was 11.7% (382 of 3262 trials contained one or more undesirable selections before the target bin was selected). Figure 8a shows the mean error rates. *Selection method* had an effect on error rate ($\chi^2(5,N=12) = 30.44$, $p < 0.0001$) and pairwise comparisons showed that DownUp caused significantly fewer errors than BackTap, LiftOff and Pierce. No other pairwise comparisons were significant. DownUp's low error rate is due to its twofold accuracy requirement: 3cm down and 3cm up in the same bin without veering into an adjacent one.

We found error rates of 6.7, 8.7, 13.5, and 17.1% for 68, 38, 26, and 20mm bins, respectively, and there was also a significant effect of *width* on error rate ($\chi^2(3,N=12) = 28.30$, $p < 0.0001$). Pairwise comparisons showed that 68mm bins caused fewer errors than both 26 and 20mm bins and that 38mm bins caused lower error rates than 20mm bins. There was no difference between the two largest and between the two smallest bin sizes. When comparing how the *selection methods* performed at each *width* (Figure 8b) we found a significant difference at each *width*, but only significant post-hoc pairwise comparisons at the smallest width, where DownUp caused fewer errors than Pierce and LiftOff.

Trial time: Trial times were positively skewed and we performed a logarithmic transformation (which resulted in distributions close to normal) before analyzing the data.

Selection method and *width* had significant effects on trial time ($F_{5,55} = 28.6$, $p < 0.0001$, $\eta^2 = 0.72$ resp. $F_{3,33} = 210.9$, $p < 0.0001$, $\eta^2 = 0.95$). Across selection methods, the geometric mean trial times (i.e., the antilog of the mean of the log-transformed data) ranged from 1.9s for the largest 68mm bins to 3.1s for the smallest 20mm bins. As a result of the increased accuracy demand, trial times increased by about 15% for each decrement in bin width. Post-hoc pairwise comparisons showed that all bin sizes differed.

Figure 8c shows the geometric means for each *selection method*. Post-hoc pairwise comparisons showed that DownUp was slower than all other selection methods and that Tap was faster than Pierce and LiftOff. There were no other differences between the methods.

The significant *method*×*width* interaction ($F_{15,165} = 2.8$, $p < 0.01$, $\eta^2 = 0.16$) plotted in Figure 8d. Except for DownUp, all methods performed about equally well at 68 and 38mm bins. With 26mm bins though, we see marked peaks for Pierce and DownUp, and moderate, similar increases in the other methods. Only with the smallest bins do BackTap and LiftOff lose ground against Tap and Dwell.

Preference ratings: According to overall preference, 9 of 12 participants rated Tap to be the best, two preferred the LiftOff method and one favored BackTap.

Summary – Session A

The results indicate that performance – in particular, errors– degrades significantly after the 38mm bin size (9 bins). We suggest that for AD-Binning, targets should not be any smaller than this size. While Tap and Dwell appear to have the least errors and a trend toward faster selection times, these may not be practical in all applications. For example, a dwell may conflict with object browsing, and Tap should only be restricted to a specific on-screen target. We continue our exploration with BackTap and LiftOff as the on-device and off-device selection methods. The same participants were recruited for Session B, providing a certain level of expertise with AD-Binning.

Session B – Space Discretization & Binning Locations

As indicated above, prior studies [38,63] have suggested an accuracy trade-off in mid-air pointing with targets distant from a reference, in our case the edge of the device. This led to the evaluation of different around-device space division or discretization methods.

The Uniform discretization (Figure 9a) divides the available space into nine equally sized bins of 37.78mm. In the Distance Dependent discretization (Figure 9b) the inner bin is 27.2mm wide and the following bins are allotted an additional multiple of 2.64mm according to their position from the inner bin. Thus, the outer bin, which is located eight positions away, is 27.2+8×2.64=48.32mm wide

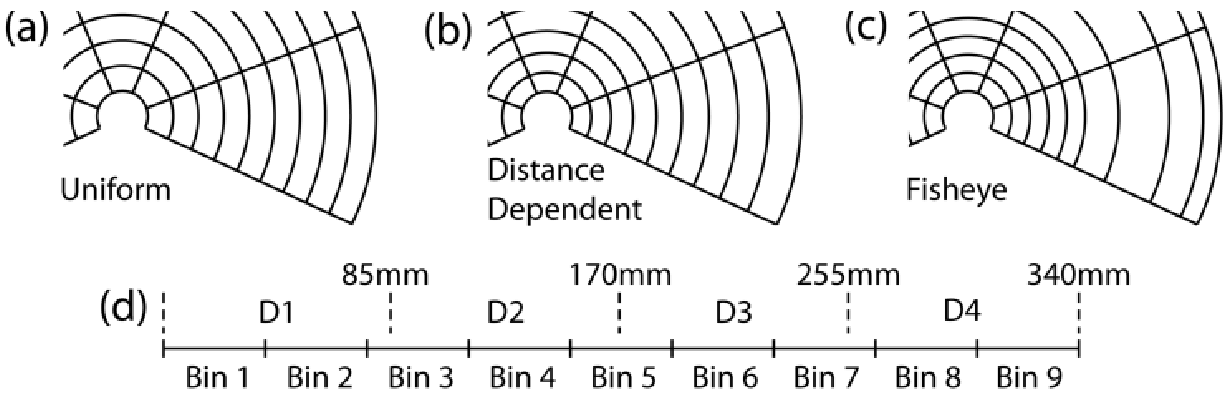

Figure 9: Discretization techniques and distance mapping.

We also included a fisheye discretization technique (Figure 9c) that uses a hysteresis function similar to [173,216] to dynamically add and remove extra space on both sides of each bin (except the first and last bins). The active bin expands to 75.28mm, its two neighbors expand to 50.28mm, and the remaining bins are 27.36mm wide. We controlled distance by dividing the available radial distance of 340mm in equal distance ranges, D1, D2, D3 and D4 (Figure 9d). A random number within the desired range was drawn and the bin at this distance was set as the next target, belonging to the corresponding distance range.

Task and experimental design

All task procedures were the same as in Session A.

Session B used a 3×2×5×4 within-subjects design for the factors *discretization* (Uniform, Distance dependent, Fisheye), *selection method* (BackTap, LiftOff), *sector* (Left, Top-Left, Top, Top-Right, Right), and *distance* (D1, D2, D3, D4). Participants performed 360 trials:

three repetitions for each combination of factor levels. We counterbalanced on discretization technique and half of the participants started with BackTap first. Participants completed 20 random practice trials and then 40 timed trials with each combination of discretization and selection methods.

Results

Error rate: BackTap had a significantly lower mean error rate than LiftOff (7.4% vs. 10.2%, Wilcoxon test: $Z = -2.7$, $p < 0.01$). We also found significant effects for *discretization* ($\chi^2(2,N=12) = 7.2$, $p < 0.05$) and *distance* ($\chi^2(3,N=12) = 15.3$, $p < 0.01$), but not for *sector*. Pairwise comparisons showed that the Fisheye, with a mean error rate of 6.6%, caused significantly fewer errors than both the Uniform and Distance dependent discretizations (error rates: 9.6% and 10.2%, respectively). Pairwise comparisons between distances showed that bins in distance range D1 caused more errors than bins in D2. There were no other pairwise differences.

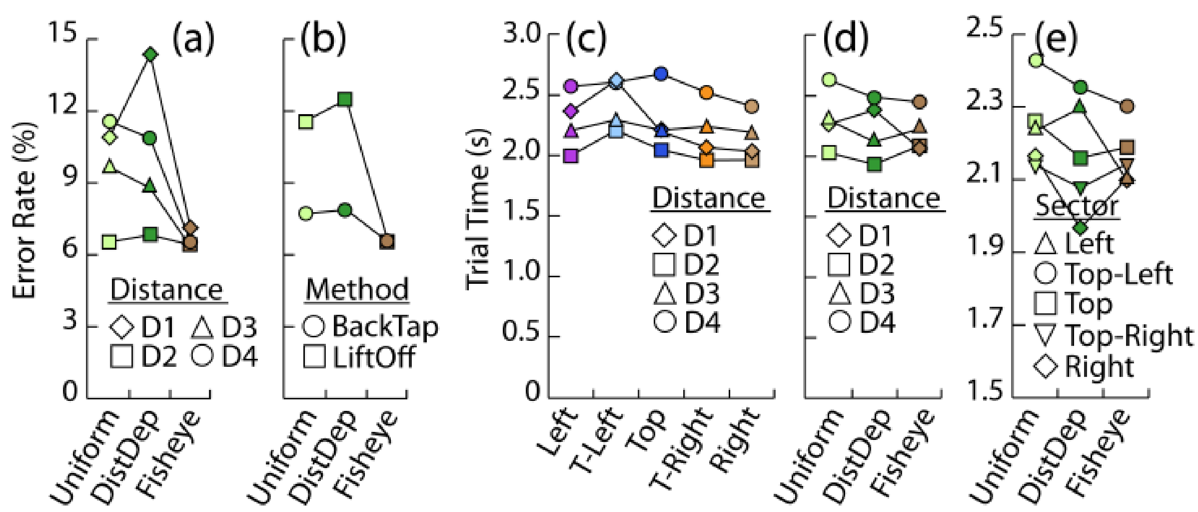

Figure 10: (a) and (b) Interaction effects, error rate. (c), (d), and (e) Interaction effects, trial time.

Interestingly, the Fisheye discretization reveals an overall equalizing effect over all distances. In the other two discretizations (Figure 10a) performance in D4 and D1 degraded, possibly due to reduced motor accuracy and smaller arc lengths, in the far and close bins, respectively. The Fisheye discretization also evened out the error rates between selection methods. The overall higher error rate with LiftOff is a result of poor performance when combined with Uniform and Distance Dependent discretization (Figure 10b). When extra space is added to the 'current' bin, as in the Fisheye, LiftOff performs as well as BackTap.

Trial time: As in Session A, trial times were positively skewed and we applied a logarithmic transform (with distributions close to normal) before analyzing the data.

The geometric mean trial time was 2.19s for BackTap and 2.20s for LiftOff. Across the two *selection methods*, the geometric means for the three discretization techniques were 2.24, 2.17 and 2.17s for the Uniform, Distance dependent and Fisheye, respectively. We did not find main effects for *selection method* or *discretization*, but there was a main effect for *sector* ($F_{4,44} = 8.0$, $p < 0.0001$, $\eta^2 = 0.42$) and for *distance* ($F_{3,33} = 33.6$, $p < 0.0001$, $\eta^2 = 0.75$).

Post-hoc pairwise comparisons between distances showed that bins in D4, with a geometric mean trial time of 2.49s, were significantly slower to select than bins located elsewhere. Bins in D2 were the fastest (1.97s). There were no differences between bins in D1 or D3 (2.18 vs. 2.16s).

Post-hoc pairwise comparisons between sectors showed that the Right and Top-Right sectors, with geometric means of 2.08s and 2.12s, respectively, were faster than the Top-Left sector which was the slowest at 2.36s. There were no other differences between any other sectors. A significant *sector*×*distance* interaction ($F_{12,132} = 3.9$, $p < 0.0001$, $\eta^2 = 0.26$) (Figure 10c), identifies D1 and D2 having marked peaks as the main sources for the overall poor performance in the Top-Left sector. We attribute these problems to screen occlusion: presumably, keeping the wrist at a natural angle when targeting Top-Left bins close to the device causes the hand and lower arm to occlude parts of the screen and the visual feedback provided by the overview.

We also observed a *discretization*×*distance* effect ($F_{6,66} = 7.4$, $p < 0.0001$, $\eta^2 = 0.40$, Figure 10d). As with errors, the Fisheye had an equalizing effect on trial time. It is notable that trial times for bins close to the device (D1) drop as a result of the Fisheye expansion. Comparing Distance Dependent to Uniform discretization reveals a clear negative effect of removing space from D1-bins (confirming Session A's result that bins should be ≥ 38mm).

The significant *discretization*×*sector* effect ($F_{8,88} = 2.7$, $p < 0.05$, $\eta^2 = 0.19$, Figure 10e), reveals that the Fisheye also equalized performance between sectors. It reduced selection times in the slow Top-Left sector, but also in the Left sector. It is also notable that the Distance dependent discretization improved performance in the Right sector.

Preference ratings: Eleven participants rated Fisheye as the preferred technique and one rated Distance Dependent as the best. Seven rated the Uniform discretization as their least preferred technique.

Summary – Session B

We observe that the Fisheye discretization had an overall equalizing effect on error rates and trial times, across selection method, distance, and sector. This technique takes the advantages of having extra space added to the both side of a bin. Our following experiments use the Fisheye for dividing the around-device space. Unexpectedly, selecting targets in the closest distances was less accurate and less efficient. Due to constrained movements with crossing arms, areas left and top-left of the device are generally more cumbersome.

Study 2 – Binning & Retrieval

Binning items could conceivably be done at any time. The user could quickly place an application icon, contact entry or web-bookmark in a system wide bin-collection for long-term storage and fast access. Binning could also be application dependent and serve more short-term purposes, such as browsing the results from a query or to manage a sub-set of items of temporary interest (e.g., yesterday's emails). In this study we compare binning techniques that provide varying degrees of user-control:

- *Automatic* provides no user control on item placement. The system assigns each item to an empty bin. The assignment can be random or based on item properties (e.g., a name, time stamp, color). Items are binned in a batch, either initiated by the user (e.g., by shaking the device) or automatically triggered through a query interface.
- *Tap-and-Bin* allows full user control. The user picks items, one by one, tapping their on-screen representations, and then, guided by the cursor in the on-screen

overview, moves the hand to the desired AD-bin to ‘drop’ it using a LiftOff gesture. Tap-and-Bin may be time-consuming with many items but facilitates individual placement strategies for improved recall. The direct acquaintance with each item in combination with the following arm movement may also help develop valuable proprioceptive memory linkages.

º *Flick-and-Bin* provides semi-automatic binning. The user indicates the desired AD-sector for a particular item by flicking it in the corresponding direction, and the system then automatically bins the item in the sector’s first empty bin. The on-screen overview provides dynamic sector highlighting during the flick. The automatic ‘first-empty’ strategy makes Flick-and-Bin fast at the expense of user control. A more elaborate version could map flick-distance to bin-distance for full user control.

Participants, Task and Experimental Design

Twelve daily computer and touch screen users (3 female) aged 20 to 39 years participated. Five had participated in study 1. Participation lasted approximately 30 minutes (including short breaks and practice).

Phase 1 of a trial consists of binning multiple icons (6 or 12); Phase 2 involves retrieving three of them. In conditions with Automatic binning, the participant taps a screen button to trigger it. With Tap-and-Bin and with Flick-and-Bin the participant taps a start button, the next icon to the bin in the trial is displayed on the screen, and timing starts. With Tap-and-Bin the participant taps the icon, moves the hand into AD-space, then bins the icon in an empty bin using a LiftOff, and timing ends. With Flick-and-Bin, the participant flicks the icon towards a sector with an empty bin and timing ends.

The on-screen overview, where empty bins are yellow and occupied are blue, provides dynamic feedback throughout the binning activity. Flick-and-Bin forces items to be binned in the directed sector with inner items filled first. Tap-and-Bin provides the most flexibility in terms of item placement.

A dialog box announces Phase 2 when all icons are binned. Dismissing the dialog box displays three random icons from Phase 1 for 10 seconds as a preparation for the upcoming three retrievals. Showing items prior to retrieval is representative of a real task where users know ahead of time what items they are looking for (such as during a search task). When the three icons disappear, timing begins, and the participant starts the first retrieval. As the retrieving finger moves beyond the screen border, the overview indicates its current location with a red marker. The bin content is shown next to the overview. When the correct bin has been found, the retrieval (and timing) ends with a LiftOff in the corresponding bin. After retrieving all three target icons, the binning Phase of the next trial starts.

The study used a 3×2×3 within-subjects design for factors *technique* (Automatic, Flick-and-Bin, Tap-and-Bin), *set size* (6 or 12 icons to bin), and *retrieval* (first, second, third in each trial). With three retrievals per trial, three trial repetitions with each technique and set size combination each, participant performed a total of 54 retrievals. The order of *technique* was counterbalanced between participants and set sizes were presented in a random order for each technique.

We used the Fisheye discretization and five sectors with three bins each. The Automatic binning prioritized inner bins and filled these starting from the right sector towards the left. Participants had two practice trials with each technique. Icons were randomly chosen from a set of 180 similarly styled images. No icon appeared in two consecutive trials.

Results

Binning phase

Binning strategies: after the study, all participants indicated that they tried to bin items strategically. With Tap-and-Bin and Flick-and-Bin, most participants categorized items (e.g., 'eatables', 'computer stuff', 'red ones', etc.) in sectors (participants did not know in advance what items to bin). As expected, it was easier to apply this strategy to six items than with twelve. Participants placed items in inner bins before the outer ones. Trials with six items provided more flexibility regarding bin choice but participants clearly avoided using the Left and Top-Left sectors (Figure 11a and b).

Binning time: As expected, Automatic binning took no placement time. Participants spent on average 2.9s to bin an item with Flick-and-Bin and slightly longer, 3.3s, when using Tap-and-Bin. With a mean trial time of 2.2s for LiftOff selections in study 1 (i.e., the same movement and gesture required by Tap-and-Bin), we see a strategizing overhead of 1.1s for Tap-and-Bin. Allegedly participants used the same strategies for both techniques but the flick gesture in Flick-and-Bin took on average 0.4s.

Retrieval phase

Error rates: In 42 of the 648 collected trials (6.5%) participants made at least one, and at most four, erroneous selections before the prompted item was selected. With all three placement techniques using LiftOff as the selection method, we found no difference in error rates between techniques. There was also no difference in error rates between *set sizes* or *retrievals*.

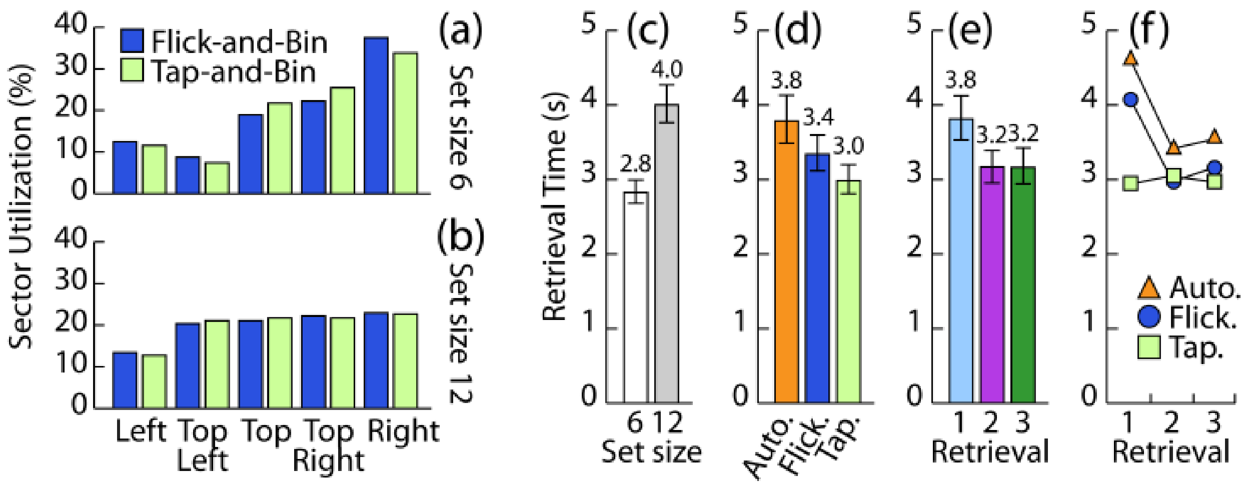

Figure 11: (a) and (b) Sector utilization. (c), (d), (e) and (f) Geometric mean retrieval times. Horizontal axis in (e) and (f) represents the first, second and third item retrieval. Error bars: 95% CI.

Retrieval time: Retrieval times were positively skewed and we applied a logarithmic transform (with distributions close to normal) before analyzing the data.

The geometric mean retrieval times for each *technique*, *set size* and *retrieval* are shown in Figure 11c, d and e. we found significant main effects for all factors (*set size*: $F_{1,11} = 100.2$, $p < 0.0001$, $\eta^2 = 0.90$; *technique*: $F_{2,22} = 7.3$, $p < 0.01$, $\eta^2 = 0.40$; *retrieval*: $F_{2,22} = 8.3$, $p < 0.01$, $\eta^2 = 0.43$).

As expected, the larger *set size* required more searching than the small set size, and thus took longer time. Post-hoc analyses between *techniques* showed that Tap-and-Bin was significantly faster than Automatic and that there were no other pairwise differences. Post-hoc analyses between *retrievals* showed that the first retrieval was significantly slower than the other two. The second and third retrievals did not differ.

The significant *placement technique*×*retrieval* interaction ($F_{4,44} = 5.9$, $p < 0.001$, $\eta^2 = 0.34$, Figure 11f) reveals that the overall advantage of Tap-and-Bin (21% vs. Automatic, 12% vs. Flick-and-Bin) is mainly a result of exceptionally fast first retrievals. With previous placement analysis showing that participants did not make use of the possibility to leave the inner bins empty with Tap-and-Bin, it is particularly interesting to note the large difference between Tap-and-Bin and Flick-and-Bin in the first retrieval. With no inner bins empty, the only difference between the two techniques is the amount of physical activity required to do the binning, a short flick for Flick-and-Bin, moving the arm and a LiftOff gesture for Tap-and-Bin. Apparently, the greater physical activity needed for Tap-and-Bin fostered spatial memory. In the first retrieval with Flick-and-Bin and Automatic, participants had to rely more on the visual overview and search. The position information participants gained during this first search was then utilized in later retrievals to improve performance to Tap-and-Bin's level.

Summary

Our results suggest that the overhead involved in manually binning items as in Tap-and-Bin is compensated by improved retrieval times due to enhanced spatial encoding. Spatial enforcement of item locations is also present while searching for items in AD-Bins: retrieval performance improved after having selected the first item in Automatic and Flick-and-Bin, as participants mentally recorded positions of subsequent items to retrieve.

Results from the above studies suggest that AD-Binning can facilitate selection of reasonably large items (study 1A), where errors and selection times across distance and sectors can be equalized using a space discretization technique such as the Fisheye (study 1B). Furthermore, exploring AD-Binning space enhances spatial encoding of item positions around the device (study 2). These results inform the selection of suitable design parameters for an efficient AD-Binning technique.

Study 3 – Analytic Task

With knowledge of how the various design parameters influence performance from the previous studies, we next demonstrate and evaluate a practical AD-Binning usage scenario for browsing information content.

Participants, Task and Experimental Design

Twelve daily computer and touch screen users (3 female) aged 18 to 35 years participated. Two were new and had not participated in any previous study. Participation lasted approximately 45 minutes (including breaks and practice).

The task simulates a frequent situation where the user has queried a system for information. In this case, a geographic tourist portal for hotel reservations where the query results are displayed on a map. Issuing the query can result in items being placed automatically in around-device space, and ready for retrieval.

Figure 12: (a) On-screen interface with 10 hotels at low density. (b) AD-Bins with 15 hotels at high density.

A trial starts with the screen displaying a prompt to search for the cheapest n-star hotel (n=number from 1-5). After reading the text, the participant taps a start button and trial time starts. The next screen shows a city map with a set of circular markers (ø 5mm) representing various hotels. We place the search criteria at the top of the screen as a constant recourse (Figure 12a). In the 'on-screen' condition, the price and rating (number of stars) for a hotel are displayed in a callout box that opens when the marker is tapped (Figure 12a). The information box is closed with a tap on the map or on another hotel marker. When the participant believes they found the hotel satisfying the search criteria, the trial ends with a tap on the button (1×0.7cm) in the callout box. If correct, the trial time stops and the text prompt for the next trial is displayed. If incorrect, an error message pops up which blocks further input for one second before it automatically fades away. After that, the search for the correct hotel can continue. Panning and zooming are fully enabled.

With AD-Binning, we use the automatic placement method such that 'proxies' to hotel markers are placed in random AD-bins. The AD-space is divided into five sectors, with a total of 5, 10 or 15 bins depending on the condition (see below), and uses the Fisheye discretization. The participant browses hotels in off-screen space by moving the index finger between bins. At bin-entry, the corresponding hotel marker is highlighted and the hotel information is shown next to the bin-overview at the bottom of the screen (Figure 12b). To select a hotel, the user performs a LiftOff inside the desired bin. The trial prompt, timing and error notifications work as previously described. On-screen panning and zooming are fully enabled.

The study used a 2×2×3 within-subjects design for factors *interface* (on-screen, AD-Binning), marker *density* (low, high), and *number of items* (5, 10, 15). Participants performed five repetitions for each combination of factor levels, for a total of 60 timed trials per participant (five practice trials were given per interface). Combinations of density and number of items were presented in random order within each interface. Six participants started with AD-Bins, six with on-screen browsing.

In low-density conditions, all hotel markers were positioned at random positions within 1.974cm of the map/screen center. The high-density conditions used 0.987cm. At least two hotels with the requested number of stars existed in each trial. Hotels, prices, stars, and marker positions were otherwise completely randomized.

Results

Error rate

In 90 of the 720 collected trials (12.5%) participants made at least one, and at most five, erroneous selection before finding and selecting the correct hotel (44 trials with on-screen, 46 with AD-Bins). Neither *interface* nor *density* influenced the error rate (Wilcoxon tests), but *number of items* did ($\chi^2(2,N=12) = 10.8$, $p < 0.01$). Post-hoc analyses showed differences between 15 and 5 items (46 vs. 19 trials) and between 15 and 10 items (46 vs. 25 trials). Naturally, with more items to manage and to compare, the risk of making a mistake increases.

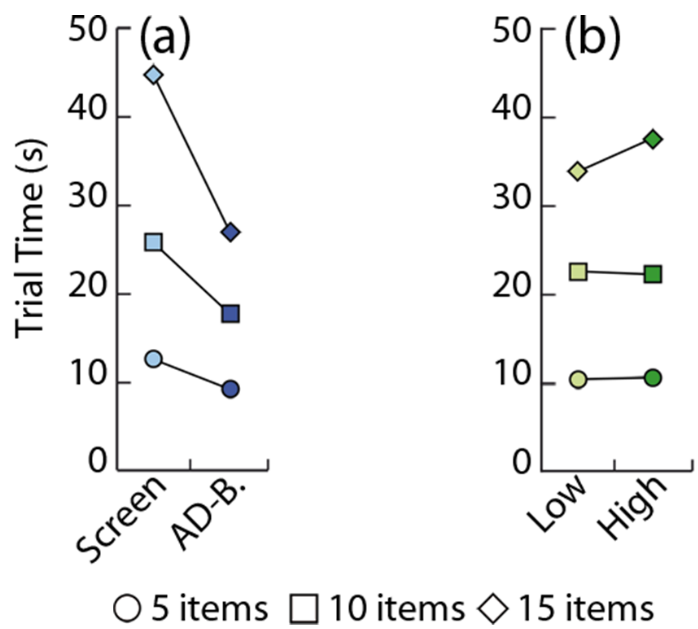

Figure 13: (a) numbers of items x interface interaction. (b) numbers of items x density interaction.

Trial time

AD-Binning, with a mean trial time of 17.5s (s.d. 9.3), was significantly faster ($F_{1,11} = 43.9$, $p < 0.00001$, $\eta^2 = 0.80$) than the on-screen interface with a mean trial time of 27.6s (s.d. 17.2). There was no difference between high and low *density* conditions but *number of items* significantly influenced trial time (Greenhouse-Geisser corrected, $F_{1.12,12.31} = 164.6$, $p < 0.00001$, $\eta^2 = 0.94$) with post-hoc comparisons showing differences between all factor levels.

Number of items interacted with both *interface* ($F_{2,22} = 24.5$, $p < 0.00001$, $\eta^2 = 0.69$) and *density* ($F_{2,22} = 3.7$, $p < 0.05$, $\eta^2 = 0.25$) (Figure 13a and b). Time savings with AD-Binning increased disproportionally with the number of items: 27.9% with five items, 31.9% with ten, and 40.2% with 15 items. Overall, AD-Binning was 35.8% faster. Across *interface*, there were no differences between low and high *density* in conditions with 5 and 10 items but performance deteriorated with 15 items at high density (Figure 13b). Presumably, this was mostly caused by the increased need for elaborate pan/zooming and closing of the callout actions in the on-screen condition (however, we note that there was no significant 3-way interaction).

Preference ratings

Three of twelve participants preferred the on-screen interface, eight preferred the AD-Binning and one was undecided. The most frequent reason for preferring AD-Binning was how it helped to recall the rough location (sector/'quadrant') of the 'current best answer'. Additionally, participants mentioned that on-screen interface forces them to frequently switch between the list and detailed views. It makes them slower in exploring on-device content.

Summary

In comparison to on-screen input, AD-Binning reduces information browsing time for three reasons. First, AD-Binning is in 'browse' mode by default. Retrieving object information involves hovering or sliding the finger between bins. In contrast, on-screen consumes at least two steps: i) tap on an icon to pull-up information, and ii) tap again to close the callout box or to retrieve information from another marker. Second, AD-Binning target sizes can be significantly larger than those on-screen. In the example, queries with 5 items used five bins, each with a larger space than when the query had 15 items. Finally, participants exploited spatial abilities with AD-Binning. They would cache in memory the best bin location satisfying the query criteria and update in memory this bin location only when the next best item was available. While this happened with automatic placement, in a full manual placement reliance on spatial memory would be even stronger as indicated from results of study 2.

Discussion

We summarize the main findings and present them as around-device binning guidelines, discuss other applications that can benefit from the guidelines and conclude with some limitations of our investigation.

Design Considerations

The results offer the following guidelines to designers of interfaces based on the concept of AD-Binning:

- *Input range:* An interaction space extending 40cm beyond each side of the device is suitable when around-device interaction is focused on a horizontal plane defined by the device. A radial division and partitioning in sectors allows for a comfortable reach.
- *Target Size:* Use the largest targets possible, and targets should not be much smaller than 4cm across. With a radial bin arrangement, interactions close to the device cause a higher number of errors, as bins are smaller there.
- *Ideal interaction regions:* Prioritize interactions on the same side as the dominant pointing hand, as users intuitively avoid interaction on the non-dominant side to avoid occluding on-screen visual guidance.
- *Around-device space division:* Fisheye discretization can suitably divide around-device space to provide equally efficient access to all content on the device and reduces accuracy requirements.
- *Selection methods:* To trigger a selection, on-screen and off-screen methods can be equally effective. Designers can choose a selection method based on task. Finger lift-off is possible in both mid-air and when the device is resting on a surface.
- *Placement methods:* Promote spatial learning through direct 'physical contact' with off-screen space. The extra time and effort required to manually place items in around-device space pays off in item browsing and retrieval activities. Rapid binning is also possible with automatic placement methods, which can be triggered through a query.

Applications

Some obvious applications for AD-Binning including photo storage (organized in sectors denoting date, event or any other semantic information), storing and retrieving items from contact lists and bookmarking items of interest when browsing a website. We also envisioned longer term applications where the user can capitalize on proprioceptive memory linkages developed over time to access content across applications or regularly issued commands, similar to CommandMaps [167]. With further development, AD-Binning could also apply to the following applications:

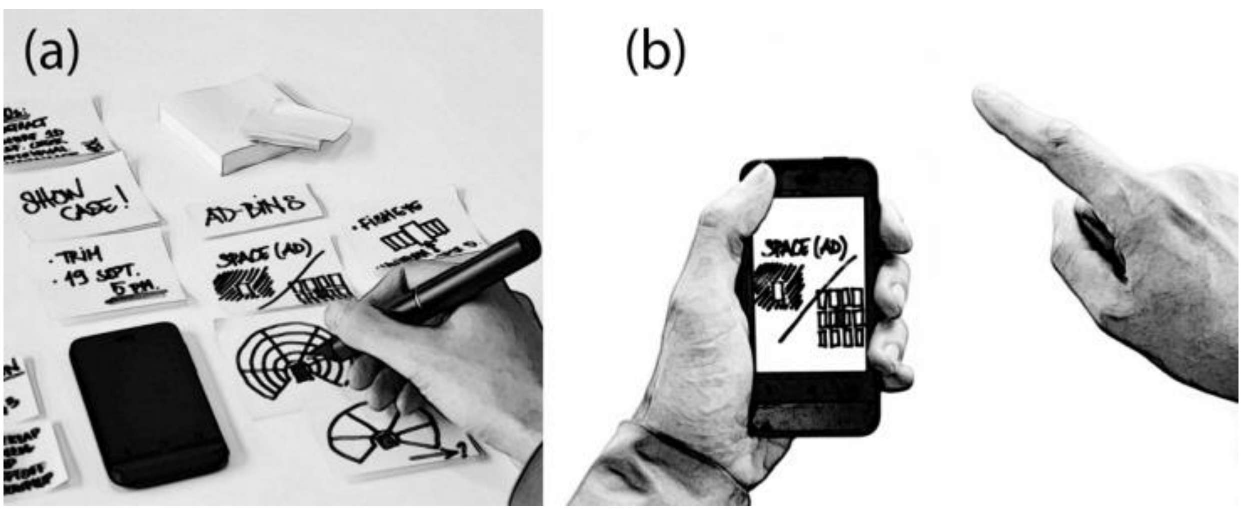

Figure 14: AD-notes. (a) Creating, arranging and storing physical notes. (b) Browsing notes stored in AD-space.

Item retrieval based on item organization. Our implementation of AD-Binning did not consider specific ordering of items. However, datasets have inherent structures that AD-Binning interfaces can leverage. For example, items could be sorted based on price, alphabetically or chronologically. In an email client, items can be placed in chronological order in around-device space. This can allow the user to retrieve items immediately based on their previously developed knowledge of AD-Binning item organization.

Mixing physical and virtual bins. AD-Binning could also be used in mixed physical and virtual workspace scenarios. By tracking the position of a digital pen, AD-Binning could facilitate note taking and brainstorming scenarios where ideas and sketches are made on physical notes arranged around the device (a). Committing the final note content and position stores the note in the corresponding AD-space for later retrieval or browsing (Figure 14b).

Summary

In this chapter, we have presented the design and evaluation of AD-Binning, a novel user interface for future small-screen mobile devices that will be able to sense finger movements in their vicinity. With AD-Binning the user can off-load screen items from the small screen into the larger off-screen space around the device. AD-Binning was mainly designed to support the user in analytic scenarios that require intensive browsing and comparisons between many alternatives, such as long query result lists or other information intensive situations where exploration is necessary before a decision is made. Such tasks can be laborious to perform using the interactions provided by small touch screens. With AD-Binning the user can efficiently store, browse and retrieve content through direct interactions in the space around the device.

Chapter 4: Acceptance Studies of Around-Device Gestures in and for Public Settings

Results from the previous chapter confirm that using the surrounding space for accessing virtual content is a promising alternative to the traditional touch screen. However, the user studies mentioned in the previous chapter are conducted in laboratory settings and very little is known about users' and spectators' comfort levels using these innovative interactions while using it in public spaces, such as in a shopping mall or on a bus. This chapter explores users' and spectators' attitudes about using around-device gestures in public.

The chapter contains three studies. The first study explores how users' comfort levels vary when they perform gestures in different regions around the device and at different distances from the device. We also explore whether users' perceptions about AD-gestures are related to personality traits, e.g., measures of extroversion. In second study in this chapter focus on the size and duration of gestures and investigates how these parameters affect users' comfort. Finally we explore spectators point of view and examine reactions when observing someone else using AD-gestures.

In the studies on this chapter we elicit participant impressions through questionnaires. In particular we use Rico and Brewster's [162] 'audience-and-location' axes to determine levels of "social acceptability". We ask participants to state in front of whom and in what locations they would feel comfortable using AD-gestures. To determine the influence of gesture parameters, we ask participants to indicate how comfortable ('very comfortable',

‘comfortable’, ‘neutral’, ‘uncomfortable’, ‘very uncomfortable’) they felt when performing the gestures. We acknowledge that the central usage of the word ‘comfort’ may leave room for diverse and individual interpretations. However, we carefully instructed participants on the intended and desired interpretation, and we explicitly asked participants to relate ‘comfort’ to social and mental aspects rather than physical ones. Furthermore, the terminology finds support in previous studies on social acceptance of gestures for mobile use [142,162,163,200,201] in which both authors and participants have used similar wording.

In these studies, we also use simple gestures that consisted of ‘drawing’ one-digit numbers in the air around the device. Such abstract drawing-gestures are easy to perform, are context free, and reduce the likelihood that participants’ responses are influenced by any uncontrolled factors such as previous experiences, cultural background, and associations to interface tasks or functionality.

In contrast to all prior studies on AD-input and the majority of earlier studies on user acceptance of gestures for mobile devices, the studies are conducted in public places. We believe this is important since it provides participants with firsthand usage experiences before making their judgements, unlike in video-based surveys where participants are asked to imagine future use and possible feelings.

Study 4 – Comfortable Region and Distance

Our first study on social acceptability was conducted in a shopping mall. Without being informed about the exact purpose of the study, participants were asked to perform a set of AD-gestures in a busy entrance zone of the mall. Participants stood and held the smartphone in the non-dominant hand at a natural viewing distance.

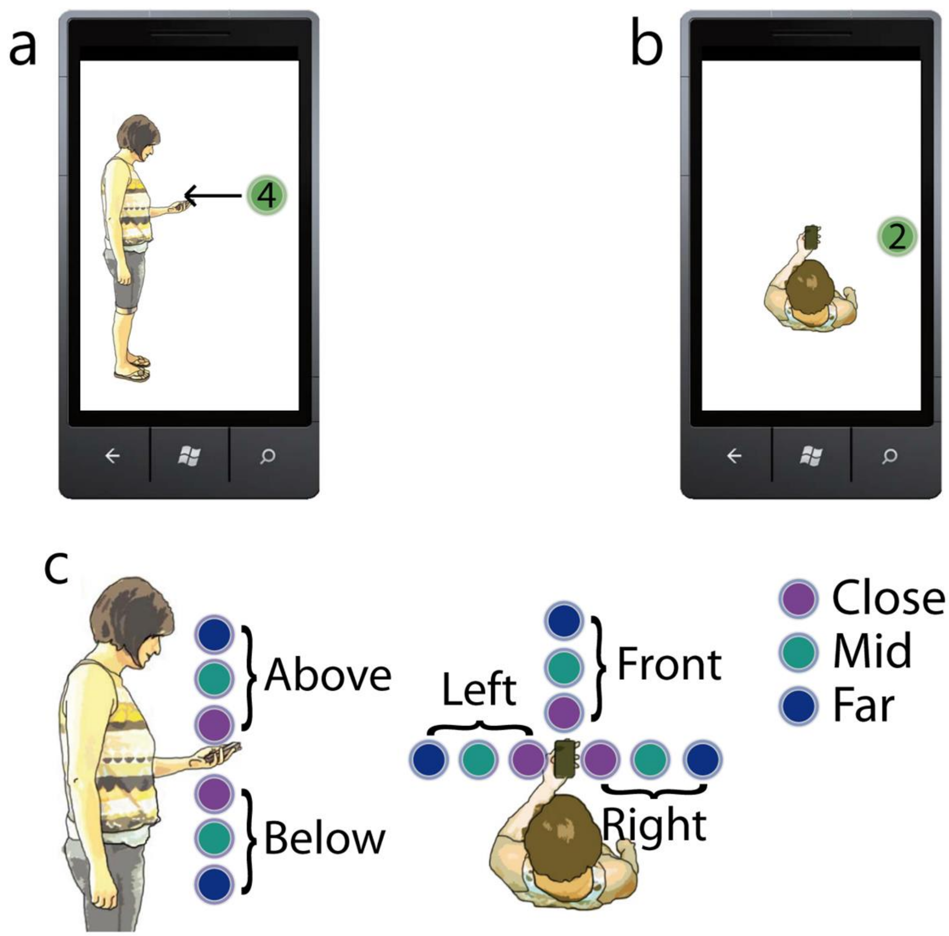

Figure 15: Example task prompts for (a) a gesture in the Above region at the Close distance and for (b) a gesture in the Right region at the Far distance. (c) Gesture distances and regions.

A set of 45 image guided gestures were solicited. The images were viewed in full-screen using the device's default image browser. Each image showed a position in the air around the device and a one-digit number to 'draw' at the indicated in-air position. (Figure 15a and b) shows two example images.

As illustrated in Figure 15c, five different drawing regions were used: *Above* and *Below* the device, to the *Left* and to the *Right* of the device, and in *Front* of the device. Each region was divided into three distance ranges, measured from the device: *Close*, *Mid*, and *Far*. *Close* corresponds to the area 0 to 15cm away from the device, *Mid* to 15 to 30cm from the device, and *Far* to more than 30cm away from the device. The furthest distance roughly corresponds to the maximum comfortable reaching range around a handheld device [71]. Five regions and three distances yield 15 around-device positions prompted during the study for each participant. In practice trials, we explained how to interpret the 15 different gesture positions shown in the task images and to "anchor" the positions in relation to body parts (such as the face, chest, and shoulder).

The experiment was self-paced and participants were instructed to work through the images and to draw the prompted numbers at a moderate speed. The next image in the set was loaded with a flick gesture on the touchscreen.

The image set was divided into three sub-sets of 15 images, with one image for each of the 15 gesture positions. The order of gesture positions and the prompted number to draw in the air was randomized within the three sub-sets. All participants used the same image set (and image sequence) in two rounds, for a total of 90 gestures.

Q1) On a scale from 1 to 6, what was your overall impression/emotion during the task?

1☐ 2☐ 3☐ 4☐ 5☐ 6☐

I hated it, it felt terribly awkward — I enjoyed it, it felt comfortable

Q2) Imagine your mobile phone could be operated through similar gestures as those you recently performed. Now, **in front of whom** do you think you would **feel comfortable** using such gestures? Select **one or more** items from the list below.

☐ I would **not** feel comfortable using them even when alone

or

☐ when alone ☐ in front of my partner ☐ in front of friends

☐ in front of family ☐ in front of colleagues ☐ in front of strangers

Q3) Now, **in which locations** do you think you would **feel comfortable** using such gestures? Select **one or more** items from the list below.

☐ I would **not** feel comfortable using them no matter where I am

or

☐ at home ☐ on the sidewalk ☐ in a pub, café, or restaurant

☐ in a shop ☐ in a museum ☐ as a passenger on a bus or train

☐ at my workplace

Q4) In the figures to the right, write a number from 1 to 5 in **each of the circles** to indicate how **comfortable** you felt performing a gesture at that particular position.

1=very comfortable 4=uncomfortable

2=comfortable 5=very uncomfortable

3=neutral

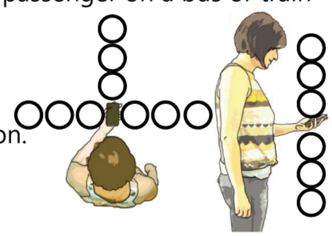

Figure 16: Study 4 questionnaire

After completing the two rounds of gestures the participant was debriefed and asked to answer four questions regarding (Q1) the overall impression/emotion during the task, (Q2) in front of whom and (Q3) in what locations he/she would feel comfortable using AD-gestures, and (Q4) how comfortable he/she felt when performing the gestures in the various in-air positions. Figure 16 shows the questionnaire. We instructed the participant to answer the questions, to interpret the central word 'comfort' from a social perspective and to ignore issues related to physical comfort and practicability. The participant also completed a Big Five Personality Test [137], which is a standard psychology test that assesses a person's personality on five broad dimensions: conscientiousness, openness, extraversion, agreeableness, and neuroticism. We used the test available at http://www.outofservice.com. We were primarily interested in identifying whether scores on extraversion correlated with perceptions of AD-gestures.

Eighteen right-handed smartphone owners (6 female) aged between 24 and 51 years (mean 31.1 years, s.d. 6.6) participated. Participation lasted roughly 30 minutes.

Results

Question 1: Only four participants indicated that their impressions/emotions during the task were more negative than positive by selecting a rating of 3. The other fourteen participants indicated having had more of a positive impression/emotion during the task: ten gave a rating of 4, three gave a rating of 5, and one participant indicated enjoyment/comfort (rating 6).

Question 2 and 3: No participant completely rejected the idea of using AD-gestures by stating that he/she would not feel comfortable using gestures even when alone. Only one stated that he would only feel comfortable using the gestures if alone and at home. Sixteen participants indicated they would be comfortable doing the gestures in at least one of the non-private settings (i.e., when not at home and when not alone). One participant thought he would feel comfortable using AD-gestures in all locations and in front of all audiences listed in the questions.

To analyze the answers to Question 2 and 3, we established an acceptance rate for each given audience and location by calculating the percentage of participants who selected each audience/location in their answers. As visible in Figure 17, the more familiar audiences, family, partner, and friends, were accepted by most participants. Only 6 of 18 participants indicated they would comfortably use AD-gestures in front of colleagues and strangers. A Cochran's Q test showed a significant difference between the audiences ($\chi^2(5,N=18) = 46.9$, $p < 0.001$). Post-hoc McNemar tests (Bonferroni: α-levels from 0.05 to 0.003) showed that the acceptance rates for the least familiar audiences, colleagues, and strangers, were significantly lower than the rates for the other audiences.

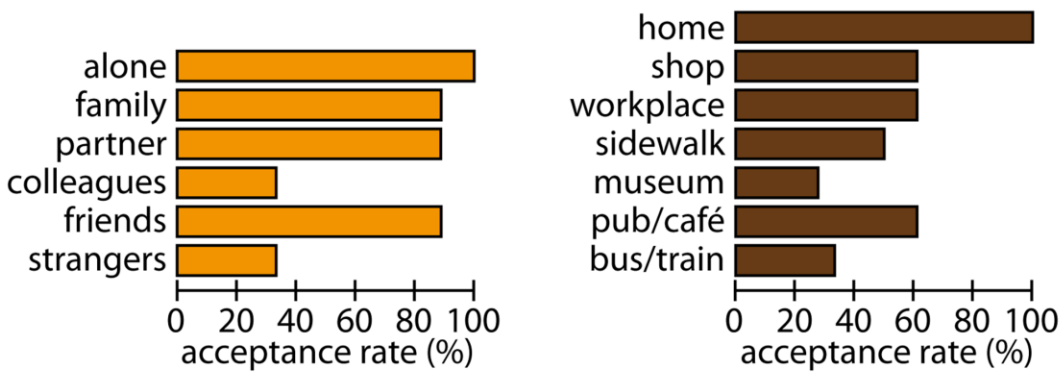

Figure 17: Acceptance rates for audiences and locations.

Also the location influenced the willingness to use AD-gestures (Cochran's Q test: $\chi^2(6,N=18) = 27.2$, $p < 0.001$). However, the results are slightly more controversial: acceptance rates of 50% were obtained for four locations (shop, workplace, sidewalk, and pub/café). All participants indicated they would feel comfortable using AD-gestures at home. Post-hoc pairwise McNemar tests (Bonferroni: α-levels from 0.05 to 0.002) showed that the rate for home was significantly higher than for the two most rejected locations, museum and bus/train (with rates of 28% and 33% respectively). No other pairwise comparisons were significant.

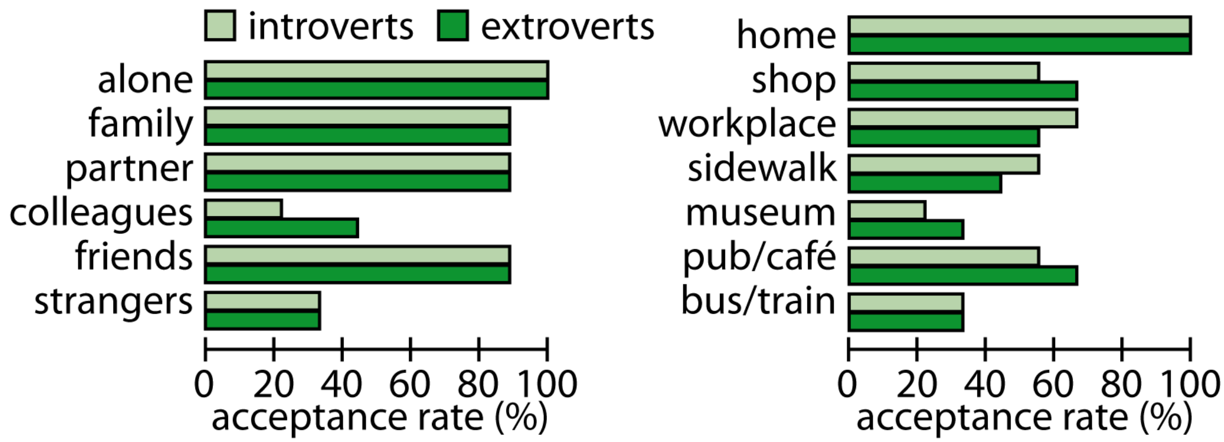

Figure 18: Acceptance rates for audiences and locations split by introverted and extroverted users.

We also examined whether there was a connection between acceptance rates and participants' extraversion personality trait. Our main focus was on this personality trait as people with higher extroversion percentiles are more sociable, friendly and enthusiastic about new things. On the Big-5 test, nine of the 18 participants had an extraversion percentile score below 50, and nine a score above 50. That is, 50% of the participants were less extraverted than 50% of all persons (over 10,000 persons) that have completed the online Big-5 test service we used. As visible in Figure 18, introverts and extroverts provided similar ratings for most audiences and locations. We did not find any significant differences in the ratings. Additionally, we examined the connection between users rating and other personality traits (i.e., Openness to Experience/Intellect, Conscientiousness, Agreeableness and Neuroticism). Similar to the extroversion results, we didn't find any significant difference between user ratings and the traits.

We can conclude that the majority of the participants were quite open to the idea of AD-gestures. Only one participant indicated a strong hesitance to perform AD-gestures in public or in front of someone else. All other participants responded that they would feel comfortable using AD-gestures outside the privacy of their home: in four of the six non-private locations 50% or more of the participants indicated that they would feel comfortable using AD-gestures. Overall, our results confirm Rico and Brewster's [162] results, which also showed that both audience and location are important factors that influence the willingness to use gestures.

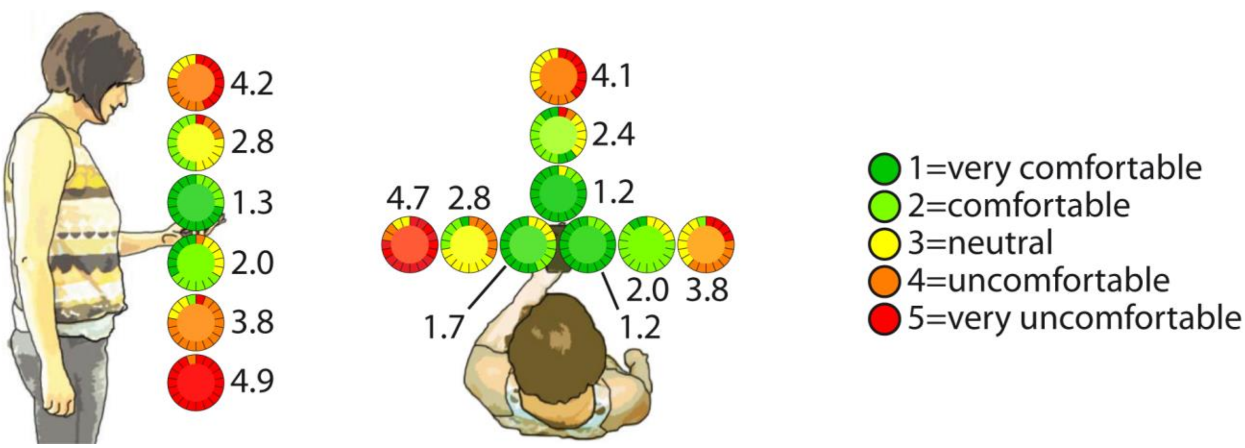

Figure 19: Comfort-ratings of the 15 gesture positions (segments around each middle circle show individual ratings, best viewed in color).

Given the low acceptance rate for strangers in the audience category, the fairly high acceptance rates for the shop, sidewalk, and pub/café – where one would expect to be seen by strangers – are somewhat surprising. However, we assume that participants indeed focused on the various locations and the circumstances that are typical for these. For example, in a pub an AD-gesture could be more easily disguised, e.g., under a table, than in a bus with a passenger sitting close by. We also suspect that acceptance depends on the frequency, size and duration of the gestures are used, their size and duration. Particularly for locations such as in a crowded bus or on a busy sidewalk where large or lengthy gestures are likely to be perceived as more inappropriate than small or quick ones. We evaluate these factors in Study 5.

Question 4: Figure 19 shows the mean rating for each of the 15 gesture positions that were rated by participants according to how comfortable they felt when doing gestures in these positions. In all five regions, the position furthest away from the device had the worst rating. Most participants stated they felt either uncomfortable or very uncomfortable doing AD-gestures far away from the device. The majority of participants indicated they felt comfortable or very comfortable when gesturing at a close distance from the device, no matter what region they used.

A Friedman test showed significant differences across the five regions. Analysis of average rating for the far, mid, and close distances (with mean rating of 4.3, 2.8, and 1.5, respectively) yield very high significance, $\chi^2(2,N=18) = 36.0$, $p < 0.001$. Post-hoc Wilcoxon tests (Bonferroni: α-levels from 0.05 to 0.016) showed differences for all three pairwise comparisons.

We also found significant differences in the ratings for the five regions ($\chi^2(4,N=18) = 36.0$, $p < 0.0001$). The mean rating for each region was: 2.8 (above), 3.6 (below), 3.1 (left), 2.6 (front), and 2.3 (right). Post-hoc Wilcoxon tests (Bon-ferroni: α-levels from 0.05 to 0.005) showed significant differences between the left and right region, and between the below region and each of the above, right and front regions. No other comparisons were significant.

Again, as with acceptance rates for audiences and locations, we did not find any differences in the distance ratings or region ratings between introverted and extroverted persons.

We can conclude that the distance between the device and the region in which AD-gestures are made strongly influence how comfortable users feel performing AD-gestures in a public setting. The level of comfort depends on the position of the gestures: gestures below the device evoke feelings of discomfort (more so the further away they are from the device) and gestures to the right of the device are preferable (note that all the participants were right-handed, we assume the results regarding the left and right regions should be mirrored for left-handed users).

Summary

Our results suggest that we could expect the majority of future around-device gesture users to have a neutral feeling when they use AD-gestures in public (Q1), but that how comfortable they feel using the gestures will depend on where and in front of whom the gestures are used (Q2 and Q3). The results also suggest that most users think that AD-gestures are compatible with many public settings, but that the acceptance for some settings is quite divergent. The results do not show that acceptance is related to the extraversion personality trait. Furthermore, the results (Q4) show that, generally, users feel more comfortable when gesturing within 30cm from the device (i.e., distances corresponding to the Close and Mid distances in the study).

Study 5 – Comfortable Gesture Size and Duration

With the knowledge that most of the participants in the previous study showed a neutral attitude towards using AD-gestures in public and that none completely rejected the idea of public use, we conducted this study to investigate how the size and duration of AD-gestures affect users' attitudes. Since Study 4 showed that there was no relationship between users' extraversion personality trait and how they perceived using AD-gestures in public, we decided not to use the Big-5 test.

For Study 5 we used the busy main entrance hall of the local university. The task and materials were similar to those used in Study 4. A set of images guided the participant through the task. As in Study 4, the images prompted one-digit numbers to be 'drawn' in the air at a specific in-air position around the device. A Silverlight Windows Phone application displayed the task images and randomized the image sequences for each new participant. One task image is shown in Figure 20. A task counter was shown in the top right corner, and a label in the top left corner indicated to the participant for how long he/she was required to draw the prompted number. When the 'start' button was pressed, the timer in the bottom right corner of the screen started. The participant was asked to re-draw the digit in the indicated location as long as the current task screen was shown. The next task screen was loaded when the timer reached the prompted duration.

We used small and large gestures. Small gestures were required to cover an area of about 15×15cm, large gestures 30×30cm. These sizes roughly correspond to half of the full distance of the preferred gesture distances defined by Study 4 results (Close and Mid). Two small gestures were prompted in each of the *Left*, *Front*, and *Right* regions (corresponding to distance *Close* and *Mid* in Study 4) and one large gesture was prompted in each of the three regions. We used three gesture durations: 3, 6, and 9 seconds (typical AD-gesture durations reported in the literature [71,99,120]). In total 27 task images were used. The used combinations of gesture size, duration and location are shown in Figure 20. Each combination was repeated twice, for a total of 54 images. After completing the 54 gestures the participant was debriefed and asked to complete the questionnaire shown in Figure 20. As in Study 4, we instructed the participant to answer these questions, to interpret the central word 'comfort' from a social perspective, and to ignore issues related to physical comfort and practicability. Participation lasted around 25 min. Eighteen right-handed smartphone owners (3 female), aged 21 to 32 years (mean 26.1, s.d. 3.6) participated. Five had participated in Study 4.

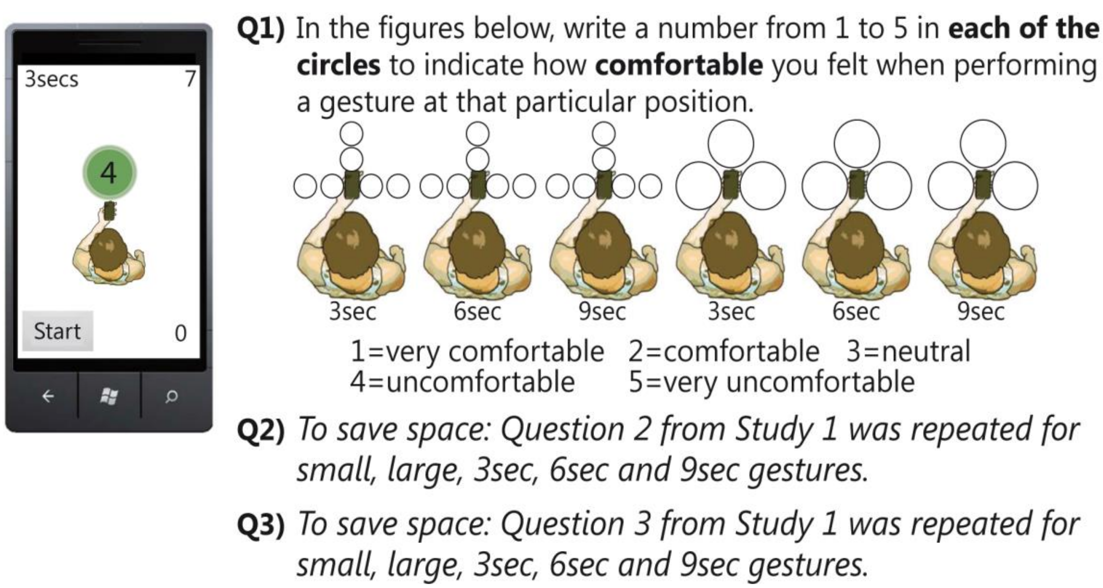

Figure 20: Left: task screen. Right: Study 2 questionnaire.

Results

Question 1: Figure 21 shows the mean rating for each of the 27 size/duration/location combinations rated by participants according to how comfortable they felt when performing these gestures. Although participants were asked to provide twice as many ratings for small than for large gestures (small gestures were performed twice in each region), we chose to make a comparison for guidance. The mean comfort-rating was 3.0 for small gestures and 3.5 for large gestures. A Wilcoxon test showed a significant difference ($Z = -2.9$, $p < 0.01$). Smaller gestures felt more comfortable, understandably because these are likely to attract less attention.

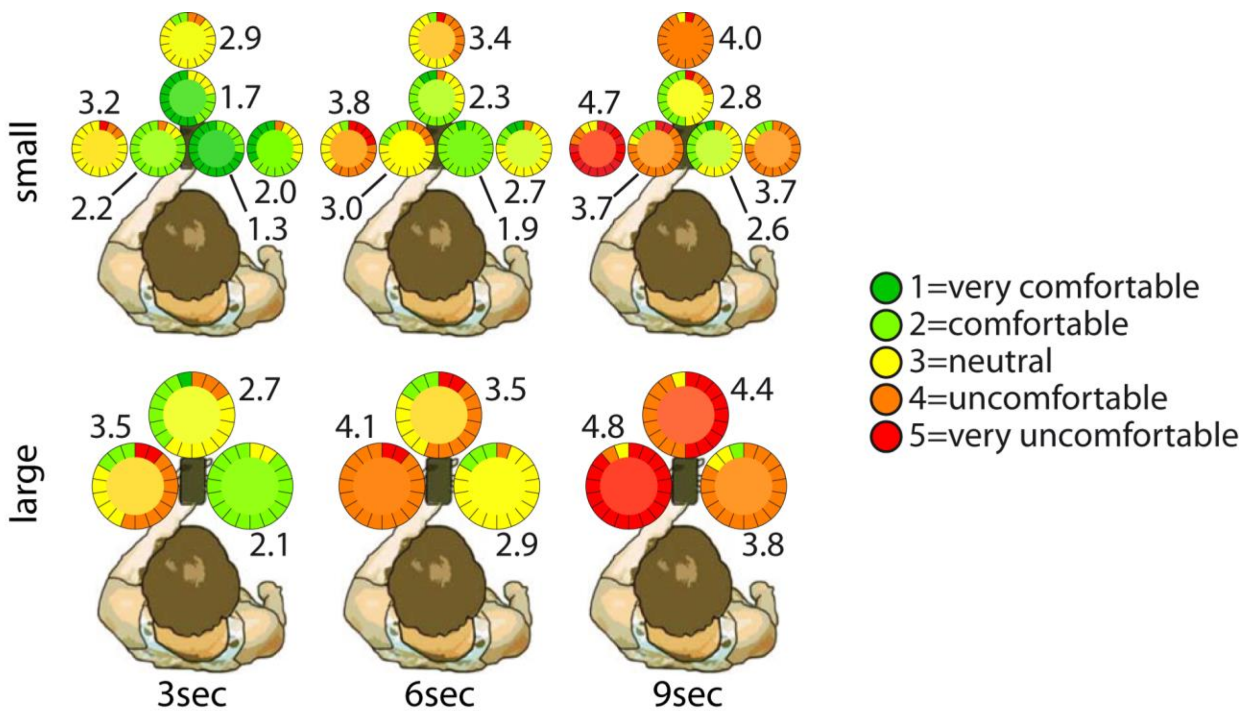

Figure 21: Comfort-ratings for all gesture size-duration combinations (segments around each middle circle show individual ratings, best viewed in color).

The average comfort-rating was 2.4 for 3sec gestures and 3.1 and 3.8 for 6sec and 9sec gestures, respectively. A Friedman test showed differences between the three gesture durations ($\chi^2(2,N=18) = 34.5$, $p < 0.0001$) and post-hoc Wilcoxon tests (Bonferroni: α-levels from 0.05 to 0.016) showed differences for all pairwise comparisons. We also found significant differences among the three regions (Friedman: $\chi^2(2,N=18) = 24.4$, $p < 0.0001$) with post-hoc Wilcoxon tests (Bonferroni: α-levels from 0.05 to 0.016) showing that the right region, with a mean rating of 2.6, was significantly different from both the front (mean 3.1) and the left (mean 3.7). Also left and front differed. Again note that all participants were right-handed and assume the results would be mirrored for left-handed users.

In Figure 21, we see an interesting interplay between position, size and duration indicating that the drawbacks of large gestures can be compensated for if they are done in a favourable location and if they are quick (e.g., to the right/3sec). Likewise, a small 3sec gesture in a less favourable region (e.g., far away in the left or front region) is rated similarly to a large 6sec gesture in the preferred right region. We also note how small 3sec gestures are consistently rated about 0.5 points higher than in Study 4 at the corresponding gesture positions (close and mid distance, compare middle part of Figure 19). Since small 3sec gestures take longer than the gestures in Study 4, where participants only had to draw one quick digit, the higher ratings in Study 5 are reasonable and confirm the robustness of the rating-based approach. We can conclude that both gesture size and duration have a significant influence on how comfortable users feel when performing AD-gestures in public. Most users indicated that they felt comfortable or neutral using small gestures and were less comfortable with large ones.

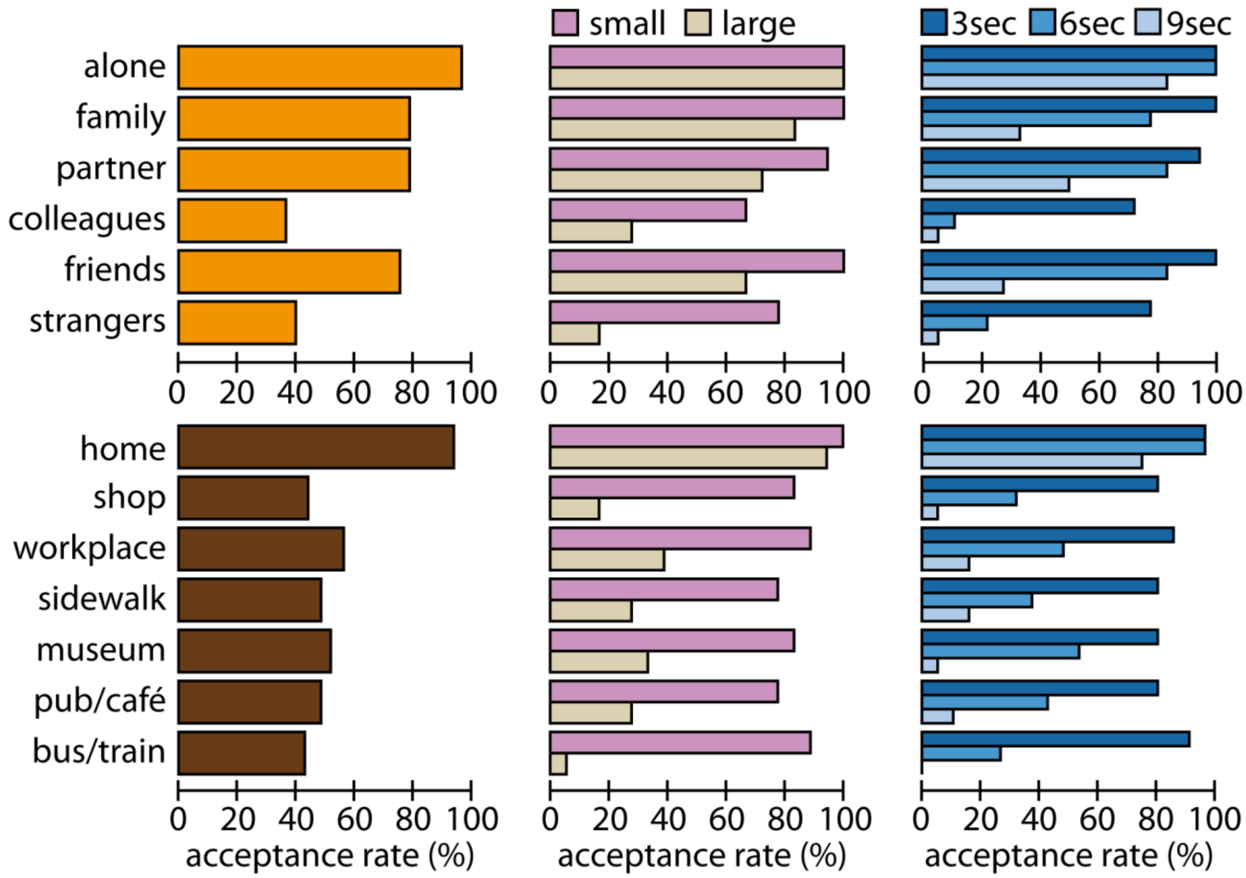

Figure 22: Acceptance rates for audiences (top) and locations (bottom). Rates aggregated across size and duration (left), split by size (middle), and split by duration (right).

Questions 2 and 3: We analyzed acceptance rates for audiences and locations aggregated across gesture sizes and durations. The results are shown in the left part of Figure 22. A Friedman test showed a significant difference between audiences ($\chi^2(5,N=18) = 59.8$, $p < 0.001$) and post-hoc Wilcoxon tests (Bonferroni: α-level 0.05 to 0.003) showed that the least desired audiences – colleagues and strangers – differed from all other audiences, as in Study 4. The rates for the audiences were similar to those in Study 4 (±10%).

A Friedman test showed differences among locations ($\chi^2(5,N=18) = 40.1$, $p < 0.001$) and post-hoc Wilcoxon tests (Bonferroni: α-level 0.05 to 0.002) showed that home differed from all other locations. In Study 4, home only differed from museum and bus/train.

The ratings for gesture sizes and gesture durations are shown in the middle and right parts of Figure 22. Clearly, for private settings (alone and home) size and duration had little or no effect. For familiar audiences (family, partner, and friends) size and duration were more important. For the least familiar audiences, colleagues and strangers, we see pronounced differences with the rates for large, 6sec and 9sec gestures around 25% or below. We also see pronounced differences for all non-private locations. In particular, we note the great difference between 3sec and 9sec gestures in the location ratings: 3sec gestures have an acceptance rate close to, or above, 80% for all locations whereas 9sec gestures have rates below 10% (ignoring the private home setting). We conclude that, indeed, both gesture size and gesture duration have a great influence on how comfortable users feel performing AD-gestures in public places (and to some extent even at home).

Summary

To our knowledge, considering unique gesture features has not been explored in prior acceptability studies. The results of Study 5 further confirm the need to examine the acceptability of gesture features separately. Gesture size and duration both impact the acceptability of AD-interactions. We note that acceptance drops rapidly after the 6-second mark. Furthermore, for all locations and audiences (except home and alone), larger gestures are seen as being less acceptable.

Q1) What were you thinking watching me gesturing around my phone, the way I just did? Select **one or more** items from the list below.

- ☐ I was wondering what you were doing
- ☐ I did not think much about it
- ☐ I thought it was annoying / disturbing
- ☐ I thought it looked fancy / interesting
- ☐ I thought “what a weird behaviour
- ☐ I thought it looked stupid / strang
- ☐ I thought the movements were inappropriate
- ☐ I thought / my impression was:______________________________

Q2) Imagine you could operate your mobile phone using gestures in the air (similar to what I was doing). Now, **in front of whom** do you think you would **feel comfortable** using such gestures? Select **one or more** items from the list below.

- ☐ I would **not** feel comfortable using them even when alone

or

- ☐ when alone
- ☐ in front of my partner
- ☐ in front of friends
- ☐ in front of family
- ☐ in front of colleagues
- ☐ in front of strangers

Q3) Now, **in which locations** do you think you would **feel comfortable** using such gestures? Select **one or more** items from the list below.

- ☐ I would **not** feel comfortable using them no matter where I am

or

- ☐ at home
- ☐ on the sidewalk
- ☐ in a pub, café, or restaurant
- ☐ in a shop
- ☐ in a museum
- ☐ as a passenger on a bus or train
- ☐ at my workplace

Figure 23: Study 6 questionnaire.

Study 6 – Spectators Attitude about AD-gestures

Alongside users’ attitudes towards AD-interactions, the reactions of persons who have watch someone using this interaction is also of interest. We call these observers “spectators”. We are unaware of any previous work on mobile gestures that investigates spectators’ reactions to the public use of gestural interaction. Alongside probing spectators’ reactions, it is also interesting to ask spectators how they think they would feel to use AD-interactions, i.e., their perspective on being a user rather than spectator.

One of the authors acted as an AD-gesture user in five different locations: in a commuter train, in a café, in a library, in a restaurant, and at a birthday party. The author worked through the same image set as was used in the previous study until someone's attention was captured. This spectator was then asked to answer the questions shown in Figure 23. Answers were collected from 24 spectators aged between 17 and 43 years (mean 26.7, s.d. 9.6). Eleven were female and all but one owned a smartphone.

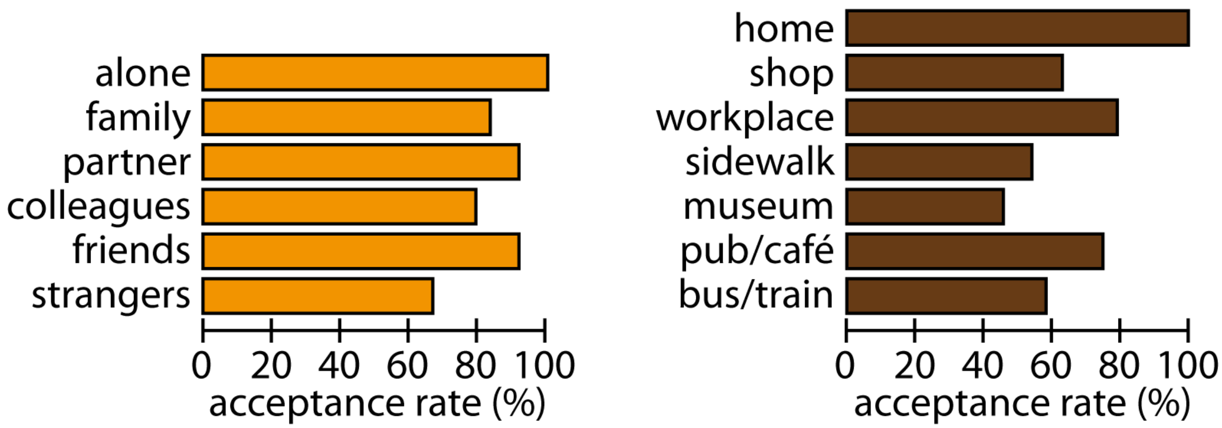

Figure 24: Acceptance rates for audiences and locations.

Results

Question 1: In Question 1 spectators were asked to select one or more statements to describe his/her thoughts when watching the AD-gestures. Twelve spectators (50%) indicated that they became curious, wondering what the user was doing. Twelve indicated that they did not think much about what they had seen and two commented that it looked "cool". One thought that it looked "fancy". Only five spectators thought it was a weird behaviour and one thought it looked stupid or strange. No one thought it was annoying or disturbing. These initial reactions were given by the spectators before they were informed about the purpose of AD-gestures. This suggests that most spectators perceived the gestures in a neutral or curious way. Very few perceived the gestures as something negative or disturbing.

Question 2 and 3: Spectators' acceptance rates, generated from answers to Question 2 and 3 (in front of whom and in what locations they thought they would feel comfortable using AD-gestures), are shown in Figure 24. As with participants in Study 4, all spectators answered that they would comfortably use AD-gestures when alone and 92% said they would feel comfortable using the gestures in front of their partner and friends. The acceptance rate for family, colleagues and strangers amount to 83%, 79%, and 67%, respectively. We found a significant difference between the audiences (Cochran's Q test: $\chi^2(5,N=24) = 15.4$, $p < 0.01$), but post-hoc pairwise McNemar tests with the conservative Bonferroni correction (α-level 0.05 to 0.003) showed no significant differences among pairs of audiences.

When compared to the acceptance rates in Study 4 where Question 2 was answered after a firsthand experience of performing AD-gestures in a public setting, we see markedly higher rates in Study 6 for the least familiar and most critical audiences, colleagues (79% vs. 33%) and strangers (67% vs. 33%). The results for the more familiar audiences are similar in the previous two studies. We also find higher acceptance rates in Study 6 for most locations. Rates for home and shop were the same in both studies, and for the other locations we find higher rates in Study 6. The differences vary between four percentage points (sidewalk) and 25 percentage points (bus/train). A Cochran's Q test showed a significant overall difference among locations ($\chi^2(6,N=24) = 28.3$, $p < 0.0001$). Post-hoc McNemar tests (Bonferroni: α-

level 0.05 to 0.002) showed that the acceptance rate for home was significantly higher than for shop, sidewalk, museum, and bus/train (no other pairwise comparison was significant).

In conclusion, the results from the spectator study indicate that AD-gestures are not likely to be perceived as obtrusive (Q1). None of the spectators we asked thought the gestures were inappropriate or annoying. Indeed, many did not think much about the gesturing they had watched, 50% became interested and/or curious. Furthermore, as with participants in Study 4 and 5, most spectators were quite open to the idea of AD-gestures and thought they would feel comfortable using them in public locations and in front of strangers. However, acceptance rates were generally much higher in Study 6 than in Study 4, indicating a possible over-estimation. A likely reason for this is the absence of an actual usage experience to relate to when providing the answers.

Discussion

To conclude this chapter, we discuss the lessons we learned and insights we gained from the three studies on social acceptance. We also demonstrate how the findings can be applied to existing around-device interactions, reflect on limitations in our approach, and point to directions for future work.

AD-Input Design Considerations & Recommendations

Intuition may provide initial guidance regarding AD-input design, suggesting general directions such that a small or quick gesture is more likely to be acceptable than a large or lengthy one. However, without empirical data, it is difficult to estimate what size is small enough and how great these effects are; when does a gesture start to feel too lengthy; and to correctly predict the consequences of changes regarding such gesture parameters. Results of our experiment provide this concrete emperical data.

First, it is evident from the results that AD-gestures belong to an acceptability-continuum where a combination of several gesture properties influences user perceptions and comfort about performing the gestures in different social contexts. Results demonstrate that users are sensitive to the parameterization of the examined properties – distance from the device, input region, gesture size and gesture duration – and that small differences in parameter settings may result in large shifts on the acceptability-continuum.

The following considerations and design recommendations emerge from our exploration of AD-gesture acceptability:

- *Distance*: AD-gestures that are closer to the device are more acceptable. Results suggest a critical point approximately 30cm away from the device. Input beyond this distance is likely to be considered "socially awkward" and thus should be avoided if possible (the region from the device extending to this point is slightly smaller than the intimate space defined in studies on proxemics [68]). The critical distance is applicable for all tested regions: to the left and right, in front, above, and below the device.

- *Input region*: Results reveal a strong preference for gestures to the right and the front of the device. This suggests that AD-input designers need to consider user handedness (which should be reflected in the operation of the system) and that they should design for input to the dominant side and in the front of the device. However, the regions above, below and to the left of the device are acceptable if distance is manipulated such that gestures are positioned near the device..

• *Size*: When in public, users indicated a strong preference for small gestures. Results indicate that caution is warranted when the gesture size approaches 15×15cm, larger gestures should be avoided.

• *Duration*: Gesture duration strongly affects users' comfort levels. Even after a few seconds of AD-input users are likely to start feeling uncomfortable. Acceptance drops rapidly after the 6-second mark.

• *Gesture property interplays*: The strong interplays we found between gesture properties suggest that AD-interaction designers can achieve socially acceptable designs even when their interactions require less favourable property characteristics. For example, the negative effects of an over-sized gesture can be reduced if the input is allowed very close to the device or in a favourable region. Thus, designers are good advised to carefully consider such interplays and to examine possibilities to encounter critical features by making changes to other gesture properties.

Adaptations to AD-Interactions

Our findings can be directly applied to several existing AD-input techniques. For example, Hoverflow [120] uses a small space (5-7cm) above the device for simple interactions such as to sweep or to rotate an image. Similarly, SideSight [20] uses proximity sensors that are capable of detecting limited space (8cm) along each side of the device. In contrast, results suggest that socially acceptable AD-input space could be larger (30cm) and could be used for complex 3D gestures such as Cyclo [132] for continuous zooming. Such gestures could extend up to 6 seconds in length without impairing users' perception of comfort.

Few AD-techniques utilize the valuable – and acceptable –space below the device. For example, the AD-Binning technique [71] relies on a large 2D space, extending up to 40cm away from the device, to allow users to store, browse and retrieve contents through gestures issued within storage bins that are positioned in AD-space. Results revealed that people feel uncomfortable using far distances for AD-input. This finding diminishes the potential value of AD-Binning. However, results can suggest alternatives that still allow users to benefit from AD-Binning. Using the space above and below the device, we could reorganize bins in a layered structure in a small 3D space. This avoids large reaching distances and might improve the acceptance of AD-interaction.

Improved Methods for Acceptance Studies

Our studies included two new approaches to collecting user opinions related to social acceptance. The first consisted of teasing apart specific gesture features. Whereas in prior work, results would indicate whether a gesture is viewed as either acceptable or not, our approach is to examine unique elements of gestures. This may not be possible with all types of gestural input. However, when the interaction modality affords this, such as with AD-input, we recommend that studies tease these apart. We found that small changes in variables had a large influence on user perception. Furthermore, teasing apart gesture features may reveal new opportunities to improve the acceptability though intelligent combinations or adjustments to the individual parameters.

The second adjustment we included was to ask participants to rate their view of a gesture after having experienced using the gesture in a public setting. Prior work has relied on visual demonstrations of the studied gestures and on participants' imagination of a future usage situation. We found that having a person rate gestures without having had a firsthand usage experience resulted in much higher acceptance rates. Overly positive responses in early design phases may allude to sub-optimal designs that future users may avoid in public settings. However, more targeted methodological research endeavors are needed to systematically disentangle the effects of firsthand usage experiences in acceptability.

We also introduced a new dimension to social acceptability studies by exploring possible linkages between personality traits and user perception. Results did not reveal any relations between the extroversion trait and user perceptions. One explanation might be that people are familiar with mid-air sensing mechanisms, through systems such as gaming consoles. The finding may also be related to the small number of participants used in Study 4. We acknowledge the limitations of the Big-Five personality test. It provides one aspect of a person's traits. Additional work is needed to identify how social acceptability tests can be linked to personality types.

Summary

In this chapter, we have presented three studies that explored the acceptability of hand gesture input in 3D space around a smartphone. The studies were performed in various public locations. We surveyed users that performed such Around-Device gestures and people who passed by about their impressions. Most users and spectators answered they would use such interactions if available on their smartphone, but also indicated they would be concerned about others' reactions. Our results show that people are selective regarding in what public settings they would use gestures. Moreover, gesture properties, such as duration and distance from the device, have a great influence on how comfortable users feel when using Around-Device gestures in public. Acceptance and perceived mental comfort markedly sink if gestures are done further than 30cm away from the device or last longer than 6 seconds. Gesture size and region (e.g., on the side, above or below the device) also matter. According to our findings and study experiences, we presented recommendations for around-device input designers and suggestions about how to improve methods used in studies related to the social acceptance of novel interaction techniques.

Chapter 5: Thumbs-Up: 3D Spatial Thumb-Reachable Space for Around-Device Interaction

In the previous two chapters, we discussed approaches to improve smartphone interaction in situations where a user holds the device with one hand and interacts with the other. With the free hand, users commonly access void in-air spaces located to the right, front or left to the device. However, we found interesting results while using around device space on the left region (for right-handed users). Results from our study 1 show that positioning items to the left of the device takes longer due to the screen occlusion caused by the hand and lower arm. Additionally, study 4 reveals that gesturing left to the device is socially uncomfortable. Therefore, in this chapter, we continue our exploration to find a complementary input region (to the large around-device space) that can be easily accessible to support around-device activities. Additionally, we are interested in exploring a mid-air space that is closer to the touch screen, as this could be used in conjunction with the touch input to design new interaction capabilities.

Around-device interaction has been widely explored with in-air spaces located to the right, front or left to the device. Ways to leverage the void space just above the device which is reachable with the thumb of the hand holding the device is an unexplored interaction space. In this chapter, we present Thumbs-Up, the in-air space next to a smartphone that can be reached with the thumb on the hand holding the device. With Thumbs-Up, the in-air space could be used to access on-screen content through directly pointing with the thumb, as shown in Figure 25. Additionally, this surrounding space could be used in conjunction with current touch input to extend the input vocabulary.

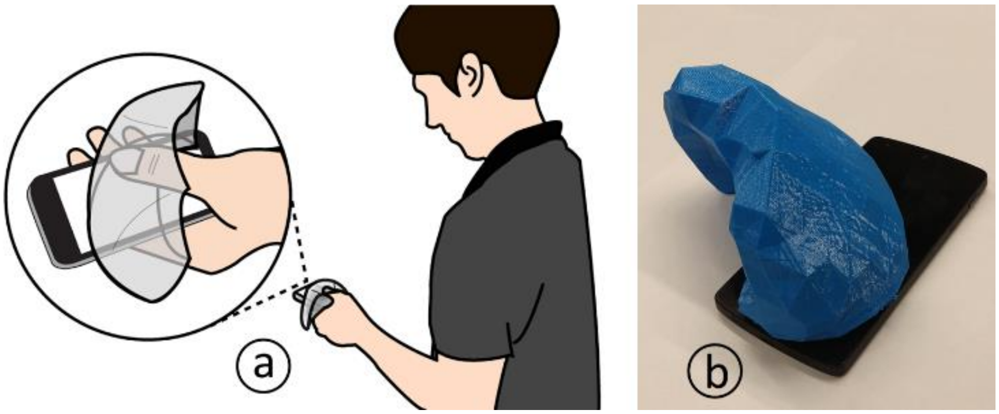

Figure 25: (a) Thumbs-Up interactions occur in the thumb-reachable in-air space around a smartphone. A 3D printed shape of this thumb-reachable space is shown in (b).

In this chapter, we start with exploring the Thumbs-Up input range. Results reveal that study participants can reach a large in-air space (up to 74mm from the touchscreen) with their thumb. With a user study, we next explore important design factors, such as selection methods and the item arrangement strategies in the Thumbs-Up space.

Study 7: Thumbs-Up Input Range

People hold and provide screen input on their smartphones in many different ways, depending on the current task (e.g., texting vs. surfing), interaction (e.g., pinch-zooming vs. scrolling), and context (e.g., sitting vs. walking and carrying a bag). We focus the common situation (out of many) where the user holds the device in the non-preferred hand and uses the index finger on the preferred hand for input. In this constellation, the thumb of the non-preferred hand, which holds the device, is mostly passive and rests along the side of the device. To see whether the passive thumb can be used for in-air input we first elicited the dimensions of the in-air volume that people can reach with the thumb when holding the device in the same hand.

We start with a study to identify the suitable thumb input range when the user holds the smartphone in the hand. Previous work [28] that has studied two-handed usage situations (where one hand holds the smartphone and the fingers on the other hand, are used for on-screen input) shows that users frequently move or rest their input fingers in the air just above the screen between on-screen interactions. This observation motivated us to distinguish between: (i) TouchSpace: the above-screen space that people use to initiate or terminate touch gestures when these are performed with the thumb; and (ii) Thumbs-Up space: the in-air space around the device that is reachable with the thumb on the hand that is holding the device.

Participants and Apparatus

Twelve right-handed smartphone owners participated (three female, mean age 24.3 years, s.d. 5.8). All participants preferred to hold and interact with the phone in portrait mode. Participants' thumbs were on average 105mm long (from the carpometacarpal joint to the tip of the thumb). On average, their palm circumference, hand size (from the tip of middle finger to bottom of the palm), and hand span measured 183mm, 184mm, and 200mm, respectively.

We used a Vicon MX system to track participants' thumb movements around a Samsung Mini S4 (screen size: 4.3", dimensions: 4.91×2.41×0.35"). We placed tracking markers on the smartphone and on a 3D-printed ring to track participants' thumbs (Figure 26). A Unity 5.0 application logged thumb movements in relation to the smartphone.

Tasks

Participants were instructed to hold the phone in their left/right hand with the left-bottom/right-bottom corner of the phone close to the centre of the palm (Figure 26). Participants performed two tasks:

Figure 26: (a) Map navigation and (b) Space filling task.

(i) Map navigation task: We asked participants to navigate a map using flick gestures with their thumb. We showed participants two familiar locations in the city in the Google Maps application and asked them to navigate from the first to the second location. To perform this kind of task, a user has to repeatedly tap the screen and then flick to pan the map. After a panning action, the user needs to readjust the thumb to start the next panning action. We included this task to identify the TouchSpace zone above the screen – the in-air space where the thumb moves after an on-screen thumb-operation.

(ii) Space filling task: We asked participants to repeatedly move the thumb in mid-air, from left-to-right and right-to-left above the screen and thereby gradually increasing the distance between the top of the thumb and the screen. We asked participants to do so until the top of the thumb had reached the maximum distance that could be managed without having to adjust the position of the device in the hand. These repetitive in-air movements generated a large thumb-reachable volume, which we refer to as Thumbs-Up space.

Participants performed the tasks standing in a room equipped with the motion tracking system. We asked participants to imagine that one of their hands is holding a coffee mug and the only the other hand is available for on-screen interactions. All participants performed the two tasks with the right and with the left thumb. The tasks were performed three times with each thumb.

Results

TouchSpace: We found that during the map navigation task, participants moved the left thumb up to a maximum of 22mm above the screen (average 18mm, s.d. 2.7mm). The right thumb was moved to a maximum of 23mm above the screen (average 19mm, s.d. 2.3mm). Accordingly, we reserve this TouchSpace for on-screen gestures only (Figure 27a and b).

Thumbs-Up space, height: The recordings from the space filling task revealed that participants could comfortably reach a maximum of 64mm above the screen with the left thumb (average 57mm, s.d. 5.8mm) and a maximum of 74mm (average 63mm, s.d. 9.6mm) with the right thumb. Accordingly, we can consider using this space (22 to 57mm above the screen for the left-hand thumb and 23 to 63mm for the right-hand thumb) for Thumbs-Up interactions.

Thumbs-Up space, width: When regarding the phone's bottom edge as the horizontal axis, participants moved the left thumb sideways within an arc spanning from 10° to 150°. With the right thumb, the movements were within an arc spanning from and 0° and 150°. With both thumbs, the corresponding arc length decreased as the thumb's height-distance from the screen increased. Our results indicate that people can comfortably reach a relatively large in-air region above and beside the device with the thumbs. Figure 27 c, d, e, and f visualize the corresponding accessible in-air space for the left and right thumbs.

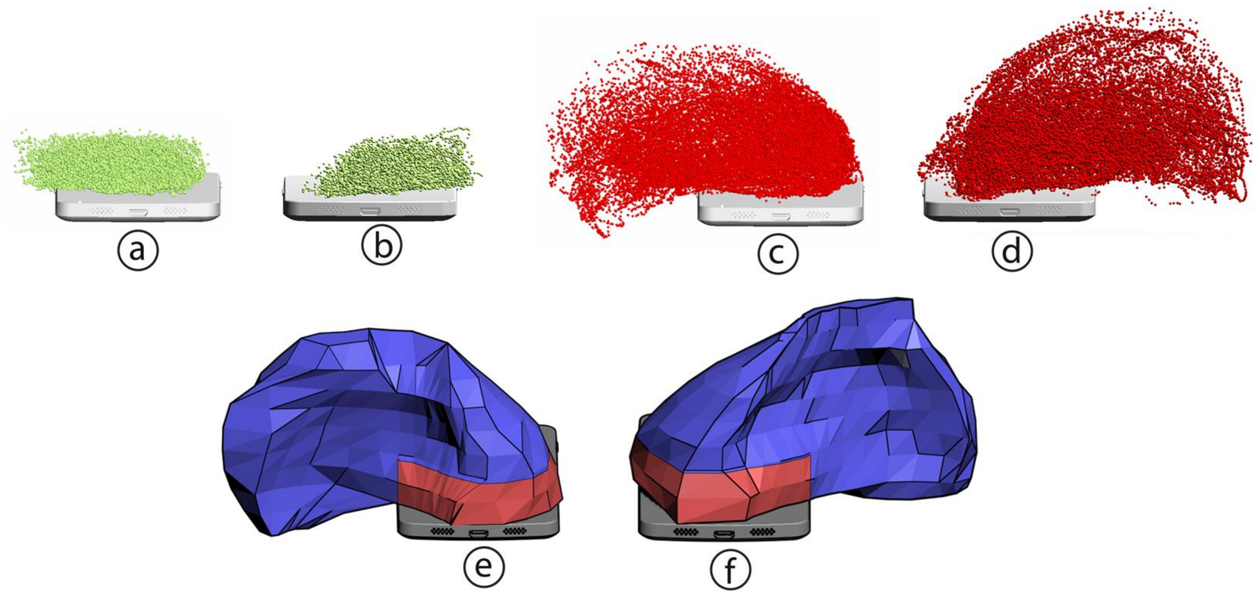

Figure 27: TouchSpace (a) the left and (b) the right thumb. Thumbs-Up space with (c) the left and (d) the right thumb. In-air thumb-reachable space with (e) the left and (f) the right thumb. The red areas indicate TouchSpace and the blue areas represent Thumbs-Up space.

We note that the in-air space for the right thumb is slightly larger than for the left thumb. On average it spans up to 64mm above the device with the left thumb and up to 74mm with the right thumb. We also observe that the region for the right thumb is larger in the horizontal direction than for the left thumb. We attribute these differences to the fact that all participants were right handed and frequently use their right thumb for one-handed smartphone interaction. This frequent usage provides them the flexibility to reach regions which are less intuitive/comfortable with the left thumb. A minor concern worth mentioning is the less diverse participants used in our study. However, we believe that more participants with balanced gender and handedness would further ascertain our findings.

Study 8: Selection and Item Placement Style on Thumbs-Up Space

After having defined the accessible in-air thumbs-up region, we explored ways to utilize this in-air space. We envision this in-air thumb-space as being useful for triggering commands and for storing, selecting and browsing information items (similar to Hasan et al.'s [71] AD-Binning concept). Accordingly, we explored ways to arrange items (i.e., 'in-air buttons') within the accessible thumb-space and methods to select such in-air items. After pilot testing with different numbers of items and ways to arrange items inside the thumb-space we found that horizontally arrange a maximum of 10 to 12 wedge-shaped items, in either two or three layers (Figure 28a and b), seems reasonable. Through pilots, we also arrived at three promising selection methods: Dwell, Tilt, and Touch (Figure 29a, b, and c).

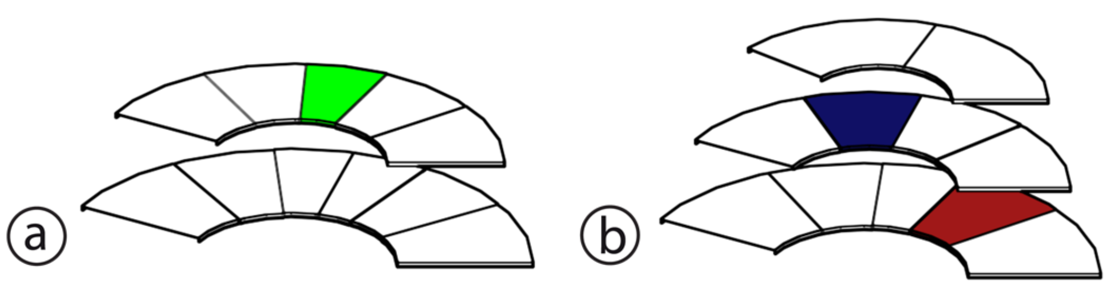

Figure 28. a) On screen visualization with a colored cursor to indicate the position of the thumb inside the in-air region in relation to the items, arranged in two and b) three layers.

To trigger a selection with Dwell, the user keeps the thumb still for 600ms within the desired item's region. To select with Tilt, the user positions the thumb within the item's desired region and then quickly tilts (or rolls) the device sideways (we use a relative 30° threshold). The Touch method involves moving the thumb to the desired in-air region and then quickly tapping down with the thumb, anywhere, on the screen. We use a threshold to resolve situations where the thumb passes through undesired items on its way to the screen (e.g., when selecting an item in an upper layer): we ignore any items which the thumb visited for less than 250ms and treat the most recent item with a visit time greater than 250ms as the selected target.

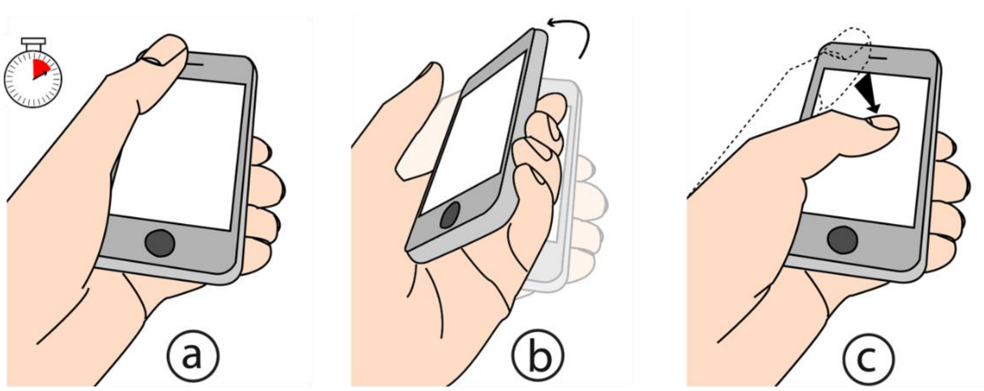

Figure 29. Selection methods. (a) Dwell: requires users to keep their finger over an item for a 600ms to trigger a selection; (b) Tilt: rotating the device clockwise or counter-clockwise direction with a certain angle triggers selection; and (c) Touch: a quick tap on the screen from the mid-air item position invokes the selection.

In our work, we assume that robust finger tracking in 3D space will be possible with future smartphones (as indicated by strong efforts in both industries, e.g., [127,156,184] and academia, e.g., [20,70,120]). We used a Vicon MX system to emulate a device with around-device tracking capabilities in our studies. The system tracks markers attached to a Google Nexus 5 phone (4.95-inch screen, 1080×1920 pixels) and to the user's left and right thumbs and right-hand index finger. Position data is sent from the Vicon system to an application (Unity + Android) on the smartphone, which interprets the data as user input and reacts with corresponding output.

To inform our design choices for the AirPanes design, we studied users' performance with the three selection methods and the two different item arrangements in a 3×2 within-subjects experiment. Twelve right-handed smart-phone owners participated (3 female; 9 male; no participant had previously participated in any pilot or study; mean age 25.2 years, s.d. 5.5).

Task

A start button in the middle and a 1.5×1.5cm visualization of the thumb-space in the top-right corner of the screen are displayed at the trial start. A tap with the tracked left thumb on the button starts timing. The target item for the trial is colored red in the visualization and a blue position cursor indicates the current location of the thumb within the thumb-space. The cursor turns green when the thumb enters the target item (Figure 28a). The timing for the trial ends and the start screen for the next trial is loaded when the system detects a correctly issued selection from within the target item. An error is recorded (timing continues) if the participant performs a selection action from a non-target item.

All participants performed two series of six blocks of trials with each of the three selection methods, one series with the 2-Layer and one series with the 3-Layer arrangement. Each block contained one trial for each of the 11 item positions (presented in random order within a block). Participants were divided into six pairs, one pair for each of the six possible presentation orders of selection methods. One participant in each pair always started with the series for the 2-Layer arrangement, one started with the 3-Layer arrangement. Each participant performed a total of 396 timed trials: 3 methods × 2 arrangements × 6 blocks × 11 item positions. A study session lasted 45 min (including instructions, practice, and breaks).

Results

Error trials: In 253 of the 4,752 trials (5.32%) participants issued one or more correct selection actions in a non-target item before correctly selecting the target item. In 173 trials (3.64%) one erroneous selection was made, in 80 trials (1.68%) more than one erroneous selection was made, at most 5 in a trial. In total, 5,115 correct selection actions were registered, on average 1.08 per trial. Of the 5,115 selections, 363 selections were on a non-target item: 99 with Touch, 130 with Dwell, and 134 with Tilt, resulting in an overall error rate of 7.10% (Touch 5.88%, Dwell 7.47%, Tilt 7.80%). A Friedman Test showed no significant difference between the three techniques and a Wilcoxon Signed-Rank test showed no difference between the two arrangements (192 erroneous selections with the 2-Layer layout, 171 with 3-Layer layout).

Figure 30 shows the number of erroneous selections and error rates for each item position in the 2-Layer and 3-Layer layouts (averaged across technique). Although the item size was slightly smaller in the middle of the layers, we see a trend with more erroneous selections for the slightly larger items at the right and left side. A possible reason for this is that the outer items require stretching the thumb more. Possible solutions to resolve this issue could be to either use a fisheye space discretization [71,173] that provides more space to the currently 'touched' item or to use a more extreme angle dependent discretization to assign even more space to outer items.

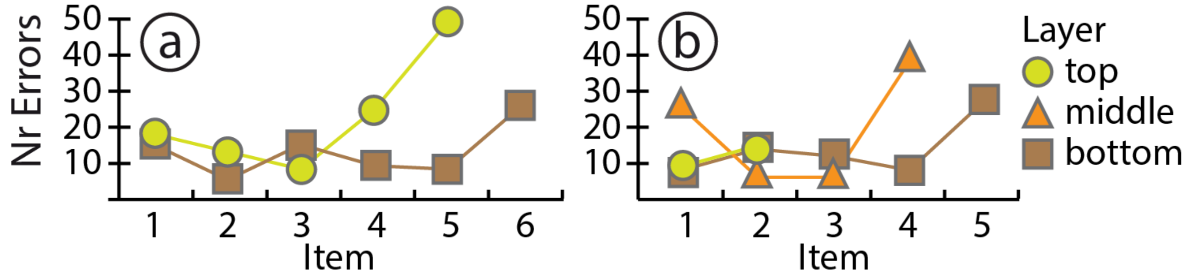

Figure 30. Number of errors for item positions from left to right in each layer with (a) 2-Layer and (b) 3-Layer layout.

Selection time: We only use error-free trials (4,499) to analyze selection time. We take the median trial time for each participant × technique × layout × position combination (i.e., using the median of one to six trials for each participant in each combination, depending on how many error trials the participant did in that particular combination). This time data was right skewed and we performed a logarithmic transformation (which resulted in a distribution close to normal) before analyzing the data.

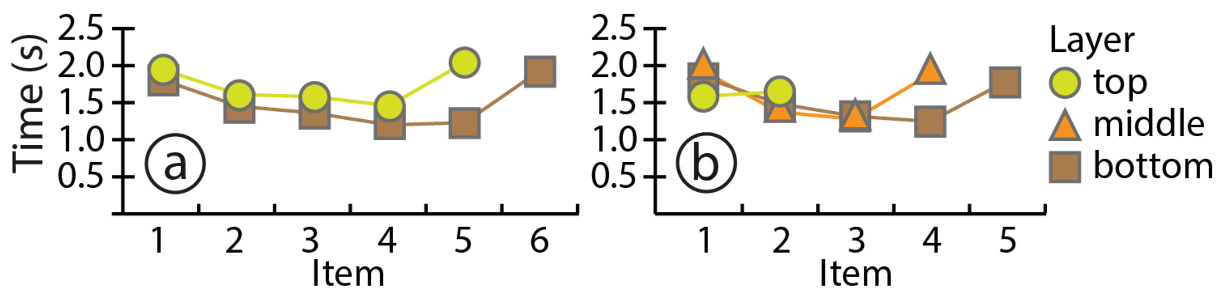

Figure 31. Geometric mean time for item positions from left to right in each layer with (a) 2-Layer and (b) 3-Layer layout.

The geometric mean selection time (i.e., the antilog of the mean of the log-transformed data) was similar for the three techniques: Touch 1.53s, Tilt 1.56ms, and Dwell 1.60ms. There was also no marked difference between the two layouts: 2-Layer 1.58s, 3-Layer 1.56s. A two-way RM-ANOVA (*technique*, *layout*) showed no significant main or interaction effect. Figure 31 shows the geometric mean selection time for each item position in the two layouts. Two separate one-way ANOVAs, one for each layout, showed that there was a statistically significant effect of item position on selection time (2-Layer: $F_{10,121} = 9.305$, $p < 0.0001$, $\eta^2 = 0.44$; 3-Layer: $F_{10,121} = 6.668$, $p < 0.0001$, $\eta^2 = 0.36$). Again, as with error trials, we see a trend with more difficulties with the outer items. Bonferroni adjusted post-hoc pairwise comparisons confirmed: the inner items, Item 3, 4, and 5, in the bottom layer in the 2-Layer layout were faster to select than the two outmost items in both layers; Item 4 in the bottom layer in the 3-Layer layout was faster to select than the two outmost items in the middle layer (all p's < 0.001).

Summary

We conclude that regarding both errors and selection time, the three selection methods and the two arrangements score about equally well. The somewhat more error and slightly longer selection times for item positions at the outer ends in a layer indicate a trend suggesting that item positions at the far ends tend to be slightly more troublesome to select than items in the middle. The results regarding the three selection methods are promising and suggest that all of our three methods are suitable for in-air thumb selections, which gives the designer the flexibility to choose method as appropriate according to other factors, such as usage context or application type.

Design Recommendations

Our investigation offers the following recommendations regarding using thumb-reachable space for around-device interactions:

Thumb-space usage: When complemented with in-air index finger input, the in-air thumb-space is well suited for input. It works best suited for short interactions involving a limited number of items. We suggest using this space to provide access to ten or fewer items, such as frequently used phone contacts or to quickly load recently visited websites.

Non-tilt based pane layout: Several participants in Study 8 reported having felt less comfortable with the tilt-based pane layout style as it required rotating the smartphone to switch between sets of panes. This movement also caused a disturbing change in viewing angel.

With insights about viable item arrangements and selection methods for in-air thumb space, we move on and explore how to use this space to support a full application that usage multiple around-device spaces.

Summary

In this chapter, we have presented Thumbs-Up space, the in-air space that is reachable with the thumb of the hand holding a smartphone. We found that, on average, the in-air space spans up to 64mm and 74mm above the device with the left thumb and right thumb respectively. With a user study, we have also examined several factors, such as selection mechanism and item placement strategies that are crucial designing interfaces with this space. Our results revealed promising selection and item placement techniques that are suitable for interacting with the Thumbs-Up space.

Chapter 6: AirPanes: Two-Handed Around-Device Interaction for Pane Switching on Smartphones

Results from our previous studies revealed that using around-device space for accessing items is faster than using traditional touch screen. We also found that users welcomed the idea of using this space with mid-air gestures in different social contexts. While such results laid the foundation for using in-air input for accessing mobile content, a complete mobile application that deploys and benefits from such an input modality has yet to be demonstrated. Accordingly, the final goal of this thesis is to explore how in-air input can enhance user performance in more complete mobile scenarios in which they are required to interact with different features of an application to make a decision. In particular, we focus on the ability to facilitate a complex analytic task with around-device interaction that requires users to cycle through a number of panes/views of an application (such as a filter pane for refining search criteria, a list pane for browsing filtered results, or a bookmarked pane to revisit bookmarked items) before making a final purchase decision.

In this chapter, we propose AirPanes, a novel strategy that leverages around-device space to structure views/panes of a mobile application and facilitates users to interact with them to carry out analytic tasks on a mobile device. We pick mobile commerce (i.e., m-commerce) applications as a scenario for designing AirPanes as millions of users use these services to purchase products online [139]. We demonstrate the benefits of AirPanes in a scenario where the user browses products, applies filters, inspects result lists, and bookmarks interesting items before making a final purchase.

In this chapter, we first distill the interface components common in interfaces of m-commerce services. Our initial investigation indicates that m-commerce interfaces use multiple panes to support major functionalities, such as filtering, browsing results and adding to favorites. We use this knowledge to identify how best to organize panes in mid-air around the device. We optimize design parameters for AirPane interactions necessary for a complete application, i.e., to select items, to examine details, to interact with filter controls, and to switch between views. We then evaluate AirPanes against a user interface representative of m-commerce app interfaces and find that AirPanes is on average 50% more efficient than the common touch interaction.

Background

We selected an m-commerce application to showcase the use of around-device interactions. Millions of consumers purchase products online using their smartphone. The number of online purchases from smartphones has steadily increased over the past years and the growing trend is expected to continue [139,140]. An m-commerce scenario is also suitable for our demonstration purpose as it involves a complex high-level analytic task (making a purchase decision) which requires the user to perform numerous subtasks, e.g., searching product categories, applying filters to narrow down searches, inspecting alternative products, adding products to a shortlist, and browsing and comparing shortlisted items before making the final purchase decision.

We surveyed numerous interfaces of major m-commerce services to learn about common user interface (UI) structures and functionality. We consulted Alexa traffic ranking [3] to find the 250 most visited websites worldwide. From these, we listed the top 24 websites with a retail service. Two companies were listed multiple times: Amazon with eight, and eBay with three different top-level domains. We decided to only include the highest ranked version for each (the other versions use the same UI structure). Accordingly, we arrived at 24-7-2 = 15 distinct services: Amazon[7], BestBuy[17], Walmart[196], Ebay[36], FlipKart[46], Alibaba[5], Tmall[190], Taobao[187], AliExpress[6], Homedepot[189], IKEA[92], Target[188], Rakuten[159], Jingdong Mall[98] and ETSY[44] also listed in Figure 32.

We used a Google Nexus 5 smartphone (4.95-inch screen, 1080×1920 pixels) and the Google Chrome browser to analyze the websites (all were responsive websites in which the content presentation depends on the client's display size). We also installed the corresponding mobile app versions from Google Play [55] (accordingly, we analyzed 15×2 interfaces). We focused our analysis on the provided shopping-related functionality and on how this functionality is organized within each interface.

All studied interfaces use similar UI mechanisms and UI structures to provide the necessary functionality, as summarized in Figure 32a. A text field is used for product search. Filter functionality is accessed through a "filter" button (most often positioned close to the search functionality) which opens a separate pane that displays the provided filter options (e.g., product categories, brand, price, color, and customer rating). Search results (filtered or unfiltered) are displayed below the search box in a scrollable overview list (Figure 32b), or in a grid. A thumbnail image and general product information are provided for each item in the list. A tap on an item opens up a new pane which shows further details about the product (Figure 32c). A tap on a "back" button displays the overview list again. From the detail pane, the user can add the item to the shopping cart (and/or favorite list) with a tap on an "add" button.

In informal tests, performed during our survey, we quickly observed that this UI structure forces users to repeatedly switch back and forth between different panes, e.g., between the main view and option view when applying and changing filter options, and between the overview list and detail views to access full information about the listed items.

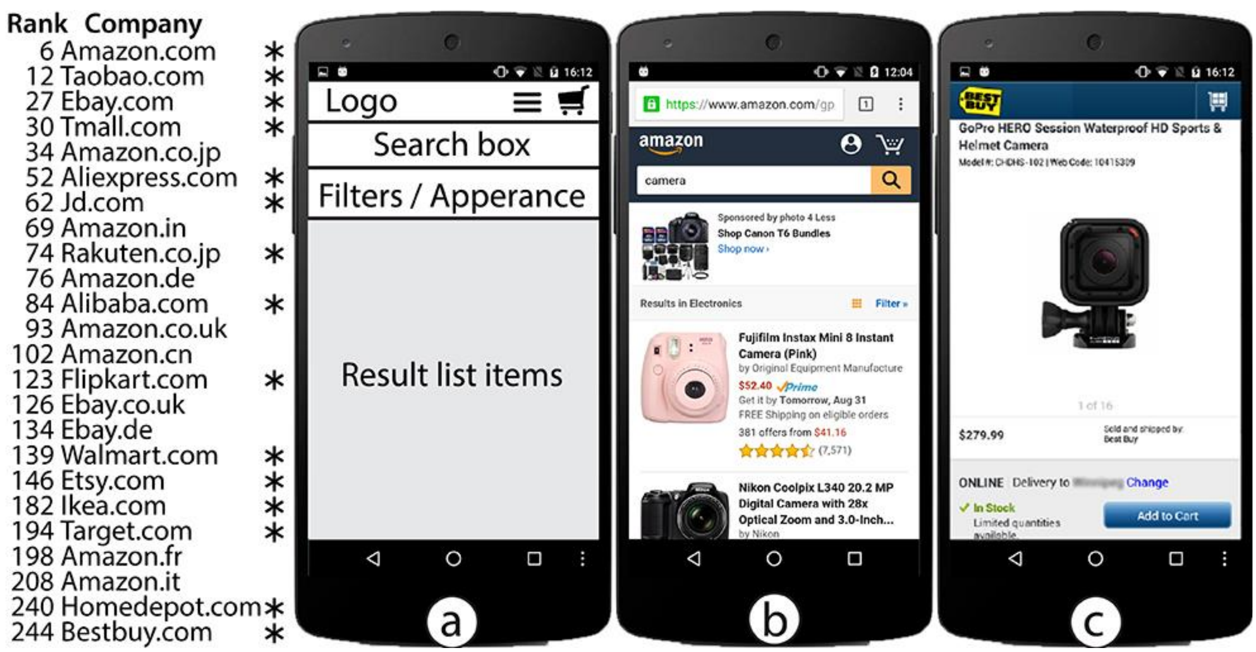

Figure 32: The 24 most visited m-commerce services (top 15 unique providers marked with a *). a) General UI structure for the surveyed interfaces. b) example responsive website's overview list and c) a detail view in a mobile app.

We note that m-commerce interfaces are not the only mobile apps using a multiple-pane structure, which imposes frequent view switching. We also observe this UI structure (and extensive view switching) in many other types of interfaces that are designed for smartphones. Examples include email clients (with separate panes to show the folder hierarchy, lists of email headers, and the corresponding email text), online video repositories (with separate panes for e.g., query purposes, result lists, and video viewing), messaging apps (with separate panes for e.g., contact lists, contact details, lists of conversation threads, and to enter messages), and interfaces to manage personal photo or music collections.

With the insight that all major m-commerce providers' interfaces use a multi-pane UI structure, which requires extensive view switching (and that the functionality of many other smartphone interfaces are organized in a similar way), we decided to focus our exploration on how to utilize around-device space for view switching purposes. Furthermore, given that no earlier project has presented in-air interactions that involve using in-air space reachable via both the left and the right hand, we were also interested in exploring those purposes for which we could capitalize on two-handed around-device input.

Study 9: AirPane Layout and Pane Switching

As we noted in our survey on m-commerce interfaces, the small display size used on smartphones forces designers to structure their interfaces in a multiple-pane. With such a structure the user has to repeatedly tap on small UI buttons (or on other interface elements, such as entries in scrollable lists) to switch views. Our AirPane approach is based on off-loading panes into around-device space. With panes residing in-air, the user can quickly switch between panes by simply moving the in-air input finger inside the desired pane's in-air area – and the pane's content is displayed on the screen. We also envision that the user interacts with pane content directly from in-air space using finger gestures.

We explored five different ways to arrange four panes in around-device space and the use of an in-air pinch-gesture to select pane content.

Stacked layout (Figure 33a): The four panes are stacked on top of each other to the right of the smartphone, two above the device, two below the device. Each pane's bottom-left x-y position aligns with the device's bottom-right x-y position. If the user moves the device, the panes follow to maintain their position relative to the device. With previous work [2] showing that social acceptability related concerns are raised for around-device gestures taking place beyond 30cm from the device, we limit our panes to a 30×30cm area. The screen's 1080×1776 pixels are mapped to 1080×1776 in-air 'pixels', each approx. 0.028×0.017cm large. The two middle pane are each 15cm high; one of our earlier studies showed this gives the user enough vertical space to perform in-air gestures inside a pane without accidentally entering (and so switching to) an adjacent pane. The top and the bottom panes extend infinitely in the upward respectively downward directions.

Tiled layout (Figure 33b): The four panes share the space used for a single pane in the Stacked layout. This requires less in-air hand movements for pane switching, but more precise movements when interacting with pane content. Each pane measures 15×15cm and extends infinitely upward and downward. The corresponding in-air 'pixel' measures approx. 0.014×0.008cm.

Tilt-Stacked layout (Figure 33c): Similar to the Stacked layout, each pane uses a 30×30cm area (the top pane extends infinitely upwards, the bottom pane downwards, in-air 'pixels' measure approx. 0.028×0.017cm). The four panes are grouped in two pairs, only one pair is accessible at a time. When the device is tilted at an angle less than 45°, the user can switch between the first pair of panes by moving the finger vertically above or below the device. When the device is titled at an angle more than 45°, the other pair is accessible. This provides a larger area for each pane and requires less vertical movements than with the Stacked layout, but introduces a tilting action.

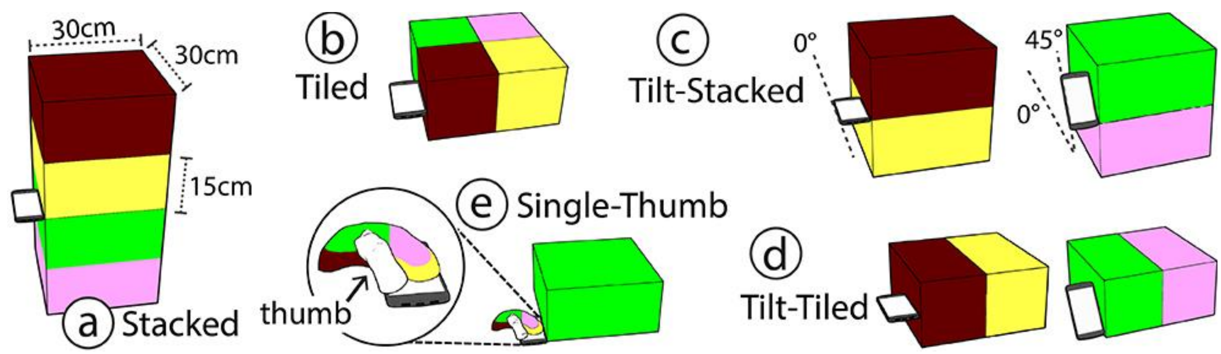

Figure 33. Five AirPane layouts.

Tilt-Tiled layout (Figure 33d): Again, panes are grouped in pairs and a tilt angel is used to switch between pairs. Each pane uses a 15×30cm area (extended infinitely upward and downward, in-air 'pixels' measure approx. 0.014×0.017cm). Horizontal in-air movements are used to switch between panes in a pair.

Single-Thumb layout (Figure 33e): Only one 30×30cm pane (extending infinitely upwards and downwards, in-air 'pixels' measure approx. 0.028×0.017cm) is accessible at a time beside the device. The user switches between panes with in-air thumb selections using the left thumb. The in-air thumb-space is divided into four approximately equally sized regions arranged in a two layers (cf. inset Figure 33e). Moving the thumb into another area and selecting the Touch method from Study 8 switches to the selected pane.

To inform our design choices regarding pane switching for our m-commerce demonstration interface, we studied 15 right-handed smartphone owners' performance with the five AirPane layouts (2 female, no participant had previously participated in any pilot or earlier study; mean age 21.1, s.d. 1.8). We used a low-level task where participants had to switch between four panes and select items in the panes.

Task

A trial consists of selecting first a blue and then a yellow 'target item', positioned in two different panes. The other two panes contain one red distractor item each. The squared items are positioned at random positions inside the pane. The blue target is labeled "1", the yellow target is labeled "2". The panes are numbered 1 to 4 (to provide a clear visual identifier for each pane). A black circular cursor on the screen provides feedback about the position of the index finger inside the active in-air pane. Figure 34a and b visualize a selection of the yellow target item; when the cursor enters the target, the target is highlighted in green. With the cursor inside the target, the participant separates the pinched thumb and index finger by at least 1cm to invoke a selection. After successfully selecting the first target the participant proceeds to find (i.e., switch to) the pane with the yellow target. Successfully selecting the yellow target ends timing and shows the start screen for the next trial (Figure 34c). An on-screen tap on the "start" button starts timing for the next trial. Selections in distractor items are ignored. Selections outside items or a selection of the yellow target before the blue target are also ignored, but the trial is marked as an error trial. Marked trials are re-queued at a random position among unfinished trials within a block of trials.

The in-air space available for an item inside a pane depends on the used pane layout. This suggests a possible trade-off between the ease and speed with which a user can switch panes and the ease and speed with which in-air items can be selected. Accordingly, we used three different on-screen item sizes to investigate such a possible trade-off: 100×100, 200×200, and 300×300 pixels.

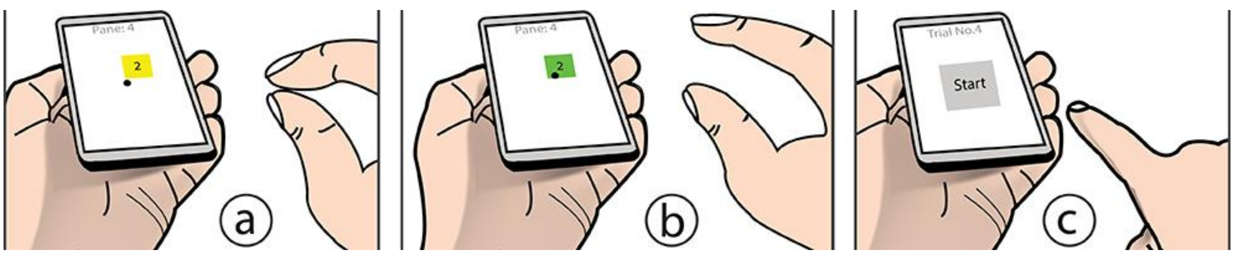

Figure 34. a) A participant moves the 'pinched' in-air fingers to steer the on-screen cursor over the target. b) Releasing the pinch with the cursor over the target selects it. c) Start screen.

Each participant completed a block of five (error free) trials for each of the 5×3 layout-size combinations. Participants had six practice trials with each layout before they started with the first block of timed trials. The presentation order of the five layouts was counterbalanced between participants and the three item sizes were presented in a random order within each layout. We instructed participants to finish each trial as quickly and accurately as possible. Participation lasted 45 min (including instruction, practice trials, and breaks).

Results

Error trials: 135 trials were marked with an error. These trials were about equally distributed across participants, layouts and target sizes.

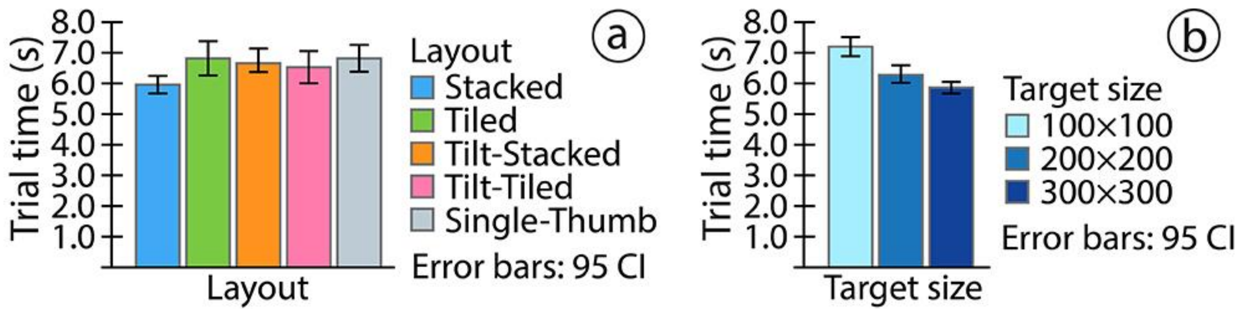

Figure 35. Geometric mean trial times, a) layout, b) target size.

Trial time: For the error free trials we use the median trial time for each layout × target size combination for each participant (i.e., 15 trials). This data was right skewed and we performed a logarithmic transformation (which resulted in a distribution close to normal) before analyzing the data. Figure 35 shows the geometric mean selection time for the five layouts and for the three target sizes. A two-way RM-ANOVA showed significant main effects for both layout ($F_{4,56} = 2.70$, $p < 0.05$, $\eta^2 = 0.16$) and target size ($F_{2,28} = 50.7$, $p < 0.0001$, $\eta^2 = 0.78$) but not a significant interaction effect. Bonferroni corrected post-hoc pairwise comparisons of layouts showed that the Stacked layout (5.96s) was significantly faster than all other layouts (all p's < 0.005), which did not differ and were 10.7% slower (or more): Tiled 6.85s, Tilt-Stacked 6.64s, Tilt-Tiled 6.60s, Single-Thumb 6.68s.

As expected, large items are faster to select than smaller (Figure 35b, 100×100 7.40s, 200×200 6.44s, and 300×300 5.97s). Bonferroni corrected post-hoc pairwise comparisons of the three target sizes showed that each pairwise comparison was significant (all p's < 0.016) with the larger target in each comparison being faster than the smaller.

Summary: Our results suggest that the Stacked layout is fast since no secondary activity is needed to switch to another pane, such as tilting the device or moving a second hand, as with Single-Thumb. A vertical hand movement is enough with Stacked. Many participants commented on the convenience and ease with moving the in-air hand in vertical directions, i.e., up and down for switching panes, compared to moving the hand toward or away from the body. A few participants mentioned that they had slight difficulties getting used to tilting the device, especially in early trials, which may explain the somewhat inferior performance with the tilt-based layouts. Our results also indicate that the trial time is strongly influenced by target size. Naturally, large items are easier and faster to select than small items. The non-significant layout×size interaction effect indicates a consistency between all layouts, with no layout suffering more (or less) due to too small in-air targets.

With insights about effective switching mechanisms and layout of panes in around-device space we now describe our m-commerce interface that utilizes AirPanes and two-handed interaction. We then present our final study in this thesis where we compared users' performance with the AirPanes interface to their performance with a classic touch-based m-commerce interface in a typical m-commerce scenario.

Study 10: Using AirPanes in an m-commerce interface

To recap, our analysis of popular m-commerce interfaces revealed five central features: 1) product search, 2) overview list of the search results, 3) detail view for the products in the search results, 4) filter function, and 5) bookmarking (temporary storage) of interesting products. We incorporated these features in our application using in-air panes. As conventional m-commerce interfaces, we use an on-screen text field and a button for the user to issue a product query. For example, the user wants to buy a camera. Now, instead of displaying overview information for each camera in the result set in a scrollable list on the screen, we off-load the result items onto a 30×30cm in-air pane to the right of the smartphone. We use the best performing Stacked layout from Study 9, and provide access to filter functionality in a second 30×30cm pane below the results pane. The user can switch between these two panes by moving the in-air hand up and down, crossing an imaginary horizontal plane (as defined by the smartphone's touch-surface).

As soon as the in-air finger enters the results pane with camera items, the user sees full information about one of the camera items on the screen, as depicted in Figure 36a. What camera is shown depends on the position of the in-air finger. The camera items in the results pane are arranged in a n×m matrix with equally sized cells filling out the pane. The size of the cells depends on the number of items in the result set. An on-screen visualization with a blue cursor (inset Figure 36a) provides feedback about the in-air finger's current position within the pane and the matrix.

When the in-air finger enters the filter pane, below the results pane, the user sees filter options on the screen, as depicted in Figure 36b. The filter pane uses a Tiled layout that divides the pane in four 15×15cm large sub-panes, one for four different filter categories (Brand, User ranking, Pixels, Price). Again, the finger's position within the pane is visualized in an on-screen overview (inset Figure 36b) and the current sub-pane is marked with a blue border. The available filter options in each sub-pane are arranged top-to-bottom in the in-air sub-pane (as on the screen). Each option in a category spans horizontally across the whole in-air sub-pane (15cm), the number of options determine the 'height' of the in-air area used for each option (e.g., with five options, each in-air area measures 15×3cm). If more than five filter options are available within a filter category, as with the "Brand" category in Figure 36b, the can access these options by moving the in-air finger below the filter pane into a new 30×30 pane where the additional options are available.

The user issues a pinch gesture to select filter options. The pinch gesture is also used when browsing result items inside the top-most in-air pane. With a pinch inside a camera item, the item is added to a list of favorites for later quick access and further inspection and a check mark is shown on the screen (a second pinch inside the same item removes it from the list and the check mark from the screen). The list of bookmarked items can be inspected using the in-air thumb-space (as in Study 7). To activate the thumb-space the user moves the right hand outside any in-air pane (e.g., moving the hand to the right thigh) and lifts the left thumb into thumb-space. Favorite items are arranged in layers inside thumb-space, as Studied in Study 8. The number of layers and the size of the items are dynamically adjusted according to the number of items that has been added to the list. When the thumb enters an item the corresponding camera information is displayed on the screen. An on-screen visualization informs about the current position within thumb-space, as shown in Figure 36c.

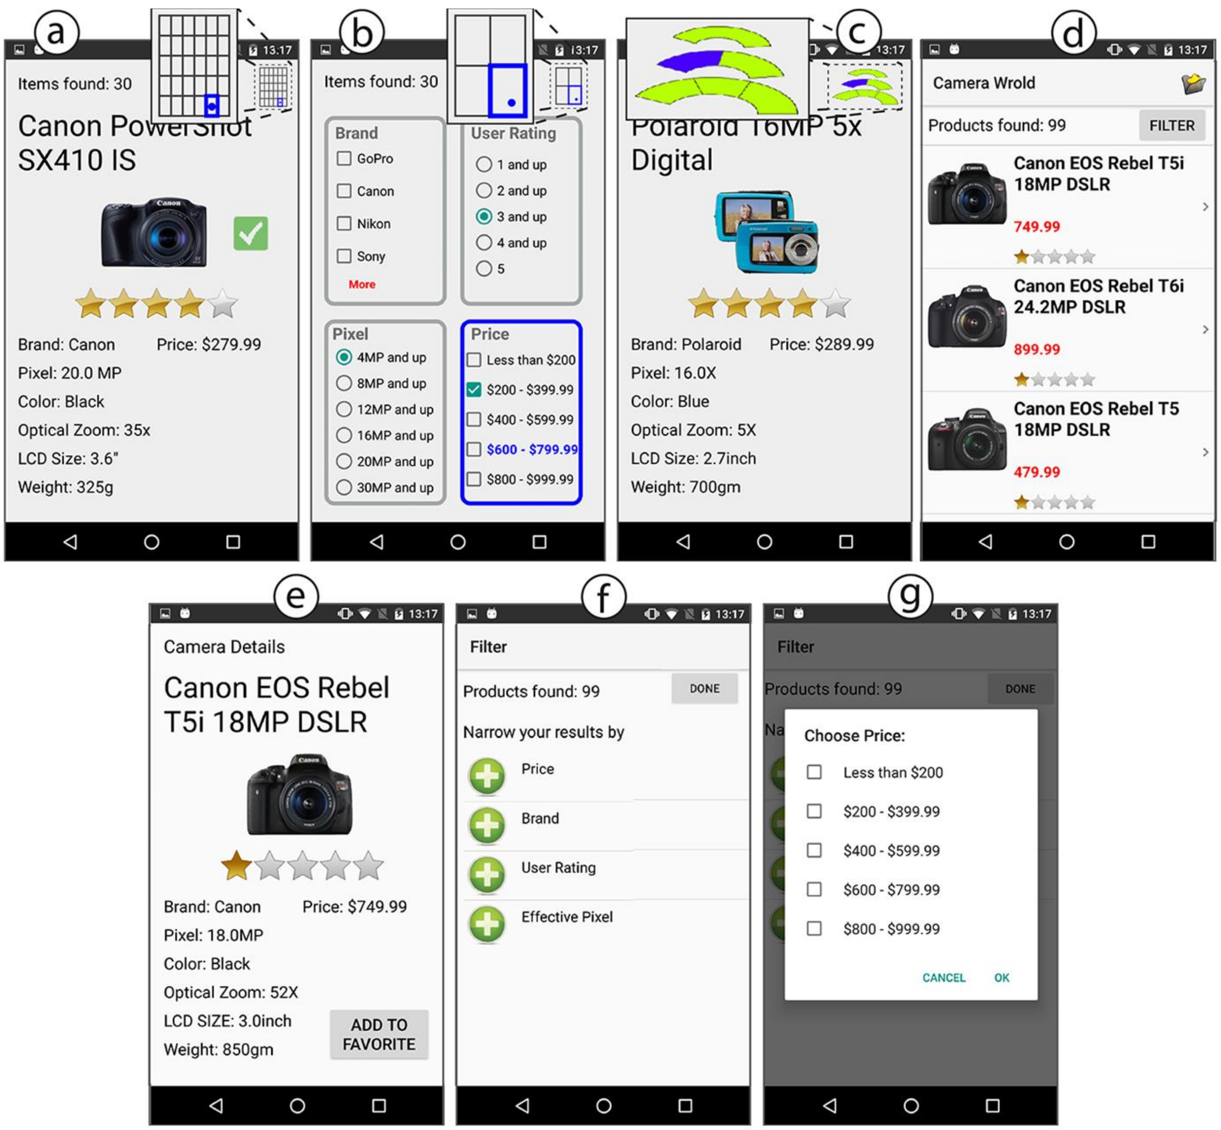

Figure 36. AirPanes: a) product details, b) filter options, and c) product detail for a bookmarked item. Classic touch interface: d) a scrollable overview list of search results, e) a view for product details, f) a view with filters, and g) a modal view with filter options.

The thumb-space is deactivated when the user moves the right hand into any in-air pane to the right of the smartphone. The results pane which is the topmost pane extends infinitely upwards. The filter pane (with its sub-panes) in the middle is 15cm high and the bottom-most pane with additional filter options extends infinitely downwards.

We evaluated the efficiency of our AirPanes interface against a touch-based m-commerce interface. The touch interface provides the same functionality as the AirPanes version. The functionality is organized in multiple views, as typical for popular m-commerce interface and shown in Figure 36. Search results are presented in a scrollable list with overview information for each item (Figure 36d). A tap on an item shows its detail view (Figure 36e), where the user can add the item to the favorite list. A check mark signals if the item is already in the favourite list (as in Figure 36a) and the “add” button is substituted by a “remove” button. From the view with the result list the user can tap on a “Filter” button to narrow down the result by first selecting a filter category in the filter view (Figure 36f), and then select desired filter options in a second view (Figure 36g). Access to the list of bookmarked favorites is provided from the view with the result list by tapping the “favorite” button in the top-right corner of the screen. The favorites are presented in a scrollable list with overview information for each item, exactly as in the main view with the search result. A tap on an item shows its detail view. From all views, a tap on the back button in Android’s navigation bar at the bottom of the screen switches back to the previous view.

Twelve right-handed male smartphone owners (age 23.2 years, s.d. 3.2) participated. All were new to the concept of around-device interactions and none had participated in any previous pilot or earlier study within our AirPanes project.

Task

The study task covers typical steps taken when searching for a camera to buy online. First, a text prompt is displayed on an external monitor, e.g., "Apply the following filters: Price between $200 and $400, user rating 3 and higher, then find the lightest camera with a weight between 400g and 500g" and "Apply the following filters: Brand: Nikon and Sony, Pixels: 10MP and up, then find the camera with the greatest optical zoom between 25X and 40X". The task prompt always includes two filters in varying combinations, one criteria range (*between* value X and Y for either optical zoom, LCD size, or weight), and one superlative (smallest, greatest, lightest, or heaviest). After reading the prompt (which remains visible throughout the trial), the participant taps a "start" button on the smartphone screen, this starts the trial timer. We exclude the entering of a search query from the task since this is done in the same way in the two interfaces. Instead, the task starts with a default set of 99 cameras. Now the participant needs to apply the two requested filters. After applying the filters either 10, 20, or 30 cameras remain in the results set (the number varies randomly between trials). The participant can now start inspecting the product details of the remaining cameras and can bookmark candidate cameras to "favorites" that correspond to the selection criteria (e.g., weight between 400g and 500g, 4 to 7 cameras per trial). After this, the participant accesses the favourite functionality to find the one and only camera that matches the complete task prompt. When the participant believes having found the correct camera, the trial ends with a tap on a "buy" button (when using the touch interface) or with a thumb-tap from the corresponding item inside thumb-up space (when using AirPanes, cf. Study 7). If it is the correct camera trial time stops and the task prompt for the next trial is displayed on the external display. If it was incorrect, the screen flashes in red and the participant can continue the search for the correct camera.

Each participant performed nine trials with each interface (three trials per result set size (10, 20, 30) in random order). Six participants started with AirPanes, six with the touch interface. We demonstrated the two interfaces and showed how the filter functionality and the bookmarking feature would help them to speed up their searches. We instructed participants to try to finish each trial as quickly as possible. Each participant had two practice trials with each interface before starting with the timed trials. Participation lasted about 1 hour (including instructions, practice, and breaks).

Results

We analyze the trial time and the time participants spent with each of the three sub-activities: applying filters (*filter time*), inspecting items and bookmarking items in the filtered set of cameras (*results time*), and inspecting and selecting from the bookmarked favorites (*favorite time*). For each separate time measure we use the median time for each interface × results size combination for each participant. The data was right skewed and we performed logarithmic transformations (which resulted in a distribution close to normal for all four-time measures) before analyzing each time measure.

Trial time: The trial time is shown in Figure 37a. The geometric mean trial time was 57.3s with AirPanes and 116.1s with the touch interface. The difference corresponds to 50.6%. A two-way RM-ANOVA (interface, results size) showed a significant effect for interface ($F_{1,11} = 222.60$, $p < 0.0001$, $\eta^2 = 0.95$) and for results size ($F_{2,22} = 48.64$, $p < 0.0001$, $\eta^2 = 0.82$), but no significant interaction effect. The trial time increased with increasing results size: 63.2s with 10 items, 86.3s with 20 items, and 99.7s with 30 items. Bonferroni corrected post-hoc pairwise comparisons between results sizes showed a significant difference for each pair (all p's < 0.016).

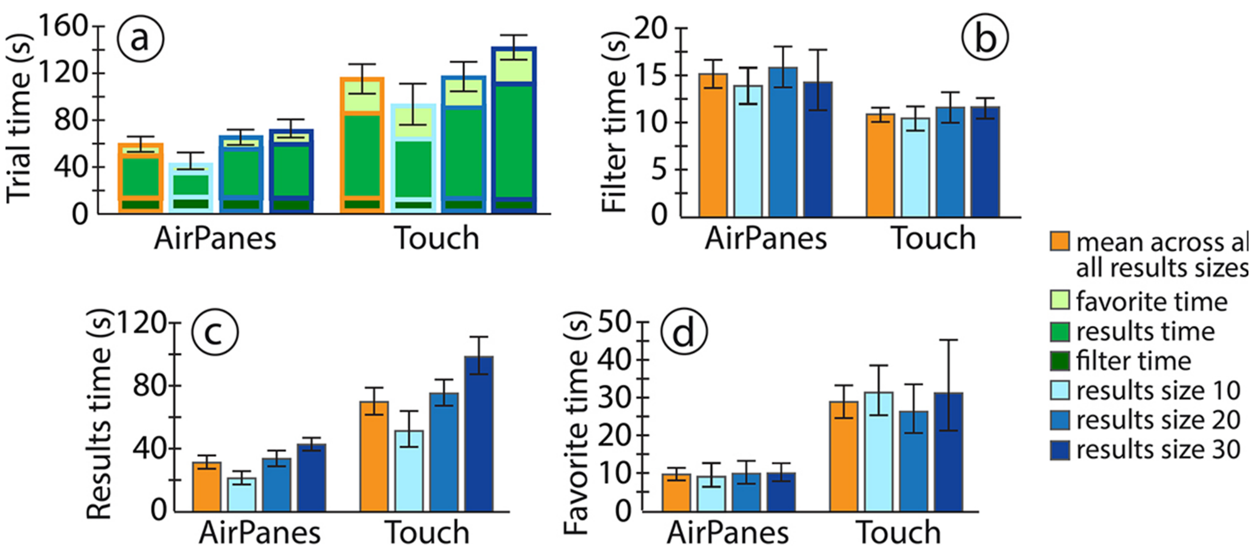

Figure 37. Geometric mean times. a) Trial time, b) filter time, c) results time, and d) favorite time. Error bars: 95 CI.

We turn to our analyses of the three separate sub-activities the participants performed to find the explanation for the great overall difference between the two interfaces.

Filter time: The time spent with the first sub-activity is shown in Figure 37b. A two-way RM ANOVA showed that across the three result sizes participants needed significantly more time ($F_{1,11} = 22.80$, $p < 0.001$, $\eta^2 = 0.68$) to apply the filters using AirPanes (15.1s) than with the touch interface (11.0s). As expected, there was no significant effect for results size (results size 10: 11.1s, result size 20: 14.6s, results size 30: 13.0s) or a significant interaction effect. Recall that the results size was relevant only after the filters had been applied. Accordingly, the time difference (3.9s) between the two interfaces mainly comes from selecting the filter controls, which was more challenging (and unfamiliar) with in-air pinches using AirPanes.

Results time: The time spent in the second sub-activity, inspecting items, and bookmarking items in the filtered set of cameras, is shown in Figure 37c. Overall, across the three results sizes, participants needed significantly ($F_{1,11} = 157.4$, $p < 0.0001$, $\eta^2 = 0.94$) more time for this activity when using the touch interface (71.3s) than when using AirPanes (30.8s). There was also a significant effect for results size ($F_{2,22} = 50.8$, $p < 0.0001$, $\eta^2 = 0.82$). Comprehensibly, with more items to inspect the longer it takes (results size 10: 35.8s, results size 20: 53.7s, results size 30: 69.5s). Bonferroni adjusted post-hoc comparisons showed that all three results sizes differed (all p's < 0.016). The relative increase in results time with increasing results size was similar for both interfaces (i.e., the RM-ANOVA did not show a significant interface × results size interaction). We attribute the slow inspection and bookmarking time with the touch interface to the frequent (and tedious) switching between items in the result list and their corresponding detail views to find the necessary information. With

AirPanes on the other hand, no switching is needed. Instead, the user does only need to move the in-air finger a short distance to enter a new in-air item and so call in its detail information on the screen.

Favorites time: The time spent in the third sub-activity, selecting from the bookmarked items, is shown in Figure 37d. Overall, across the three result sizes, participants were significantly faster ($F_{1,11} = 59.2$, $p < 0.0001$, $\eta^2 = 0.84$) selecting from the favorite set when using AirPanes (9.7s) than when using the touch interface (28.7s). As with the filtering activity, results size had no significant effect on how much time was needed to identify and select the target item from the favorites (results size 10: 17.4s, results size 20: 19.8s, results size 30: 20.8s). Again, as when involved with inspecting the filtered set of cameras, we attribute the disadvantage with the touch interface to the frequent view switching that is required to going back and forth between camera overviews and details. With AirPanes, only a small movement of the in-air thumb is necessary to view information about a new item.

Summary

Our results show that in comparison to touch input, in-air interactions with AirPanes reduces browsing time by taking the advantage of an in-air 'hover' state. With AirPanes, detailed item information can be inspected without having to perform an action (i.e., tapping) to open the detail view. This makes AirPanes an efficient browsing interface, which does not require frequent switching between views, as in many touch-based interfaces (where the user has to tap on an item in a list, open its details and then tap again to switch back to the list). Our results also demonstrate that users can capitalize on this in-air 'hover' state when using their thumb.

Discussion

In the light of these promising results, we discuss design recommendations, potential issues regarding the integration of AirPanes in smartphone applications and present future directions to extend our work.

Design Recommendations

Our investigation offers the following recommendations regarding around-device interactions:

Nested pane layout: In our AirPanes design, we used nested in-air panes to provide access to the four filter categories and their options. This nesting strategy increases the number of panes that can be used in an application. We recommend designers to adopt nesting instead of stacking several planes on top of each other which requires larger vertical movements for switching between panes.

Fine and coarse task assignment: We observed that the in-air thumb has limited movement flexibility and preciseness when the user holds the smartphone in the same hand.

Therefore, we suggest designers to only use the thumb-space for coarse-grained interactions (e.g., making selections).

Considerations for Integration of AirPanes in Applications

AirPanes scalability: We tested AirPanes with a limited number items in a pane (e.g., a maximum of 30 cameras in the results pane). However, the design can be extended to accommodate a larger number of panes. For instance, a UI scrolling strategy could be adopted where moving the index finger to a certain in-air location would show new panes. Moreover, an AirPanes application could incorporate other strategies, such as pagination, to divide large sets of in-air items into sub-sets which could be triggered in thumb-space.

Pane organization styles: Our AirPanes implementation considers a mixed pane organization, the panes with filter options and the results pane are stacked. The panes with filter categories were tiled. We envision that AirPanes can work with many alternative pane organizations. For example, AirPanes could be designed for one-handed use using only stacked or tiled layout. In these ways, all the panes would be accessible with the index finger on the free hand. Alternatively, a pane organization could leverage two-handed usage by assigning the thumb-space for triggering panes that are accessed with the other hand as shown in Figure 38.

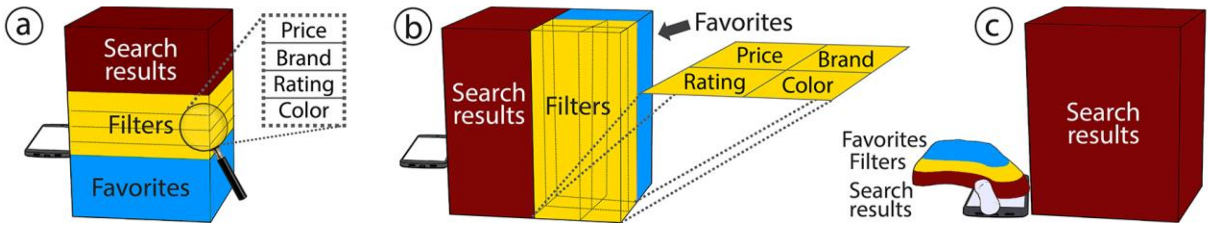

Figure 38. In-air pane organization: a) and b) for one-handed interaction with nested panes, c) for two-handed interaction using in-air thumb space for pane switching.

Generalization of AirPanes for other Applications: In a recent study Carrascal et al. [22] found that people use their smartphones for the following most frequent purposes: social networking, searching and browsing, SMS/texting, phone, audio, and email. We explored two applications for each of these purposes (i.e., a total of 10 applications) to learn about the generalizability of AirPanes to other application categories. We observed that due to the limited display size, frequent pane switching is common for the investigated applications. For instance, results in search applications (e.g., Google [53] or Yahoo [207]) are most often displayed with headlines and snippets and the user has to open the results, one by one, which involves frequent view switching. Since AirPanes reduces the need for frequent switching through its in-air 'hover' state, we believe most applications that require view switching would benefit from an AirPanes design.

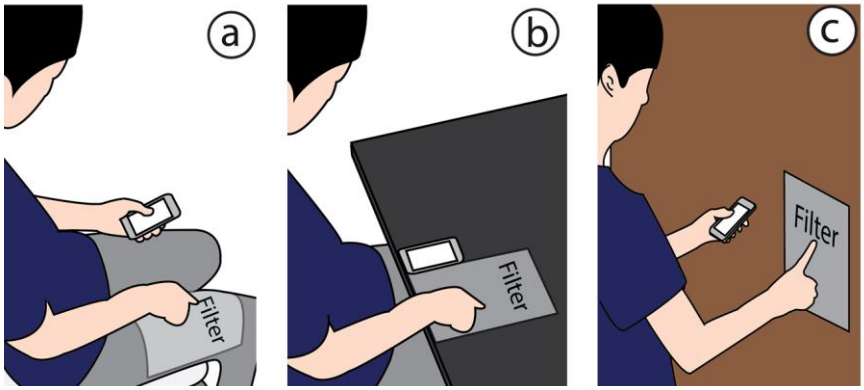

Figure 39. AirPanes can leverage a nearby surface for haptic feedback, e.g., (a) a knee, (b) table, or (c) wall.

Users Context: We explored AirPanes with a task where the user was standing. However, AirPanes could also be used while sitting, or when the smartphone is laying on a table or any other surface. In such cases, AirPanes could use the surrounding physical surface to leverage haptic feedback into mixed physical and in-air interactions. For instance, AirPanes could use the surrounding surface (e.g., use the user's thigh, the table or the wall as shown in Figure 39) for applying filters, and then use the in-air space for browsing results.

Summary

We have presented AirPanes, a novel technique that utilizes in-air spaces to organize multiple panes that are commonly seen in smartphone applications. We demonstrated the value of AirPanes for carrying out a complex task in an m-commerce application that requires users to frequently switch between multiple in-air panes, such as filters, results, and favourites to make a purchase decision. Through two initial studies, we identified properties of various design factors relevant to AirPanes. Our final study confirmed that AirPanes facilitates complex analytic tasks by reducing task time by 50% compared to a standard touch screen interface. Overall, we believe that AirPanes is the first successful step in using in-air spaces for a complete mobile application and so pushes the boundary of current around-device interactions.

Chapter 7: Conclusion

Searching and browsing information on mobile devices is popular as it allows access to digital content anytime, anywhere. However, mobile devices have numerous limitations, such as small screens and limited input modalities (i.e., touch). Information exploration tasks on these devices thus involve many minute operations, such as flicking through screens, and opening and closing items of interest. This results in a less than optimal information browsing experience.

This thesis focused on proposing and developing solutions to address these limitations to facilitate users with information exploration on mobile devices. More specifically, my thesis has demonstrated utilizing around-device space as an effective medium to overcome the limitations of information exploration on small devices. we demonstrated the effectiveness of using around-device space for m-commerce analytic tasks that often require that users quickly browse and compare the breadth of available choices before making a decision. With ten user studies, we explored the design space for around-device interface and how such interfaces can co-exist with on-screen interactions.

In the following sections, we first provide summaries of our findings based on results presented in chapter 3 -6 and discuss limitations that can lead to potential future research directions. We next discuss on the future opportunities that the thesis opens up for the imminent researchers in the around-device interaction field.

Summary

In this thesis, we showed that in-air space around a mobile device can effectively be use for performing complex analytic tasks that require browsing through large information space. To achieved that goal, we started with exploring whether the devices' peripheral space could be used to interact with on-screen items. In chapter 3, we proposed an interactive method, Around-Device Binning or AD-Binning that allows users to directly store and retrieve virtual content in a 2D virtual interactive plane parallel to the device. Interactions with this virtual plane around the mobile device can be carried out by hovering the index finger around the mobile device. Using a systematic design process, we explored various factors related to designing the around-device space for content exploration. A final study was conducted to evaluate AD-Binning with practical usage scenarios where users were required to browse content for making a decision such as finding the cheapest hotel from a map. Our results revealed that participants were faster-exploring information with AD-Binning than with traditional on-screen touch techniques.

AD-Binning demonstrated the advantages of extending mobile devices' input by leveraging around-device space. However, there was uncertainty about users' perceptions regarding the social acceptance of around-device gestures (AD-Gestures) such as pointing to an around-device location to access AD-Binning content. In chapter 4, we explored how comfortable users feel performing AD-gestures in public places, such as a shopping mall or restaurant. We started with examining the influence of fundamental AD-Gesture features on users' level of comfort, such as the distance from the device, gesture size and gesture duration. We found that small or quick gestures close to the device are more likely to be acceptable than a large or lengthy one performed far from the device. Additionally, results revealed that users are selective regarding in what public settings, such as location (e.g., home, shop, workplace, bus/train) and audience (e.g., alone, family, partner, colleagues or stranger) they would use gestures. Also, results confirmed that 3D space around a device is socially acceptable for interacting with around-device content. AD-Binning only utilizes a 2D plane for interacting with virtual content. Inspired by this results, we next explored how to use 3D space around-device space to interact with mobile applications to facilitating complex analytic tasks in a mobile application scenario.

Prior work on around-device interaction has mostly considered the in-air spaces (i.e., right, front and left regions) around a mobile device. No attention has been paid how to use the thumb to access the thumb-reachable in-air space just above the device. In chapter 5, we presented Thumbs-Up, the in-air space next to a mobile device that can be reached with the thumb on the hand holding the device. With Thumbs-Up, the in-air space could be used to access on-screen content through directly pointing with the thumb. Additionally, this surrounding space could be used in conjunction with other around-device spaces to extend the in-air input vocabulary. We explored several parameters related to designing the thumb-reachable in-air space for accessing in-air content. With insights about designing Thumbs-Up space, we next explored how to combine this space for supporting two-handed in-air interaction for analytic design making tasks with around-device input.

Given the limited visual output space, mobile app designers often structure the interface into multiple views or windows. For instance, purchasing a product from an e-commerce application requires frequent switches between filter, results, and cart views to make a purchase decision. This task is commonly time-consuming on mobile devices because of these frequent view switches. In chapter 6, we designed an interactive method, AirPanes, a technique where the user accesses and switches between information and functionality views – or panes – that are located in the air around the device. With AirPanes, a user can access different views by pointing to different regions around the mobile device (e.g., top or bottom) to access the panes. We investigated whether browsing large information spaces with AirPanes facilitates analytic tasks that require users to log items and re-inspect them to make a decision. Our results revealed that AirPanes is faster than tradition touch interfaces for exploring a large amount of data on mobile devices.

Overall, these explorations revealed the potential of using in-air around-device space for supporting complex analytic tasks that require users to make decisions while browsing through large information space. Additionally, results from the user studies confirmed that around-device space is an promising alternate input to the standard touch technique that can be integrated to the next-generation mobile devices.

Limitations

Though our explorations have revealed enormous potential of using in-air space, however, there are several limitations that could potentially be investigated in future. In this section, we highlight these limitations and briefly discuss how to address these to facilitate complex analytic tasks with around-device interactions on future mobile devices.

In chapter 3, we used a commercial tracking solution (i.e., Vicon motion capture system) to emulate the around-device environment. Our developed prototype, AD-Binning relies on a robust tracking mechanism (Vicon cameras) for around-device interaction and in 3D space. Further experimentation is needed for determining suitable design parameters for devices equipped with new sensors that track fingers in off-screen space. Results reveal that around-device space is best discretized into forty-five bins (based on 5 sectors and the smallest target size). To augment this space, items can be stacked on top of one another in 3D which we described in a later chapter (chapter 6). The results are also dependent on an overview, which consumes space on the screen. Additional work is needed to identify whether such visual guidance can be eliminated after repeated use in a given task and application. Finally, the automatic placement strategy inserts items in bins in a random manner. More robust layout mechanisms are needed to provide for an efficient organization of around-device items. For example, in the map application, items in one area could be assigned to corresponding relative regions in off-screen space.

Additionally, we only considered AD-Binning for analytic tasks. Future work could consider extrapolating the results to other forms of tasks in around-device space, such as selecting commands, bridging between physical items around the device and AD-Binning, and coupling around-device input with on-screen interaction. Fruthermore, the impact of mobility, such as walking or on the bus, on AD-Binning's interface could be another potential direction to explore in future.

In chapter 4, we acknowledge the limited methodological support for the central use of perceived 'mental comfort' as a predictor of social acceptance. However, as numerous previous study designers [142,162,163,200,201] and many of their participants have used a similar terminology, our choice was not a farfetched one. We also recognize that user acceptance and social acceptance are multifaceted concepts, by far not limited to the perceived or expected levels of mental comfort [34,142]. In the studies, we focused on social settings and ignored important cultural factors, such as participants' cultural background. The studies were mainly conducted in Canada with persons living there. Little is known about how cultural aspects influences user perceptions about, and the social acceptance of novel interaction techniques [201]. Accordingly, and with all the participants living in a western culture, we are wary of generalizing the results to non-western users and cultures. We suspect that examining culture-dependent differences of technology adoption and social acceptance would be a challenging but very fruitful path for future work. Additionally, we are also wary of assuming the results apply to other age groups as most of our 60 participants in studies 4-6 were 25 to 35 years old (mean 27.8, s.d. 7.6).

In chapter 5, we introduced the idea of using thumb movement in the Thumbs-reachable in-air space. Though current smartphones have the necessary capabilities to detect thumb movements in the air around the device, attaching external sensors (e.g., an omnidirectional mirror on top of the smartphone's camera [11] or wearing a magnetic ring on the thumb) could be used to track the thumb in Thumbs-Up space. Further investigation is required to make mobile devices self-contained to track such in-air movement. Thumbs-Up interactions require a trigger mechanism to activate the Thumbs-Up space. A tap at the back of the smartphone (which produces a vibrational signal that is easily detected by the on-board accelerometer) or a quick swipe with the thumb back and forth across the screen bezel could be used. Alternatively, a press on a physical button or a tap on a special screen button could activate/deactivate the Thumbs-Up space. Furthermore, Thumbs-Up space allows the user to explore on-device items by moving the thumb. A selection is required when the user wants to invoke an item and put it into focus for more details. Specific finger movements such as a rapid thumb raise in Thumbs-Up space could be used to trigger the selection. We consider in-air thumb-space as a complementary input region. In our future work will focus on exploring this space for one-handed use cases such as performing two-finger gestures with this in-air space, which is yet to explore.

In chapter 6, we presented AirPanes, which demands accurate and robust finger detection around the smartphone. Although current smartphones do not provide such features, a recent study, Song et al. [179] showed that the smartphone's camera can be used to recognize in-air hand movement but in limited directions. Further explorations are needed to make devices capable of precise around-device hand tracking. We considered the smartphone to be in 'AirPanes' mode by default. However, an explicit mode switching mechanism must be developed to activate and deactivate around-device input to counter accidental in-air gesture events. This can be done, for example, with a physical button that the user presses to activate AirPanes mode. Though the current smartphone design allows holding of the device with one hand, the flat shape of the phone prohibits users from reaching the top regions with the thumb. We are also interested in investigating how to change the device's shape in order to better match the space that can be reachable with the thumb.

Future Work

My research opens the possibility of using the surrounding space of mobile device to facilitate complex analytic tasks that involve interacting with large number of on-device content. Results also revealed that performing such tasks with mid-air gestures is more efficient than using touch input. My future research goal is to build on these results by designing efficient mobile interfaces to facilitate around-device interaction for many mobile or wearable application. Equipped with my research expertise and knowledge that I gathered in my thesis, I will (i) explore robust and reliable sensing methods for self-contained mobile devices to detect hand gestures; (ii) investigate novel approaches to provide feedback for virtual content just beside the screen; and (iii) build new around-device display techniques to extend the current mobile devices display capabilities.

Around-device sensing

On mobile devices, vision-based systems (e.g., Surround-See) provide basic around-device input and gesture detection in real-time. However, this approach provides low accuracy and limited depth sensing capabilities warranting further investigation of suitable around-device sensing methods. A major step to my future goals is to investigate precise and reliable sensing methods that would recognize around-device gestures and finger inputs without relying on an external tracking mechanism. Possible technical approaches to achieve this goal could be to (i) use mechanical waves (e.g., Sonar) propagated from and captured by mobile devices for detection; (ii) use advance imaging technologies (e.g., thermal imaging) on mobile device to capture emitted infrared radiation from hand; or (iii) explore advanced sensors (e.g., 3D range sensor) to detect hand and finger inputs. Building a reliable and robust sensing method would allow me to provide effective around-device interaction capabilities and eventually allow around-device interaction to become mainstream.

Around-device feedback

Providing feedback about the items that are placed around the device is challenging due to limited screen space and on-board sensing capabilities. A common approach to representing such off-screen information is to use visual cues (e.g., Overview [71] and EdgeRadar [64]) which consume a significant portion screen space on the limited mobile display. These cues impose clutter and do not scale effectively to large amounts of content which make information exploration tedious. Additionally, switching between the workspace and the visual cues imposes additional cognitive load and breaks the seamless interaction. Therefore, we need to have effective feedback mechanisms to get rapid and accurate information about around-device content. In my future work, I will explore various ways to provide haptic feedback about the off-screen content. This form of feedback could be generated by using an electromagnetic field or high audio frequency. Building robust prototypes with mid-air feedback will provide an opportunity to not only gauge the effectiveness of my around-device interactions but also help me to transfer such interactions into mainstream mobile and wearable devices.

Around-device display

While smartphones have become indispensable in our everyday activities, they still have limitations in displaying and interacting with a large amount of information. As discussed before, mobile devices interaction space could be extended by augmenting a virtual interactive plane that extends their physical form. However, very little is known about how to extend the display space for presenting more information. My future research will explore around-device display where I aim to display off-screen information directly in mid-air around the mobile device. Examples of how an around-device display could be used include a map on a mobile device that extends its physical boundary and displays geographic information (e.g., regions, objects) that are spatially located around the mobile device. I will investigate a variety of technical solutions, such as using nano-projection or holographic displays for mobile platforms to display off-screen information.

Final words

In summary, my research took a novel approach to push the boundary of mobile interaction technologies and techniques. I strongly believe that contributions of this thesis will result in a paradigm shift in how users interact with information on mobile and wearable devices in the future. The outcomes of my research will open new paradigms that could replace or complement existing mobile interactions.

References

1. 12 Mobile Statistics to Get You Amped for 2015. 2016. Retrieved June 30, 2016 from http://www.highervisibility.com/blog/12-mobile-statistics-to-get-you-amped-for-2015/
2. David Ahlström, Khalad Hasan, and Pourang Irani. 2014. Are You Comfortable Doing That?: Acceptance Studies of Around-device Gestures in and for Public Settings. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services* (MobileHCI '14), 193–202. https://doi.org/10.1145/2628363.2628381
3. Alexa Top 500 Global Sites. 2016. Retrieved November 16, 2016 from http://www.alexa.com/topsites
4. Jason Alexander, Mark T. Marshall, and Sriram Subramanian. 2011. Adding Haptic Feedback to Mobile Tv. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '11), 1975–1980. https://doi.org/10.1145/1979742.1979899
5. Alibaba. 2016. Retrieved November 30, 2016 from https://www.alibaba.com
6. Aliexpress. 2016. Retrieved November 30, 2016 from //www.aliexpress.com
7. Amazon.com: Online Shopping for Electronics, Apparel, Computers, Books, DVDs & more. 2016. Retrieved November 30, 2016 from https://www.amazon.com/
8. Daniel Ashbrook, Patrick Baudisch, and Sean White. 2011. Nenya: Subtle and Eyes-free Mobile Input with a Magnetically-tracked Finger Ring. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '11), 2043–2046. https://doi.org/10.1145/1978942.1979238
9. Oscar Kin-Chung Au, Xiaojun Su, and Rynson W.H. Lau. 2014. LinearDragger: A Linear Selector for One-finger Target Acquisition. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), 2607–2616. https://doi.org/10.1145/2556288.2557096
10. Daniel Avrahami, Jacob O. Wobbrock, and Shahram Izadi. 2011. Portico: Tangible Interaction on and Around a Tablet. In *Proceedings of the 24th Annual ACM*

Symposium on User Interface Software and Technology (UIST '11), 347–356. https://doi.org/10.1145/2047196.2047241

11. Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2008. ILoveSketch: As-natural-as-possible Sketching System for Creating 3D Curve Models. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (UIST '08), 151–160. https://doi.org/10.1145/1449715.1449740
12. Ravin Balakrishnan and Ken Hinckley. 2000. Symmetric Bimanual Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '00), 33–40. https://doi.org/10.1145/332040.332404
13. Patrick Baudisch and Gerry Chu. 2009. Back-of-device Interaction Allows Creating Very Small Touch Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '09), 1923–1932. https://doi.org/10.1145/1518701.1518995
14. Patrick Baudisch and Ruth Rosenholtz. 2003. Halo: A Technique for Visualizing Off-screen Objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '03), 481–488. https://doi.org/10.1145/642611.642695
15. Patrick Baudisch, Xing Xie, Chong Wang, and Wei-Ying Ma. 2004. Collapse-to-zoom: Viewing Web Pages on Small Screen Devices by Interactively Removing Irrelevant Content. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology* (UIST '04), 91–94. https://doi.org/10.1145/1029632.1029647
16. Joanna Bergstrom-Lehtovirta and Antti Oulasvirta. 2014. Modeling the Functional Area of the Thumb on Mobile Touchscreen Surfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), 1991–2000. https://doi.org/10.1145/2556288.2557354
17. Best Buy: Expert Service. Unbeatable Price. 2016. Retrieved November 30, 2016 from http://www.bestbuy.com/
18. Xiaojun Bi, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2011. Magic Desk: Bringing Multi-touch Surfaces into Desktop Work. In *Proceedings of the SIGCHI*

Conference on Human Factors in Computing Systems (CHI '11), 2511–2520. https://doi.org/10.1145/1978942.1979309

19. Stefano Burigat, Luca Chittaro, and Andrea Vianello. 2012. Dynamic Visualization of Large Numbers of Off-screen Objects on Mobile Devices: An Experimental Comparison of Wedge and Overview+Detail. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services* (MobileHCI '12), 93–102. https://doi.org/10.1145/2371574.2371590
20. Alex Butler, Shahram Izadi, and Steve Hodges. 2008. SideSight: Multi-"Touch" Interaction Around Small Devices. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (UIST '08), 201–204. https://doi.org/10.1145/1449715.1449746
21. W. Buxton and B. Myers. 1986. A Study in Two-handed Input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '86), 321–326. https://doi.org/10.1145/22627.22390
22. Juan Pablo Carrascal and Karen Church. 2015. An In-Situ Study of Mobile App & Mobile Search Interactions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), 2739–2748. https://doi.org/10.1145/2702123.2702486
23. Tom Carter, Sue Ann Seah, Benjamin Long, Bruce Drinkwater, and Sriram Subramanian. 2013. UltraHaptics: Multi-point Mid-air Haptic Feedback for Touch Surfaces. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13), 505–514. https://doi.org/10.1145/2501988.2502018
24. Didier Casalta, Yves Guiard, and Michel Beaudouin-Lafon. 1999. Evaluating Two-handed Input Techniques: Rectangle Editing and Navigation. In *CHI '99 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '99), 236–237. https://doi.org/10.1145/632716.632862
25. Youli Chang, Sehi L'Yi, Kyle Koh, and Jinwook Seo. 2015. Understanding Users' Touch Behavior on Large Mobile Touch-Screens and Assisted Targeting by Tilting Gesture.

In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), 1499–1508. https://doi.org/10.1145/2702123.2702425

26. Wei-Hung Chen. 2015. Blowatch: Blowable and Hands-free Interaction for Smartwatches. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (CHI EA '15), 103–108. https://doi.org/10.1145/2702613.2726961
27. Xiang "Anthony" Chen, Julia Schwarz, Chris Harrison, Jennifer Mankoff, and Scott Hudson. 2014. Around-body Interaction: Sensing & Interaction Techniques for Proprioception-enhanced Input with Mobile Devices. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services* (MobileHCI '14), 287–290. https://doi.org/10.1145/2628363.2628402
28. Xiang "Anthony" Chen, Julia Schwarz, Chris Harrison, Jennifer Mankoff, and Scott E. Hudson. 2014. Air+Touch: Interweaving Touch & In-air Gestures. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), 519–525. https://doi.org/10.1145/2642918.2647392
29. Alton Y. K. Chua, Radhika Shenoy Balkunje, and Dion Hoe-Lian Goh. 2011. Fulfilling Mobile Information Needs: A Study on the Use of Mobile Phones. In *Proceedings of the 5th International Conference on Ubiquitous Information Management and Communication* (ICUIMC '11), 92:1–92:7. https://doi.org/10.1145/1968613.1968721
30. Karen Church and Nuria Oliver. 2011. Understanding Mobile Web and Mobile Search Use in Today's Dynamic Mobile Landscape. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services* (MobileHCI '11), 67–76. https://doi.org/10.1145/2037373.2037385
31. Karen Church and Barry Smyth. 2009. Understanding the Intent Behind Mobile Information Needs. In *Proceedings of the 14th International Conference on Intelligent User Interfaces* (IUI '09), 247–256. https://doi.org/10.1145/1502650.1502686
32. Karen Church, Barry Smyth, and Mark T. Keane. 2006. Evaluating Interfaces for Intelligent Mobile Search. In *Proceedings of the 2006 International Cross-disciplinary*

Workshop on Web Accessibility (W4A): Building the Mobile Web: Rediscovering Accessibility? (W4A '06), 69–78. https://doi.org/10.1145/1133219.1133232
33. A. Cockburn, P. Quinn, C. Gutwin, G. Ramos, and J. Looser. 2011. Air Pointing: Design and Evaluation of Spatial Target Acquisition with and Without Visual Feedback. *Int. J. Hum.-Comput. Stud.* 69, 6: 401–414. https://doi.org/10.1016/j.ijhcs.2011.02.005
34. Fred D. Davis, Richard P. Bagozzi, and Paul R. Warshaw. 1989. User Acceptance of Computer Technology: A Comparison of Two Theoretical Models. *Manage. Sci.* 35, 8: 982–1003. https://doi.org/10.1287/mnsc.35.8.982
35. Bruno R. De Araùjo, Géry Casiez, and Joaquim A. Jorge. 2012. Mockup Builder: Direct 3D Modeling on and Above the Surface in a Continuous Interaction Space. In *Proceedings of Graphics Interface 2012* (GI '12), 173–180. Retrieved November 16, 2016 from http://dl.acm.org/citation.cfm?id=2305276.2305305
36. eBay | Electronics, Cars, Fashion, Collectibles, Coupons and More. 2016. Retrieved November 30, 2016 from http://www.ebay.com/
37. Darren Edge and Alan F. Blackwell. 2009. Bimanual Tangible Interaction with Mobile Phones. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction* (TEI '09), 131–136. https://doi.org/10.1145/1517664.1517697
38. Barrett Ens, David Ahlström, Andy Cockburn, and Pourang Irani. 2011. Characterizing User Performance with Assisted Direct Off-screen Pointing. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services* (MobileHCI '11), 485–494. https://doi.org/10.1145/2037373.2037445
39. Barrett Ens, Ahmad Byagowi, Teng Han, Juan David Hincapié-Ramos, and Pourang Irani. 2016. Combining Ring Input with Hand Tracking for Precise, Natural Interaction with Spatial Analytic Interfaces. In *Proceedings of the 2016 Symposium on Spatial User Interaction* (SUI '16), 99–102. https://doi.org/10.1145/2983310.2985757
40. Barrett Ens, Rasit Eskicioglu, and Pourang Irani. 2011. Visually Augmented Interfaces for Co-located Mobile Collaboration. In *Distributed User Interfaces*, José A. Gallud, Ricardo Tesoriero and Victor M. R. Penichet (eds.). Springer London, 169–176.

Retrieved June 30, 2016 from http://link.springer.com/chapter/10.1007/978-1-4471-2271-5_19

41. Barrett Ens, Tovi Grossman, Fraser Anderson, Justin Matejka, and George Fitzmaurice. 2015. Candid Interaction: Revealing Hidden Mobile and Wearable Computing Activities. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology* (UIST '15), 467–476. https://doi.org/10.1145/2807442.2807449
42. Barrett M. Ens, Rory Finnegan, and Pourang P. Irani. 2014. The Personal Cockpit: A Spatial Interface for Effective Task Switching on Head-worn Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), 3171–3180. https://doi.org/10.1145/2556288.2557058
43. Barrett Ens, Eyal Ofek, Neil Bruce, and Pourang Irani. 2015. Spatial Constancy of Surface-Embedded Layouts Across Multiple Environments. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction* (SUI '15), 65–68. https://doi.org/10.1145/2788940.2788954
44. Etsy. 2016. Retrieved November 30, 2016 from https://www.etsy.com/ca/
45. Jackson Feijó Filho, Wilson Prata, and Thiago Valle. 2015. Advances on Breathing Based Text Input for Mobile Devices. In *Universal Access in Human-Computer Interaction. Access to Today's Technologies*, Margherita Antona and Constantine Stephanidis (eds.). Springer International Publishing, 279–287. Retrieved November 25, 2016 from http://link.springer.com/chapter/10.1007/978-3-319-20678-3_27
46. Flipkart.com. 2016. Retrieved November 30, 2016 from https://www.flipkart.com
47. Euan Freeman, Stephen Brewster, and Vuokko Lantz. 2014. Tactile Feedback for Above-Device Gesture Interfaces: Adding Touch to Touchless Interactions. In *Proceedings of the 16th International Conference on Multimodal Interaction* (ICMI '14), 419–426. https://doi.org/10.1145/2663204.2663280
48. Euan Freeman, Stephen Brewster, and Vuokko Lantz. 2015. Interactive Light Feedback: Illuminating Above-Device Gesture Interfaces. In *Human-Computer Interaction – INTERACT 2015*. Springer International Publishing, 478–481. Retrieved

November 30, 2016 from http://link.springer.com/chapter/10.1007/978-3-319-22723-8_42

49. Euan Freeman, Stephen Brewster, and Vuokko Lantz. 2016. Do That, There: An Interaction Technique for Addressing In-Air Gesture Systems. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), 2319–2331. https://doi.org/10.1145/2858036.2858308
50. GestIC Technology - GestIC® Technology | Touch and Input Sensing | Microchip Technology Inc. 2016. Retrieved November 18, 2016 from http://www.microchip.com/design-centers/touch-input-sensing/gestic-technology/overview
51. Verena Giller, Rudolf Melcher, Johann Schrammel, Reinhard Sefelin, and Manfred Tscheligi. 2003. Usability Evaluations for Multi-device Application Development Three Example Studies. In *Human-Computer Interaction with Mobile Devices and Services*, Luca Chittaro (ed.). Springer Berlin Heidelberg, 302–316. Retrieved June 30, 2016 from http://link.springer.com/chapter/10.1007/978-3-540-45233-1_22
52. Tiago Gonçalves, Ana Paula Afonso, Maria Beatriz Carmo, and Paulo Pombinho. 2013. Comparison of Off-screen Visualization Techniques with Representation of Relevance on Mobile Devices. In *Proceedings of the 27th International BCS Human Computer Interaction Conference* (BCS-HCI '13), 9:1–9:10. Retrieved June 30, 2016 from http://dl.acm.org/citation.cfm?id=2578048.2578061
53. Google. 2016. Retrieved November 16, 2016 from https://www.google.ca/
54. Google and Nielsen: How Mobile is the "Path to Purchase" | WompMobile Blog | Adaptive Mobile Websites |. 2016. Retrieved December 3, 2016 from http://www.wompmobile.com/blog/google-and-nielsen-how-mobile-is-the-path-to-purchase/
55. Google Play. 2016. Retrieved November 16, 2016 from https://play.google.com/store
56. Jens Grubert, Matthias Heinisch, Aaron Quigley, and Dieter Schmalstieg. 2015. MultiFi: Multi Fidelity Interaction with Displays On and Around the Body. In

Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15), 3933–3942. https://doi.org/10.1145/2702123.2702331

57. Jens Grubert, Matthias Kranz, and Aaron Quigley. 2015. Design and Technology Challenges for Body Proximate Display Ecosystems. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct* (MobileHCI '15), 951–954. https://doi.org/10.1145/2786567.2794310
58. Jens Grubert, Matthias Kranz, and Aaron Quigley. 2016. Challenges in Mobile Multi-Device Ecosystems. *ResearchGate*. Retrieved December 1, 2016 from https://www.researchgate.net/publication/303521197_Challenges_in_Mobile_Multi-Device_Ecosystems
59. Jens Grubert, Eyal Ofek, Michel Pahud, Matthias Kranz, and Dieter Schmalstieg. 2016. GlassHands: Interaction Around Unmodified Mobile Devices Using Sunglasses. In *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces* (ISS '16), 215–224. https://doi.org/10.1145/2992154.2992162
60. Yves Guiard. 1987. Asymmetric Division of Labor in Human Skilled Bimanual Action. *Journal of Motor Behavior* 19, 4: 486–517. https://doi.org/10.1080/00222895.1987.10735426
61. Sidhant Gupta, Dan Morris, Shwetak N. Patel, and Desney Tan. 2013. AirWave: Non-contact Haptic Feedback Using Air Vortex Rings. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (UbiComp '13), 419–428. https://doi.org/10.1145/2493432.2493463
62. Sean Gustafson, Patrick Baudisch, Carl Gutwin, and Pourang Irani. 2008. Wedge: Clutter-free Visualization of Off-screen Locations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), 787–796. https://doi.org/10.1145/1357054.1357179
63. Sean Gustafson, Daniel Bierwirth, and Patrick Baudisch. 2010. Imaginary Interfaces: Spatial Interaction with Empty Hands and Without Visual Feedback. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology* (UIST '10), 3–12. https://doi.org/10.1145/1866029.1866033

64. Sean G. Gustafson and Pourang P. Irani. 2007. Comparing Visualizations for Tracking Off-screen Moving Targets. In *CHI '07 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '07), 2399–2404. https://doi.org/10.1145/1240866.1241014
65. Sean G. Gustafson, Bernhard Rabe, and Patrick M. Baudisch. 2013. Understanding Palm-based Imaginary Interfaces: The Role of Visual and Tactile Cues when Browsing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13), 889–898. https://doi.org/10.1145/2470654.2466114
66. Sean Gustafson, Christian Holz, and Patrick Baudisch. 2011. Imaginary Phone: Learning Imaginary Interfaces by Transferring Spatial Memory from a Familiar Device. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology* (UIST '11), 283–292. https://doi.org/10.1145/2047196.2047233
67. Hiroyuki Hakoda, Takuro Kuribara, Keigo Shima, Buntarou Shizuki, and Jiro Tanaka. 2015. AirFlip: A Double Crossing In-Air Gesture Using Boundary Surfaces of Hover Zone for Mobile Devices. In *Human-Computer Interaction: Interaction Technologies*. Springer International Publishing, 44–53. Retrieved November 25, 2016 from http://link.springer.com/chapter/10.1007/978-3-319-20916-6_5
68. Edward T. Hall. 1990. *The Hidden Dimension*. Anchor, New York.
69. Chris Harrison and Scott E. Hudson. 2008. Scratch Input: Creating Large, Inexpensive, Unpowered and Mobile Finger Input Surfaces. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (UIST '08), 205–208. https://doi.org/10.1145/1449715.1449747
70. Chris Harrison and Scott E. Hudson. 2009. Abracadabra: Wireless, High-precision, and Unpowered Finger Input for Very Small Mobile Devices. In *Proceedings of the 22Nd Annual ACM Symposium on User Interface Software and Technology* (UIST '09), 121–124. https://doi.org/10.1145/1622176.1622199
71. Khalad Hasan, David Ahlström, and Pourang Irani. 2013. Ad-binning: Leveraging Around Device Space for Storing, Browsing and Retrieving Mobile Device Content. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13), 899–908. https://doi.org/10.1145/2470654.2466115

72. Khalad Hasan, David Ahlström, and Pourang Irani. 2015. SAMMI: A Spatially-Aware Multi-Mobile Interface for Analytic Map Navigation Tasks. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '15), 36–45. https://doi.org/10.1145/2785830.2785850
73. Khalad Hasan, David Ahlström, and Pourang P. Irani. 2015. Comparing Direct Off-Screen Pointing, Peephole, and Flick & Pinch Interaction for Map Navigation. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction* (SUI '15), 99–102. https://doi.org/10.1145/2788940.2788957
74. Khalad Hasan, Xing-Dong Yang, Andrea Bunt, and Pourang Irani. 2012. A-coord Input: Coordinating Auxiliary Input Streams for Augmenting Contextual Pen-based Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '12), 805–814. https://doi.org/10.1145/2207676.2208519
75. Khalad Hasan, Xing-Dong Yang, Hai-Ning Liang, and Pourang Irani. 2012. How to Position the Cursor?: An Exploration of Absolute and Relative Cursor Positioning for Back-of-device Input. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services* (MobileHCI '12), 103–112. https://doi.org/10.1145/2371574.2371591
76. Tomi Heimonen. 2009. Information Needs and Practices of Active Mobile Internet Users. In *Proceedings of the 6th International Conference on Mobile Technology, Application & Systems* (Mobility '09), 50:1–50:8. https://doi.org/10.1145/1710035.1710085
77. Annika M. Hinze, Carole Chang, and David M. Nichols. 2010. Contextual Queries Express Mobile Information Needs. In *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services* (MobileHCI '10), 327–336. https://doi.org/10.1145/1851600.1851658
78. Nambu Hirotaka. 2003. Reassessing Current Cell Phone Designs: Using Thumb Input Effectively. In *CHI '03 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '03), 938–939. https://doi.org/10.1145/765891.766081
79. T. Hoshi. 2011. Development of aerial-input and aerial-tactile-feedback system. In *2011 IEEE World Haptics Conference*, 569–573. https://doi.org/10.1109/WHC.2011.5945548

80. Takayuki Hoshi, Takayuki Iwamoto, and Hiroyuki Shinoda. 2009. Non-contact tactile sensation synthesized by ultrasound transducers. In *World Haptics 2009 - Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 256–260. https://doi.org/10.1109/WHC.2009.4810900
81. Takayuki Hoshi, Masafumi Takahashi, Kei Nakatsuma, and Hiroyuki Shinoda. 2009. Touchable Holography. In *ACM SIGGRAPH 2009 Emerging Technologies* (SIGGRAPH '09), 23:1–23:1. https://doi.org/10.1145/1597956.1597979
82. T. Hoshi and Y. Nishiyama. 2010. Observations of airflow arising from Airborne Ultrasound Tactile Display. In *Proceedings of SICE Annual Conference 2010*, 384–385.
83. T. Hoshi and H. Shinoda. 2006. A Sensitive Skin Based on Touch-Area-Evaluating Tactile Elements. In *2006 14th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 89–94. https://doi.org/10.1109/HAPTIC.2006.1627069
84. T. Hoshi and H. Shinoda. 2007. Free-Form Tactile Sensor Using 3-Dimensional Shape Capture Sheet. In *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)*, 403–408. https://doi.org/10.1109/WHC.2007.56
85. T. Hoshi, M. Takahashi, T. Iwamoto, and H. Shinoda. 2010. Noncontact Tactile Display Based on Radiation Pressure of Airborne Ultrasound. *IEEE Transactions on Haptics* 3, 3: 155–165. https://doi.org/10.1109/TOH.2010.4
86. Zahid Hossain, Khalad Hasan, Hai-Ning Liang, and Pourang Irani. 2012. EdgeSplit: Facilitating the Selection of Off-screen Objects. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services* (MobileHCI '12), 79–82. https://doi.org/10.1145/2371574.2371588
87. Wolfgang Hürst and Philipp Merkle. 2008. One-handed Mobile Video Browsing. In *Proceedings of the 1st International Conference on Designing Interactive User Experiences for TV and Video* (UXTV '08), 169–178. https://doi.org/10.1145/1453805.1453839

88. Sungjae Hwang, Myungwook Ahn, and Kwang-yun Wohn. 2013. MagGetz: Customizable Passive Tangible Controllers on and Around Conventional Mobile Devices. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13), 411–416. https://doi.org/10.1145/2501988.2501991
89. Sungjae Hwang, Andrea Bianchi, Myungwook Ahn, and Kwangyun Wohn. 2013. MagPen: Magnetically Driven Pen Interactions on and Around Conventional Smartphones. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services* (MobileHCI '13), 412–415. https://doi.org/10.1145/2493190.2493194
90. Sungjae Hwang, Dongchul Kim, Sang-won Leigh, and Kwang-yun Wohn. 2013. NailSense: Fingertip Force As a New Input Modality. In *Proceedings of the Adjunct Publication of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13 Adjunct), 63–64. https://doi.org/10.1145/2508468.2514711
91. Tuuli Hyvärinen, Anne Kaikkonen, and Mika Hiltunen. 2005. Placing Links in Mobile Banking Application. In *Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices & Services* (MobileHCI '05), 63–68. https://doi.org/10.1145/1085777.1085788
92. IKEA.com - International homepage. 2016. Retrieved November 30, 2016 from http://www.ikea.com/
93. Increasing the Appeal of Mobile Tv Using Haptic Feedback Acm Classification Keywords - Semantic Scholar. 2016. Retrieved December 1, 2016 from /paper/Increasing-the-Appeal-of-Mobile-Tv-Using-Haptic-Alexander-Marshall/f80d0da45ad928964120c31761da7580f8ac7814
94. Pourang Irani, Carl Gutwin, and Xing Dong Yang. 2006. Improving Selection of Off-screen Targets with Hopping. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '06), 299–308. https://doi.org/10.1145/1124772.1124818
95. Takayuki Iwamoto, Mari Tatezono, and Hiroyuki Shinoda. 2008. Non-contact Method for Producing Tactile Sensation Using Airborne Ultrasound. In *Haptics: Perception,*

Devices and Scenarios. Springer Berlin Heidelberg, 504–513. Retrieved December 1, 2016 from http://link.springer.com/chapter/10.1007/978-3-540-69057-3_64

96. Dominik Jäckle, Bum Chul Kwon, and Daniel A. Keim. 2015. Off-Screen Visualization Perspectives: Tasks and Challenges. Retrieved November 21, 2016 from http://bib.dbvis.de/uploadedFiles/offscreen_perspectives_jaeckle_vds15.pdf
97. Robert J. K. Jacob, Linda E. Sibert, Daniel C. McFarlane, and M. Preston Mullen Jr. 1994. Integrality and Separability of Input Devices. *ACM Trans. Comput.-Hum. Interact.* 1, 1: 3–26. https://doi.org/10.1145/174630.174631
98. Jingdong (JD.COM) - integrated online shopping preferred - genuine low-cost, quality assurance, distribution in a timely manner, easy shopping! 2016. Retrieved November 30, 2016 from https://www.jd.com/
99. Brett Jones, Rajinder Sodhi, David Forsyth, Brian Bailey, and Giuliano Maciocci. 2012. Around Device Interaction for Multiscale Navigation. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services* (MobileHCI '12), 83–92. https://doi.org/10.1145/2371574.2371589
100. Matt Jones, George Buchanan, and Harold Thimbleby. 2002. Sorting Out Searching on Small Screen Devices. In *Human Computer Interaction with Mobile Devices*, Fabio Paternò (ed.). Springer Berlin Heidelberg, 81–94. Retrieved June 30, 2016 from http://link.springer.com/chapter/10.1007/3-540-45756-9_8
101. Steve Jones, Matt Jones, and Shaleen Deo_andA2. 2004. Using Keyphrases As Search Result Surrogates on Small Screen Devices. *Personal Ubiquitous Comput.* 8, 1: 55–68. https://doi.org/10.1007/s00779-004-0258-y
102. Raine Kajastila and Tapio Lokki. 2013. Eyes-free interaction with free-hand gestures and auditory menus. *International Journal of Human-Computer Studies* 71, 5: 627–640. https://doi.org/10.1016/j.ijhcs.2012.11.003
103. Maryam Kamvar, Melanie Kellar, Rajan Patel, and Ya Xu. 2009. Computers and Iphones and Mobile Phones, Oh My!: A Logs-based Comparison of Search Users on Different Devices. In *Proceedings of the 18th International Conference on World Wide Web* (WWW '09), 801–810. https://doi.org/10.1145/1526709.1526817

104. Shaun K. Kane, Daniel Avrahami, Jacob O. Wobbrock, Beverly Harrison, Adam D. Rea, Matthai Philipose, and Anthony LaMarca. 2009. Bonfire: A Nomadic System for Hybrid Laptop-tabletop Interaction. In *Proceedings of the 22Nd Annual ACM Symposium on User Interface Software and Technology* (UIST '09), 129–138. https://doi.org/10.1145/1622176.1622202

105. Amy K. Karlson and Benjamin B. Bederson. 2007. ThumbSpace: Generalized One-handed Input for Touchscreen-based Mobile Devices. In *Proceedings of the 11th IFIP TC 13 International Conference on Human-computer Interaction* (INTERACT'07), 324–338. Retrieved November 16, 2016 from http://dl.acm.org/citation.cfm?id=1776994.1777034

106. Amy K. Karlson and Benjamin B. Bederson. 2008. One-handed Touchscreen Input for Legacy Applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), 1399–1408. https://doi.org/10.1145/1357054.1357274

107. Amy K. Karlson, Benjamin B. Bederson, and John SanGiovanni. 2005. AppLens and launchTile: Two Designs for One-handed Thumb Use on Small Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '05), 201–210. https://doi.org/10.1145/1054972.1055001

108. Bonifaz Kaufmann and David Ahlström. 2013. Studying Spatial Memory and Map Navigation Performance on Projector Phones with Peephole Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13), 3173–3176. https://doi.org/10.1145/2470654.2466434

109. Hamed Ketabdar, Amirhossein Jahanbekam, Kamer Ali Yuksel, Tobias Hirsch, and Amin Haji Abolhassani. 2011. MagiMusic: Using Embedded Compass (Magnetic) Sensor for Touch-less Gesture Based Interaction with Digital Music Instruments in Mobile Devices. In *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction* (TEI '11), 241–244. https://doi.org/10.1145/1935701.1935749

110. Hamed Ketabdar, Peyman Moghadam, Babak Naderi, and Mehran Roshandel. 2012. Magnetic Signatures in Air for Mobile Devices. In *Proceedings of the 14th International*

Conference on Human-computer Interaction with Mobile Devices and Services Companion (MobileHCI '12), 185–188. https://doi.org/10.1145/2371664.2371705

111. Hamed Ketabdar, Mehran Roshandel, and Kamer Ali Yüksel. 2010. MagiWrite: Towards Touchless Digit Entry Using 3D Space Around Mobile Devices. In *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services* (MobileHCI '10), 443–446. https://doi.org/10.1145/1851600.1851701
112. Hamed Ketabdar, Mehran Roshandel, and Kamer Ali Yüksel. 2010. Towards Using Embedded Magnetic Field Sensor for Around Mobile Device 3D Interaction. In *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services* (MobileHCI '10), 153–156. https://doi.org/10.1145/1851600.1851626
113. Hamed Ketabdar and Kamer Ali Yüksel. 2010. Smart Ring: Controlling Call Alert Functionality Based on Audio and Movement Analysis. In *Proceedings of the 15th International Conference on Intelligent User Interfaces* (IUI '10), 415–416. https://doi.org/10.1145/1719970.1720049
114. Hamed Ketabdar, Kamer Ali Yüksel, and Mehran Roshandel. 2010. MagiTact: Interaction with Mobile Devices Based on Compass (Magnetic) Sensor. In *Proceedings of the 15th International Conference on Intelligent User Interfaces* (IUI '10), 413–414. https://doi.org/10.1145/1719970.1720048
115. J. Kim, J. He, K. Lyons, and T. Starner. 2007. The Gesture Watch: A Wireless Contact-free Gesture based Wrist Interface. In *2007 11th IEEE International Symposium on Wearable Computers*, 15–22. https://doi.org/10.1109/ISWC.2007.4373770
116. Sunjun Kim, Jihyun Yu, and Geehyuk Lee. 2012. Interaction Techniques for Unreachable Objects on the Touchscreen. In *Proceedings of the 24th Australian Computer-Human Interaction Conference* (OzCHI '12), 295–298. https://doi.org/10.1145/2414536.2414585
117. Kinect - Windows app development. 2016. Retrieved November 16, 2016 from https://developer.microsoft.com/en-us/windows/kinect

118. Jarrod Knibbe, Diego Martinez Plasencia, Christopher Bainbridge, Chee-Kin Chan, Jiawei Wu, Thomas Cable, Hassan Munir, and David Coyle. 2014. Extending Interaction for Smart Watches: Enabling Bimanual Around Device Control. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '14), 1891–1896. https://doi.org/10.1145/2559206.2581315
119. Sven Kratz, Patrick Chiu, and Maribeth Back. 2013. PointPose: Finger Pose Estimation for Touch Input on Mobile Devices Using a Depth Sensor. In *Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces* (ITS '13), 223–230. https://doi.org/10.1145/2512349.2512824
120. Sven Kratz and Michael Rohs. 2009. Hoverflow: Exploring Around-device Interaction with IR Distance Sensors. In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '09), 42:1–42:4. https://doi.org/10.1145/1613858.1613912
121. Sven Kratz, Michael Rohs, Dennis Guse, Jörg Müller, Gilles Bailly, and Michael Nischt. 2012. PalmSpace: Continuous Around-device Gestures vs. Multitouch for 3D Rotation Tasks on Mobile Devices. In *Proceedings of the International Working Conference on Advanced Visual Interfaces* (AVI '12), 181–188. https://doi.org/10.1145/2254556.2254590
122. Gordon Kurtenbach and William Buxton. 1994. User Learning and Performance with Marking Menus. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '94), 258–264. https://doi.org/10.1145/191666.191759
123. Gordon P. Kurtenbach, Abigail J. Sellen, and William A. S. Buxton. 1993. An Empirical Evaluation of Some Articulatory and Cognitive Aspects of Marking Menus. *Hum.-Comput. Interact.* 8, 1: 1–23. https://doi.org/10.1207/s15327051hci0801_1
124. Jianwei Lai and Dongsong Zhang. 2014. ExtendedThumb: A Motion-based Virtual Thumb for Improving One-handed Target Acquisition on Touch-screen Mobile Devices. In *Proceedings of the Extended Abstracts of the 32Nd Annual ACM Conference on Human Factors in Computing Systems* (CHI EA '14), 1825–1830. https://doi.org/10.1145/2559206.2581158

125. Gierad Laput, Robert Xiao, Xiang "Anthony" Chen, Scott E. Hudson, and Chris Harrison. 2014. Skin Buttons: Cheap, Small, Low-powered and Clickable Fixed-icon Laser Projectors. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), 389–394. https://doi.org/10.1145/2642918.2647356
126. Mathieu Le Goc, Stuart Taylor, Shahram Izadi, and Cem Keskin. 2014. A Low-cost Transparent Electric Field Sensor for 3D Interaction on Mobile Devices. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems* (CHI '14), 3167–3170. https://doi.org/10.1145/2556288.2557331
127. Leap Motion. 2016. Retrieved November 16, 2016 from https://www.leapmotion.com/
128. Andrea Leganchuk, Shumin Zhai, and William Buxton. 1998. Manual and Cognitive Benefits of Two-handed Input: An Experimental Study. *ACM Trans. Comput.-Hum. Interact.* 5, 4: 326–359. https://doi.org/10.1145/300520.300522
129. Frank Chun Yat Li, David Dearman, and Khai N. Truong. 2009. Virtual Shelves: Interactions with Orientation Aware Devices. In *Proceedings of the 22Nd Annual ACM Symposium on User Interface Software and Technology* (UIST '09), 125–128. https://doi.org/10.1145/1622176.1622200
130. I. Scott MacKenzie. 1992. Fitts' Law As a Research and Design Tool in Human-computer Interaction. *Hum.-Comput. Interact.* 7, 1: 91–139. https://doi.org/10.1207/s15327051hci0701_3
131. I. Scott MacKenzie and Yves Guiard. 2001. The Two-handed Desktop Interface: Are We There Yet? In *CHI '01 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '01), 351–352. https://doi.org/10.1145/634067.634275
132. Sylvain Malacria, Eric Lecolinet, and Yves Guiard. 2010. Clutch-free Panning and Integrated Pan-zoom Control on Touch-sensitive Surfaces: The Cyclostar Approach. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '10), 2615–2624. https://doi.org/10.1145/1753326.1753724
133. Man-Systems Integration Standards (MSIS). 2016. Retrieved June 30, 2016 from http://msis.jsc.nasa.gov/

134. Market Research on What People Watch, Listen To and Buy | Nielsen. 2016. Retrieved December 3, 2016 from http://www.nielsen.com/ca/en.html

135. Diego Martinez Plasencia, Edward Joyce, and Sriram Subramanian. 2014. MisTable: Reach-through Personal Screens for Tabletops. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), 3493–3502. https://doi.org/10.1145/2556288.2557325

136. Maurice R. Masliah and Paul Milgram. 2000. Measuring the Allocation of Control in a 6 Degree-of-freedom Docking Experiment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '00), 25–32. https://doi.org/10.1145/332040.332403

137. Robert R. McCrae and Oliver P. John. 1992. An Introduction to the Five-Factor Model and Its Applications. *Journal of Personality* 60, 2: 175–215. https://doi.org/10.1111/j.1467-6494.1992.tb00970.x

138. Ross McLachlan and Stephen Brewster. 2015. Bimanual Input for Tablet Devices with Pressure and Multi-Touch Gestures. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '15), 547–556. https://doi.org/10.1145/2785830.2785878

139. Mobile eCommerce Statistics 2016 | Rise in Percentage of Mobile mCommerce Shopping. 2016. Retrieved November 16, 2016 from http://www.outerboxdesign.com/web-design-articles/mobile-ecommerce-statistics

140. Mobile marketing statistics 2016. 2016. Retrieved June 30, 2016 from http://www.smartinsights.com/mobile-marketing/mobile-marketing-analytics/mobile-marketing-statistics/

141. Mobile Now Exceeds PC: The Biggest Shift Since the Internet Began | Search Engine Watch. 2016. Retrieved June 30, 2016 from https://searchenginewatch.com/sew/opinion/2353616/mobile-now-exceeds-pc-the-biggest-shift-since-the-internet-began

142. Calkin S. Montero, Jason Alexander, Mark T. Marshall, and Sriram Subramanian. 2010. Would You Do That?: Understanding Social Acceptance of Gestural Interfaces. In

Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '10), 275–278. https://doi.org/10.1145/1851600.1851647

143. Heiko Müller, Andreas Löcken, Wilko Heuten, and Susanne Boll. 2014. Sparkle: An Ambient Light Display for Dynamic Off-screen Points of Interest. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational* (NordiCHI '14), 51–60. https://doi.org/10.1145/2639189.2639205

144. Masahiro Nakamura, Go Inaba, Jun Tamaoki, Kazuhito Shiratori, and Junichi Hoshino. 2006. Mounting and Application of Bubble Display System: Bubble Cosmos. In *Proceedings of the 2006 ACM SIGCHI International Conference on Advances in Computer Entertainment Technology* (ACE '06). https://doi.org/10.1145/1178823.1178879

145. Kei Nakatsuma, Hiroyuki Shinoda, Yasutoshi Makino, Katsunari Sato, and Takashi Maeno. 2011. Touch Interface on Back of the Hand. In *ACM SIGGRAPH 2011 Posters* (SIGGRAPH '11), 39:1–39:1. https://doi.org/10.1145/2037715.2037760

146. Rajalakshmi Nandakumar, Vikram Iyer, Desney Tan, and Shyamnath Gollakota. 2016. FingerIO: Using Active Sonar for Fine-Grained Finger Tracking. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), 1515–1525. https://doi.org/10.1145/2858036.2858580

147. Takehiro Niikura, Yuki Hirobe, Alvaro Cassinelli, Yoshihiro Watanabe, Takashi Komuro, and Masatoshi Ishikawa. 2010. In-air Typing Interface for Mobile Devices with Vibration Feedback. In *ACM SIGGRAPH 2010 Emerging Technologies* (SIGGRAPH '10), 15:1–15:1. https://doi.org/10.1145/1836821.1836836

148. Stina Nylander, Terés Lundquist, and Andreas Brännström. 2009. At Home and with Computer Access: Why and Where People Use Cell Phones to Access the Internet. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '09), 1639–1642. https://doi.org/10.1145/1518701.1518951

149. Marianna Obrist, Sue Ann Seah, and Sriram Subramanian. 2013. Talking About Tactile Experiences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13), 1659–1668. https://doi.org/10.1145/2470654.2466220

150. OptiTrack: Motion Capture Systems. 2016. Retrieved June 30, 2016 from https://www.optitrack.com/
151. Anna Ostberg and Nada Matic. 2015. Hover Cursor: Improving Touchscreen Acquisition Of Small Targets With Hover-enabled Pre-selection. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (CHI EA '15), 1723–1728. https://doi.org/10.1145/2702613.2732903
152. Michel Pahud, Ken Hinckley, Shamsi Iqbal, Abigail Sellen, and Bill Buxton. 2013. Toward Compound Navigation Tasks on Mobiles via Spatial Manipulation. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services* (MobileHCI '13), 113–122. https://doi.org/10.1145/2493190.2493210
153. Pekka Parhi, Amy K. Karlson, and Benjamin B. Bederson. 2006. Target Size Study for One-handed Thumb Use on Small Touchscreen Devices. In *Proceedings of the 8th Conference on Human-computer Interaction with Mobile Devices and Services* (MobileHCI '06), 203–210. https://doi.org/10.1145/1152215.1152260
154. Shwetak N. Patel and Gregory D. Abowd. 2007. Blui: Low-cost Localized Blowable User Interfaces. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology* (UIST '07), 217–220. https://doi.org/10.1145/1294211.1294250
155. R. L. Potter, L. J. Weldon, and B. Shneiderman. 1988. Improving the Accuracy of Touch Screens: An Experimental Evaluation of Three Strategies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '88), 27–32. https://doi.org/10.1145/57167.57171
156. Project Soli. 2016. Retrieved November 16, 2016 from https://atap.google.com/soli/
157. Qian Qin, Michael Rohs, and Sven Kratz. 2011. Dynamic Ambient Lighting for Mobile Devices. In *Proceedings of the 24th Annual ACM Symposium Adjunct on User Interface Software and Technology* (UIST '11 Adjunct), 51–52. https://doi.org/10.1145/2046396.2046418

158.I. Rakkolainen. 2007. How Feasible Are Star Wars Mid-air Displays. In *Information Visualization, 2007. IV '07. 11th International Conference*, 935–942. https://doi.org/10.1109/IV.2007.64

159.Rakuten.com. 2016. Retrieved November 30, 2016 from http://www.rakuten.com/

160. Gonzalo Ramos, Matthew Boulos, and Ravin Balakrishnan. 2004. Pressure Widgets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '04), 487–494. https://doi.org/10.1145/985692.985754

161. Gabriel Reyes, Dingtian Zhang, Sarthak Ghosh, Pratik Shah, Jason Wu, Aman Parnami, Bailey Bercik, Thad Starner, Gregory D. Abowd, and W. Keith Edwards. 2016. Whoosh: Non-voice Acoustics for Low-cost, Hands-free, and Rapid Input on Smartwatches. In *Proceedings of the 2016 ACM International Symposium on Wearable Computers* (ISWC '16), 120–127. https://doi.org/10.1145/2971763.2971765

162.Julie Rico and Stephen Brewster. 2010. Usable Gestures for Mobile Interfaces: Evaluating Social Acceptability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '10), 887–896. https://doi.org/10.1145/1753326.1753458

163.Sami Ronkainen, Jonna Häkkilä, Saana Kaleva, Ashley Colley, and Jukka Linjama. 2007. Tap Input As an Embedded Interaction Method for Mobile Devices. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction* (TEI '07), 263–270. https://doi.org/10.1145/1226969.1227023

164. Anne Roudaut, Stéphane Huot, and Eric Lecolinet. 2008. TapTap and MagStick: Improving One-handed Target Acquisition on Small Touch-screens. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (AVI '08), 146–153. https://doi.org/10.1145/1385569.1385594

165.Deepak Ranjan Sahoo, Takuto Nakamura, Asier Marzo, Themis Omirou, Michihiro Asakawa, and Sriram Subramanian. 2016. JOLED: A Mid-air Display Based on Electrostatic Rotation of Levitated Janus Objects. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (UIST '16), 437–448. https://doi.org/10.1145/2984511.2984549

166. Samsung Canada | Mobile | TV | Home Appliances. 2016. Retrieved November 18, 2016 from http://www.samsung.com/ca/home/

167. Joey Scarr, Andy Cockburn, Carl Gutwin, and Andrea Bunt. 2012. Improving Command Selection with CommandMaps. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '12), 257–266. https://doi.org/10.1145/2207676.2207713

168. Stefano Scheggi, Leonardo Meli, Claudio Pacchierotti, and Domenico Prattichizzo. 2015. Touch the Virtual Reality: Using the Leap Motion Controller for Hand Tracking and Wearable Tactile Devices for Immersive Haptic Rendering. In *ACM SIGGRAPH 2015 Posters* (SIGGRAPH '15), 31:1–31:1. https://doi.org/10.1145/2787626.2792651

169. Paul Schmieder, John Hosking, Andrew Luxton-Reilly, and Beryl Plimmer. 2013. Thumbs Up: 3D Gesture Input on Mobile Phones Using the Front Facing Camera. In *Human-Computer Interaction – INTERACT 2013*, Paula Kotzé, Gary Marsden, Gitte Lindgaard, Janet Wesson and Marco Winckler (eds.). Springer Berlin Heidelberg, 318–336. Retrieved September 21, 2016 from http://link.springer.com.uml.idm.oclc.org/chapter/10.1007/978-3-642-40480-1_20

170. Sue Ann Seah, Diego Martinez Plasencia, Peter D. Bennett, Abhijit Karnik, Vlad Stefan Otrocol, Jarrod Knibbe, Andy Cockburn, and Sriram Subramanian. 2014. SensaBubble: A Chrono-sensory Mid-air Display of Sight and Smell. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems* (CHI '14), 2863–2872. https://doi.org/10.1145/2556288.2557087

171. Karsten Seipp and Kate Devlin. 2014. Backpat: Improving One-handed Touchscreen Operation by Patting the Back of the Device. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '14), 555–558. https://doi.org/10.1145/2559206.2574811

172. Marcos Serrano, Khalad Hasan, Barrett Ens, Xing-Dong Yang, and Pourang Irani. 2015. Smartwatches + Head-Worn Displays: the "New" Smartphone. In *CHI 2015 Workshop on Mobile Collocated Interactions*. Retrieved November 28, 2016 from http://oatao.univ-toulouse.fr/15288/

173. Kang Shi, Pourang Irani, Sean Gustafson, and Sriram Subramanian. 2008. PressureFish: A Method to Improve Control of Discrete Pressure-based Input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), 1295–1298. https://doi.org/10.1145/1357054.1357256
174. Sujan Shrestha. 2007. Mobile Web Browsing: Usability Study. In *Proceedings of the 4th International Conference on Mobile Technology, Applications, and Systems and the 1st International Symposium on Computer Human Interaction in Mobile Technology* (Mobility '07), 187–194. https://doi.org/10.1145/1378063.1378094
175. Rajinder Sodhi, Matthew Glisson, and Ivan Poupyrev. 2013. AIREAL: Tactile Gaming Experiences in Free Air. In *ACM SIGGRAPH 2013 Emerging Technologies* (SIGGRAPH '13), 2:1–2:1. https://doi.org/10.1145/2503368.2503370
176. Rajinder Sodhi, Ivan Poupyrev, Matthew Glisson, and Ali Israr. 2013. AIREAL: Interactive Tactile Experiences in Free Air. *ACM Trans. Graph.* 32, 4: 134:1–134:10. https://doi.org/10.1145/2461912.2462007
177. SoftKinetic - 3D Vision Leader. 2016. Retrieved November 16, 2016 from https://www.softkinetic.com/
178. Timothy Sohn, Kevin A. Li, William G. Griswold, and James D. Hollan. 2008. A Diary Study of Mobile Information Needs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), 433–442. https://doi.org/10.1145/1357054.1357125
179. Jie Song, Gábor Sörös, Fabrizio Pece, Sean Ryan Fanello, Shahram Izadi, Cem Keskin, and Otmar Hilliges. 2014. In-air Gestures Around Unmodified Mobile Devices. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), 319–329. https://doi.org/10.1145/2642918.2647373
180. Yang Song, Hao Ma, Hongning Wang, and Kuansan Wang. 2013. Exploring and Exploiting User Search Behavior on Mobile and Tablet Devices to Improve Search Relevance. In *Proceedings of the 22Nd International Conference on World Wide Web* (WWW '13), 1201–1212. https://doi.org/10.1145/2488388.2488493

181. Martin Spindler, Martin Schuessler, Marcel Martsch, and Raimund Dachselt. 2014. Pinch-drag-flick vs. Spatial Input: Rethinking Zoom & Pan on Mobile Displays. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems* (CHI '14), 1113–1122. https://doi.org/10.1145/2556288.2557028
182. Adam J. Sporka, Sri H. Kurniawan, Murni Mahmud, and Pavel Slavík. 2006. Non-speech Input and Speech Recognition for Real-time Control of Computer Games. In *Proceedings of the 8th International ACM SIGACCESS Conference on Computers and Accessibility* (Assets '06), 213–220. https://doi.org/10.1145/1168987.1169023
183. Christian Steins, Sean Gustafson, Christian Holz, and Patrick Baudisch. 2013. Imaginary Devices: Gesture-based Interaction Mimicking Traditional Input Devices. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services* (MobileHCI '13), 123–126. https://doi.org/10.1145/2493190.2493208
184. Structure Sensor - 3D scanning, augmented reality, and more for mobile devices. 2016. Retrieved November 16, 2016 from http://structure.io/
185. Kazuki Takashima, Nana Shinshi, and Yoshifumi Kitamura. 2015. Exploring Boundless Scroll by Extending Motor Space. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '15), 557–566. https://doi.org/10.1145/2785830.2785884
186. Tango. 2016. Retrieved November 17, 2016 from https://get.google.com/tango/
187. Taobao. 2016. Retrieved November 30, 2016 from https://world.taobao.com/
188. Target : Expect More. Pay Less. 2016. Retrieved November 30, 2016 from http://intl.target.com/
189. The Home Depot Canada. 2016. Retrieved November 30, 2016 from https://www.homedepot.ca/en/home.html
190. tmall.com. 2016. Retrieved November 30, 2016 from https://www.tmall.com/
191. Ryoko Ueoka, Mami Yamaguchi, and Yuka Sato. 2016. Interactive Cheek Haptic Display with Air Vortex Rings for Stress Modification. In *Proceedings of the 2016 CHI*

Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16), 1766–1771. https://doi.org/10.1145/2851581.2892299

192. Wouter Van Vlaenderen, Jens Brulmans, Jo Vermeulen, and Johannes Schöning. 2015. WatchMe: A Novel Input Method Combining a Smartwatch and Bimanual Interaction. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (CHI EA '15), 2091–2095. https://doi.org/10.1145/2702613.2732789

193. VICON. Motion Capture Systems. *VICON*. Retrieved November 16, 2016 from http://www.vicon.com

194. Visbox, Inc. 2016. Retrieved November 21, 2016 from http://www.visbox.com/

195. Daniel Vogel and Patrick Baudisch. 2007. Shift: A Technique for Operating Pen-based Interfaces Using Touch. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '07), 657–666. https://doi.org/10.1145/1240624.1240727

196. Walmart.com. 2016. Retrieved November 30, 2016 from https://www.walmart.com/

197. QianYing Wang, Tony Hsieh, and Andreas Paepcke. 2009. Piles across space: Breaking the real-estate barrier on small-display devices. *International Journal of Human-Computer Studies* 67, 4: 349–365. https://doi.org/10.1016/j.ijhcs.2008.10.004

198. Saiwen Wang, Jie Song, Jaime Lien, Ivan Poupyrev, and Otmar Hilliges. 2016. Interacting with Soli: Exploring Fine-Grained Dynamic Gesture Recognition in the Radio-Frequency Spectrum. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (UIST '16), 851–860. https://doi.org/10.1145/2984511.2984565

199. Hongyi Wen, Julian Ramos Rojas, and Anind K. Dey. 2016. Serendipity: Finger Gesture Recognition Using an Off-the-Shelf Smartwatch. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), 3847–3851. https://doi.org/10.1145/2858036.2858466

200. Julie R. Wiliamson, Andrew Crossan, and Stephen Brewster. 2011. Multimodal Mobile Interactions: Usability Studies in Real World Settings. In *Proceedings of the*

13th International Conference on Multimodal Interfaces (ICMI '11), 361–368. https://doi.org/10.1145/2070481.2070551

201.Julie R. Williamson, Stephen Brewster, and Rama Vennelakanti. 2013. Mo!Games: Evaluating Mobile Gestures in the Wild. In *Proceedings of the 15th ACM on International Conference on Multimodal Interaction* (ICMI '13), 173–180. https://doi.org/10.1145/2522848.2522874

202. Graham Wilson, Thomas Carter, Sriram Subramanian, and Stephen A. Brewster. 2014. Perception of Ultrasonic Haptic Feedback on the Hand: Localisation and Apparent Motion. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems* (CHI '14), 1133–1142. https://doi.org/10.1145/2556288.2557033

203. Anusha Withana, Roshan Peiris, Nipuna Samarasekara, and Suranga Nanayakkara. 2015. zSense: Enabling Shallow Depth Gesture Recognition for Greater Input Expressivity on Smart Wearables. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), 3661–3670. https://doi.org/10.1145/2702123.2702371

204. Robert Xiao, Greg Lew, James Marsanico, Divya Hariharan, Scott Hudson, and Chris Harrison. 2014. Toffee: Enabling Ad Hoc, Around-device Interaction with Acoustic Time-of-arrival Correlation. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services* (MobileHCI '14), 67–76. https://doi.org/10.1145/2628363.2628383

205. Xiangye Xiao, Qiong Luo, Dan Hong, Hongbo Fu, Xing Xie, and Wei-Ying Ma. 2009. Browsing on Small Displays by Transforming Web Pages into Hierarchically Structured Subpages. *ACM Trans. Web* 3, 1: 4:1–4:36. https://doi.org/10.1145/1462148.1462152

206. Wenchang Xu, Chun Yu, Jie Liu, and Yuanchun Shi. 2015. RegionalSliding. *Pervasive Mob. Comput.* 17, PA: 63–78. https://doi.org/10.1016/j.pmcj.2014.02.005

207. Yahoo. 2016. Retrieved November 16, 2016 from //ca.yahoo.com/

208. Xing-Dong Yang, Khalad Hasan, Neil Bruce, and Pourang Irani. 2013. Surround-see: Enabling Peripheral Vision on Smartphones During Active Use. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13), 291–300. https://doi.org/10.1145/2501988.2502049

209. Xing-Dong Yang, Pourang Irani, Pierre Boulanger, and Walter Bischof. 2009. One-handed Behind-the-display Cursor Input on Mobile Devices. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '09), 4501–4506. https://doi.org/10.1145/1520340.1520690

210. Xing-Dong Yang, Edward Mak, Pourang Irani, and Walter F. Bischof. 2009. Dual-Surface Input: Augmenting One-handed Interaction with Coordinated Front and Behind-the-screen Input. In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '09), 5:1–5:10. https://doi.org/10.1145/1613858.1613865

211. Koji Yatani, Kurt Partridge, Marshall Bern, and Mark W. Newman. 2008. Escape: A Target Selection Technique Using Visually-cued Gestures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), 285–294. https://doi.org/10.1145/1357054.1357104

212. Hyunjin Yoo, Jungwon Yoon, and Hyunsoo Ji. 2015. Index Finger Zone: Study on Touchable Area Expandability Using Thumb and Index Finger. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct* (MobileHCI '15), 803–810. https://doi.org/10.1145/2786567.2793704

213. Polle T. Zellweger, Jock D. Mackinlay, Lance Good, Mark Stefik, and Patrick Baudisch. 2003. City Lights: Contextual Views in Minimal Space. In *CHI '03 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '03), 838–839. https://doi.org/10.1145/765891.766022

214. Chen Zhao, Ke-Yu Chen, Md Tanvir Islam Aumi, Shwetak Patel, and Matthew S. Reynolds. 2014. SideSwipe: Detecting In-air Gestures Around Mobile Devices Using Actual GSM Signal. In *Proceedings of the 27th Annual ACM Symposium on User*

Interface Software and Technology (UIST '14), 527–534.
https://doi.org/10.1145/2642918.2647380

215. Shengdong Zhao and Ravin Balakrishnan. 2004. Simple vs. Compound Mark Hierarchical Marking Menus. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology* (UIST '04), 33–42. https://doi.org/10.1145/1029632.1029639

216. Shengdong Zhao, Pierre Dragicevic, Mark Chignell, Ravin Balakrishnan, and Patrick Baudisch. 2007. Earpod: Eyes-free Menu Selection Using Touch Input and Reactive Audio Feedback. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '07), 1395–1404. https://doi.org/10.1145/1240624.1240836

217. Junhan Zhou, Yang Zhang, Gierad Laput, and Chris Harrison. 2016. AuraSense: Enabling Expressive Around-Smartwatch Interactions with Electric Field Sensing. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (UIST '16), 81–86. https://doi.org/10.1145/2984511.2984568

Appendix A

ASSOCIATION FOR COMPUTING MACHINERY, INC. LICENSE TERMS AND CONDITIONS

Nov 30, 2016

This Agreement between Mohammad K Hasan ("You") and Association for Computing Machinery, Inc. ("Association for Computing Machinery, Inc.") consists of your license details and the terms and conditions provided by Association for Computing Machinery, Inc. and Copyright Clearance Center.

All payments must be made in full to CCC. For payment instructions, please see information listed at the bottom of this form.

License Number	3998861478774
License date	Nov 30, 2016
Licensed Content Publisher	Association for Computing Machinery, Inc.
Licensed Content Publication	Proceedings
Licensed Content Title	Ad-binning: leveraging around device space for storing, browsing and retrieving mobile device content
Licensed Content Author	Khalad Hasan, et al
Licensed Content Date	Apr 27, 2013
Type of Use	Thesis/Dissertation
Requestor type	Author of this ACM article
Is reuse in the author's own new work?	Yes
Format	Print and electronic
Portion	Full article
Will you be translating?	No
Order reference number	
Title of your thesis/dissertation	Around-Device Interaction for Exploring Large Information Spaces on Mobile Devices
Expected completion date	Apr 2017
Estimated size (pages)	200
Requestor Location	Mohammad K Hasan

Winnipeg, MB R3T3M2
Canada
Attn: Mohammad K Hasan

Billing Type	Credit Card
Credit card info	Master Card ending in 0908
Credit card expiration	10/2020
Total	10.75 CAD

Terms and Conditions

Rightslink Terms and Conditions for ACM Material

1. The publisher of this copyrighted material is Association for Computing Machinery, Inc. (ACM). By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at).
2. ACM reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.
3. ACM hereby grants to licensee a non-exclusive license to use or republish this ACM-copyrighted material* in secondary works (especially for commercial distribution) with the stipulation that consent of the lead author has been obtained independently. Unless otherwise stipulated in a license, grants are for one-time use in a single edition of the work, only with a maximum distribution equal to the number that you identified in the licensing process. Any additional form of republication must be specified according to the terms included at the time of licensing.

*Please note that ACM cannot grant republication or distribution licenses for embedded third-party material. You must confirm the ownership of figures, drawings and artwork prior to use.

4. Any form of republication or redistribution must be used within 180 days from the date stated on the license and any electronic posting is limited to a period of six months unless an extended term is selected during the licensing process. Separate subsidiary and subsequent republication licenses must be purchased to redistribute copyrighted material on an extranet. These licenses may be exercised anywhere in the world.
5. Licensee may not alter or modify the material in any manner (except that you may use, within the scope of the license granted, one or more excerpts from the copyrighted material, provided that the process of excerpting does not alter the meaning of the material or in any way reflect negatively on the publisher or any writer of the material).
6. Licensee must include the following copyright and permission notice in connection with any reproduction of the licensed material: "[Citation] © YEAR Association for Computing Machinery, Inc. Reprinted by permission." Include the article DOI as a link to the definitive version in the ACM Digital Library. Example: Charles, L. "How to Improve Digital Rights Management," Communications of the ACM, Vol. 51:12, © 2008 ACM, Inc. http://doi.acm.org/10.1145/nnnnnn.nnnnnn (where nnnnnn.nnnnnn is replaced by the actual number).

7. Translation of the material in any language requires an explicit license identified during the licensing process. Due to the error-prone nature of language translations, Licensee must include the following copyright and permission notice and disclaimer in connection with any reproduction of the licensed material in translation: "This translation is a derivative of ACM-copyrighted material. ACM did not prepare this translation and does not guarantee that it is an accurate copy of the originally published work. The original intellectual property contained in this work remains the property of ACM."
8. You may exercise the rights licensed immediately upon issuance of the license at the end of the licensing transaction, provided that you have disclosed complete and accurate details of your proposed use. No license is finally effective unless and until full payment is received from you (either by CCC or ACM) as provided in CCC's Billing and Payment terms and conditions.
9. If full payment is not received within 90 days from the grant of license transaction, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted.
10. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.
11. ACM makes no representations or warranties with respect to the licensed material and adopts on its own behalf the limitations and disclaimers established by CCC on its behalf in its Billing and Payment terms and conditions for this licensing transaction.
12. You hereby indemnify and agree to hold harmless ACM and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.
13. This license is personal to the requestor and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.
14. This license may not be amended except in a writing signed by both parties (or, in the case of ACM, by CCC on its behalf).
15. ACM hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and ACM (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.
16. This license transaction shall be governed by and construed in accordance with the laws of New York State. You hereby agree to submit to the jurisdiction of the federal and state courts located in New York for purposes of resolving any disputes that may arise in connection with this licensing transaction.
17. There are additional terms and conditions, established by Copyright Clearance Center, Inc. ("CCC") as the administrator of this licensing service that relate to billing and payment

for licenses provided through this service. Those terms and conditions apply to each transaction as if they were restated here. As a user of this service, you agreed to those terms and conditions at the time that you established your account, and you may see them again at any time at http://myaccount.copyright.com

18. Thesis/Dissertation: This type of use requires only the minimum administrative fee. It is not a fee for permission. Further reuse of ACM content, by ProQuest/UMI or other document delivery providers, or in republication requires a separate permission license and fee. Commercial resellers of your dissertation containing this article must acquire a separate license.

Special Terms:

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

ASSOCIATION FOR COMPUTING MACHINERY, INC. LICENSE TERMS AND CONDITIONS

Nov 30, 2016

This Agreement between Mohammad K Hasan ("You") and Association for Computing Machinery, Inc. ("Association for Computing Machinery, Inc.") consists of your license details and the terms and conditions provided by Association for Computing Machinery, Inc. and Copyright Clearance Center.

All payments must be made in full to CCC. For payment instructions, please see information listed at the bottom of this form.

License Number	3998870171175
License date	Nov 30, 2016
Licensed Content Publisher	Association for Computing Machinery, Inc.
Licensed Content Publication	Proceedings
Licensed Content Title	Are you comfortable doing that?: acceptance studies of around-device gestures in and for public settings
Licensed Content Author	David Ahlström, et al
Licensed Content Date	Sep 23, 2014
Type of Use	Thesis/Dissertation
Requestor type	Author of this ACM article
Is reuse in the author's own new work?	Yes
Format	Print and electronic
Portion	Full article
Will you be translating?	No
Order reference number	
Title of your thesis/dissertation	Around-Device Interaction for Exploring Large Information Spaces on Mobile Devices
Expected completion date	Apr 2017
Estimated size (pages)	200
Requestor Location	Mohammad K Hasan Winnipeg, MB R3T3M2 Canada Attn: Mohammad K Hasan
Billing Type	Credit Card
Credit card info	Master Card ending in 0908

Credit card expiration	10/2020
Total	10.75 CAD

Terms and Conditions

Rightslink Terms and Conditions for ACM Material

1. The publisher of this copyrighted material is Association for Computing Machinery, Inc. (ACM). By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at).
2. ACM reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.
3. ACM hereby grants to licensee a non-exclusive license to use or republish this ACM-copyrighted material* in secondary works (especially for commercial distribution) with the stipulation that consent of the lead author has been obtained independently. Unless otherwise stipulated in a license, grants are for one-time use in a single edition of the work, only with a maximum distribution equal to the number that you identified in the licensing process. Any additional form of republication must be specified according to the terms included at the time of licensing.
*Please note that ACM cannot grant republication or distribution licenses for embedded third-party material. You must confirm the ownership of figures, drawings and artwork prior to use.
4. Any form of republication or redistribution must be used within 180 days from the date stated on the license and any electronic posting is limited to a period of six months unless an extended term is selected during the licensing process. Separate subsidiary and subsequent republication licenses must be purchased to redistribute copyrighted material on an extranet. These licenses may be exercised anywhere in the world.
5. Licensee may not alter or modify the material in any manner (except that you may use, within the scope of the license granted, one or more excerpts from the copyrighted material, provided that the process of excerpting does not alter the meaning of the material or in any way reflect negatively on the publisher or any writer of the material).
6. Licensee must include the following copyright and permission notice in connection with any reproduction of the licensed material: "[Citation] © YEAR Association for Computing Machinery, Inc. Reprinted by permission." Include the article DOI as a link to the definitive version in the ACM Digital Library. Example: Charles, L. "How to Improve Digital Rights Management," Communications of the ACM, Vol. 51:12, © 2008 ACM, Inc. http://doi.acm.org/10.1145/nnnnnn.nnnnnn (where nnnnnn.nnnnnn is replaced by the actual number).
7. Translation of the material in any language requires an explicit license identified during the licensing process. Due to the error-prone nature of language translations, Licensee must include the following copyright and permission notice and disclaimer in connection with any reproduction of the licensed material in translation: "This translation is a derivative of ACM-copyrighted material. ACM did not prepare this translation and does not guarantee

that it is an accurate copy of the originally published work. The original intellectual property contained in this work remains the property of ACM."

8. You may exercise the rights licensed immediately upon issuance of the license at the end of the licensing transaction, provided that you have disclosed complete and accurate details of your proposed use. No license is finally effective unless and until full payment is received from you (either by CCC or ACM) as provided in CCC's Billing and Payment terms and conditions.

9. If full payment is not received within 90 days from the grant of license transaction, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted.

10. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.

11. ACM makes no representations or warranties with respect to the licensed material and adopts on its own behalf the limitations and disclaimers established by CCC on its behalf in its Billing and Payment terms and conditions for this licensing transaction.

12. You hereby indemnify and agree to hold harmless ACM and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

13. This license is personal to the requestor and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.

14. This license may not be amended except in a writing signed by both parties (or, in the case of ACM, by CCC on its behalf).

15. ACM hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and ACM (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.

16. This license transaction shall be governed by and construed in accordance with the laws of New York State. You hereby agree to submit to the jurisdiction of the federal and state courts located in New York for purposes of resolving any disputes that may arise in connection with this licensing transaction.

17. There are additional terms and conditions, established by Copyright Clearance Center, Inc. ("CCC") as the administrator of this licensing service that relate to billing and payment for licenses provided through this service. Those terms and conditions apply to each transaction as if they were restated here. As a user of this service, you agreed to those terms and conditions at the time that you established your account, and you may see them again at any time at http://myaccount.copyright.com

ASSOCIATION FOR COMPUTING MACHINERY, INC. LICENSE TERMS AND CONDITIONS

Nov 30, 2016

This Agreement between Mohammad K Hasan ("You") and Association for Computing Machinery, Inc. ("Association for Computing Machinery, Inc.") consists of your license details and the terms and conditions provided by Association for Computing Machinery, Inc. and Copyright Clearance Center.

All payments must be made in full to CCC. For payment instructions, please see information listed at the bottom of this form.

License Number	3998870298405
License date	Nov 30, 2016
Licensed Content Publisher	Association for Computing Machinery, Inc.
Licensed Content Publication	Proceedings
Licensed Content Title	Thumbs-Up: 3D Spatial Thumb-Reachable Space for One-Handed Thumb Interaction on Smartphones
Licensed Content Author	Khalad Hasan, et al
Licensed Content Date	Oct 15, 2016
Type of Use	Thesis/Dissertation
Requestor type	Author of this ACM article
Is reuse in the author's own new work?	Yes
Format	Print and electronic
Portion	Full article
Will you be translating?	No
Order reference number	
Title of your thesis/dissertation	Around-Device Interaction for Exploring Large Information Spaces on Mobile Devices
Expected completion date	Apr 2017
Estimated size (pages)	200
Requestor Location	Mohammad K Hasan Winnipeg, MB R3T3M2 Canada Attn: Mohammad K Hasan
Billing Type	Credit Card
Credit card info	Master Card ending in 0908

Credit card expiration	10/2020
Total	10.75 CAD

Terms and Conditions

Rightslink Terms and Conditions for ACM Material

1. The publisher of this copyrighted material is Association for Computing Machinery, Inc. (ACM). By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at).
2. ACM reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.
3. ACM hereby grants to licensee a non-exclusive license to use or republish this ACM-copyrighted material* in secondary works (especially for commercial distribution) with the stipulation that consent of the lead author has been obtained independently. Unless otherwise stipulated in a license, grants are for one-time use in a single edition of the work, only with a maximum distribution equal to the number that you identified in the licensing process. Any additional form of republication must be specified according to the terms included at the time of licensing.
*Please note that ACM cannot grant republication or distribution licenses for embedded third-party material. You must confirm the ownership of figures, drawings and artwork prior to use.
4. Any form of republication or redistribution must be used within 180 days from the date stated on the license and any electronic posting is limited to a period of six months unless an extended term is selected during the licensing process. Separate subsidiary and subsequent republication licenses must be purchased to redistribute copyrighted material on an extranet. These licenses may be exercised anywhere in the world.
5. Licensee may not alter or modify the material in any manner (except that you may use, within the scope of the license granted, one or more excerpts from the copyrighted material, provided that the process of excerpting does not alter the meaning of the material or in any way reflect negatively on the publisher or any writer of the material).
6. Licensee must include the following copyright and permission notice in connection with any reproduction of the licensed material: "[Citation] © YEAR Association for Computing Machinery, Inc. Reprinted by permission." Include the article DOI as a link to the definitive version in the ACM Digital Library. Example: Charles, L. "How to Improve Digital Rights Management," Communications of the ACM, Vol. 51:12, © 2008 ACM, Inc. http://doi.acm.org/10.1145/nnnnnn.nnnnnn (where nnnnnn.nnnnnn is replaced by the actual number).
7. Translation of the material in any language requires an explicit license identified during the licensing process. Due to the error-prone nature of language translations, Licensee must include the following copyright and permission notice and disclaimer in connection with any reproduction of the licensed material in translation: "This translation is a derivative of ACM-copyrighted material. ACM did not prepare this translation and does not guarantee

that it is an accurate copy of the originally published work. The original intellectual property contained in this work remains the property of ACM."
8. You may exercise the rights licensed immediately upon issuance of the license at the end of the licensing transaction, provided that you have disclosed complete and accurate details of your proposed use. No license is finally effective unless and until full payment is received from you (either by CCC or ACM) as provided in CCC's Billing and Payment terms and conditions.
9. If full payment is not received within 90 days from the grant of license transaction, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted.
10. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.
11. ACM makes no representations or warranties with respect to the licensed material and adopts on its own behalf the limitations and disclaimers established by CCC on its behalf in its Billing and Payment terms and conditions for this licensing transaction.
12. You hereby indemnify and agree to hold harmless ACM and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.
13. This license is personal to the requestor and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.
14. This license may not be amended except in a writing signed by both parties (or, in the case of ACM, by CCC on its behalf).
15. ACM hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and ACM (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.
16. This license transaction shall be governed by and construed in accordance with the laws of New York State. You hereby agree to submit to the jurisdiction of the federal and state courts located in New York for purposes of resolving any disputes that may arise in connection with this licensing transaction.
17. There are additional terms and conditions, established by Copyright Clearance Center, Inc. ("CCC") as the administrator of this licensing service that relate to billing and payment for licenses provided through this service. Those terms and conditions apply to each transaction as if they were restated here. As a user of this service, you agreed to those terms and conditions at the time that you established your account, and you may see them again at any time at http://myaccount.copyright.com