

Maptimizer: Using Optimization to Tailor Tactile Maps to Users Needs

Megan Hofmann
Human Computer
Interaction Institute,
Carnegie Mellon University
Pittsburgh, PA

Scott E. Hudson
Human Computer
Interaction Institute,
Carnegie Mellon University
Pittsburgh, PA

Kelly Mack, Jessica Birchfield, Jerry Cao,
Autumn Hughes, Shriya Kurpad, Kathryn J Lum,
Emily Warnock, Anat Caspi
Paul G. Allen
School of Computer Science,
University of Washington
Seattle, WA

Jennifer Mankoff
Paul G. Allen
School of Computer Science,
University of Washington
Seattle, WA

ABSTRACT

Tactile maps can help people with visual impairments navigate and familiarize themselves with unfamiliar locations. Ideally, tactile maps are created by considering an individual’s unique needs and abilities because of their limited space for representation. However, significant customization is not supported by existing tools for generating tactile maps. We present Maptimizer which generates tactile maps that are customized users’ preferences and requirements. Maptimizer uses a two stage optimization process to pair representations with geographic information and tune those representations to present that information more clearly. In a user study, Maptimizer helped participants more successfully and efficiently identify locations of interest in unknown areas. These results demonstrate the utility of optimization techniques and generative design in complex accessibility domains.

1 INTRODUCTION

Tactile Maps display information about a geographic location using a set of raised tactile features that can support navigation by people with visual impairments. When used alongside other navigation aids, a tactile map can help the user better understand a geographic space and the relationship between different geographic features. For instance, a tactile map could help a user identify a geographic feature that can support them in orienting themselves, such as a loud fountain at the center of a University courtyard. However, to better support navigation, tactile maps must be more widely available and include information that is tailored to the user’s specific needs, their abilities, and the location that the map represents.

The availability of tactile maps is limited by the inaccessibility of tools for generating tactile maps [11] and the lack of customization accessible to users [31]. The increasing availability of consumer 3D printers offers an opportunity to produce complex tactile maps automatically, without the aid of a sighted cartographer. A variety of approaches have been explored (e.g., [11, 17, 18, 31]), each focusing on designing different tactile representations of geographic features which can be automatically sourced from online repositories like OpenStreetMap [25]. In a few cases, visually-impaired

users are afforded some customizing opportunities ranging from setting the location and scale of a map [19] to individually tailoring how different geographic features are represented [31].

We argue that the dichotomy between non-customizable and overly customizable but burdensome interfaces leaves users struggling to create maps that better fit their needs. On the one hand, users often require more customization options to ensure that they can easily read the map and that it contains the information that is most important to them. On the other hand, making a wide range of parameters customizable creates a complex design task. Ideally, the user would be able to rapidly iterate and create a variety of maps to meet their needs. However, such iterations are largely inaccessible because 3D printing maps is time consuming and usually requires the assistance of a sighted person to operate the machine. Many users may not consider this extensive design task worth their efforts. Further, the user alone may lack critical contextual information that will make the map more usable. This is especially true when the map displays a location unfamiliar to the user—a common case. We require a tool that offers a variety of customization opportunities but can also rapidly iterate over a variety of tactile map designs and incorporate context that is unavailable to the user.

We offer an alternative approach to tactile map customization, rooted in principles of Ability-Based Design [32]. We contribute *Mapimizer*, an optimization-based generative design tool for creating highly-customized tactile maps. Maptimizer gathers user’s preferences from a screen-readable web application and uses their preferences to generate a tactile map that optimizes the communicability of the tactile representations and the informativeness of the geographic features embedded in the map, while minimizing attention costs such as clutter and information-density. When creating a tactile map, Maptimizer adjusts how geographic features are represented based on what the user prefers and can best distinguish given their abilities. Further, Maptimizer optimizes maps based on contextual information about the location. We use optimization methods (e.g., linear programming [33], Ant Colony Optimization [9]) to explore the wide space of tactile maps for a design that best fits the users needs at a given location.

Following a review of related work, we present Maptimizer’s method for creating tactile maps, user-interface, and optimization algorithm. We then present a user-evaluation which demonstrates that Maptimizer’s optimized tactile maps support the identification of new locations over tactile maps that are customized entirely by the user or offer no customization parameters. Our study with six blind/low vision individuals reveals opportunities to further improve tactile maps using our optimization-based approach.

2 RELATED WORK

2.1 3D Printing Tactile Graphics

Recently, researchers have leveraged the emergence of consumer-grade 3D printing to produce new forms of assistive technology. In particular, a variety of new methods for creating tactile graphics for people who have visual impairments have been introduced in a range of domains (e.g., educational models [5, 28], screen navigation [20, 34], appliance overlays [13], and tactile maps [11, 17, 31]). Work that has focused on developing accessible interfaces for creating these tactile graphics (e.g., [5, 13, 31]) tend to simplify the modeling process to a small set of parameters that can be set by the user. For example, Brown and Hurst generate tactile graphics that represent math equations (e.g., bar charts, line graphs) by having users fill in the equation and some specifications about the line thickness of the graphic, and the system generates a tactile version of the equation plotted [5]. Taylor et al. [31] use a similar approach to generate tactile maps by having the user input a specific geographic region and map preferences, and the system generates a corresponding tactile map for 3D printing. As the number of parameters users must set increases (e.g., only a few for math equations [5], potentially dozens for tactile maps), the design task becomes complex and cumbersome. Indeed, Taylor et al’s users sometimes struggled to narrow down the map to the set of information they needed while creating an easy to read result. Just like any other designer, a visually-impaired user may need to iterate over a variety of parameters to create the best tactile graphic. Unfortunately, iteration over 3D models is not screen reader accessible, and little work has explored ways to support modeling by people with visual impairments [29].

2.2 Tactile Maps for Navigation

Tactile maps are a critical type of tactile graphic worth designing. Previous work has demonstrated the potential benefits and challenges of using tactile maps as a way-finding and orientating tool for people who are blind or have low vision. Studies have shown that context-awareness of hazards and the immediacy of a changing environment affects the ability of a visually-impaired person to orient and navigate confidently [4]. While access to rich information raises the comfort level associated with independent travel and orienting, it is important to determine which presented information will be most helpful for the user to avoid unnecessary information overload [14]. Providing users who have visual impairments with navigational information requires consideration of their specific needs and personalization to their interests. The MoBIC Travel Aid is particularly relevant, as it considers both the general needs of visually-impaired travellers and specific user preferences, drawing extensively from user feedback and involvement [26].

Several researchers have tackled the challenges of making tactile maps more useful, easier to fabricate, or inexpensive. For example, Taylor et al’s “TactileMaps.net” system can generate 3D printable maps with a small set of user-customizable options [31]. They note the challenge of enabling advanced customizations without making the interface too complex or inaccessible to screen readers. Gotzelmann et al. expanded on the use of 3D printing with *LucentMaps* which integrated 3D printed, capacitive-sensing materials into the tactile maps to create audio labels [11]. The audio feedback provided real world distances as well as more specific details of places like accessibility features, stores, or entertainment features. Similarly, Taylor et al. created tactile maps which would be placed over the phone in a 3D printed case, both of which had buttons that afford interaction with the map [30]. While *TactileMaps.net* [31] and *LucentMaps* [11] only produce 2.5D maps, Holloway et al. designed maps with more complex 3D icons that mirrored the real life objects they represented [17, 18]. These icons were easily recognizable and reportedly helped with route planning and mental model strength of the location. Each of these tools focus on new ways of producing tactile maps, however little focus is centered on the task of designing a customized map.

2.3 Optimization and Ability Based Design

One approach to streamlining the tactile graphic design process is to apply optimization methods, which have been applied in other accessibility contexts. For example, Guo et al. used computer vision techniques to create 3D printed tactile labels for appliances (e.g., microwaves) [13]. In the separate domain of grip design, Chen et al. defined a parameterized space of common assistive gripping models and related them to different gripping styles for people with mobility impairments [7]. Beyond accessibility, optimization methods are widely used in other digital fabrication domains (e.g., [8, 16, 35]).

Optimization methods are often applied in accessibility domains when an Ability-Based Design approach is followed [32]. This approach is particularly important because it enables designers to consider people with multiple disabilities rather than treating people as having only one disability [15]. Ability-based design emphasizes the value of customization. In particular, it calls for technologies to adapt to the abilities of its users through automatic customization. While this approach has inspired significant work in user-interface adaptation (e.g., [6, 10, 23, 24]), as well as in personalized accessible routing [2], there is little cross-over into digital fabrication research. Each of these tools rely on models of the user’s abilities and use optimization methods to adapt complex interfaces to those needs and abilities. Digital fabrication and design processes share many similarities to interface design, but no work has applied a similar approach as Gajos et al. [10, 23], Mott et al. [24], or Carter et al. [6].

2.4 Optimizing Map User Interfaces

Optimization techniques are often used to render digital maps. For example, Agrawala and Stolte [1] use a simulated annealing optimization method to simplify how routing information is displayed to a driver so that it is easy to understand and minimizes how much attention a driver needs to pay to a map while driving. Grabler et al. [12] also applied simulated annealing to design tourist maps

of city that present the most informative pieces of information to individual users, rather than presenting the same information to all users. This demonstrates the use of these optimization techniques to create customized maps that meet the needs of individual users. Similarly, Lee et al. [21, 22], use optimization techniques to select what features are displayed to the user, prioritizing features that provide the most important contextual information during route navigation. Individually, these systems present optimization techniques that focus on three key characteristics of maps: how communicative the representations of information are [12], how informative information displayed is [21, 22], and how much of the user’s attention is drawn from their primary task (e.g., driving) to the map itself [6]. These characteristics are critical to the design of high-quality customized tactile maps.

3 MAPTIMIZER

Maptimizer enables people with visual impairments to independently generate a 3D printable tactile map that is uniquely tailored to their preferences. In particular, we consider two types of preferences that will affect the quality and usability of the tactile maps. First, a user’s preferences may affect what types of *geographic features* are presented in the map. Beyond simply showing roadways and buildings, a user may need specific information about points of interest, accessible infrastructure (e.g. stairs, ramps, tactile markers, audible and tactile traffic signals), general infrastructure (e.g., walkways, entrances to buildings), and/ or amenities (e.g., benches, restrooms). Second, users may have preferences for how this information is represented (i.e., its *representation*), such as the types of symbols used to represent information and the size of those symbols. However, designing a tactile map requires more than a user’s preferences. The information they prefer must be supported by sufficient contextual information (e.g., features critical for orientation) to navigate effectively, and showing too much information can clutter the map making it difficult to read. *Maptimizer* is a generative design tool that uses optimization methods to generate maps that consider these different requirements and prioritize the end-user’s preferences. We use a two stage optimization process that first uses linear-programming to pair communicative representations to informative geographic features based on users’ preferences and second tunes parameters of those representations (e.g., the depth of a feature) to increase their communicability and reduce clutter on the map.

In *Maptimizer*, tactile maps are generated by pairing geographic features to representations of those features. To avoid user confusion, each representation can only be used for one geographic feature and each geographic feature can only be shown with one representation. Representations will generate a component of the tactile map 3D model. For instance, a *path* representation will create a series of raised lines and a *location* representation will create a peg that sticks out of the map to mark that location. Exemplar representations are presented in Figure 1. These representations are used to present three types of *geographic features*: regions (e.g., buildings, water ways), paths (e.g., roads, sidewalks, bus routes), or locations (e.g., points of interests, benches, tactile pavement). For example, a raised path could represent roads or footpaths but cannot represent a set of benches or buildings.

3.1 Geographic Features


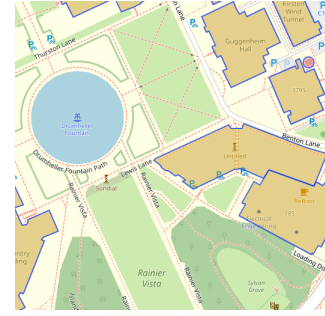

We gather information about geographic features from the OpenStreetMap API [25]. Which geographic features are available depends on the location and data available through OpenStreetMap. Generally, most areas have data about major architectural features (e.g., roads, buildings), and ecological features (e.g., green spaces, water ways). Some areas may also include details about amenities or accessibility features (e.g. accessible traffic signals, curb-cuts, tactile pavement). In many locations, the types of information available can be significantly limited and are being addressed by related work [3, 27]. Our goal is to collect a wide variety of information from OpenStreetMap and enable users to determine which pieces of information are most critical to them. They may consider points of interests, routes they intend to take, or accessibility features that will help them plan a visit to the area.

There are two important attributes of a geographic feature: the user’s preference for that feature and the area of the map taken up by the feature. Throughout the paper, we will denote a user’s preference for a geographic feature, g , as $u(g)$. This value is provided through a user interface and falls on a scale from 0 (not important) to 10 (most important). We denote the area of a geographic feature as $Area(g)$. The area of a region-feature is the area of polygons that make up the region (e.g., the area covered by a set of buildings or the area of a waterway). Path-features are multiple connected lines, we approximate the area as the length of these lines multiplied by a width of 1mm. Similarly, location features are a set of coordinates in the map so we approximate their area as the count of locations in the feature. This essentially treats each location as a dot on the map that has an area of 1mm^2 . Note that we do not consider sizing-parameters such as the size of pegs that mark each location or the width of roads. These are accounted for by the size of representations (e.g., their depth).

3.2 Representations

Regions, paths, and locations can each be represented in a variety of ways. Locations are represented by differently shaped pegs that stick out of the map at the corresponding coordinates. We tested a variety of peg shapes based on related work [11, 17, 31] and a pilot test with a blind researcher and selected three types of pegs for the system used in our user study: pyramids, domes, and cubes. Pyramids are triangular pegs that have a triangular pyramid point on top. Domes are cylindrical pegs with a half sphere on top. Cube pegs are square pegs with a flat top (Figure 1a). Paths are represented by raised solid-lines, dashed-lines (Figure 1b), or dotted-lines. For regions, the simplest representation is to emboss a set of polygons shaped like that region (Figure 1c). To differentiate more regions (e.g., buildings vs water), textures can be added on top of these embossed polygons. Based on our pilot tests, we developed three textures that emboss small triangles, circles, or squares in a grid pattern across the texture. In total, we include ten representations: four regions, three paths, three locations. Our set of representations is extensible, and this set ensured a wide variety of map designs for our user evaluation.


These representations are individually parameterized. Each representation has a depth parameter, d_r , which determines how much it is raised out of the map base. Path and location representations,




(a) Locations are represented with differently shaped pegs.



(b) Paths are a series of raised dashed lines.



(c) Building Regions represented by raised polygons on the tactile map.



(d) Combinations of different representations are added onto a base-map to create a complex tactile map.

Figure 1: A tactile map is produced by pairing geographic features to tactile representations of those features. There are three types of geography features which can be related to specific representations: (a) locations, (b) paths, (c) regions. Combining different representations creates a full tactile map (d).

additionally, have a width parameter, w_r , which determines the width of the path or peg. These parameters will determine how much space is taken up by a represented geographic feature in the final tactile map. This, in turn, effects how easy different geographic features are to identify and distinguish. Larger representations (i.e., with greater depths and widths) will stand out of the map more because they take up more space. Correspondingly, smaller representations will be more difficult to find and identify. Additionally, users provide their own preference-ranking, $u(r)$, for each type of representation in the user interface on a scale from 0 to 10.

3.3 User Interface

We provide a simple screen-readable web interface that collect's participant's preferences for different ways of representing information and different geographic features (Figure 2). Based on prior work [30], the interface designed to be easy to use with a screen reader. Once the user has provided the location of their map, either an address or a coordinate, it queries the OpenStreetMap Data set for geographic features in that area. Each of these geographic features is presented in a table and the user can rank them on a scale from 0 to 10, with 10 being the most important geographic features,

Location

Map Size in MM

Map Base Depth in MM

Preference for Triangular Textures

Preference for Square Textures

Preference for Circular Textures

Geography Sets

Available Geography Sets

Geography Set	Importance	Remove Geography
Buildings	<input type="text" value="10"/>	<input type="button" value="Remove"/>
Water	<input type="text" value="3"/>	<input type="button" value="Remove"/>
Roads	<input type="text" value="4"/>	<input type="button" value="Remove"/>
Foot Paths	<input type="text" value="10"/>	<input type="button" value="Remove"/>
Picnic Tables	<input type="text" value="5"/>	<input type="button" value="Remove"/>
Building Entrances	<input type="text" value="10"/>	<input type="button" value="Remove"/>
Restrooms	<input type="text" value="10"/>	<input type="button" value="Remove"/>

(a) Maptimizer Screen Readable Interface

Your 3D printed map will have the following geographic sets represented with the described icons or representations

Buildings

The top view of buildings will be raised out of the map by 1mm.

Water

The top view of water will be raised out of the map by 1mm. From that a grid of 2 by 2 mm squares will be raised to create a bumpy texture.

Roads

The network of roads will be 3mm wide lines raised out of the map by 4mm. The lines will be dashed with 1mm dashes and 1mm breaks

Foot Paths

The network of roads will be 3mm wide lines raised out of the map by 4mm.

Building Entrances

Building entrances will be marked by a 6mm tall peg in the shape of a square.

Restrooms

Restrooms will be marked by a 6mm tall peg in the shape of a circle.

Picnic Tables

Picnic Tables entrances will be marked by a 6mm tall peg in the shape of a Triangle.

(b) Screen Readable Optimized tactile map Legend

Figure 2: Users set their preferences for different types of geographic features and how representations are made in a simple screen readable interface. Maptimizer displays the optimized tactile map’s legend and creates an STL file for 3D printing.

and 0 being not at all important. These rankings will be used in the optimization process and are denoted $u(g)$. To generate each tactile map, the user also ranks different types of representations (i.e., $u(r)$) based on their textures (e.g., square, circular, triangular). These are also ranked on a scale of 0 to 10. A screenshot of these inputs is shown in Figure 2a. These rankings (e.g., $u(g)$, $u(r)$) will be used in our optimization process to generate a tactile map that uses high-quality, preferred representations to present the most important geographic information to the user.

Once the tactile map is generated via optimization, a screen readable legend describing the map is presented to the user (Figure 2b). Each included geographic feature is labeled with a header and is followed by a brief description of its representation. Maptimizer also outputs a 3D model of the optimized tactile map to be 3D printed. For all of the maps in our study, the maps are 175 mm² squares which represents an 250 m² square area (i.e., 7mm per 10 meters). Tactile maps of this size and resolution maximized the high-quality printing space on our 3D printer bed while still generally producing high quality maps with legible representations.

3.4 Objective Function

When designing a tactile map to represent a variety of geographic features, we must consider three conflicting qualities: *communicability*, how effectively the map conveys relevant information to a particular user; *informativeness*, how valuable is the information the map conveys; and *attention-costs*, how cluttered the map is. We adopt these concepts from existing systems (e.g., [1, 12, 21, 22]) and adapt them to the domain of tactile maps.

Considered alone, each of these attributes is easy to optimize but will produce a poor performing map. For example, a tactile map that uses only a user’s preferred representations will be very communicative but may not be able to show every piece of information that is important without cluttering the map (i.e., increasing the attention cost). Then again, a sparse map with only a few key pieces of information will be less cluttered (i.e., decreasing attention costs) but may not use the representations that are easiest to read (i.e., reducing communicability) and may leave out important information (i.e., reducing informativeness). Features like communicability and informativeness are largely dependent on a user’s preferences and may even be dependent on location. Thus, representing different geographic features the same way in maps of different locations will rarely produce a high quality map. Alternatively, a sighted cartographer could consider these requirements and craft high quality customized maps, but this would be prohibitively expensive for most tactile map applications. As such, we formalize these properties (i.e., communicability, informativeness, attention-cost) to create the objective function which Maptimizer will maximize by generating unique tactile maps based on a user’s preferences and a location’s context.

For the purposes of optimization, we define a tactile map as a set of a representations paired to geographic features such that each geographic feature has one, and only one, unique representation. We denote the pairing between a geographic feature, g , and a representation, r , $p_{r \rightarrow g}$ and the set pairs in a specific tactile map as \mathbb{P} . Each type of representation has been ranked by the user in the user-interface, $u(r)$, as have all geographic features, $u(g)$. We denote the

communicability of a representation as $C(r)$, which measures how well that representation matches a users preferences and stands out of the map. We denote the informativeness of an geographic feature as $I(g)$ which measures the importance the information in a geographic feature. Finally, the attention cost of representing geographic features in a particular way is denoted $A(r, g)$ and measures how much the pairing of r to g clutters the tactile map. Generally, for each pairing of representations and geographic features we want to maximize the sum of the representation's communicability and the geographic feature's informativeness while minimizing their attention cost ($o(r, g)$, Equation 1a). We use the control variables ζ , ι , and α to weight communicability, informativeness, and attention costs in the objective function. For all maps produced for our user study we set their values to $\zeta = 1$, $\iota = 2$, and $\alpha = 1$, based on initial pilot tests. To evaluate a whole tactile map we calculate the sum, $O(\mathbb{P})$ (Equation 1b), for each pairing of representations and geographic features.

$$o(r, g) = \zeta C(r) + \iota I(g) - \alpha A(r, g) \quad (1a)$$

$$\text{Maximize: } O(\mathbb{P}) = \sum_{Pr, g \in \mathbb{P}} o(r, g) \quad (1b)$$

3.4.1 Estimating Communicability.

$$C(r) = \begin{cases} u(r) \frac{d_r}{d_{max}} & \text{if } r \text{ represents a region} \\ u(r) \frac{d_r w_r}{d_{max} w_{max}} & \text{otherwise} \end{cases} \quad (2)$$

The communicability of a representation measures how easy a representation is to distinguish from other representations in the map. Our estimation of communicability considers the user's ranking of that representation, $u(r)$, and how large their are (i.e., larger representations are easier to distinguish). We estimate size as the proportion of the size of the representation defined by its parameters (d_w, w_r) over the maximum allowed size of those parameters (d_{max}, w_{max} (Equation 2)). For each representation, these parameters can range from 1mm to 10mm. Generally, larger representations, with greater parameters, will be easier to recognize making them more communicative. We multiply this by the user's ranking from 0 to 10. If the user indicates a strong preference for a particular type of representation (e.g., setting $u(r)$ to 10), using that representation will amplify the communicability score.

3.4.2 User Ranked Information-Value.

$$I(g) = \frac{\text{Area}(g)}{\text{Map Area}} u(g) \quad (3)$$

The informativeness of a geographic feature is highly dependent on the user's information preferences which, in turn, is dependent on how they expect to use the map. Additionally, informativeness is dependent on information which the user may not be aware of. For example, a user may highly value common features such as buildings and roads but be unaware of a water feature (e.g., a lake) which is critical to understanding a location. A highly informative map includes features that provide context and that a user prefers and prioritizes features that meet both criteria.

We estimate informativeness as the product of two terms (Equation 3). The first term, approximates how critical a feature is to understanding a location based on how much of the map it takes up

(i.e., its area). We use a heuristic which assumes that if a geographic feature takes up a large portion of the map it must provide significant contextual information. We calculate this as the proportion of the area of the geographic feature over the total area of the map (i.e., 175mm by 175mm). The second term is the user's ranking of each geographic feature's value, $u(g)$, gathered from the user interface. This term gives the user the most control of their tactile maps because it has the strongest influence on what types of information will be presented in the final map.

3.4.3 Estimating Attention Costs.

$$A(r, g) = \begin{cases} \frac{d_r \text{Area}(g)}{\text{Map Volume}} & \text{if } r \text{ represents a region} \\ \frac{d_r w_r \text{Area}(g)}{\text{Map Volume}} & \text{otherwise} \end{cases} \quad (4)$$

Alone, maximizing communicability and informativeness would produce tactile maps packed with many communicative representations of informative geographic features. However, these tactile maps would be dense and cluttered and may confuse users or obscure valuable information. To penalize cluttering the tactile map, we subtract an attention cost term which estimates how much of the tactile map is covered by the representation of a specific geographic feature. To measure this we calculate the *volume* of a representation-geography pair by multiplying the sizing parameters of the representation (e.g., d_r, w_r) by the area of the geographic feature. We divide that by the *map volume* which is simply the map area multiplied by the maximum allowed depth of any representation (10 mm in all of our samples). The proportion of the represented geographic feature's volume to the total map volume will penalize larger representations and geographic features that might occlude the rest of the map. We found this simpler measure to be effective in the tactile maps produced for our user study.

3.5 Multi Stage tactile map Optimization

Given Maptimizer's objective function (Equation 1b), our optimization method generates tactile maps over two stages (Figure 3). The first stage determines what types of representations will be paired with geographic features. It does this by considering the user's preferences (e.g., $u(r), u(g)$) gathered from the Maptimizer interface. The second stage optimizes the sizing parameters of each of the representations that were selected in the first stage. Tuning these parameters makes the geographic features easier to identify (i.e., communicative) while minimizing clutter (i.e., attention-cost). Our optimization method is not guaranteed to find the globally optimal tactile map. However, based on the results of our user evaluation, we expect that the space of tactile maps has many high-quality local maxima. Maptimizer's goal is to generate a tactile map that more effectively supports a users needs than those that use the same representations of geographic features regardless of user preferences or context about the mapped location.

3.5.1 Stage 1: Representation and Geographic Feature Pairing. The first step to generating a tactile map is to identify the most informative geographic features and pair them to representations that will best communicate that information. To do this we first consider the pairings of geographic features to representations that would maximize our objective function (Equation 1b) if there were no constraints. This would be a set of pairs where every geographic

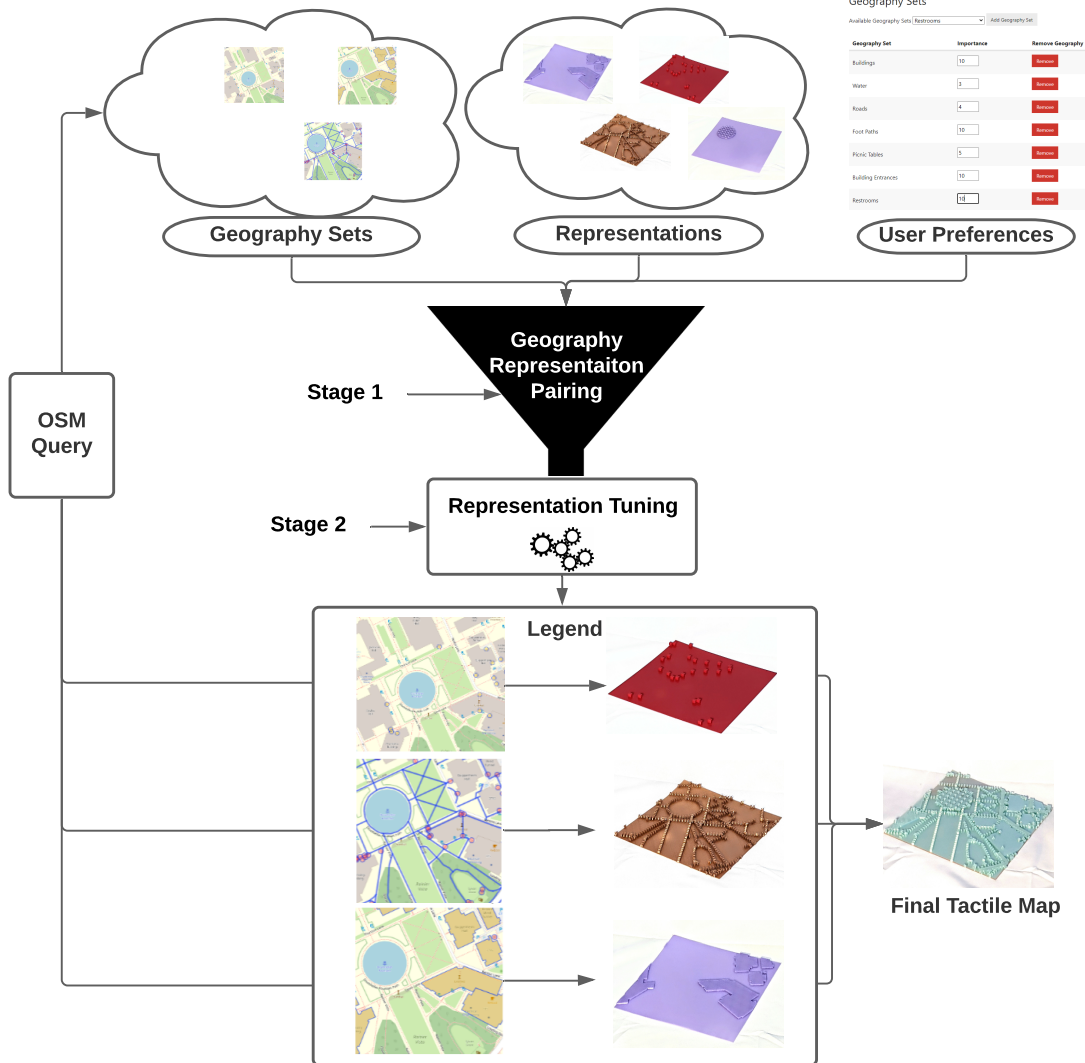


Figure 3: Maptimizer takes in a geographic feature, representation set, and user preferences in to pair representations to geographic features (stage 1). The representation parameters are tuned in stage 2. These representations take in geographic data from OpenStreetMap to combine representations into the final 3D printed map.

feature is paired once to every representation meaning that all of the most informative features are paired with all of the most communicative representations. However, this map would be unusable because there would be no way to distinguish between different features on the map. Thus, we apply three constraints. First, each representation can only be paired with one geographic feature, otherwise the user could not distinguish between the geographic features. Second, each geographic feature can only be paired to one representation so that overlapping representations are not generated. Finally, a representation can only be paired to an geographic feature if they are *compatible*. Naturally, region representations are compatible with geographic regions, path representations are compatible with sets of paths, and peg representations are compatible with locations.

We use linear-programming to maximize our objective function subject to these constraints. To do this we introduce binary-weights, $\omega_{r,g}$ into our objective function which determine if the pairing of a representation r to a geographic feature g will be included in the final tactile map. If the weight $\omega_{r,g}$ is set to 1, then the pair, $p_{r,g}$ will be included in the optimized tactile map. Otherwise, the weight must be 0, and the pair will not be included in the tactile map. The optimal set of pairings between representations and geographic features, \mathbb{P} , is a subset of all pairings, \mathbb{P}_{all} , of all representations, \mathbb{R} , to all geographic features, \mathbb{G} . To generate the set of optimal pairings, we reformulate our objective function and constraints as a solvable set of linear equations. Once we have solved for weights between 0 and 1, each weight is rounded up or down to create the final binary weights.

$$\begin{aligned} & \text{Maximize} \quad \sum_{p_{r,g} \in \mathbb{P}_{all}} \omega_{r,g} o(r, g) \\ & \text{s.t.} \quad \left\{ \begin{array}{l} \forall r \in \mathbb{R} \quad \sum_{g \in \mathbb{G}} \omega_{r,g} \leq 1 \\ \forall g \in \mathbb{G} \quad \sum_{r \in \mathbb{R}} \omega_{r,g} \leq 1 \\ \forall p_{r,g} \in \mathbb{P}_{all} \quad \omega_{r,g} > 0 \implies r \text{ is compatible with } g \end{array} \right. \quad (5) \end{aligned}$$

We formulate our pairing constraints as linear equations that isolate each pairing weight and ensure representations and geographic features are not overused. First, we constrain the weights to ensure that no representation is used by more than one geographic feature. That is, for all representations $r \in \mathbb{R}$, the sum of the pairing weight between r and each geographic feature, $g \in \mathbb{G}$, is less than or equal to one. This implies that only one or none of the pairs will be use. Similarly, we constrain the weights to ensure that no geographic feature is paired to more than one representation. That is for all geographic features, the sum of the pairing weight between the feature and each representation is less than or equal to one. Finally, for each pair between a representation and geographic feature, the corresponding weight can only be greater than zero if the representation and feature are compatible.

3.5.2 Stage 2: Representation Parameter Tuning. The first stage of optimization will determine what geographic information will be included in the map and what types of representation will be used to represent it. However, we still need to tune the sizing parameters of each representation to maximize our objective function (e.g., adjusting the width of a path). Increasing a representation's parameters will increase its communicability (Equation 2). Conversely, decreasing these parameters will decrease the attention cost (Equation 4). By making these adjustments we are making trade offs between *communicability* and *attention cost* in our objective function (Equation 1a). To solve for representation parameters that improve our objective score we use an Ant Colony Optimization method [9].

In this optimization stage we search for a tactile map that has the geography-representation pairs from the first optimization stage with new depth and width parameters for each representation. Each representation has a free depth parameter and path and location representations additionally have a free width parameter (i.e. d_r, w_r). We start our optimization process with a tactile map that uses the pairing set in the last stage and has all representation parameters set to 1mm. Over a large number of iterations, we modify this map by incrementing and decrementing different representation parameters by 1mm. We choose 1mm increments because smaller differences are not easily detected by most visually-impaired users and may be unreliable on most consumer 3D printers. Our goal is to find a tactile map that uses the pairings of geographic-information to representations from stage 1 and has representation parameters that maximize our objective function (Equation 1b).

We use Ant Colony Optimization to find the optimal tactile map by conducting a series of traversals of the search space. After a maximum number of iterations is reached (e.g. 1000), we return the highest scoring tactile map. At each step of the optimization we

decide how to modify the tactile map (i.e., what representation parameter to modify and by what increments) based on a probability determined by two factors: *desirability* and *history of success*. Desirability estimates how much we expect the tactile map to improve by taking modifying its representation's parameter. The history of success measures how much the objective score has been increased when we have taken this step in past iterations.

First, we consider the desirability, $D(\Delta, r)$, of incrementing a representation, r , by a modifier Δ . The modifier can be either -1 or 1 mm. Desirability estimates by how much better the resulting tactile map, m' , will be than the current tactile map, m . Generally, we want to encourage increasing the size of highly-ranked representations, which increases communicability, while decreasing the size of lower-ranking representations, which decreases attention costs. If Δ is positive, we estimate desirability as the proportion that the communicability score will increase by (i.e., $\zeta u(r)$) over the amount the attention cost will increase by (i.e., $\frac{\alpha \text{Area}(g)}{\text{Map Volume}}$). So, if the communicability score will increase more than the attention cost, the desirability will be greater than 1. If the increment is negative, desirability is the inverse proportion. That is, we estimate desirability as the portion that the attention cost will decrease by the reduction in the communicability score. If the attention cost decreases more than communicability the desirability will be greater than 1. Note that informativeness (Equation 3) is not effected by representation parameters and is, therefore, not included in this measure of desirability.

$$D(\Delta, r) = \begin{cases} \frac{\zeta u(r) \text{Map Volume}}{\alpha \text{Area}(g)} & \Delta_r > 0 \\ \frac{\alpha \text{Area}(g)}{\zeta u(r) \text{Map Volume}} & \Delta_r < 0 \end{cases} \quad (6)$$

Next we consider the history of success. Ant Colony optimization uses the metaphor of a colony of ants searching for a food source, in this case a high-performing tactile maps. As single ants (i.e., iterations of the optimization) take steps through the search space, they leave behind a pheromone which tells future ants that there is food in that direction. The more ants that follow a specific path, the more pheromone is left behind and the more it encourages future ants to follow the same path. We maintain a set of edges that record each step taken by the optimizer. At the start of the search process, each edge has a pheromone weight, ρ_e equal to 1. Every time an edge is crossed (i.e., we apply the same increment on a representation parameter), we increment that pheromone weight by the difference between the objective score of the tactile map before taking the step, $O(m)$, and the objective score after taking that step, $O(m')$. If the tactile map improved, the weight will increase and in future iterations we will be more likely to increment that parameter in the same way. If the map is worse, it will discourage repeating this modification.

Using these measures of desirability and a history of success we can estimate the probability of taking any edge in the search space. We set the probability of taking an edge, e , from m to m' to the desirability, $D(\Delta, r)$, times the pheromone weight, ρ_e , over the sum of the desirability and pheromone weight for all other possible edges from m .

Algorithm 1 Tunes the depth and width parameters of a tactile map to maximize the objective-function

Input: N ▷ The number of search iterations. Generally set to 1000 iterations.
Input: m_0 ▷ A tactile map with the representation geographic feature pairs set by linear programming with all representation parameters set to 1mm.
Output: m_{\max} ▷ The tactile map with tuned width and height variables that maximize the objective function.

- 1: $visited_maps \leftarrow \{m_0\}$: The set of tactile maps that have been visited by the search process.
- 2: $crossed_edges \leftarrow \{\}$: The set of edges that have been crossed in the search process.
- 3: $i \leftarrow 0$: The number of search-iterations completed.
- 4: **while** $i < N$ **do**
- 5: $m \leftarrow$ Choose a tactile map from $visited_maps$ at random with a bias towards high performing tactile maps.
- 6: $e \leftarrow$ Choose an edge from m based on the probability distribution $P(e|m)$. ▷ see Equation 7
- 7: $m' \leftarrow$ update m using Δ_r as defined by e
- 8: **if** e is not in $crossed_edges$ **then**
- 9: add m' to $visited_maps$
- 10: add e to $crossed_edges$
- 11: $\rho_e \leftarrow 0$
- 12: $\rho_e \leftarrow \rho_e + (O(m') - O(m))$
- 13: $i \leftarrow i + 1$

return highest-scoring tactile map in $visited_maps$.

$$P(e|m) = \frac{\rho_e D(\Delta, r)}{\sum_{e' \in E_m} \rho_{e'} D(\Delta', r)} \quad (7)$$

Given this probability distribution, our optimized search process executes as follows (Algorithm 1). We create a set of *visited maps*. Initially, this set only includes a tactile map with the pairings of representations to geographic features resulting from the stage 1 linear programming optimization. All representation parameters are initialized to 1mm. We also create an empty set of *crossed edges* to track the paths we will visit at each iteration. For N iterations, we first choose a tactile map, m , at random from the set of visited maps with a bias towards the one that have the highest objective scores. Next we select an edge, e from m , using the probability distribution defined by desirability and each edge’s pheromone weight (Equation 7). If we have not crossed e before, we generate the resulting tactile map, m' , and add it to the set of visited maps. We also add e to the set of crossed edges. We then update the pheromone-weight of e by adding the difference in objective scores. We repeat this until the maximum number of iterations is met and then return the highest scoring tactile map that we visited.

This search method is not guaranteed to find a globally optimal tactile map. The space of tactile maps is complex and has many good local maxima. Instead, it efficiently explores a diverse set of tactile maps in high-performing sections of the search space. As the edge-weights, ρ_e , are updated this intensifies the search over high-performing regions yielding finely tuned results from those areas. Based on the results of our user evaluation, we found that this two stage optimization process tends to produce high quality maps that support users in navigational tasks.

4 USER EVALUATION

We conducted a two part user-evaluation of the optimized tactile maps to answer two research questions. First, does the ability to customize what information is presented in a tactile map improve the experience for users? Second, does the combination of optimization with customization further improve the user’s experience?

Throughout our study we considered three tactile map conditions: 1) *standardized-maps* made with TouchMapper [19] which cannot be customized by the user, 2) *customized-maps* that only include the most important information defined by a user, and 3) *optimized-maps* which consider both a users preferences and other factors and are generated with Maptimizer. Participants were unaware of how any map was created until after they completed the study. We only told them that we had made these tactile maps and wanted their feedback. Given tactile maps of the same locations, we measure the quality of a tactile map based on: a user’s preference for different types of maps, and the user’s ability to quickly and accurately find a specific location on a map using a verbal description.

To test the quality of our *optimized-maps* compared to *customized-maps* and *standardized-maps*, we conducted a user study with six participants who identified as blind or low-vision from Seattle. Participants were asked to first complete an online survey that contained questions about 1) their demographic information, 2) previous experience with tactile maps, and 3) questions from the Maptimizer interface. The results of these surveys were used to generate customized maps with Maptimizer. Participants then met researchers at a public location for an hour long session where they used these tactile maps. Participants received a travel stipend to commute to the location and were compensated \$40 for their time. This study was approved by the second author’s institutional review board.

Two of our participants identified as men, and four identified as women. Our participants ages ranged from 28 to 72 years of age (mean of 47 years, standard deviation of 18 years). Four participants identified as blind, and two identified as having low-vision. On a scale from 1 to 5, 5 being very familiar, on average participants rated their familiarity with tactile maps as 2.66 with a standard deviation of .5. Only one participant had no prior experience with tactile maps. We found no significant effects of participants demographic information on the results of the study.

4.1 Methods

We conducted the study in two parts. The first part helped participants familiarize themselves with the tactile maps and practice describing the representations to researchers. In the second part of the study we tested whether maps that were customized or optimized affected participants' ability to identify a location on the map from a verbal description.

4.1.1 Part 1: Familiarity with Tactile Maps. We first showed participants three different 3D printed tactile maps of the same location in Seattle: a park that has roads, walkways, buildings, a lake, and some small locations of interest (e.g., picnic tables, benches, sculptures). The three maps were created for our three conditions: *standardize-maps*, *customized-maps*, and *optimized-maps*.

Standardized-Map (Figure 4a): Maps generated with TouchMapper [19] cannot be customized; all maps with TouchMapper use pair the geographic features to the same representations. We use TouchMapper as a control condition to measure the effect of customization on the user's experience.

Customized-Map (Figure 4b): For customized maps we used the six most highly ranked geographic features from the participant's survey. During pilot tests, we found that any more than six features consistently produced cluttered map that were too difficult to read. These were each mapped to the most highly-ranked, compatible representations available in Maptimizer. Standard depth and width parameters were used, rather than optimizing them. This experimental-condition helps us compare the effect of customization through the user's preferences without optimization.

Optimized-Map (Figure 4c): By feeding the participant's preferences into the Maptimizer interface, we generated optimized maps that considered both the participant's preferences and contextual information such as the size of different geographic features of the location. This is our experimental-condition that helps us compare the effects of optimization with standardized and customized maps.

We randomized the order in which these three maps were shown to the participants. Following a think aloud protocol, participants examined the maps and were free to ask questions about what different elements on the map represented. For instance, a participant could identify a rounded peg and ask what information it represented. We asked participants to verbally describe the elements of the map they wanted to know about, rather than allowing them to physically point out the feature. This think aloud approach gives us a better understanding of which representations were easiest to identify and distinguish from one another. For example, a participant could not just point out a round peg, they had to describe it as a peg (or equivalent phrase) and differentiate its shape from other pegs (e.g., squares and triangles). As participants explored the maps, they provided feedback around what they found useful or confusing with respect to both the information conveyed and the textures used to represent them. After examining all three maps, we asked them which map they preferred and why.

4.1.2 Part 2: Location Identification. Next, we wanted to examine how customization and optimization effected the participants' ability to identify a location on a map, given a verbal description. We presented users with three new maps (i.e., standardized, customized, optimized) but this time each map was of a different location in

Seattle—a university campus, a public market, a different park with a lake. We counterbalanced map-generation conditions with these locations across three groups to ensure that if one location was more difficult to understand, this would effect each map generation condition equally. The first group received an optimized university map, a customized park map, and a standardized market map. The second group received a customize market, standardized park, and optimized market. The third group received a standardized university, optimized park, and customized market. Two participants were randomly assigned to each group.

For each location, we read aloud a short verbal description of a location on the map. Features that were available in each map-condition could be used to identify this location. In the customized and optimized conditions, these features may not be present depending on whether the user happened to have a strong preference for that type of location—a sculpture on a university campus, the end of a foot path leading to a beach in a park, or an information booth in the market. Each verbal description described the location as relative to an area (e.g., buildings, water), a path (e.g., road, foot-paths), and a set of locations (e.g., benches). We read the description to the participants then handed them the map to start exploring. We would repeat the description as often as the participant requested, but did not provide any other details about the location. Just as in the first task, participants could ask what a feature of the map represented by describing it verbally.

Participants could give up their search for the location, or would announce when they believed they had found it. We recorded the time it took them to complete the task, whether or not they found the correct location, and how confident they were that they had found the correct location on a scale from 0 to 5; 0 was reserved for participants who did not identify a location and 5 indicated the highest confidence in the location they found.

4.2 Results

Based on the results summarized in Table 1, we found that optimization helped participants identify a location. Overall, participant's preferences for different types of maps (e.g., standardized, customized, optimized) were dependent on how naturally specific representations paired with specific types of features. However, none of these generation methods consistently paired information this way. When we focused participants efforts on using the maps to identify a location, participants had significantly more success with optimized maps. We suspect this is because the optimized maps provided both information the participants had ranked highly, and contextual information which was critical to identifying the location.

4.3 Participant's Preferences

The first part of the study showed that participants' preferences for different ways of generating a map varied. Overall four people preferred the optimized map and two preferred the standardized map while no one preferred the customized maps. While thinking aloud, our participants revealed insights into what made maps attractive. For instance, the standardized maps used a "kind of rippled" (P1) texture to represent water. Some participants (P1, P3, P6), intuitively knew this was water, but others assumed it was a "hilly grass, kind




Figure 4: Three maps of the same public park generated with: (a) a standardized map, (b) a user-customized map, and (c) an optimized map produced with Maptimizer.

	Standardized	Customized	Optimized
Preferred Map-Condition	2 (33%)	0 (0%)	4 (66%)
Identified correct location	2 (33%)	3 (50%)	6 (100%)
Confidence in location M (SD)	2.17 (2.04)	3.33 (2.25)	4.33 (1.03)
Time elapsed M (SD)	169.00 (112.93)	167.83 (99.72)	150.16 (111.54)

Table 1: The summary of three key statistics from the location finding task (number of people who identified the correct location, average participant confidence that the location they selected was correct, and time elapsed in the location finding task) for the three types of maps. The Maptimizer maps outperformed both alternatives in all three statistics.

of thing” (P2). Alternatively, some participants based their preferences on how easy it was for them to identify information that was important to them after learning about the area. For example, P3 preferred the optimized map “because it identifies the different items [she] needs” such as water, the entrances to buildings, and a picnic table; whereas, in the customized map “the water is not identified which would be a little confusing”. Notably, P3’s customized map did not include water because she had ranked water features lower than other features. Since the water took up such a large portion of the map, Maptimizer added the water feature to provide context.

Most of our participants associated specific geographic features with representations that evoked an image (e.g., the rippled water texture). However, neither our customization condition nor our optimization method currently takes this into consideration. For instance, P2 did not think the square texture Maptimizer assigned to the lake matched his image of water, instead describing it as “Arlington Cemetery for flees”. However, later during the location finding task, he found the information booth quickly because Maptimizer had assigned it a round peg representation. He said it reminded him of a lighthouse, “what’s the shining light coming out of this lighthouse? It’s information”. Relating specific types of representations to geographic information is a difficult task, especially as different types of information are represented creating conflicts. Exploring ways to include common representations of information in our optimization process is a promising area of future work.

4.4 Effects on Location Identification

While the number of participants is small, we found that Maptimizer’s optimized maps helped participants identify locations more accurately. When presented with optimized maps, all six participants correctly identified the described location on their map,

regardless of which group they were in. Accuracy varied when participants used standardized and customized maps. With the standardized maps, only the participants given the standardized park map could identify the correct location. Similarly, with the control maps, only half of the participants could identify the location. Two of these participants had the park map and the third had the market map. A χ^2 test shows that the maps location had a significant effect on success on location finding ($\chi^2 = 6.08, p < .05$). However, a χ^2 test demonstrates that the way a map was created (e.g., standardized, customized, optimized) also had a significant effect on participants’ success at finding the correct location ($\chi^2 = 6.08, p < .05$). Since, map location and the way a map was generated were counterbalanced, we suspect that these two effects are independent and the high success rate with optimized maps is attributed to the generation method.

5 LIMITATIONS AND FUTURE WORK

Based on our study, Maptimizer’s main limitation is the inability to associate geographic features with representations that evoke the image of that information (e.g., rippled water, light-house pegs). Instead, our optimization process considers the communicability of a representation separately from the informativeness of a geographic feature. Even with our small group of participants, there was little agreement on what representations evoked which features. For instance, participants disagreed about the intuitiveness of the rippled water. A larger study of users’ preferences may reveal more consistent patterns which could be used in the optimization process. This suggests that customization is key to designing high-quality tactile maps since there is limited space to present information and what representations present that information the best is highly dependent on individual users.

Despite such limitations of the maps, Maptimizer did improve the users' ability to identify locations, meaning that optimization may have a critical role to play in generating high-quality tactile maps. Beyond tactile maps, the process of matching customized representations to sets of information may help us to build a variety of end-user tailored tactile graphics (e.g., educational aids, children's books) or could be applied to other way-finding tools (e.g., audio descriptions of routes). In particular, the concepts such as communicability, informativeness, and attention cost we adopted from prior work (e.g., [1, 12, 21, 22]) may generalize beyond navigational aids. Future work should explore how Maptimizer's approach generalizes to other accessibility domains.

6 CONCLUSION

Tactile Maps are useful tools to support navigation by people with visual impairments. However, to make tactile maps more readily available, we need tools that can generate those maps for new locations without the support of a sighted cartographer. Beyond increasing availability, these tactile maps should be customized to meet the specific needs of an individual user. The design of each map should consider what information is most valuable to a particular person, and how information can be represented most effectively. However, requiring the user to fully define their tactile map creates a cumbersome design process which may not be worth their efforts. Instead, optimization techniques enable us to automatically adapt designs to user's specific abilities and needs.

In this paper, we have presented Maptimizer, a tool which uses optimization methods and an ability-based design [32] approach to create customized tactile maps. Our user study comparing tactile maps created with a standardized maps, customized maps, and optimized maps, demonstrates that users have strong personal preferences for how information is represented in a tactile map. However, different users will have diverse and conflicting sets of preferences which indicates the need for customization. Customization alone does not always produce maps that provide sufficient information to perform navigational tasks. Based on the results of our user study, Maptimizer's optimization method successfully incorporates context about a location and user preferences to generate more usable tactile maps.

REFERENCES

- [1] Maneesh Agrawala and Chris Stolte. 2001. Rendering Effective Route Maps: Improving Usability through Generalization. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. Association for Computing Machinery, New York, NY, USA, 241–249. <https://doi.org/10.1145/383259.383286>
- [2] N. Bolten and A. Caspi. 2021. Towards routine, city-scale accessibility metrics: Graph theoretic interpretations of pedestrian access using personalized pedestrian network analysis. *PLoS one* 16, 3 (2021), 10 pages. <https://doi.org/10.1371/journal.pone.0248399>
- [3] N. Bolten, S. Mukherjee, V. Sipeeva, A. Tanweer, and A. Caspi. 2017. A pedestrian-centered data approach for equitable access to urban infrastructure environments. *IBM Journal of Research and Development* 61, 6 (2017), 10:1–10:12. <https://doi.org/10.1145/JRD.2017.2736279>
- [4] Nicholas A. Bradley and Mark D. Dunlop. 2005. An Experimental Investigation into Wayfinding Directions for Visually Impaired People. *Personal and Ubiquitous Computing* 9, 6 (Nov 2005), 395–403. <https://doi.org/10.1007/s00779-005-0350-y>
- [5] Craig Brown and Amy Hurst. 2012. VizTouch: Automatically Generated Tactile Visualizations of Coordinate Spaces. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction* (Kingston, Ontario, Canada) (*TEI '12*). Association for Computing Machinery, New York, NY, USA, 131–138. <https://doi.org/10.1145/2148131.2148160>
- [6] Scott Carter, Amy Hurst, Jennifer Mankoff, and Jack Li. 2006. Dynamically Adapting GUIs to Diverse Input Devices. In *Proceedings of the 8th International ACM SIGACCESS Conference on Computers and Accessibility* (Portland, Oregon, USA) (*Assets '06*). Association for Computing Machinery, New York, NY, USA, 63–70. <https://doi.org/10.1145/1168987.1169000>
- [7] Xiang 'Anthony' Chen, Jeeun Kim, Jennifer Mankoff, Tovi Grossman, Stelian Coros, and Scott E. Hudson. 2016. Reprise: A Design Tool for Specifying, Generating, and Customizing 3D Printable Adaptations on Everyday Objects. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) (*UIST '16*). Association for Computing Machinery, New York, NY, USA, 29–39. <https://doi.org/10.1145/2984511.2984512>
- [8] Ruta Desai, James McCann, and Stelian Coros. 2018. Assembly-Aware Design of Printable Electromechanical Devices. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Berlin, Germany) (*UIST '18*). Association for Computing Machinery, New York, NY, USA, 457–472. <https://doi.org/10.1145/3242587.3242655>
- [9] M. Dorigo, M. Birattari, and T. Stutzle. 2006. Ant colony optimization. *IEEE Computational Intelligence Magazine* 1, 4 (2006), 28–39. <https://doi.org/10.1109/MCI.2006.329691>
- [10] Krzysztof Z. Gajos, Daniel S. Weld, and Jacob O. Wobbrock. 2010. Automatically generating personalized user interfaces with Supple. *Artif. Intell.* 174, 12 (2010), 910–950. <https://doi.org/10.1016/j.artint.2010.05.005>
- [11] Timo Götzelmann. 2016. LucentMaps: 3D Printed Audiovisual Tactile Maps for Blind and Visually Impaired People. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility* (Reno, Nevada, USA) (*ASSETS '16*). Association for Computing Machinery, New York, NY, USA, 81–90. <https://doi.org/10.1145/2982142.2982163>
- [12] Floraine Grabler, Maneesh Agrawala, Robert W. Sumner, and Mark Pauly. 2008. Automatic Generation of Tourist Maps. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 1–11. <https://doi.org/10.1145/1360612.1360699>
- [13] Anhong Guo, Jeeun Kim, Xiang 'Anthony' Chen, Tom Yeh, Scott E. Hudson, Jennifer Mankoff, and Jeffrey P. Bigham. 2017. Facade: Auto-Generating Tactile Interfaces to Appliances. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 5826–5838. <https://doi.org/10.1145/3025453.3025845>
- [14] Richard Guy and Khai Truong. 2012. CrossingGuard: Exploring Information Content in Navigation Aids for Visually Impaired Pedestrians. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 405–414. <https://doi.org/10.1145/2207676.2207733>
- [15] Megan Hofmann, Devva Kasnitz, Jennifer Mankoff, and Cynthia L Bennett. 2020. Living Disability Theory: Reflections on Access, Research, and Design. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility* (Virtual Event, Greece) (*ASSETS '20*). Association for Computing Machinery, New York, NY, USA, Article 4, 13 pages. <https://doi.org/10.1145/3373625.3416996>
- [16] Megan Hofmann, Jennifer Mankoff, and Scott E. Hudson. 2020. KnitGIST: A Programming Synthesis Toolkit for Generating Functional Machine-Knitting Textures. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (*UIST '20*). Association for Computing Machinery, New York, NY, USA, 1234–1247. <https://doi.org/10.1145/3379337.3415590>
- [17] Leona Holloway, Kim Marriott, and Matthew Butler. 2018. Accessible Maps for the Blind: Comparing 3D Printed Models with Tactile Graphics. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173772>
- [18] Leona Holloway, Kim Marriott, Matthew Butler, and Samuel Reinders. 2019. 3D Printed Maps and Icons for Inclusion: Testing in the Wild by People Who Are Blind or Have Low Vision. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility* (Pittsburgh, PA, USA) (*ASSETS '19*). Association for Computing Machinery, New York, NY, USA, 183–195. <https://doi.org/10.1145/3308561.3353790>
- [19] Samuli Kärkkäinen. 2018. Tactile Maps Easily | Touch Mapper. <https://touchmapper.org/en> [Online; accessed 25. Aug. 2021].
- [20] Shaun K. Kane, Meredith Ringel Morris, and Jacob O. Wobbrock. 2013. Touchplates: Low-Cost Tactile Overlays for Visually Impaired Touch Screen Users. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (Bellevue, Washington) (*ASSETS '13*). Association for Computing Machinery, New York, NY, USA, Article 22, 8 pages. <https://doi.org/10.1145/2513383.2513442>
- [21] Joonhwan Lee, Jodi Forlizzi, and Scott E. Hudson. 2005. *Studying the Effectiveness of MOVE: A Contextually Optimized in-Vehicle Navigation System*. Association for Computing Machinery, New York, NY, USA, 571–580. <https://doi.org/10.1145/1054972.1055051>
- [22] Joonhwan Lee, Jodi Forlizzi, Scott E. Hudson, and Soojin Jun. 2015. Use of the Backseat Driving Technique in Evaluation of a Perceptually Optimized In-Car Navigation Display. *International Journal of Human-Computer Interaction* 31, 2 (2015), 128–138. <https://doi.org/10.1080/10447318.2014.959104>

- arXiv:<https://doi.org/10.1080/10447318.2014.959104>
- [23] Louis Li and Krzysztof Z. Gajos. 2014. Adaptive Click-and-Cross: Adapting to Both Abilities and Task Improves Performance of Users with Impaired Dexterity. In *Proceedings of the 19th International Conference on Intelligent User Interfaces* (Haifa, Israel) (*IUI '14*). Association for Computing Machinery, New York, NY, USA, 299–304. <https://doi.org/10.1145/2557500.2557511>
- [24] Martez E. Mott and Jacob O. Wobbrock. 2019. Cluster Touch: Improving Touch Accuracy on Smartphones for People with Motor and Situational Impairments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3290605.3300257>
- [25] OpenStreetMap Community Contributors. 2021. OpenStreetMap. <https://www.openstreetmap.org/> [Online; accessed 25. Aug. 2021].
- [26] Helen Petrie, Valerie Johnson, Thomas Strothotte, Andreas Raab, Rainer Michel, Lars Reichert, and Axel Schalt. 1997. MoBIC: An Aid to Increase the Independent Mobility of Blind Travellers. *British Journal of Visual Impairment* 15, 2 (1997), 63–66. <https://doi.org/10.1177/026461969701500205> arXiv:<https://doi.org/10.1177/026461969701500205>
- [27] Manaswi Saha, Michael Saugstad, Hanuma Teja Maddali, Aileen Zeng, Ryan Holland, Steven Bower, Aditya Dash, Sage Chen, Anthony Li, Kotaro Hara, and Jon Froehlich. 2019. *Project Sidewalk: A Web-Based Crowdsourcing Tool for Collecting Sidewalk Accessibility Data At Scale*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3290605.3300292>
- [28] Lei Shi, Yuhang Zhao, and Shirri Azenkot. 2017. Designing Interactions for 3D Printed Models with Blind People. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility* (Baltimore, Maryland, USA) (*ASSETS '17*). Association for Computing Machinery, New York, NY, USA, 200–209. <https://doi.org/10.1145/3132525.3132549>
- [29] Alexa F. Siu, Son Kim, Joshua A. Miele, and Sean Follmer. 2019. ShapeCAD: An Accessible 3D Modelling Workflow for the Blind and Visually-Impaired Via 2.5D Shape Displays. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility* (Pittsburgh, PA, USA) (*ASSETS '19*). Association for Computing Machinery, New York, NY, USA, 342–354. <https://doi.org/10.1145/3308561.3353782>
- [30] Brandon Taylor, Anind Dey, Dan Siewiorek, and Asim Smailagic. 2016. Customizable 3D Printed Tactile Maps as Interactive Overlays. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility* (Reno, Nevada, USA) (*ASSETS '16*). Association for Computing Machinery, New York, NY, USA, 71–79. <https://doi.org/10.1145/2982142.2982167>
- [31] Brandon T. Taylor, Anind K. Dey, Dan P. Siewiorek, and Asim Smailagic. 2015. TactileMaps.Net: A Web Interface for Generating Customized 3D-Printable Tactile Maps. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers and Accessibility* (Lisbon, Portugal) (*ASSETS '15*). Association for Computing Machinery, New York, NY, USA, 427–428. <https://doi.org/10.1145/2700648.2811336>
- [32] Jacob O. Wobbrock, Shaun K. Kane, Krzysztof Z. Gajos, Susumu Harada, and Jon Froehlich. 2011. Ability-Based Design: Concept, Principles and Examples. *ACM Trans. Access. Comput.* 3, 3, Article 9 (April 2011), 27 pages. <https://doi.org/10.1145/1952383.1952384>
- [33] Eitan Zemel. 1984. An O(n) algorithm for the linear multiple choice knapsack problem and related problems. *Inf. Process. Lett.* 18, 3 (Mar 1984), 123–128. [https://doi.org/10.1016/0020-0190\(84\)90014-0](https://doi.org/10.1016/0020-0190(84)90014-0)
- [34] Xiaoyi Zhang, Tracy Tran, Yuqian Sun, Ian Culhane, Shobhit Jain, James Fogarty, and Jennifer Mankoff. 2018. Interactiles: 3D Printed Tactile Interfaces to Enhance Mobile Touchscreen Accessibility. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility* (Galway, Ireland) (*ASSETS '18*). Association for Computing Machinery, New York, NY, USA, 131–142. <https://doi.org/10.1145/3234695.3236349>
- [35] Lifeng Zhu, Weiwei Xu, John Snyder, Yang Liu, Guoping Wang, and Baining Guo. 2012. Motion-Guided Mechanical Toy Modeling. *ACM Trans. Graph.* 31, 6, Article 127 (2012), 10 pages. <https://doi.org/10.1145/2366145.2366146>