

Lightweight Relief Shearing for Enhanced Terrain Perception on Interactive Maps

Wesley Willett

Inria / University of Calgary
wesley.willett@ucalgary.ca

Bernhard Jenny

Oregon State University
jennyb@geo.oregonstate.edu

Tobias Isenberg, Pierre Dragicevic

Inria

[tobias.isenberg,pierre.dragicevic]@inria.fr

ABSTRACT

We explore interactive relief shearing, a set of non-intrusive, direct manipulation interactions that expose depth and shape information in terrain maps using ephemeral animations. Reading and interpreting topography and relief on terrain maps is an important aspect of map use, but extracting depth information from 2D maps is notoriously difficult. Modern mapping software attempts to alleviate this limitation by presenting digital terrain using 3D views. However, 3D views introduce occlusion, complicate distance estimations, and typically require more complex interactions. In contrast, our approach reveals depth information via shearing animations on 2D maps, and can be paired with existing interactions such as pan and zoom. We examine explicit, integrated, and hybrid interactions for triggering relief shearing and present a version that uses device tilt to control depth effects. Our evaluation shows that these interactive techniques improve depth perception when compared to standard 2D and perspective views.

Author Keywords

Relief shearing; terrain maps; depth perception; interaction; plan oblique relief.

ACM Classification Keywords

H.5.2. Information Interfaces and Presentation: User Interfaces—*Input devices and strategies.*

INTRODUCTION AND MOTIVATION

Topography and elevation are among the most important pieces of information on many maps, yet interpreting terrain from two-dimensional topographic maps is challenging for many users [21]. Traditionally, cartographers attempted to make terrain more accessible by using relief representations such as contour lines, shaded relief, and elevation coloring [11]. Most modern computational mapping tools still use these same representations or display terrain using 3D perspective views (e.g., Google Earth). These 3D views facilitate more immediate perception of the shape of terrain, but introduce occlusions and perspective distortions that complicate distance and elevation estimation. In addition, 3D views often require users to control the camera, which can be cumbersome and disorienting.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CHI 2015, April 18–23, 2015, Seoul, Republic of Korea.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3145-6/15/04 ...\$15.00.
<http://dx.doi.org/10.1145/2702123.2702172>

We explore *interactive relief shearing*, a set of non-intrusive, direct manipulation interactions (Figure 1) that expose depth and shape information in terrain maps using brief ephemeral animations. Based on this principle, we contribute a set of new, lightweight interaction techniques that use motion to enhance the perception of depth but remain compatible with the hardware and interaction techniques used for most 2D maps. We also describe two experiments in which we compared these techniques to existing interaction and rendering approaches. We find that relief-shearing animations improve users' ability to perceive depth, both when maps are animated automatically and when animations are connected to panning interactions. We also show that maps with interactive relief shearing help users make more accurate elevation judgments than either standard shaded relief maps or 3D perspective views.

BACKGROUND

Our work builds on prior research in relief rendering, 3D representation, and interaction.

Relief Rendering Techniques

Historically, communicating the shape and geometrical complexity of terrain in orthographic maps has been one of the biggest challenges faced by cartographers. In response, map-makers have developed a range of techniques for representing elevation including contour lines, spot heights, shaded relief, and elevation coloring. Contour lines are difficult to read compared to other techniques [19, 21], but with training it is possible to extract absolute elevations by interpolating between lines. Spot heights and soundings, meanwhile, are useful for showing the elevations of select points but they cannot portray the continuity of the terrain surface [19]. Relief shading or




Figure 1. Using *explicit* interactive relief shearing, the user grabs a point on the map (left) and drags it upward (right). The base of the terrain model stays in place, but the model is sheared so that the selected point remains under the cursor, revealing the shape of the terrain. When released, the map animates back to the original un-sheared position.

“hill shading”, a technique where terrain is shaded to simulate the highlights and shadows produced by a light source, is widely used on modern maps including many digital mapping services. This approach, pioneered by Swiss cartographers in the late 19th century and further refined and documented by cartographers like Eduard Imhof [11], uses natural visual cues to suggest the shape of the terrain. Similarly, elevation coloring (or “hypsometric tinting”) uses color gradients to visualize terrain elevation. Colors may be modulated with illumination to more clearly illustrate shaded and illuminated terrain slopes. Relief shading and elevation coloring provide only relative elevation, but are easier and faster to read than contour lines and portray the terrain as a continuous surface [4, 19, 20].

Even for contemporary mapmakers, representing terrain in a map requires compromising between clearly conveying the shape of the terrain and communicating other information. For example, colorful stepped hypsometric tints can make elevation comparisons easier, but they also make it harder to use color elsewhere on the map. Imhof recommends a number of combinations of techniques appropriate for various map scales [11, chap. 14b]. However, even if contour lines, spot heights, relief shading, and hypsometric tinting are combined well it can still be difficult for many users of 2D maps to compare elevations or to understand the overall shape of terrain.

Alternatively, elevation information can be represented by applying projections that deviate from the traditional, orthographic top-down views—such as perspective, cylindrical, or oblique projections [12]. Mapmakers may also create views that distort the underlying map to emphasize specific terrain features [6, 13]. The panorama maps often used to show ski slopes are a common example of this technique. These types of maps are widely used and preferred by many map users—largely because the sides and edges of features more closely resemble what viewers see when looking at the actual terrain. Unfortunately, distorting or projecting terrain typically introduces occlusions which obscure portions of the map and can impede distance and area judgments.

Insights from Perceptual Psychology

While the problem of conveying terrain relief has been studied for several centuries, the mechanisms of 3D shape perception are barely starting to be fully understood. We now know that humans infer 3D shape from shading, gradients of texture, contour, binocular disparity (stereopsis), and motion [25]. Cartographers have extensively used—or experimented with [3]—all these cues except for motion, which has been mostly disregarded. Yet motion is known in perceptual psychology to be one of the most important sources of depth information. Depth information from motion can be inferred quickly, in less than a second [18], although findings remain controversial as to which of these cues is the most effective [26]. The extraction of 3D structure from motion has been given different names depending on the conditions under which it occurs [28], here we refer to the general phenomenon as *kinetic depth*.

Computer-Supported Map Exploration

Compared to early hand-drawn maps, computers dramatically expanded the range of possibilities for representing terrain relief. Maps can now be automatically generated, making it eas-

ier to explore more advanced relief rendering techniques. Dynamic animated and interactive displays as well as advanced input and output hardware also provide new opportunities for map rendering, but introduce new challenges for interaction.

2D Map Displays. Most applications that display 2D maps support a set of simple interaction techniques for navigation. Typically, navigation techniques are meant to emulate the way people examine paper maps: panning moves the map around, while zooming adjusts the map’s distance to the user. On multi-touch devices, most map applications let users pan, zoom, and rotate content using similar “drag” and “pinch” interactions. However, when terrain is displayed in map applications, relief information is just as difficult to extract as from traditional paper maps and depends heavily on graphical cues like shading.

3D Map Displays. Other map applications such as Google Earth offer dynamically rendered perspective views of terrain and 3D navigation techniques. This approach has several advantages for relief perception. First, perspective rendering can reveal terrain from viewpoints that make it easier for users to understand its shape. Second, interacting with 3D maps provides kinetic depth cues that can immediately reveal complex terrain structures.

However, 3D map displays also have a number of limitations. In particular, interacting with full 3D views can be more complicated than with 2D representations because it requires the user to manage more degrees of freedom [23]. In addition, 2D interactions do not translate into unambiguous movements in 3D space and complex 3D rotations can make it complicated to obtain the desired view. As a result, camera control can often distract users from their primary task. In response, prior research on 3D navigation has suggested adding constraints to reduce the number of degrees of freedom [23, 24]. Applications like Google Earth only allow users to control heading (rotation around z) and pitch (rotation around x), while Apple Maps and the web-based Google Maps default to toggling the pitch between a few fixed inclinations. Other research has focused on devising more direct 2D-to-3D interaction mappings (e.g. [2, 22]) which, while effective, remain more complicated than traditional map interactions.

Advanced Display and Input Hardware. Over the past few decades, researchers have also explored a number of display technologies for better conveying depth [16] including stereoscopic, autostereoscopic (3D without glasses), and volumetric displays. Stereoscopic displays have made their way to cinemas and home theaters, while autostereoscopic displays are now available on mass-market mobile devices like Nintendo’s 3DS. These displays provide a vivid impression of depth and can heighten a user’s sense of immersion. However, it does not appear that viewers are able to extract more information from stereopsis than from kinetic depth [25, 26]. While volumetric displays and other spatialized techniques seemingly outperform stereoscopic displays on depth judgment tasks [9], the technology remains bulky and expensive.

In summary, prior work in cartography has largely focused on 2D techniques for representing depth. More recent advances in 3D display technologies have led to an increased focus on reproducing realistic 3D views, often resulting in more com-

Interactive & Multi-Surface Maps




Figure 2. Projections for rendering terrain. Plan oblique relief and relief shearing create the same visual result, but differ in construction.

plex interactions. Meanwhile, kinetic depth, which perceptual psychology has shown to be a key element of 3D structure extraction, has not been used to its full potential. We explore a family of simple and lightweight interaction techniques that provide depth cues using kinetic depth and remain compatible with existing interactive maps.


INTERACTIVE RELIEF SHEARING

Our approach, which we refer to as *interactive relief shearing*, uses shearing animations that enhance depth perception for terrain maps. Interactive relief shearing provides many of the benefits of 3D views—users can view terrain from alternate angles that more closely resemble what they see in the real world and the maps provide kinetic depth information that improves perception of the terrain’s shape. However, this approach also preserves many of the benefits of using 2D terrain maps—interactions remain simple and the map always returns to a standard orthographic view, making distance and area measurements easier.

Maps with *relief shearing* are identical to maps with *plan oblique relief* (Figure 2), a parallel projection technique in which the image plane is translated relative to the ground plane [12]. While Jenny and Patterson [12] coined the term *plan oblique relief* based on the concept of a translated image plane for ray-tracing, we achieve the effect by linearly shearing points in the model in x and y directions based on their elevation. Both techniques produce views in which low-elevation points remain near their original positions, but higher-elevation points are shifted laterally. This exposes the sides of terrain features and emphasizes ridges, saddles, and other terrain features that are less apparent in orthogonal 2D maps. Relief shearing is not a new invention and was used to great effect by 20th century cartographers like Erwin Raisz and Heinrich C. Berann [12]. However, constructing these maps manually was cumbersome and required considerable artistic talent, and thus the technique was not used often.

While Jenny and Patterson [12] used plan oblique relief to generate static maps, we couple shearing to user input. This allows users to interactively shear the underlying model and use kinetic depth cues to better understand the shape of the terrain. To create these views, we first overlay a 2D map image on top of a digital elevation model. When viewed from above using an orthographic projection, the map appears identical to the original image. However, users can now shear the model to expose the shape of the underlying terrain. In the simplest case, which we refer to as *explicit shearing* (Figure 1), a user grabs a point on the map and begins to drag. The selected point on the model’s surface (c) follows the cursor, while the corresponding point at the base of the terrain model (s) remains fixed. As the cursor moves, we dynamically shear the model based on the angle between these two points, so that c always stays beneath the cursor, but the base does not move.

CHI 2015, Crossings, Seoul, Korea



By grabbing and dragging at any point on the map, the user can quickly view the terrain from multiple angles. Shearing also provides motion parallax that enhances depth perception and helps users interpret the shape of the terrain. We shear the entire model (rather than locally deforming the terrain in the area of the drag) to ensure that motion can be interpreted consistently across the whole map.

Integrating Panning and Shearing

We have explored several variations of this basic interaction technique, focusing on interactions that add minimal overhead for the user and which are compatible with the kinds of maps typically used in desktop and mobile settings.

Explicit Shearing. As illustrated in Figure 1, explicit shearing varies the shear in direct response to user input. Grabbing and dragging on a point shears the map in the direction of the drag. When the map is released, the animation reverses and the point returns to its original position. Explicit shearing gives users direct control over the strength and direction of the shear and allows a user to hold terrain at a particular orientation for more detailed inspection. However, explicit shearing also adds a new interaction to the standard set (pan, zoom, and rotate) and requires an additional control or gesture to activate it. In our implementation, users can shear terrain by clicking and dragging with the right mouse button (on desktop machines) or by dragging after a long press (on touch-enabled devices). However, these interactions may conflict with existing uses of the same controls in other mapping tools.

Integrated Shearing. Because explicit shearing must be initiated directly by the user, it provides motion cues only when a user decides to shear. Moreover, it cannot be used while panning or zooming. To address this, we designed an even more lightweight approach, which we call *integrated shearing* (Figure 3) which integrates shearing with the existing panning interactions common on almost all 2D digital maps.

As before, grabbing and dragging a point on the map causes the model to shear in the direction of the cursor. However, when using integrated shearing, we also add an elastic force that pulls the *base* of the model towards the cursor. This force ensures that any time the cursor stops moving (even momentarily) the base of the model is pulled back into position beneath the selected point. The final result is identical to panning the map with the same cursor motion, but the intermediate shearing animation provides new kinetic depth cues with each interaction and helps users maintain peripheral and ongoing awareness of the terrain.

We model the elastic force as a simulated spring with $length = 0$ situated at the base of the model. One end of the spring (s_0) is anchored to the model directly below the initial press, while the other (s_1) follows the current position of the cursor (Figure 3). At each animation step, the elastic force of the spring tries to pull s_0 (and thus the entire base of the




Figure 3. Using *integrated shearing*, the user grabs a point on the map (a) and drags it (b). This shears the model to keep the selected point under the cursor, while an elastic force pulls the map back into alignment (c). The final result (d) is identical to panning the map using the same motion.

model) back into alignment with s_1 . We compute the velocity of the base at animation step t as:

$$v_t = k|s_1 - s_0| + Dv_{(t-1)}$$

where k specifies the spring constant and D specifies the drag coefficient. We then adjust the strength and direction of the shear to match the new angle between s_0 and c . While the cross-sectional diagrams in Figures 1 and 3 illustrate shearing in only one dimension, in practice, we model the spring in 2D and adjust the shearing in both x and y .

Integrated shearing produces motion that adapts smoothly to changes in panning speed and direction and intuitively feels like pulling the top of a jelly model. When released, the model elastically snaps back upright and may oscillate as it comes to rest—depending on the spring constant and drag coefficient used. In our experience, the oscillation produced by using a stronger spring and less drag provides a satisfying conclusion to panning operations, and the additional motion at the end helps persist the depth effect. These oscillations can also be triggered in response to zoom events, providing integrated depth cues as users zoom in and out.

Hybrid Explicit/Integrated Shearing. There is a clear trade-off between explicit and integrated shearing. Integrated shearing adds frequent peripheral depth information without adding any new interactions, but those depth cues are ephemeral and only last as long as the user is panning. Explicit shearing, by contrast, requires direct initiation but gives users more control over the persistence of depth cues and makes it possible to examine terrain from particular angles.

To help provide both peripheral cues and explicit manipulation of terrain, we can also combine explicit and integrated shearing into one *hybrid shearing* interaction. Hybrid shearing extends our integrated shearing model so that movement in a small radius around the selected point shears the model without panning. To achieve this, we allow the length of the spring (previously 0) to grow as long as the cursor stays inside the specified radius (Figure 4). As a result, small movements do not apply any elastic force and do not move the base. However, once the cursor leaves this area, the spring can no longer grow and further motion will cause the model to pan (as with standard integrated shearing).

Hybrid shearing supports both direct shearing and panning via one interaction. However, short pans become more difficult

and—because the shearing behavior changes depending on how far a user drags—it can be harder for a user to predict what effect their interaction will have.

Shearing via Device Tilt. Panning and shearing can also be decoupled entirely and triggered using orthogonal input modalities. For example, on a handheld mobile device, we can use the tilt of the device to control shearing. This allows users to manipulate the map using the standard set of touch interactions, but simultaneously and independently adjust the shearing by varying the device’s position and tilt. When holding the device, a user’s passive body motion can also produce faint variations in the amount of shear, providing subtle depth cues even when she is not actively manipulating the map. This effect is related to, but different from, head-coupled parallax—where head tracking is used to control a virtual window into a 3D scene [8]. Although both techniques provide kinetic depth cues that can change based on the position of the device, tilt-driven shearing can produce views (e.g., oblique views of the sides of terrain features) that are not possible with head-coupled parallax.

Because tilt interactions require the user to actively hold and reposition the output device, they are impractical in some settings (e.g., desktop PCs). Also, while shearing and pointing use different inputs, it is possible for the two to conflict. For example, tilting the device can impede target acquisition by shearing the model and moving touch targets out from under a user’s finger. In practice, however, we find that shearing by changing the device tilt tends to involve relatively small movements that are strongly coordinated with the user’s body motion, making shearing behavior quite predictable.




Figure 4. Using *hybrid shearing*, dragging within a small radius (left) shears the map and it snaps back to its original position when released (*explicit shearing* behavior). Dragging outside of this radius (right) shears and pans the map (*integrated shearing* behavior).

Shearing via Other Inputs. While we have focused on touch- and tilt-driven interactions, other input modes could also pair well with relief shearing. For example, using back-of-device interactions to shear and front-of-device interactions to pan could make it easier to perform both kinds of manipulations simultaneously. Face-tracking using on-device cameras [8] could also be an alternative to device tilt for controlling shearing on desktop devices.

Shearing Direction

When shearing is driven using direct touches or drags on the terrain, the dragged point should remain attached to the finger or cursor to preserve the impression of direct manipulation [10]. In our experience, this effect works well when a user drags on a topographically distinct high point on the map because the selected point appears to be pulled directly while the base of the map and terrain at lower elevations appear to lag behind (Figure 5-Left). This preserves the physical intuition of pulling from the top of an elastic model.

However, if a user drags from a low-elevation point, areas at high elevations will appear to move more quickly than the dragged point and may even occlude it, breaking that physical intuition. As a result, we have found it preferable to invert the shearing direction when interacting with a low point—treating the *top* of the model as fixed (rather than the base) and shearing points based on their distance below it. Low points then appear to be pulled directly, while the higher points lag behind (Figure 5-Right). This gives the sense of pulling on the same model, but from the bottom rather than the top.

In practice, the distinction between a low point and a high point can change when maps are zoomed to show small regions of terrain. For example, Figure 5 illustrates how a high alpine valley can appear to be a high point when zoomed out (as in view A), but then appear to be low point when the view is zoomed closer (view B). In our prototypes and evaluations, we address this issue by classifying each clicked point as either a local minimum or local maximum, based on the elevations of points in its immediate vicinity. We sample all elevations in a fixed radius around the selected point—performing the computation in screen space rather than model space to account for both the zoom level and the resolution of the elevation model. As Figure 5 illustrates, we assume a point is a local maximum (red) if more than half of the neighboring points are below it, and a local minimum (blue) otherwise.

IMPLEMENTATION

In order to explore the space of designs we have implemented several variants of our technique. The principal version is a Java application using OpenGL and based on software by Jenny *et al.* [14]. Our application overlays image files containing 2D maps or satellite imagery on top of digital elevation models and can be configured to use the explicit, integrated, or hybrid interaction models. We also implemented a variant for Windows tablets that uses touch input and data from the device's internal tilt sensor to control terrain shearing.

EXPERIMENTS

We conducted a series of experiments to assess the impact of our proposed interactive shearing techniques.

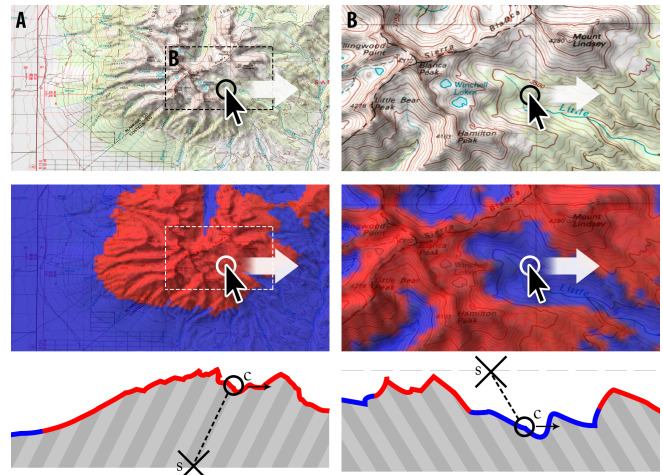


Figure 5. Shearing direction depends on the relative height of the target point. Dragging on a high point (left) shears the model in the direction of the drag using the bottom of the model as the baseline. Dragging a low point (right) uses the top of the visible model as the baseline for the shear. As illustrated here, same point can be either a local maximum (red) or local minimum (blue) depending on the zoom level.

Task Choice

Past experiments in the cartography literature have used a variety of tasks to assess participants' ability to parse and interpret terrain information. These include elevation-reading tasks in which participants must determine the height of a point, line-of-sight tasks in which participants are asked to mentally visualize the terrain between two points and determine if a clear line of sight exists, and comparison tasks where participants must compare heights, slopes, or distances between multiple points on a map (e.g., [4, 19, 20, 21]).

We use elevation comparison tasks, in which we present participants with a pair of points and ask them to determine which of the two points has the highest elevation. These tasks have been widely used in prior studies of map-reading [19, 20] and we focus on them because they provide a strong indicator of whether or not a viewer understands the shape of the underlying terrain. Moreover, elevation comparisons are an implicit building block for other more complex map-reading tasks—such as determining elevation changes along a route or predicting visibility between locations.

Experiment 1: Depth Perception

In our first experiment, we sought to understand the extent to which animated shearing impacts viewers' ability to perceive depth on a map. We compared three different conditions:

Static. This condition showed elevation using traditional 2D techniques including contour lines and relief shading.

Animated Shearing. This condition used the same 2D techniques, but added a continuous animation that sheared the terrain in a clockwise circular pattern with a period of 1.5 seconds. This animation let users see all sides of the terrain and provided continuous kinetic depth cues.

Integrated. This condition used the same 2D cues as the other conditions, but participants could interactively pan and shear maps using *integrated shearing*.

The animated condition would be impractical in most real-world situations where the constant motion would make selec-

tion and direct interaction difficult. However, it provides very strong kinetic depth cues and allowed us to assess the overall impact of shearing animations on terrain perception. The interactive integrated shearing condition, meanwhile, allowed us to test whether participants were able to effectively access depth information by panning the map. We hypothesized that:

- H1. Participants would be better able to discriminate depth when using an animated map rather than a static one.
- H2. Participants would achieve similar depth discrimination when they interactively activated shearing as they did when shearing was animated.

Procedure. We used examples of three common types of real-world maps featuring varying degrees of relief and terrain complexity. The three maps included: a very detailed Swiss topographic map of the Bernese Alps with contour lines and shading; a US National Park Service map of the Glacier National Park with shading and hypsometric tints but no contours; and an urban terrain map of San Francisco from Google Maps that included shading. We overlaid the maps on digital elevation models drawn from the SRTM THIRD dataset¹.

For the elevation comparison tasks, we generated pairs of stimuli randomly according to a set of pre-defined criteria. We discarded points that fell within the outer 5% of the map in either x or y . We also ensured that points fell in the same area of the map and that the distance between them was less than 15% of the map height or width, whichever was larger. To create a continuum of both easy and hard tasks, we selected pairs with a range of different elevation changes between the higher and lower points. When selecting pairs, we eliminated pairs whose elevations were too similar or too far apart, and selected even numbers of pairs whose elevation differences fell into five discrete bins between the minimum and maximum differences. We tuned these elevation differences through a series of internal pilots to ensure that the task was sufficiently challenging to produce a high error rate, but that the most similar elevations were still possible to disambiguate.

At the start of each task, we displayed a white disk which hid the area of map where the new points would appear. This allowed participants to center their view before the task and reduced the need to visually search for the points once they were displayed. After one second, we removed the disk and participants were able to see the two stimulus points, one on the left (red) and one on the right (blue). We instructed participants to “quickly and accurately” determine which of the two points was the highest and press a key (“1” for left, “2” for right) to indicate the higher point. To reduce confusion with the keyboard mapping, we also labeled the physical “1” and “2” keys with red and blue marks that matched the colors of the points. To reduce confusion about which of the two points was the leftmost and which was the rightmost, we avoided pairs of points which were more than 60° from horizontal.

When contour lines are available, a viewer could spend a large amount of time comparing and tracing individual lines in order to determine the higher point. Because we wanted to assess participants’ intuitive sense of the shape of the terrain

¹<http://geo.spectrumtelecom.ca/>

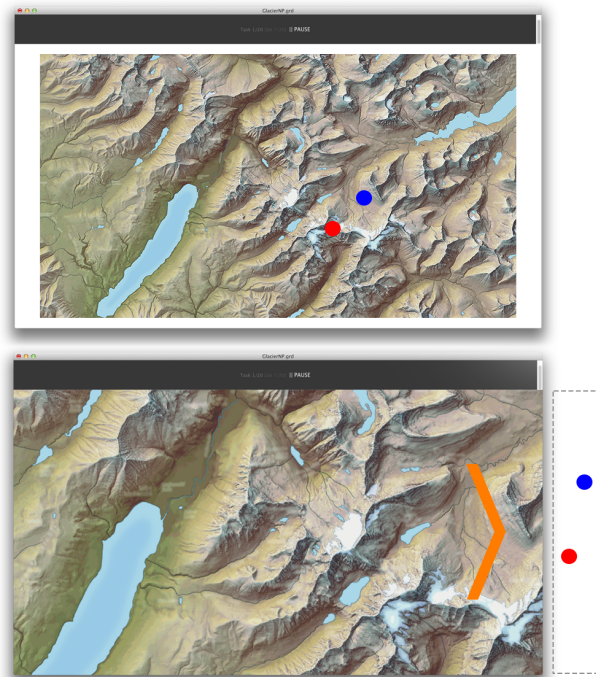


Figure 6. In Experiment 1 (top), participants could always see the entire map. In Experiment 2 (bottom) we increased the size of the map to necessitate panning, and placed all new pairs of points off-screen (simulated here) with an arrow showing their direction and distance. To improve legibility, this figure shows points at 4× their actual diameter.

and not their map reading ability, we chose to fix the amount of time available for the task and focus instead on the error rate. Therefore, we allowed participants to see each pair of points for only 4 seconds before fading out the map and points (participants could still answer after the map was hidden).

We used a within-subjects design in which each participant saw all combinations of maps and motion conditions and completed 45 trials for each (3 maps × 3 conditions × 45 tasks = 405 tasks). To help control for learning effects, we presented tasks in blocks of 15 and systematically varied the order in which maps and motion conditions were presented. Each participant received blocks of tasks in a unique, randomized order where we ensured that no two adjacent blocks used the same map or the same motion condition.

We used e-mail to recruit a set of 20 graduate student participants. Half of the participants were experts — cartographers and cartography students with prior training in map reading and experience with Geographic Information Systems (GIS). The other half had no GIS or cartography experience beyond basic map reading skills. Sixteen participants were male and four were female, with ages ranging from 23–33 ($mean = 27.8$, $SD = 2.8$). Participants self-administered the experiment in a controlled environment by following a set of printed instructions. An evaluator was present to monitor the experiment and answer questions from the participants. Each participant first completed a series of 45 training tasks (15 for each of the three motion conditions) using a fourth map not included in the actual study. Participants were encouraged to repeat the training if they felt they needed additional practice. Participants then proceeded to the main experiment—where they required




Figure 7. (Top) Mean accuracy across all participants by motion condition. Accuracy is normalized so that a score of 0 indicates a result equal to chance. (Bottom) Mean improvement in accuracy between conditions (right technique minus left technique). All error bars show 95% bootstrap confidence intervals.

an average of 36 minutes ($SD = 6$) to complete all 405 trials. Afterward, participants completed a short exit questionnaire. Each participant received a \$15 gift certificate. Participants used a 24" display at 1920 × 1080 in a standard desktop setting. They interacted with the maps using a mouse and entered their inputs using a standard keyboard.

Results. We collected a total of 8,100 elevation judgments from our 20 participants. Due to increasing concerns in a variety of research fields about the use of null hypothesis significance testing [5, 7], we analyzed these results using estimation techniques and report effect sizes with confidence intervals rather than p -value statistics. This reporting methodology is consistent with recent APA recommendations [27].

To compare the three conditions, we computed an accuracy score $a = ((\frac{\text{correct responses}}{\text{total responses}} - 0.5)/0.5)$ for each participant in each condition. This score normalizes accuracy so that a 1 indicates a perfect performance while a 0 indicates a result equal to chance (e.g., randomly guessing). We then computed means and 95% bootstrap confidence intervals [15] across all participants (Figure 7-Top). These results showed that participants tended to be better at differentiating elevations in both the animated and interactive conditions than in the static condition. Pairwise comparisons between the conditions (Figure 7-Bottom) further illustrate that both animation and integrated shearing improved users' performance over the static condition. However, we saw very little difference in performance between the animated and integrated conditions.

These results confirm both of our initial hypotheses—showing that (H1) the presence of constant motion cues from animation improves depth discrimination, and (H2) that users are able to achieve similar improvement when motion cues triggered by their panning interactions. Expert users likely benefited more from the two motion conditions than non-experts did (Figure 8). However, non-experts still showed a clear increase in accuracy in both the animated and integrated conditions.

To check for the presence of learning effects, we computed mean accuracy scores for each participant for each of the 27 task blocks, then used those scores to perform a linear regression for $error\ rate = a \times block\ number + b$. The resulting model ($a = 0.006$, 95% CI [0.000, 0.011], $b = 0.438$, 95% CI [0.388, 0.487]) confirms that the accuracy of partici-




Figure 8. Mean improvement in accuracy between conditions (right technique minus left technique) for expert and for non-expert users. Error bars show 95% bootstrap confidence intervals.




Figure 9. Mean accuracy for each map across all individuals. Accuracy is normalized so that a score of 0 indicates a result equal to chance. Error bars show 95% confidence intervals.

pants' responses did not change much as their familiarity with the maps and techniques increased.

We also computed per-participant accuracies and confidence intervals for each map used in the study (Figure 9). We found that participants' judgments were much more accurate when using the topographically simple Google map than on the more complex National Park Service map. Participants were even less accurate when using the Swiss map, which was more detailed and had considerably more diverse terrain. This ordering held for nearly every single user in the study—regardless of expertise—and confirms that we indeed selected maps that spanned a range of difficulty levels.

Experiment 2: Depth Perception with Integrated Shearing

Based on the positive results of our first experiment, we conducted a second study in which we compared relief shearing against existing 2D and perspective techniques. This second experiment used the same basic elevation comparisons as Experiment 1 but situated them in a more realistic panning task. We also used a Microsoft Surface Pro 2 tablet with an 10.6" screen instead of a desktop PC, which allowed us to include a condition that used orientation sensors to control relief shearing. We compared five different interaction techniques:

Standard Panning. A standard 2D map with no motion cues.

Integrated. The integrated shearing technique.

Top-Down Central Perspective. This condition simulated the default perspective view available in applications like Apple Maps and Google Maps. The camera is fixed to show the terrain from directly above, but the perspective distortion provides some parallax depth cues, especially for steep terrain. We used a field of view of 60°.

Inclined Central Perspective. This condition simulated the inclined perspective views available in these same commercial tools. We set the field of view to 45°, oriented the camera to 30° from vertical, and configured panning so that it translated the camera parallel to the ground plane.

Device Tilt. In this condition, users could adjust the tilt of the device to shear the map independently of their panning.

We hypothesized that:

- H3.** Integrated shearing would improve depth perception when compared to standard panning (no motion cues).
- H4.** Integrated shearing would improve depth perception when compared to top-down perspective which provides more subtle parallax cues.
- H5.** Integrated shearing would improve depth perception when compared to inclined perspective which introduces pronounced depth compression effects.

We also included the more experimental *device tilt* condition to explore whether or not participants would actively vary tilt to manipulate the view, or if depth cues from their passive motion would be sufficient to improve their depth perception.

Procedure. Because Experiment 1 showed no evidence of learning effects or interactions between maps and motion conditions, we simplified Experiment 2 by reducing the number of maps. We used the Swiss map for training tasks and the US National Park Service map for the study itself.

To create realistic panning tasks we zoomed the maps to 2× the size of the viewport. For each task, we generated a pair of points off-screen with a large arrow indicating their direction and distance (Figure 6-Bottom). Participants needed to pan the view (using touch interactions) to reveal and then compare the points. We locked the tablet in landscape mode but allowed participants to physically tilt, turn, and rest the device however they wanted. We also asked them to indicate the higher point by tapping it directly, eliminating the need for a keyboard.

We selected points using the same criteria as in Experiment 1, but always ensured that the next point pair was not in the current view. Because participants could now indicate the higher point via touch, we also removed the restriction on the angle between points. Unlike the first experiment, we did not time-limit participants, but encouraged them to answer “quickly and accurately” and use less than 10 seconds per trial.

As in Experiment 1, we used a within-subjects design where each participant completed tasks using all of the interaction techniques (5 conditions × 60 tasks = 300 tasks). We presented tasks in blocks of 20 and systematically varied the order in which they were presented. We also included a series of 5 training blocks (10 tasks each) to introduce each technique.

We recruited 14 unpaid participants (8 female, 6 male), none of whom had taken part in our first experiment. Participants’ ages ranged from 24–40 (*mean* = 30.75, *SD* = 5.77). Again, participants self-administered the experiment under the supervision of an evaluator, using a combination of printed and on-screen instructions. Participants took an average of 53 minutes (*SD* = 13) to complete the experiment.

Results. As in Experiment 1, we computed accuracy scores for each condition (Figure 10-Top) as well as pairwise differences between each condition. We found that participants had the highest accuracy scores when using *integrated shearing*. Furthermore, pairwise comparisons (Figure 10-Middle) show that participants were reliably more accurate when using *integrated shearing* than any of the other four conditions, confirming hypotheses H3–H5.

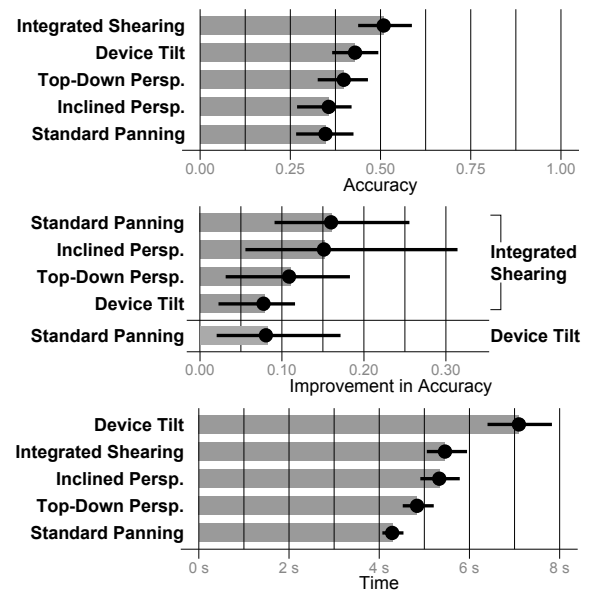


Figure 10. (Top) Mean accuracy across all participants by condition. Accuracy is normalized so that a score of 0 indicates a result equal to chance. (Middle) Mean improvement in accuracy between conditions (right technique minus left technique). (Bottom) Mean task completion times (in seconds). Error bars show 95% bootstrap confidence intervals.

Participants were also more accurate when using *device tilt* than when using standard panning, but device tilt was not clearly better than either of the perspective techniques. However, we observed that participants used device tilt much differently than the other techniques—often holding the tablet at extreme angles and attempting to align the stimulus points in order to compare them. As a result, participants tended to take substantially longer to complete tasks when using the device tilt shear (Figure 10-Bottom). Despite this, 8 of the 14 participants reported that they preferred *device tilt*, while 3 preferred *integrated shearing*, and just 1 preferred each of the other three conditions—although there was likely a novelty effect.

When kinetic depth cues are available, there may be a speed-accuracy tradeoff, since users can continue to extract information by moving the map and viewing the points from new orientations. However, with the exception of *device tilt*, participants spent only about 1 second longer in the motion conditions—even though we did not impose a time limit.

DISCUSSION

The results of these experiments indicate that using shearing to add subtle motion cues can improve users’ ability to perceive depth in maps while adding little interaction cost. In particular, we found that integrated shearing, coupled to existing pan interactions, can help users discriminate elevations more accurately than standard shaded relief maps or perspective views. Next, we describe some of the benefits and trade-offs of interactive relief shearing, and examine several aspects of the technique that merit further exploration.

Benefits and Trade-offs of Interactive Relief Shearing

Using relief shearing to impart depth information on maps has a number of potential benefits, both when compared to traditional 2D maps and to other more-complex approaches.

First, **interactive relief shearing provides kinetic depth information that enhances 3D perception.** These cues make elevation discrimination and map-reading tasks that require depth judgment easier, especially when compared to traditional 2D maps. However, if no depth judgment is required, the animations may be distracting and could make interaction and distance estimation more difficult than they would be with a 2D map.

Second, **interactive relief shearing is primarily a 2D technique and is compatible with 2D map views.** As a result, our approach still facilitates 2D interaction and 2D visual judgment tasks that may be more difficult when using perspective views, stereo displays, and other approaches that attempt to approximate full 3D. Moreover, the motion is ephemeral and maps can rapidly revert to a standard 2D view. Conversely, relief shearing may still be less effective than full 3D views for tasks that require specific viewing angles—for example examining the profile of a mountain ridge or estimating slope steepness.

Third, **interactive relief shearing requires little new interaction** and can be coupled with existing operations like pan and zoom or controlled via an orthogonal input modality. *Integrated shearing*, in particular, provides depth cues without requiring users to learn any new controls or hidden gestures. Moreover, because these interactions rely on direct manipulation, they are likely to be easy to master. Coupling shearing to pan and zoom may be particularly beneficial for search and route-finding tasks that require depth discrimination and already entail extensive map manipulation. However, more complex versions of the technique such as *hybrid shearing* may increase the difficulty of panning. Pairing the depth effect with interaction also makes it less effective than stereo displays or physical models for situations where interaction is impossible or where the user needs continuous depth perception.

Finally, because it requires only a display and 2D input, **interactive relief shearing can be used on most modern PCs and mobile devices.** As a result, our approach is currently much more accessible than ones that rely on a stereo displays, head tracking, or other physical hardware. It could be easily implemented in modern web browsers and mobile devices with 3D graphics capabilities as well as forthcoming e-paper and electrowetting displays with higher refresh rates.

Interaction and Orientation

In addition to our formal experiments, one of the authors also used the the tablet version of our tool extensively over the course of a week-long backpacking trip through the Swiss Alps. During five days on trail, the author and several colleagues used maps augmented with explicit, integrated, and animated shearing to navigate and identify terrain features.

We found that shearing maps to view the profiles of terrain features was especially useful when attempting to match landmarks on a map to the surrounding real-world terrain. However, comparing profiles usually requires reorienting the map so that the face of the terrain feature in question can be sheared upwards. While it is often possible to reorient by physically turning the tablet, combining the device tilt and compass inputs to rotate and shear the map automatically could make matching landmarks even easier.




Figure 11. Canyon maps like this one (shown in its original orientation and rotated by 180°) are especially susceptible to the terrain reversal effect. Depending on motion, lighting direction, and other cues, the canyon can appear to be either concave or convex.

Animated Shear for Non-Panning Tasks

Integrated shearing when panning appears to add useful depth cues with little additional interaction cost. However, common digital mapping tasks like examining a route or set of search results can involve relatively little panning since the map is often automatically centered on the area of interest. In these cases it may be helpful to include a short shearing animation when the map is first presented. Such animations could also be triggered to help viewers re-situate when a map is changed significantly, e.g., when new pins or layers are added.

Interestingly, continuous shearing animations (like those used in Experiment 1) may also be useful in some cases when interaction is not necessary or when the map is being referenced indirectly or from a distance. For example, during our live deployment in Switzerland we found the continuous animations useful during long train trips—where the tablet was positioned on a table and could be used as a glanceable reference for multiple users. Our experience suggests that continuous motion may be useful for enhancing depth on ambient or public map displays intended for group viewing.

The Terrain Reversal Effect

In some cases, we have found that interactively shearing maps can trigger the “terrain reversal effect” [1, 11], in which the elevations on a map appeared to reverse—with high points like peaks and ridges appearing to be low and low points like canyons and valleys appearing to be high. This kind of reversal is well-documented in the cartography literature [11] and is known to occur on shaded-relief maps, especially when the light source or orientation of the map changes (as in Figure 11). The effect is likely the result of an underlying perceptual prior whereby the perceptual system tends to perceive ambiguous objects as convex rather than concave [17]. Because the size and speed of objects in a sheared projection do not change with depth, motion cues do not unambiguously indicate which areas are closer to the viewer and which are further away. As a result, motion from relief shearing animations can have two equally valid interpretations—for example, both a canyon and a ridge with the same elevation change could have identical motion profiles, depending on the orientation of the shear.

The reversal effect is highly individual. In our studies, a few participants experienced reversals on all orthographic maps (including static 2D versions) while many others experienced no reversals at all. However, in our experience, reversals seem to be less common if shearing is triggered using direct manipulation or device tilt—possibly because the sense of direct control makes it easier to interpret the resulting motion.

CONCLUSION

Our explorations show that interactive relief shearing can improve depth perception when compared to both traditional 2D maps and 3D perspective views. Our approach is simple, lightweight, and compatible with existing interaction techniques and could thus be a compelling addition to existing web-based mapping tools and mobile devices. Relief shearing animations may also be useful for adding depth cues to maps on public displays or in other applications where understanding the shape of the terrain is important, but complex interactions or advanced hardware like stereoscopic displays is impractical. Moreover, this work illustrates how lightweight, targeted interactions can make maps (as well as height maps and other 2.5D visualizations) even more usable and accessible—providing a richer sense of space while preserving simplicity.

ACKNOWLEDGMENTS

We thank our participants, Google for a Faculty Research Award, Tom Patterson for use of his Glacier National Park map, and Brooke Marston who helped administer the experiment and provided feedback on the paper.

REFERENCES

- Bernabé-Poveda, M. A., and Çöltekin, A. Prevalence of the terrain reversal effect in satellite imagery. *International Journal of Digital Earth* (2015). To appear.
- Bowman, D. A., Kruijff, E., LaViola, Jr., J. J., and Poupyrev, I. *3D User Interfaces: Theory and Practice*. Addison-Wesley, Boston, 2005.
- Buchroithner, M. F., and Knust, C. True-3D in cartography—Current hard-and softcopy developments. In *Geospatial Visualisation*. Springer, Heidelberg/Berlin, 2013, 41–65.
- Castner, H. W., and Wheate, R. Re-assessing the role played by shaded relief in topographic scale maps. *The Cartographic Journal* 16, 2 (Dec. 1979), 77–85.
- Cumming, G. The new statistics: Why and how. *Psychological Science* 25, 1 (Jan. 2014), 7–29.
- Degener, P., and Klein, R. A variational approach for automatic generation of panoramic maps. *ACM Transactions on Graphics* 28, 1 (Jan. 2009), 2:1–2:14.
- Dragicevic, P., Chevalier, F., and Huot, S. Running an HCI experiment in multiple parallel universes. In *CHI Extended Abstracts*, ACM (New York, 2014), 607–618.
- Francone, J., and Nigay, L. Using the user's point of view for interaction on mobile devices. In *Proc. IHM*, ACM (New York, 2011), 4:1–4:8.
- Grossman, T., and Balakrishnan, R. An evaluation of depth perception on volumetric displays. In *Proc. AVI*, ACM (New York, 2006), 193–200.
- Hancock, M., Cate, T., and Carpendale, S. Sticky tools: Full 6DOF force-based interaction for multi-touch tables. In *Proc. ITS*, ACM (New York, 2009), 133–140.
- Imhof, E. *Cartographic Relief Presentation*. Esri Press, Redlands, CA, USA, 2007.
- Jenny, B., and Patterson, T. Introducing plan oblique relief. *Cartographic Perspectives*, 57 (Spring 2007), 21–40.
- Jenny, H., Jenny, B., Cartwright, W. E., and Hurni, L. Interactive local terrain deformation inspired by hand-painted panoramas. *The Cartographic Journal* 48, 1 (Feb. 2011), 11–20.
- Jenny, H., Jenny, B., and Hurni, L. Interactive design of 3D maps with progressive projection. *The Cartographic Journal* 47, 3 (Aug. 2010), 211–221.
- Kirby, K. N., and Gerlanc, D. BootES: An R package for bootstrap confidence intervals on effect sizes. *Behavior Research Methods* 45, 4 (Dec. 2013), 905–927.
- Knust, C., and Buchroithner, M. F. Principles and terminology of true-3D geovisualisation. *The Cartographic Journal* 51, 3 (Aug. 2014), 191–202.
- Langer, M. S., and Bühlhoff, H. H. A prior for global convexity in local shape-from-shading. *Perception* 30, 4 (Jan. 2001), 403–410.
- Loomis, J. M., and Eby, D. W. Perceiving structure from motion: Failure of shape constancy. In *Proc. ICCV*, IEEE Computer Society (Los Alamitos, 1988), 383–391.
- Phillips, R. J., Lucia, A., and Skelton, N. Some objective tests of the legibility of relief maps. *The Cartographic Journal* 12, 1 (June 1975), 39–46.
- Potash, L. M., Farrell, J. P., and Jeffrey, T. S. A technique for assessing map relief legibility. *The Cartographic Journal* 15, 1 (June 1978), 28–35.
- Rapp, D. N., Culpepper, S. A., Kirkby, K., and Morin, P. Fostering students' comprehension of topographic maps. *Journal of Geoscience Education* 55, 1 (Jan. 2007), 5–16.
- Reisman, J. L., Davidson, P. L., and Han, J. Y. A screen-space formulation for 2D and 3D direct manipulation. In *Proc. UIST*, ACM (New York, 2009), 69–78.
- Shneiderman, B. Why not make interfaces better than 3D reality? *IEEE Computer Graphics and Applications* 23, 6 (Nov./Dec. 2003), 12–15.
- Smith, G., Stuerzlinger, W., Salzman, T., Watson, B., and Buchanan, J. 3D scene manipulation with 2D devices and constraints. In *Proc. Graphics Interface*, CIPS (Toronto, 2001), 135–142.
- Todd, J. T. The visual perception of 3D shape. *Trends in Cognitive Sciences* 8, 3 (Mar. 2004), 115–121.
- Todd, J. T., and Norman, J. F. The visual perception of 3-D shape from multiple cues: Are observers capable of perceiving metric structure? *Perception & Psychophysics* 65, 1 (Jan. 2003), 31–47.
- VandenBos, G. R., Ed. *Publication Manual of the American Psychological Assn.*, 6th ed. APA, Wash. D.C., 2009.
- Vezzani, S., Kramer, P., and Bressan, P. Stereokinetic effect, kinetic depth effect, and structure from motion. In *The Oxford Handbook of Perceptual Organization*. Oxford University Press, Oxford, UK, 2014. In press.