

MistForm: Adaptive Shape Changing Fog Screens

Yutaka Tokuda, Mohd Adili Norasikin, Sriram Subramanian, Diego Martinez Plasencia

School of Engineering and Informatics, University of Sussex, Brighton, United Kingdom

{y.tokuda, m.norasikin, sriram, d.martinez-plasencia}@sussex.ac.uk



Figure 1. (A) MistForm is a shape-changing fog screen, enabling single user and two user interaction (B) with 2D or 3D content. (C) We describe a computational approach and projection algorithm which work together to minimize artifacts (shape distortion and uneven brightness) related to projecting on a shape changing fog display (note parallel edges and brightness in the checkerboard).

ABSTRACT

We present MistForm, a shape changing fog display that can support one or two users interacting with either 2D or 3D content. Mistform combines affordances from both shape changing interfaces and mid-air displays. For example, a concave display can maintain content in comfortable reach for a single user, while a convex shape can support several users engaged on individual tasks. MistForm also enables unique interaction possibilities by exploiting the synergies between shape changing interfaces and mid-air fog displays. For instance, moving the screen will affect the brightness and blurriness of the screen at specific locations around the display, creating spaces with similar (collaboration) or different visibility (personalized content). We describe the design of MistForm and analyse its inherent challenges, such as image distortion and uneven brightness on dynamic curved surfaces. We provide a machine learning approach to characterize the shape of the screen and a rendering algorithm to remove aberrations. We finally explore novel interactive possibilities and reflect on their potential and limitations.

Author Keywords

3D displays; Shape Changing displays; Non-solid diffusers.

ACM Classification Keywords

H.5.1. [Multimedia Information Systems] Artificial, augmented and virtual realities

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 2017, May 06–11, 2017, Denver, CO, USA
© 2017 ACM. ISBN 978-1-4503-4655-9/17/05...\$15.00
DOI: <http://dx.doi.org/10.1145/3025453.3025608>

INTRODUCTION

We present MistForm, a shape-changing display featuring a fog display surface and mechanical actuation of the fog manifold. This allows the displacement of a continuous fog surface in a range of 18 cm. MistForm retains affordances of both shape-changing interfaces (i.e. change the position of each point of the display) and immaterial screens (i.e. reach through capabilities and directional light scattering). For instance, a convex shape (Figure 1(A)) can cover user's field of view, optimize its concavity to allow comfortable hand reach and allow for interaction in front and behind the display (i.e. through the fog).

However, it is the interplay between shape changing and mid-air affordances that enables unique 2D and 3D interaction capabilities. First, moving parts of the fog display will change the angle between the observer and the light projected, useful to attenuate common issues in fog displays (uneven brightness and blending of neighboring pixels). Secondly, MistForm can transform these issues into controllable features. Using shape to control uneven brightness and blending can create regions with optimum visibility to several users in front of it (i.e. shared interaction) or disjoint regions where each user can see one region clearly while the others are dim and blurry (e.g. to support personal views or perspective corrected multiuser 3D, as in Figure 1(B)). Finally, adjusting shape to keep the fog surface within the depth of field of the user's eyes features MistForm as the first display with an adjustable zone of comfort [41], allowing direct hand interaction with 3D content in much bigger volumes than typical planar 3D displays.

MistForm's unique ability to use shape to control fog's visual properties also introduces specific challenges. Projected content will suffer from both uneven brightness (i.e. implicit to fog displays) and geometric distortion (i.e. due to projecting on a changing curved surface), especially significant for perspective corrected 3D content.

In the first part of the paper, we focus on addressing the engineering challenges related to creating shape changing fog screens. We first contribute the design of our immaterial shape changing display, MistForm. Next, we identify the challenges related to projecting on these screens and formulate a description of the problem and approach which is scalable for larger display sizes and higher shape resolution (i.e. more actuators). Concretely, we use a data driven approach to estimate the 3D shape of the fog screen based on the input of their linear actuators, and evaluate it using our prototype. Finally, we contribute a projection algorithm to correct image distortion (see Figure 1(C)) by taking shape and fog brightness profile into account, together with other relevant factors (i.e. projector properties, user's location), and report the reduction in distortion achieved.

In the second part of the paper, we contribute an exploration of the interaction possibilities enabled by shape changing fog screens, using our prototype to demonstrate key scenarios. We particularly focus on scenarios exploiting the interplay between shape changing displays and fog displays. We finish the paper with a discussion on the future possibilities and limitations of this type of systems.

RELATED WORK

Our shape changing fog display explores the synergy and engineering problems of integrating shape changing displays and immaterial mid-air displays. In this section, we review previous works on these fields.

Shape Changing Displays

Visions such as Kinetic Interactions [28], Organic User Interfaces [7] or Radical Atoms [14] have fueled research on shape changing interfaces in HCI.

When focusing on display systems, FEELEX [15] was one of the first to feature computer control of bits of the display surface using mechanical linear actuators. Other systems such as Lumen [33], InForm [10] or ShapeClip [12] have adapted this concept to different sizes and formats. Surfex [6] relied in shape memory alloys embedded in foam, the Actuated Workbench [27] or BubbleWrap [2] used magnets, while Takashima et al. [42] used flat panels mounted on Roomba robots, to mention just some approaches. Works from Coelho and Zigelbaum [7] or Rasmussen et al. [36] explore this design space and summarize practical approaches, and Roudaut et al. [37] describe shape resolution under ten quantifiable features.

However, most previous works focus on solid display elements, with only a few systems describing mid-air or permeable systems. ZeroN [20] features a single mid-air element relying on electromagnets. Omirou et al. [25] use ultrasound levitation to control several elements and Sahoo et al. [39] also demonstrated controlled rotation of the particles. As per mid-air permeable systems, Sahoo et al. [38] use charged fog constrained between transparent electrodes to create a small shape changing display surface, at the expense of user's reachability. Lam et al. [17] described a

two-dimensional array of small, flat fog emitters which can create small fog displays at discrete positions across a tabletop, but continuity of the screen can only be ensured across display elements in the same line. A second approach by Lam et al. [18] mounted fog emitters on moving guides, allowing for continuous displacement, but not for continuity across the display. Thus content must be kept small (within the limits of one emitter).

Immaterial Mid-air Displays

Immaterial mid-air displays form an image plane in space, avoiding accommodation conflicts in the proximities of that plane (i.e. in front and behind), and enabling reach through interactions. Two main approaches are available:

Light converging optics

Light converging optical elements (i.e. optical combiners, concave mirrors, convex and Fresnel lenses) have been often used to create a mid-air image by presenting a stigmatic pair of points of the image at symmetric positions with respect to the optical elements [3]. Approaches using optical combiners (e.g. half silvered mirrors) between the user and the image [11, 19, 26, 30, 32] hinder reachability and interaction. Retro-reflective imaging [43, 46] allows for reachable floating images by combining off-the-shelf beam splitters and retro-reflectors. Instead, HaptoMime [23] uses Aerial Imaging Plate (AIP) providing similar affordances. Vermeer [4] uses a swept volume 3D display at the focal point of a concave mirror and forms the volumetric 3D content in a small area above the mirror. Although such light converging methods are effective to present a floating image, when real objects (e.g. fingers) occlude the light converging elements, the real objects always appear in front of the floating image causing incorrect occlusions and eye fatigue.

Light scattering particle

This approach uses particles floating in the air to scatter projected light. Water drop particles [1, 8], fog particles [16, 31, 34] and airborne particles [29, 40] have been used as scattering mediums. Such light scattering pattern depends on the size and shape of the particles used. Most relevant to this research, Mie scattering (i.e. from spherical particles whose diameter is around or larger than wavelength of incident lights, such as in a fog screen) diffuses light directionally towards the projection axis. This produces uneven brightness on the images (i.e. very bright when looking straight on the projector, decaying to the sides, as angle increases). This issue has been reported [35] and corrected in previous works [31]. However, this directional scattering can also offer interesting affordances, to create face to face [24] or multi-view displays [45]. We exploit these in this paper, focusing on their interplay with display's changes in shape, which remain unexplored.

MISTFORM: SHAPE CHANGING FOG SCREEN

MistForm creates a continuous and permeable display surface of 84 cm x 56 cm, which can be moved forward and backwards up to 18 cm, enabling a range of shapes to be created. Mistform's working volume is designed to cover the




Figure 2. Overview of the main elements in MistForm. The actuator module uses a PVC flexible pipe (10cm diameter) and an array of 5 linear actuators. A fan array creates the airflow around the pipe and into the filter module, to create a shape changing fog display with working volume W:84xD:18xH:59 cm space for direct 3D interaction an average (or shorter) adult user can reach, as detailed later in the paper. It can also be used to support collaborative interaction for two users in front of it.

The design of MistForm is an adapted version of the one proposed by Kataoka et al. [16], using a laminar flow of fog constrained among curtains of air to avoid turbulence and maintain the consistency of the display surface. We followed the guidelines proposed by Martinez et al. [31] (projector above the display, fog flowing down) as our interactive scenarios also make use of reach through interactions. We also chose to use a long-throw projector, placed at 2 m from the display. Firstly, this minimizes the angle between projector rays and the observer in front of the display, which allows a better use of the directional light scattering provided by fog screen (i.e. smaller angles provide more brightness). Secondly, a longer throw allows projected content to remain in focus in the 18 cm where our fog screen can move.

As a main difference to previous approaches, our display includes a flexible fog manifold and five linear actuators, displacing the manifold in a range of 18 cm, to enable the shape changing fog surface. This moving manifold requires an additional display space of 9 cm to each side to protect the working space with an additional laminar flow of air.

We finally use an off-the-shelf 3D projector and shutter glasses, to display stereo content, Kinect v1 for hand tracking and interaction (i.e. its laser projector works even through the fog) and OptiTrack for head tracking and perspective corrected rendering.

Modular design

Figure 2 shows the modules in our shape changing fog screen in detail: the air curtain module, the manifold actuator module and the flow filter module. The air curtain module consists of a 7 x 3 array of 12 cm ARX DC ceramic bearing fans (291.6 m³/h), covering the entire display area (84 cm x

36 cm). This module blows air around the flexible pipe and into the flow filter module, to produce a laminar flow of fog.

The screen actuator module consists of five 12V stepper motors and a Ø 100mm flexible PVC ducting pipe. We made a Ø 2mm hole array at the bottom of the flexible pipe and placed it on top of the flow filter module. The pipe is fixed at both ends and is transversely actuated by five stepper motors by linear motion shafts. This allows a displacement of 18 cm at the max speed of 40 cm/s. Mechanical switches at the end of the shafts' range (home positions), initialize them and correct drift during operation.

The fog distribution system delivers fog to the flexible manifold, using a fog machine and a 12 cm DC circulation fan. The fog distribution system controls the fog density and velocity, and influences pipe flexibility (i.e. air pressure changes the stiffness of the fog pipe).

The flow filter module consists of multiple layers of Ø 6 mm vertical fluted plastic. This fluted array filters the wake turbulence created when the flow from the air curtain module blows around the circular pipe, creating a laminar air curtain around the fog screen. This combination of modules results in a dynamic fog screen moving inside a volume of W:84 cm x D:18 cm x H:59 cm. However, this arrangement allows for 3D content to be projected without vergence-vs-accommodation conflict [13] within a volume of W:84 cm x D:30 cm x H:59 cm (MistForm's working volume), as detailed in our exploration of interaction possibilities.

Challenges & General Approach

MistForm offers interesting affordances, such as maintaining the display surface within ergonomic reach for the user, better support for free hand 3D interaction or support for 3D multiuser scenarios using proxemics cues.

However, projection onto dynamic curved surfaces inherently introduces distortion in the shape of the contents. This, and the uneven brightness distribution typical in fog screens (i.e. Mie scattering) can seriously hinder its use.

We contribute a projection algorithm to overcome these limitations. The algorithm computes the contents as they should be seen from the users' point of view (i.e. eye position). Then, for each pixel projected on the curved display surface, the algorithm determines the color that should be mapped to that pixel using the 3D position of the point on the screen and the location of the user's eyes.

Although a full description of the algorithm is provided later in the paper, the explanation above helps illustrate how our algorithm requires accurate registration of all users' eyes, projector and the 3D shape of the display itself.

While the first two are common in VR and projection mapping systems, the real-time characterization of the display shape remains a specific challenge of MistForm and one that we address in this paper, as described in the following section.

SHAPE RECONSTRUCTION

This section provides a deeper description of our reconstruction approach, to aid replication by other researchers.

Modeling of pipe shapes

The deformation of our fog screen is driven by the displacements of our five actuators and the two fixed ends of the display. Our choice of stepper motors allows for controlled positioning of the display, as long as the actuators speed and torque limits are met.

According to these criteria, it should be possible to create a direct mapping between the number of steps taken from our homing position (s), steps per revolution ($S=200$), number of teeth in our gear ($T=11$) and the pitch of our raft ($P=1.9$ cm), and the location x_i of our servo as shown in Eq. (1).

$$x_i = \frac{S_i}{S} \cdot T \cdot P \quad (1)$$

These five inputs are defined as a vector $\bar{x} \in \mathbb{R}^5$. Our goal is to estimate the 3D position \mathbf{v} of each point of our display surface, given a specific input vector \bar{x} . To retrieve these points we projected a sparse set of 15x8 points evenly distributed across the display volume.

This allows us to use a data driven approach, formulating our goal as the computation of the position of these projected points, $V = \{\mathbf{v}_1, \dots, \mathbf{v}_{R \times C}\} \subseteq \mathbb{R}^d$, where $R=8$, $C=15$ and $N=R \cdot C$. Later on, spline fitting will allow us to interpolate intermediate 3D points.

For our explanations, we will make use of homogeneous coordinates (i.e. 3D points as $\mathbf{v}_i(\mathbf{x}, \mathbf{y}, \mathbf{z}, 1) \in \mathbb{R}^4$), homogeneous 4x4 matrices and right hand systems of reference. This will ease our later explanation of our projection algorithm. For some explanations, we will vectorize $V = \{\mathbf{v}_i\}_{i=1}^N$ into a column vector $\mathbf{V} \in \mathbb{R}^{4N}$, representing the sparse shape to compute for a given $\bar{x} \in \mathbb{R}^5$ (i.e. display state). To describe a set C of these variables, we will label them as $V_c = \{\mathbf{v}_i^c\}_{i=1}^N$, \mathbf{V}_c and \bar{x}_c , with $c \in C$.

Ground Truth Capture

Our data driven approach to reconstruct the display shape requires the acquisition of training data. We then use part of this data to correlate the display state to screen shape, and the rest to later validate the correctness of our model.

Preparing for data acquisition

We used a stereoscopic camera (Stereo Labs ZED) to capture 25 images of a checkerboard of known size (10x7tiles, 2.4cm side) at different positions and orientations in front of the camera. We then used the approach by Zhang et al. [47] to retrieve each camera's intrinsic and distortion parameters and their position relative to the checkerboard in each image. We then used Levenberg–Marquardt optimization method [22] to compute the relative position from one camera to another. We finally used a checkerboard placed at a known position, to retrieve the position of the camera pair relative to MistForm (see M_{screen}^{cam} in Figure 3(A)).




Figure 3. Sampling ground truth data: we projected an array of 15x8 points on the display and used a calibrated stereo camera (i.e. intrinsics and extrinsics) to retrieve 1023 samples.

Dataset construction

When we project our static pattern of 15x8 points (see Figure 3(B)), the setup above allows us to use stereo triangulation to detect their 3D coordinates $\mathbf{w}_i(\mathbf{x}, \mathbf{y}, \mathbf{z}, 1) \in \mathbb{R}^4$, relative to the stereo camera. Points are then transformed to MistForm's space using our camera extrinsics as $\mathbf{v}_i = M_{screen}^{cam} \cdot \mathbf{w}_i$.

We collected data from $J=1023$ random shapes, recording the actuators' input (\bar{x}) and three samples of our 15x8 projected points. We created our ground truth V as the average of these three samples, creating our fixed dataset $\mathbf{D} = \{\{\bar{x}_j, V_j\}_{j=1}^J\}$, that we will use for training (subset \mathbf{M} , with 800 samples) and evaluation (subset \mathbf{E} with 223 samples) to assess the quality of our model.

Data driven Shape reconstruction

Regularized Least Square Linear Regression Model

To predict shape V from the five actuators' input (\bar{x}), we use a linear regression model: $V = \bar{x}^T \mathbf{W}$ where $\mathbf{W} \in \mathbb{R}^{5 \times N}$.

Our five actuators move linearly along the Z axis and our projected pattern (15 horizontal dots) assures at least 3 or 4 points fall between two actuators, giving enough resolution to reconstruct the shape of the pipe between two actuators. Very little variation is expected along the Y axis, as our laminar filter module creates a thin straight layer of fog flowing downwards. Therefore, we can reasonably assume the linear relationship between V and actuator's input (\bar{x}).

Let $\mathbf{X} = [\bar{x}_1, \dots, \bar{x}_M]^T \in \mathbb{R}^{M \times 5}$ and $\mathbf{Y} = [V_1^T, \dots, V_M^T] \in \mathbb{R}^{M \times N}$. Applying *Regularized Least Square* method, we can calculate the optimized value of \mathbf{W} as follows:

$$\mathbf{W}^* = \arg \min_{\mathbf{W}} \|\mathbf{Y} - \mathbf{X}\mathbf{W}\|_2^2 + \lambda \|\mathbf{W}\|_2^2 \quad (2)$$

where $\mathbf{W} \in \mathbb{R}^{5 \times N}$ is the linear regression weight matrix and λ is a regularization parameter to penalize large, complex \mathbf{W} value. The solution has the following closed form, where $\mathbf{I} \in \mathbb{R}^{5 \times 5}$ is the identity matrix:

$$\mathbf{W}^* = (\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I})^{-1} \mathbf{X}^T \mathbf{Y} \quad (3)$$

To account for missing data in our set \mathbf{M} , we relied on our assumption that points remain constant across Y (i.e. fog flows downwards). Thus, we assign missing Y values to the average Y value of the corresponding row, and missing X and Z values to the average X and Z values of the column. After the imputation process, we use this complete training data matrix to solve Eq. (3), producing a matrix \mathbf{W} .

Matrix W allows us to estimate the position of a sparse set of 15x8 points on our fog screen at run time. The continuous shape (intermediate points) is interpolated using cubic splines. Our 800 shapes of training data is robust enough to consider drift effects (i.e. missed steps) for moderate periods of time, but we still use the mechanical switches to reset position when actuators reach the end of their range.

Shape reconstruction evaluation

We evaluated our shape reconstruction model using our evaluation set E (223 shapes). Here we report the accuracy achieved and analyze the causes of error by visualizing the error distribution on the fog screen at the XY plane.

To gain insight on the error introduced by the turbulence of the fog itself, independently of the shape, we first measured the deviation between the three measurements we took from each point, for each given shape. For a fixed point, we averaged its deviation across all the different shapes, to provide an estimate of the flow turbulence at that point. Figure 4 (left) shows the application of this process to all points across the display (i.e. display turbulence). The average deviation across all points is 0.94 cm, with a minimum deviation of 0.21 cm to the top of the display and up to 3.8 cm to the bottom right. The five black arrows mark X positions of our actuators, to evaluate their effects.

We can observe higher error patterns at the left and right sides caused by: a) friction between the moving laminar flow (both fog and protective air) and outer non-moving air; and b) blending between neighboring pixels [35], which made measurements less accurate. High error values are also present to the bottom of the fog screen, as a result of flow becoming more turbulent (and nonlinear) as it travels away from the flow filter module. It is worth noting that, while the top of the display shows less turbulence (smaller deviation), there is a point of high error in the top center position. This position matches the location of our projector through the display and we believe the bright spot around the projector could make the detection of projected points around that area less reliable.

Figure 4 (right) shows the average error from our regression model across the display. The global average error of the model is 1.2 cm, with a maximum error of 1.84 cm to the bottom of the display around $X=-20$ cm. A minimum error of 0.41 cm is located around the top center. We found this average error acceptable as it lies in the order of magnitude of the inherent turbulence (0.94 cm) and the thickness of our fog screen (~ 2.0 cm).

The error is also more evenly distributed than in the previous configuration, but it still shows the same tendencies we observed from the Figure 4(left), showing higher values at the bottom and both sides of the fog screen. The transition from laminar to turbulent flow is more noticeable below $Y=-45$ cm, again leading to higher errors in the shape estimation.




Figure 4. Estimation of the turbulence across the display (left) and reconstruction error from our regression model (right). Black arrows indicate step motor X positions.

There are unexpected high error bands around $X = -20$ cm (between the first and second stepper motor), not observed in Figure 4(left), indicating a lower predictability of that part of the pipe. After post-hoc investigation of the pipe and the step motors connection, we found that the length of the flexible pipe between these stepper motors was slightly longer than others sections, which resulted in a more flexible (and less predictable) part of the display.

Shape control: from intended shape to actuators' input

Our *Regression* model allows us to, given the actuators input, compute the resulting display shape. However, in most situations (as exemplified later in our examples), we will need the reverse process: to compute what the input must be to produce a desired shape.

Unfortunately, our method does not support inverse regression (i.e. not invertible matrix). We implemented a simple gradient descent method which, given a set of input points, finds the actuator's input that results in a shape minimizing error (i.e. sum of the squares of the distances). We use Eq. (1) to produce the first estimation for actuators position and resulting shape, which we refine iteratively. At each step we compute the gradient for each actuator (i.e. how much a change of one step affects error), and iterate for the actuator maximizing gradient (minimizing error).

PROJECTING ON MISTFORM

The model described above allows us to control the 3D shape of our fog surface. However, as introduced earlier, observers' location and projector's parameters are also required to correct shape distortion and uneven brightness.

We first describe the process followed to calibrate our projector and then provide a detailed description of our rendering process. We finally provide an evaluation, assessing the final round trip reprojection error a user would perceive, as a result of inaccuracies in our calibration and 3D reconstruction method.

Projector calibration

We used a projector-camera system (inFocus In116A and Kinect v1) to determine their intrinsic and extrinsic properties. We first determined projector's intrinsics and its position relative to our depth camera and then computed the position of our camera relative to MistForm.

Projector's intrinsics and extrinsics relative to Kinect were computed by using a sparse dataset of points mapping projector 2D coordinates to Kinect 3D coordinates (fixed 9x6 checkerboard projected onto a flat white board, collecting a total of 594 points). We used the approach by Zhang et al. [47] to retrieve our projector's extrinsics (M_{kinect}^{proj}) and intrinsics, and transformed the intrinsics matrix¹ into the equivalent projection matrix (P_{proj}).

We determined the position of Kinect relative to MistForm (M_{screen}^{kinect}) using a checkerboard, as we did with our stereo camera before. This allows us to compute the position of the projector to MistForm as $M_{screen}^{proj} = M_{screen}^{kinect} \cdot M_{kinect}^{proj}$.

Rendering algorithm

MistForm's rendering algorithm makes use of several stages, explained in the following subsections. Our current implementation of MistForm with a single 3D projector can provide support for stereoscopic rendering for two users (exemplified later in the paper), but the same algorithm can apply to glasses-free multi-projector arrays supporting several users. Also, our brightness compensation makes use of the attenuation profile of fog. However, using other profiles (e.g. BRDF of cotton fabric) would allow the algorithm to be reused for other shape changing displays.

Rendering user specific views

For a variable number of N users, we produce a set of 2·N observer views E , one for each eye (Figure 5 illustrates an example for one user, for simplicity). In the case of a multi-projector array, one observer view would be computed per projector. Let $e_i \in E$ be an observer view, M_{screen}^{ei} its position relative to MistForm and P_{ei} the projection matrix that determines what the eye would see through MistForm's front panel (defined by corners $TL(-0.45, 0, 0.18)$ and $BR(0.45, 0.56, 0.18)$).




Figure 5. Example of rendering algorithm for a user that computes the views for each user's eyes through MistForm's front panel. Each pixel F projected on the screen is reprojected towards the observer (F'), to determine its color.

¹ <http://ksimek.github.io/2012/08/13/introduction/>

We then render the 3D scene into a texture t_i , capturing the observer's view of the 3D scene from its position. This produces a total set of textures $T = \bigcup_{i=1}^{2N} \{t_i\}$, with the views of all N users around the display.

Geometric Compensation

Compensation of geometric distortion in each texture t_i can be faced as a projection mapping problem. Using our regression model, we create a 3D model of the current shape of the fog surface, and render it from the projector's perspective (i.e. using P_{proj} and M_{screen}^{proj}), which allows us to control the projection on each point of the fog surface.

We created a custom fragment shader that reprojects each fragment (i.e. point on the fog screen) to observer's space and then retrieves the colour that should be projected on that point from t_i . Let f_{NDC} be the coordinates of the screen point in projector's NDC space. Its mapping to observers' NDC coordinates (f_o) can be computed as shown in Eq.(4).

$$f_o = P_{ei} \cdot (M_{screen}^{ei})^{-1} \cdot M_{screen}^{proj} \cdot (P_{proj})^{-1} \cdot f_{NDC} \quad (4)$$

The mapping from observer's NDC to UV texture coordinates can then be easily computed as in Eq.(5).

$$(u, v)_{ti} = (0.5 \cdot (f_o.x + 1), 0.5 \cdot (1 - f_o.y)) \quad (5)$$

It must be noted that all matrices (and their inverse) involved in this reprojection algorithm are constant across all pixels in a frame and can be precomputed once per frame. Thus, only matrix multiplications are required, incurring on minimal performance hit.

Brightness Compensation

We follow the approach described by Martinez et al. [31] to correct uneven brightness (see Figure 6), but instead of using a sparse attenuation mask and interpolation, we apply it on a per pixel basis in the fragment shader. We do this by encoding the attenuation profile (i.e. mapping $\cos \alpha$ to the inverse of the brightness distribution) as a precomputed texture. This is then used by our fragment shader to correct brightness on each pixel, based on the dot product between the projected ray and the observer (see angle α in Figure 5).

Evaluation

To demonstrate our algorithm can correct the projection distortion induced by the shape of the screen, we displayed several virtual spheres on 50 different fog shapes (Figure 7, left), and analyzed their reprojection error (i.e. distance between where each sphere should be seen and where it was actually detected). To do this, we used a calibrated camera at 70 cm in front of the display (i.e. position of an observer's eye). The spheres were distributed in three planes of 11x6 spheres, across MistForm's volume (see *enhanced working volume for 3D free-hand interaction*, later). We measured reprojection error using three alternatives:

- *No Correction*: This mode illustrates the effect of not taking the shape curvature into account. As a best effort in this case, we compute perspective corrected views as seen through the plane $Z=0$ (i.e. centre of the volume).




Figure 6. Example image without correction (A) and with correction (B).

- *Naïve Spline Reconstruction*: We produced a naïve reconstruction model based on cubic splines. We used the displacement of our control points (Eq. (1)) and the position of the two fixed ends of the pipe as control points, to approximate the shape of the pipe, and simply extruded this curve downwards (i.e. fog flowing downwards), to estimate the shape of the display. The projection algorithm proposed was used, the only change was the reconstruction algorithm.
- *Regression Model Reconstruction*: We used our regression based reconstruction method and projection algorithm, as explained throughout the paper.

Figure 7 (right) shows the results obtained from each of these three configurations. Unsurprisingly, the Naïve approach yields very negative results, with an average reprojection error of 12.34 ± 8.9 px (mean and standard deviation). Results only remain relatively low at the top centre of the display, where the camera is aligned to the projector, yielding minimum distortion. Considering our camera location and extrinsic parameters, this can translate in an error in the range of 0.49-3.02 cm. The high deviation in error across the display also implies noticeable distortion as content is moved across different points of the display.

Naïve Spline and *Regression* models resulted in average errors of 6.36 ± 3.72 px and 3.79 ± 2.2 px, respectively. Error distribution follows our previous observations, being relatively higher to the sides and bottom of the display. The *Spline* model shows higher errors in the right side of the display. This can indicate parts of the display with a more

complex dynamic behavior of the pipe, which does not adapt to a *Spline* model defined by the position of the actuators alone. In this part of the display, the additional resolution in our model (i.e. we projected 3-4 points between each two actuators) allows our model to capture these behaviours, resulting in a more stable error distribution. These models result in average errors of 0.91 cm for the *Spline* model and 0.55 cm for *Regression*. This difference might not seem enough to justify our more complex reconstruction approach, however, we believe the results from our model could be subject to a ceiling effect due to the inherent turbulence in the fog. Besides, the higher deviation in error in the *Spline* model (error in the range of 0.37 – 1.43 cm) results in a less stable correction as contents are moved across the display. Our model provides more stable results (0.22 – 0.85 cm).

EXPLORING INTERACTION POSSIBILITIES

The type of displays proposed in this paper inherit affordances both from shape changing and fog displays. However it's the combination of both that enables new interaction possibilities, both for 2D and 3D content. We first identify these novel affordances and then discuss the opportunities they raise for 2D and 3D interaction.

Identifying affordances

Blending, uneven brightness and shape changes:

Fog screens introduce visual artefacts due to directional light scattering (uneven brightness) and volumetric scattering (blending pixels) of small particles [31, 34]. Blending to neighbouring pixels is related to the thickness of the fog display. Light starts scattering as it travels through the fog curtain, which will make it visible in other areas/pixels of the display (see red area in Figure 8(A)). Thus blending will be minimal where projector's light rays are perpendicular to the fog surface and increase to the sides.

Uneven brightness is also related to the angle between the projector rays and the observer, making the display brighter between the observer and the projector and decreasing to the sides (yellow represents brightness in Figure 8(A)).

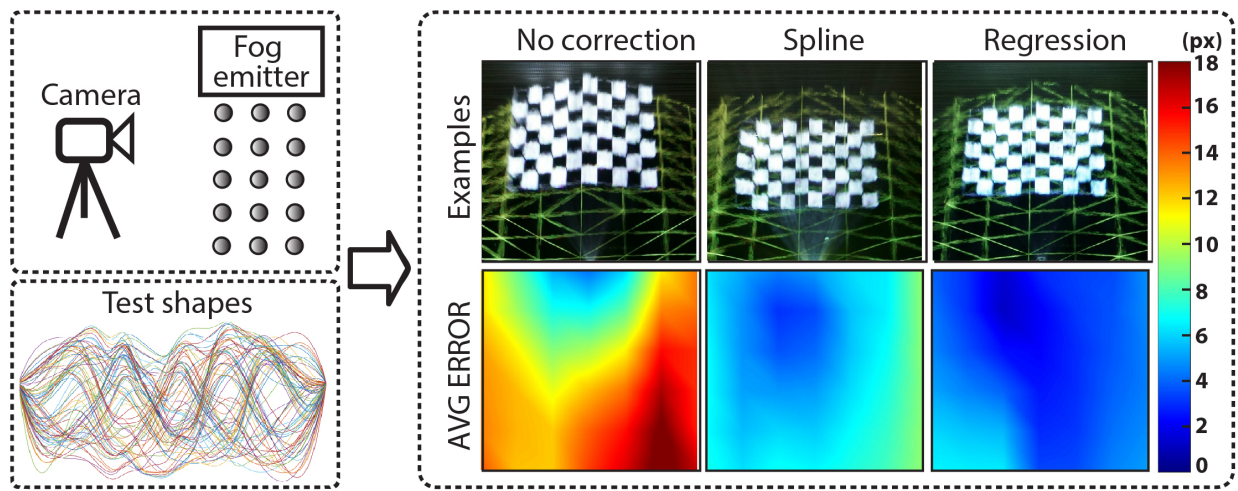


Figure 7. We projected virtual spheres distributed across MistForm's working space (top left), using 50 different shapes (bottom left). We report the reprojection error for three different approaches: no correction, Naïve Spline correction and using our model.




Figure 8. (A) Interplay between the position of the display, brightness and blending with neighbouring pixels. Both effects become more apparent when the display moves close to the user (B), and decrease with distance (C)

In a system such as MistForm (with the user standing at an interactive range to the display and the projector far away) small displacements of the screen can have a drastic influence on these effects. As shown in Figure 8(B), when MistForm is at its closest position to a user, both uneven brightness and blending become more significant than with the display 18 cm further to the back (Figure 8(C)).

MistForm's shape change capabilities not only attenuate such common issues but also turn issues into controllable features, which enable zones of shared visibility (Fig. 9(C)), personalized 2D (Fig. 9(D) and Fig. 10(B)) and 3D (Fig. 12) views, as we explain in the following next sections.

Zone of Comfort (ZoC), fog screens and shape changes:

Auto-stereoscopic, multi-view and, in general, any planar 3D displays only allow small working volumes, due to the interplay between vergence-accommodation conflict [13] and the limited depth of field (DOF) of the human eye.

When seeing 3D content in these displays, our eyes converge on the 3D object, but they must focus on the display surface (i.e. vergence-vs-accommodation). Given the DOF of our eyes (± 0.3 dioptries (DPT)), only real objects within 0.3DPT of the display surface will also be in focus, which defines the display's *zone of comfort* (ZoC) [41]. Thus, if we focus at a display within our arms' reach (e.g. 50 cm ~ 2 DPT), we will

only see our hand in focus in the 7cm in front of the display (43 cm ~ 2.3 DPT), which means we will only be able to interact comfortably with 3D contents inside those 7 cm.

In comparison, MistForm offers two advantages. First, being immaterial, our hands can access the space behind the display surface. In the example above this would give access to an extra 9 cm (59 cm ~ 1.7 DPT) resulting in a ZoC of 16cm. Secondly, by moving the fog surface, the location of this bigger ZoC can be dynamically adjusted. Tracking user's focus (e.g. gaze or hand trackers), the *dynamic ZoC* can move to keep both hands and 3D content within the ± 0.3 DPT range, allowing comfortable interaction (no accommodation issues). Contents outside, would still be correctly visible (using our algorithm), but the *adjustable ZoC* should move towards them if we also want to interact with them. Our implementation (later in the paper) describes how this can be used to increase the working volume from 7cm up to 30 cm.

Interacting with 2D content:

When not using 3D glasses, MistForm can be used to display 2D content on its display surface. The *geometric compensation* is not needed here, but the display's shape reconstruction will still be needed to wrap contents correctly (e.g. make the curvature of a dialog window match the actual curvature of the screen) and to allow user's input to be mapped to content correctly (e.g. map a finger touch to a small button).

Exploiting the relationship between *brightness*, *blending* and *shape*, MistForm's projection surface can adapt to the position of users in front of it. For a single user, a concave shape can cover the space his arms can reach comfortably (Figure 9(A)). The central part of the shape would minimize artefacts, enabling a main space for interaction. The space to the periphery, would suffer both from blending and lower brightness, being more appropriate for holding peripheral content. Adapting the shape as the user moves (Figure 9(B)) can keep contents within reach or even at constant positions relative to the user (central or peripheral), as a reduced version of the personal Cockpit [9].

If a second user joins, this affordance can be applied to reinforce proxemics interactions between the users [5]. For collaborating users (see Figure 9(C)), the display should be




Figure 9. Examples for 2D interaction. (A) A convex shape produces a central region with optimum visibility, with bigger distortion to the sides. (B) This shape can adapt to user's displacements in front of the display. (C) Concave shapes offer support for close collaboration, while a triangular shape creates regions for personal interaction (D).




Figure 10. Convex shapes adapt well for single user scenarios (A). Triangular shapes create disjoint spaces for two users (B)

placed to the back of their o-space (but still within the limits of the user's reach). This will assure brightness is more even for both observers, while retaining the spatial features of the F-formation.

In contrast, if users break apart to focus on individual tasks (e.g. separation, divergent users' views), the display could progressively vary shape from a concave to an almost triangular shape (see Figure 9(D)), creating a differentiated surface for each user. Unlike the previous configuration, this would result in very different angles between each point of the screen and the observers. While the space in front of each user offers good visibility, their visibility of the other user's space would suffer from high brightness attenuation and blending, which both work together to create personal spaces for each user (see Figure 10(B)).

Interacting with 3D content

Enhanced working volume for free-hand 3D interaction

MistForm was designed relying on the possibility to dynamically adjust the *ZoC* of the display, by adjusting the position and shape of the fog display.

Our implementation offers a total working volume of up to 30cm, covering the space where an average male adult user can interact comfortably. We used the ergonomic metric RULA [21] and computed minimum (i.e. upper arm at -10° ; and lower arm at 100°) and maximum interaction distances (i.e. upper arm at 45° ; and lower arm at 80°). Using average anthropometric measurements (i.e. upper and lower arm length of 36.88 and 35.92 cm), this results in a working volume of 28.8cm and 55.4cm, which aligns with estimates by other researchers (e.g. 0.4-0.8 arm's length as in [44]).




Figure 11. Our actuator's range (18cm) is enough to dynamically extend the display's zone of comfort to all the range an average male user can reach comfortably, by adjusting the position of the fog screen.




Figure 12. Triangular shapes provide users with independent 3D views, while minimizing visibility of the other user's view.

Figure 11(A) and (B) illustrate how, for a user standing at 32 cm from our display, the range of our linear actuators (18cm) is enough to cover this working volume. In this situation, for any 3D content shown inside the working volume, it is possible to place the diffuser surface within 0.3 DPT distance, assuring that both the object and the user's hands interacting with it stay within the *ZoC* (dotted line).

In our implementations, we use user's hand position to detect the intent of his/her interaction. The bounding box of the objects within 7cm to the hand are fed to our shape control algorithm to automatically determine the best display position. This binary condition can cause sudden changes of shape when hand movement makes an object fall out of the range, and progressive approaches should be used instead.

Support for multiuser scenarios

The dependency between *brightness*, *blending* and *shape* allows MistForm to deliver different 3D views to its users, within some limits.

Multiview support is depicted in Figure 12, and it is based on the use of triangular shapes. As explained in our multiuser 2D scenarios, this configuration maximizes the difference in perceived brightness between the two regions of the display and was used as a way to enforce disjoint working spaces.

Here we use this same property to allow for different perspectives of the same shared 3D object, presented in the o-space between the users. Each user will see a bright image of the 3D object from his/her perspective, while only getting a residual image (less bright and distorted) from the other user's perspective.

Our approach cannot allow for full multiuser 3D support (e.g. contents spanning across the whole screen) and can only dedicate specific parts of the display to each user. However this can still be interesting to examine common objects of interest, it can be supported with a conventional 3D projector and can be combined with personalized 2D contents to support mixed focus scenarios as those in [31].

DISCUSSION

MistForm is the first system to explore the novel interactive possibilities that raise from combining: a) the inherent properties of fog as a display medium; and b) the ability to dynamically track and move the immaterial screen to affect its visual properties. Our exploration shows how this can offer exciting possibilities. Adjusting the shape of the display can be used to reduce blending and uneven brightness,

known issues in fog displays. Besides, using shape changes can transform these issues into useful controllable features. We explore this approach, identifying display configurations to optimize visibility to all users, or to produce personalized regions to specific users. Combining these patterns with proxemics or ergonomics considerations and specific use cases, have allowed us to identify interesting usage scenarios.

Moving the fog surface enables a *dynamic/movable zone of comfort (ZoC)*. This has allowed us to produce a prototype which (using a relatively small actuation range of 18 cm) provides a working volume over four times bigger than an alternative planar 3D display and can cover the comfortable interactive range of an adult user. *MistForm's dynamic ZoC* can offer a flexible alternative to other display techniques in scenarios where 3D interaction with our real hands, real tools or semiotic gestures to co-workers are important. For instance, in a 3D CAD application, a designer using a CAVE would only be able to interact comfortably with 3D contents inside a few centimeters in front of the display. Proxies (virtual hands, pointers) allow for interaction beyond this space, but at the expense of the expressiveness of the users' real hands (eyes focus on the CAVE wall, not on the hands). If the designer needs to be aware of a client's feedback (e.g. client pointing to a part of a 3D object), the designer will be forced to switch focus between the display surface (to see the content) and the client's hand (to see the gesture). This can hinder collaboration and add unnecessary visual fatigue (shifting focus between hands and display even if both are virtually close). VR solutions can attenuate these issues (no shifts in focus), but direct facial expressions are occluded and collaboration is limited by the expressiveness of the mediators (avatars) used.

Benefits also apply in training scenarios like Surgery, where tools (e.g. scalpel) are central to the task. *MistForm* allows hands, scalpel and (virtual) organs to remain in focus. Awareness of their surrounding (no HMDs) would allow trainees to look and engage with other users as in the real world (nurses in an operation theater).

The approach and methods described, and their evaluation using our *MistForm* prototype, should also serve as a basis for other researchers exploring this space, helping them solve the most demanding challenges involved.

The prototype presented, however, is not free of limitations. Our naïve choice of actuators and fans is noisy (66dBA at a user's interaction area), which can hinder interaction, moreover when considering collaborative scenarios (i.e. talk to peers while collaborating).

Using an aerosol heavier than air (e.g. cooling down fog, or using dry ice) could be first venue to explore. Our fog aerosol is lighter than air, which forces us to use a relatively faster flow of air to achieve the desired display height. A heavier aerosol would perform with a slower flow, allowing the use of less powerful (and less noisy) blowers.

A better choice of linear actuators could also reduce noise even further, and could also improve *MistForm's* ability to adjust its Zone of Comfort (currently limited to 40 cm/s), making it suitable for tasks requiring faster hand motion.

Our choice of active shutter glasses also affects perceived brightness, which could become an issue if bigger form factors are explored (i.e. brightness would become even lower to the sides, and shutter glasses would further decrease brightness). Although this is not an issue in our system, this choice of 3D technology was simply a matter of convenience and we would recommend interested practitioners to use alternative technologies, such as passive glasses. This would involve the calibration of a second projector, but the methods and algorithms described here can still be applied.

Using multiple projectors is another interesting line of exploration. As demonstrated by Yagi et al. [45], a multi-view display could be created using a dense array of projectors, removing the need for 3D glasses. Besides, as these rely on using the brighter parts of the scattering function (i.e. users see the pixels from the projector that is directly aligned to their eyes), this would allow for the use of less bright projectors and still make better use of their brightness profile. Moving the display surface relative to the user could also help align the views to the user's eyes, avoiding crosstalk between views.

CONCLUSION

MistForm is a novel display technology combining the affordances of shape changing and fog displays. We presented a suitable design and identified challenges related to this technology, addressing the estimation of the display shapes created and a projection algorithm to remove optical artefacts (geometric and brightness compensation). We evaluated both our reconstruction and projection algorithms, showing that our approach can successfully minimize distortion due to projecting on changing surfaces.

We then explored the novel possibilities enabled by combining the shape changing and mid-air nature of our system. For example, *MistForm* can use shape changes to affect how brightness and blending affects different areas of the display, and we show how this can be used to support different scenarios involving single or multiple users.

We also show how moving the fog surface can create a dynamic zone of comfort for the display, enabling free hand interaction over bigger volumes without vergence vs accommodation issues. We believe these features illustrate the potential of *MistForm* to enable new forms of interaction and collaboration.

ACKNOWLEDGEMENT

This work is supported by the European Research Council (Starting Grant Agreement 278576) under the Seventh Framework Programme and Universiti Teknikal Malaysia Melaka (UTeM) and Ministry of Higher Education Malaysia (MoHE). We thank Luis Jose Berna for technical support and Luis Veloso for taking the images and videos.

REFERENCES

1. Peter C. Barnum, Srinivasa G. Narasimhan, and Takeo Kanade. 2010. A Multi-Layered Display with Water Drops. *ACM Transactions on Graphics (TOG)* 297, 4 (2010), 1–7. DOI:http://doi.org/10.1145/1833351.1778813
2. Olivier Bau, Uros Petrevski, and Wendy Mackay. 2009. BubbleWrap: A Textile-Based Electromagnetic Haptic Display. In *Extended Abstracts on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 3607–3612. DOI:http://doi.org/10.1145/1520340.1520542
3. Oliver Bimber. 2004. Combining Optical Holograms with Interactive Computer Graphics. *Computer* 37, 1 (2004), 85–91. DOI:http://doi.org/10.1109/MC.2004.1260730
4. Alex Butler, Otmar Hilliges, Shahram Izadi, Steve Hodges, David Molyneaux, David Kim, and Danny Kong. 2011. Vermeer: Direct Interaction with a 360° Viewable 3D Display. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST '11)*. ACM, New York, NY, USA, 569–576. DOI:http://doi.org/10.1145/2047196.2047271
5. T. Matthew Ciolek and Adam Kendon. 1980. Environment and the Spatial Arrangement of Conversational Encounters. *Sociological Inquiry* 50, 3–4 (1980), 237–271. DOI:http://doi.org/10.1111/j.1475-682X.1980.tb00022.x
6. Marcelo Coelho, Hiroshi Ishii, and Pattie Maes. 2008. Surfex: a programmable surface for the design of tangible interfaces. In *Extended Abstracts on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 3429–3434. DOI:http://doi.org/http://doi.acm.org/10.1145/1358628.1358869
7. Marcelo Coelho and Jamie Zigelbaum. 2011. Shape-changing Interfaces. *Personal and Ubiquitous Computing* 15, 2 (2011), 161–173. DOI:http://doi.org/10.1007/s00779-010-0311-y
8. Shin-ichiro Eitoku, Kunihiro Nishimura, Tomohiro Tanikawa, and Michitaka Hirose. 2009. Study on Design of Controllable Particle Display Using Water Drops Suitable for Light Environment. In *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology (VRST '09)*. ACM, New York, NY, USA, 23–26. DOI:http://doi.org/10.1145/1643928.1643936
9. Barrett Ens, Rory Finnegan, and Pourang Irani. 2014. The Personal Cockpit: A Spatial Interface for Effective Task Switching on Head-worn Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 3171–3180. DOI:http://doi.org/10.1145/2556288.2557058
10. Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii. 2013. inFORM: Dynamic Physical Affordances and Constraints through Shape and Object Actuation. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, New York, NY, USA, 417–426. DOI:http://doi.org/10.1145/2501988.2502032
11. Martin Hachet, Benoit Bossavit, Aurélie Cohé, and Jean-Baptiste de la Rivière. 2011. Toucheo: Multitouch and Stereo Combined in a Seamless Workspace. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST '11)*. ACM, New York, NY, USA, 587–592. DOI:http://doi.org/10.1145/2047196.2047273
12. John Hardy, Christian Weichel, Faisal Taher, John Vidler, and Jason Alexander. 2015. ShapeClip: Towards Rapid Prototyping with Shape-Changing Displays for Designers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 19–28. DOI:http://doi.org/10.1145/2702123.2702599
13. David M. Hoffman, Ahna R. Girshick, Kurt Akeley, and Martin S. Banks. 2015. Vergence–accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of Vision* 8, 3 (2015), 1–30. DOI:http://doi.org/10.1167/8.3.33.Introduction
14. Hiroshi Ishii, Dávid Lakatos, Leonardo Bonanni, and Jean-Baptiste Labruné. 2012. Radical Atoms: Beyond Tangible Bits, Toward Transformable Materials. *Interactions XIX*, 38–51. DOI:http://doi.org/10.1145/2065327.2065337
15. Hiroo Iwata, Hiroaki Yano, Fumitaka Nakaizumi, and Ryo Kawamura. 2001. Project FEELEX: Adding Haptic Surface to Graphics. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. ACM, New York, NY, USA, 469–476. DOI:http://doi.org/10.1145/383259.383314
16. Akira Kataoka and Yasuhiro Kasahara. 1993. A Method and apparatus for a fog screen and image-forming method using the same. US Patent 5,270,752. (14 December 1993), Filed December 4, 1992.
17. Miu-Ling Lam, Bin Chen, Kit-Yung Lam, and Yaozhun Huang. 2014. 3D Fog Display using Parallel Linear Motion Platforms. In *International Conference on Virtual Systems & Multimedia (VSMM)*. IEEE, 234–237. DOI:http://doi.org/10.1109/VSMM.2014.7136689

18. Miu Ling Lam, Bin Chen, and Yaozhong Huang. 2015. A Novel Volumetric Display using Fog Emitter Matrix. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 4452–4457. DOI:http://doi.org/10.1109/ICRA.2015.7139815
19. Jinha Lee, Alex Olwal, Hiroshi Ishii, and Cati Boulanger. 2013. SpaceTop: Integrating 2D and Spatial 3D Interactions in a See-through Desktop Environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 189–192. DOI:http://doi.org/10.1145/2470654.2470680
20. Jinha Lee, Rehmi Post, and Hiroshi Ishii. 2011. ZeroN: Mid-Air Tangible Interaction Enabled by Computer Controlled Magnetic Levitation. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST '11)*. ACM, New York, NY, USA, 327–366. DOI:http://doi.org/10.1145/2047196.2047239
21. McAtamne Lynn and E. Nigel Corlett. 1993. RULA: a survey method for the investigation of work-related upper limb disorders. *Applied ergonomics* 24, 2 (1993), 91–99.
22. Donald W. Marquardt. 1963. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *Journal of the Society for Industrial and Applied Mathematics* 11, 2 (1963), 431–441. DOI:http://doi.org/10.1137/0111030
23. Yasuaki Monnai, Keisuke Hasegawa, Masahiro Fujiwara, Kazuma Yoshino, Seki Inoue, and Hiroyuki Shinoda. 2014. HaptoMime: Mid-Air Haptic Interaction with a Floating Virtual Screen. In *Proceedings of the 27th annual ACM symposium on User interface software and technology (UIST '14)*. ACM, New York, NY, USA, 663–667. DOI:http://doi.org/10.1145/2642918.2647407
24. Alex Olwal, Stephen DiVerdi, Ismo Rakkolainen, and Tobias Höllerer. 2008. Consigalo: Multi-user Face-to-face Interaction on Immaterial Displays. In *Proceedings of the 2nd international conference on INtelligent TEchnologies for interactive enterTAINment (INTETAIN '08)*. ICST, 8. DOI:http://doi.org/10.4108/ICST.INTETAIN2008.2485
25. Themis Omirou, Asier Marzo Perez, Sriram Subramanian, and Anne Roudaut. 2016. Floating Charts: Data Plotting using Free-Floating Acoustically Levitated Representations. In *IEEE Symposium on 3D User Interfaces (3DUI)*. 187–190. DOI:http://doi.org/10.1109/3DUI.2016.7460051
26. Hilliges Otmar, David Kim, Shahram Izadi, Malte Weiss, and Andrew D. Wilson. 2012. HoloDesk: Direct 3D Interactions with a Situated See-Through Display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 2421–2430. DOI:http://doi.org/10.1145/2207676.2208405
27. Gian Pangaro, Maynes-Aminzade Dan, and Hiroshi Ishii. 2003. The Actuated Workbench: Computer-Controlled Actuation in Tabletop Tangible Interfaces. In *Proceedings of the 15th annual ACM Symposium on User Interface Software and Technology (UIST '02)*. ACM, New York, NY, USA, 181–190. DOI:http://doi.org/http://doi.acm.org/10.1145/571985.572011
28. Amanda Parkes, Ivan Poupyrev, and Hiroshi Ishii. 2008. Designing Kinetic Interactions for Organic User Interfaces. *Communications of the ACM* 51, 6 (2008), 58–65. DOI:http://doi.org/10.1145/1349026.1349039
29. Kenneth Perlin and Jefferson Y. Han. 2006. Volumetric display with dust as the participating medium. US Patent 6,997,558. (14 February 2006), Filed December 5, 2003.
30. Diego Martinez Plasencia, Florent Berthaut, Abhijit Karnik, and Sriram Subramanian. 2014. Through the Combining Glass. In *Proceedings of the 27th annual ACM symposium on User interface software and technology (UIST '14)*. ACM, New York, NY, USA, 341–350. DOI:http://doi.org/10.1145/2642918.2647351
31. Diego Martinez Plasencia, Edward Joyce, and Sriram Subramanian. 2014. MisTable: Reach-through Personal Screens for Tabletops. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3493–3502. DOI:http://doi.org/10.1145/2556288.2557325
32. Timothy Poston and Luis Serra. 1994. The Virtual Workbench: Dextrous VR. In *Proceedings of the Virtual Reality and Software Technology Conference (VRST '94)*. World Scientific, River Edge, NJ, USA, 111–121.
33. Ivan Poupyrev, Tatsushi Nashida, Shigeaki Maruyama, Jun Rekimoto, and Yasufumi Yamaji. 2004. Lumen: Interactive Visual and Shape Display for Calm Computing. In *ACM SIGGRAPH 2004 Emerging technologies (SIGGRAPH '04)*. ACM, New York, NY, USA, 17. DOI:http://doi.org/10.1145/1186155.1186173
34. Ismo Rakkolainen. 2008. Mid-air Displays Enabling Novel User Interfaces. In *Proceedings of the 1st ACM International Workshop on Semantic Ambient Media Experiences (SAME '08)*. ACM, New York, NY, USA, 25–30. DOI:http://doi.org/10.1613/jair.301
35. Ismo Rakkolainen. 2008. Measurements and Experiments of the Immaterial Virtual Reality Display. In *3DTV Conference: The True Vision-Capture, Transmission and Display of 3D Video*.

- IEEE, 37–40.
DOI:<http://doi.org/10.1109/3DTV.2008.4547802>
36. Majken K. Rasmussen, Esben W. Pedersen, Marianne G. Petersen, and Kasper Hornbæk. 2012. Shape-Changing Interfaces: A Review of the Design Space and Open Research Questions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 735–744. DOI:<http://doi.org/10.1145/2207676.2207781>
 37. Anne Roudaut, Abhijit Karnik, Markus Löchtfeld, and Sriram Subramanian. 2013. Morphees: Toward High “Shape Resolution” in Self-Actuated Flexible Mobile Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 593–602. DOI:<http://doi.org/10.1145/2470654.2470738>
 38. Deepak Ranjan Sahoo, Diego Martinez Plasencia, and Sriram Subramanian. 2015. Control of Non-Solid Diffusers by Electrostatic Charging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI'15)*. ACM, New York, NY, USA, 11–14. DOI:<http://doi.org/10.1145/2702123.2702363>
 39. Deepak Ranjan Sahoo, Nakamura Takuto, Asier Marzo, Themis Omirou, Michihiro Asakawa, and Sriram Subramanian. 2016. JOLED: A Mid-Air Display Based on Electrostatic Rotation of Levitated Janus Objects. In *29th ACM User Interface Software and Technology Symposium (UIST '16)*. ACM, New York, NY, USA, .
 40. Stefan Schneegass, Florian Alt, Jürgen Scheible, and Albrecht Schmidt. 2014. Midair Displays: Concept and First Experiences with Free-Floating Pervasive Displays. In *Proceedings of The International Symposium on Pervasive Displays (PerDis'14)*. ACM, New York, NY, USA, 27–31. DOI:<http://doi.org/10.1145/2611009.2611013>
 41. Takashi Shibata, Joohwan Kim, David M. Hoffman, and Martin S. Banks. 2011. The zone of comfort: Predicting visual discomfort with stereo displays. *Journal of vision* 11, 8 (2011), 1–29. DOI:<http://doi.org/10.1167/11.8.11>
 42. Kazuki Takashima, Takafumi Oyama, Yusuke Asari, Ehud Sharlin, Saul Greenberg, and Yoshifumi Kitamura. 2016. Study and Design of a Shape-Shifting Wall Display. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. ACM, New York, NY, USA, 796–806. DOI:<http://doi.org/10.1145/2901790.2901892>
 43. Yutaka Tokuda, Atsushi Hiyama, Michitaka Hirose, and Hirotugu Yamamoto. 2015. R2D2 w / AIRR: Real time & Real space Double-Layered Display with Aerial Imaging by Retro-Reflection. In *SIGGRAPH Asia 2015 Emerging Technologies (SA'15)*. ACM, New York, NY, USA, 20. DOI:<http://doi.org/10.1145/2818466.2818484>
 44. Dimitar Valkov, Frank Steinicke, Gerd Bruder, and Klaus Hinrichs. 2011. 2D Touching of 3D Stereoscopic Objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 1353–1362. DOI:<http://doi.org/10.1145/1978942.1979142>
 45. Asuka Yagi, Masataka Imura, Yoshihiro Kuroda, and Osamu Oshiro. 2011. 360-Degree Fog Projection Interactive Display. In *SIGGRAPH Asia 2011 Emerging Technologies (SA '11)*. ACM, New York, NY, USA, 19. DOI:<http://doi.org/10.1145/2073370.2073388>
 46. Hirotugu Yamamoto, Yuka Tomiyama, and Shiro Suyama. 2014. Floating aerial LED signage based on aerial imaging by retro-reflection (AIRR). *Optics Express* 22, 22 (2014), 26919–26924. DOI:<http://doi.org/10.1364/OE.22.026919>
 47. Zhengyou Zhang. 2000. A Flexible New Technique for Camera Calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 11 (2000), 1330–1334. DOI:<http://doi.org/10.1109/34.888718>