

Interactive Design of Gallery Walls via Mixed Reality

Haikun Huang
George Mason University
Fairfax, Virginia, United States
hhuang25@gmu.edu

Yuxuan Zhang
Wayfair Inc.
Boston, Massachusetts, United States
yuzhang@wayfair.com

Tomer Weiss
New Jersey Institute of Technology
Newark, New Jersey, United States
tweiss@njit.edu

Rebecca W Perry
Wayfair Inc.
Boston, Massachusetts, United States
reperry@wayfair.com

Lap-Fai Yu
George Mason University
Fairfax, Virginia, United States
craigyu@gmu.edu

Abstract—We present a novel interactive design tool that allows users to create and visualize gallery walls via a mixed reality device. To use our tool, a user selects a wall to decorate and chooses a focal art item. Our tool then helps the user complete their design by optionally recommending additional art items or automatically completing both the selection and placement of additional art items. Our tool holistically considers common design criteria such as alignment, color, and style compatibility in the synthesis of a gallery wall. Through a mixed reality device, such as a Magic Leap One headset, the user can instantly visualize the gallery wall design in situ and can interactively modify the design in collaboration with our tool’s suggestion engine. We describe the suggestion engine and its adaptability to users with different design goals. We also evaluate our mixed-reality-based tool for creating gallery wall designs and compare it with a 2D interface, providing insights for devising mixed reality interior design applications.

Index Terms—Design Interfaces; Mixed Reality; Spatial Computing

I. INTRODUCTION

The advent of mixed reality devices (e.g., Microsoft HoloLens, Magic Leap One) gives rise to new and exciting opportunities for spatial computing. The superior immersive visualization and interaction experience provided by these devices promises to change the way interior design is performed as they allow users to instantly preview and modify designs in real spaces.

Interior design has historically been a costly and time-intensive process. The conventional design process involves contemplating fabric swatches and inspirational photos, as well as talking to a designer. A professional designer may make use of 3D modeling software to preview a design on screen through sophisticated manual operations. The designer’s client must then mentally translate what they see on the 2D screen to how the design may look in a real living space. Without convenient means for visualizing and modifying designs, the design process can be tedious and nonintuitive. Such limitations restrict the ability of general users to engage in interior design even though they may have creative ideas.

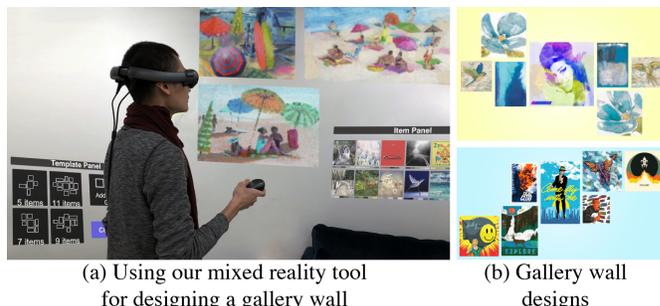


Fig. 1: (a) A user wearing a Magic Leap One headset designs a gallery wall using our tool. The figure shows what the user sees in mixed reality while designing: the control panels and his gallery wall design overlaid on the real wall. (b) Some gallery walls designed by users with our tool.

In our work, we attempt to address these challenges by devising an interior design tool leveraging the visualization and interaction capabilities of the latest consumer-grade mixed reality devices. Since interior design is a very broad area, as an early attempt to investigate such design applications based on mixed reality, we particularly focus on the design of gallery walls, which are common for decorating interior spaces such as living rooms, hotel lobbies, galleries, etc. Figure 1 shows a user designing a gallery wall using our tool.

A **gallery wall** refers to a cluster of wall **art items** artistically arranged on a wall. A gallery wall commonly contains a **focal item** near the center of the arrangement that sets the tone for the overall design. The other items are called **auxiliary items** and are placed around and with reference to the focal item. The auxiliary items are generally compatible with the focal item in terms of color and style. These definitions follow the conventions used by designers in creating a gallery wall design using a traditional workflow.

Our mixed reality design tool is suitable for novice users. Users are able to directly visualize how the gallery wall design will look on the real wall while interactively creating the

design in mixed reality. The visualization and user-interaction components of our tool keeps the user in the loop of the design process, allowing quick exploration of desirable designs through trial-and-error.

Furthermore, by comparing the color and style compatibility between different art items, our suggestion engine lets the user quickly browse through many desirable design suggestions generated by our tool in mixed reality in real time, hence avoiding the manual and mental efforts in browsing through a large database of wall art items.

We make the following major contributions in this work:

- Based on interviews with professional designers, we devise a computational approach for facilitating and automating the design of gallery walls, which enables a novel mixed reality interactive design tool.
- We demonstrate how mixed reality technology, which bridges the gap between real-world scene knowledge and design suggestions computed in a virtual setting, can be adopted for interior design. In our case, we particularly demonstrate how such an approach can be applied for designing gallery walls.
- We conduct experiments to evaluate the user experience and performance of using our novel tool for gallery wall design. We also conduct a perceptual study to evaluate the quality of the gallery wall designs created by users with our tool.

We believe these contributions will inspire future research in creating mixed reality interfaces for interior design.

II. RELATED WORK

To the best of our knowledge, there is no existing work on using mixed or augmented reality for gallery wall design. Regardless, we briefly review the existing research work and commercial tools relevant to our problem domain.

A. Extended Reality for Interior Design

Companies have been exploring the use of virtual, augmented, or mixed reality technologies for creating and visualizing interior design. Matterport uses a 3D camera to capture the color and depth of real-world living spaces, which can be visualized in 3D by users wearing a virtual reality headset. Such an approach finds promising applications for virtual real estate tours. On the other hand, roOomy provides virtual staging services, enabling users to see previews of interior designs via augmented reality devices showing virtual furniture objects overlaid on a real scene. Furniture retailers such as Wayfair also develop virtual and augmented reality experiences with capabilities such as customizing the design of outdoor spaces with furnishings and décor.

It is noteworthy that several companies provide web interfaces or mobile applications for designing gallery walls. For example, Shutterfly allows users to upload their photos and arrange them using preset layouts provided by the company. Art.com provides a mobile application that allows users to select individual art items or preconfigured gallery wall layouts

from their large wall art collection and visualize them via augmented reality on a mobile device.

On the research side, Zhang et al. [1] proposed an approach to add furniture items and relight scenes on a RGBD-enabled tablet. Yue et al. [2] developed a mixed reality system which allows users to efficiently edit a scene for applications such as room redecoration. Virtual content needs to be adapted to fit the current scene. Nuernberger et al. [3] devised a technique to align virtual content with the real world. Chae et al. [4] proposed a space manipulation technique for placing distant objects by dynamically squeezing surrounding spaces in augmented reality. Lindlbauer and Wilson [5] showed how to manipulate space and time in augmented reality.

Compared to the existing approaches, our tool not only uses extended reality technologies for visualizing gallery wall designs, but also reasons about the spatial and color compatibility. Our tool facilitates the design process by suggesting desirable combinations and placements of art items, taking the color of the wall into account.

B. Automated Layout Design

Recently, layout design automation has received much research attention. There are previous efforts on automatic 2D graphics layout design [6], [7], poster design [8], website design [9], magazine covers [10], photo collages [11], [12], etc. We focus on reviewing automatic scene layout design works.

Merrell et al. [13] proposed an interactive tool for furniture layout design based on interior design guidelines, while Yu et al. [14] devised an optimization framework for automatic furniture layout design. Fisher et al. [15] characterized structural relationships in scenes based on graph kernels and later proposed an example-based approach [16] for synthesizing 3D object arrangements for modeling partial scenes. More recently, Wang et al. [17] applied deep convolutional priors for indoor scene synthesis, while Weiss et al. [18] proposed a physics-based approach for fast and scalable furniture layout synthesis. The aforementioned works are mostly focused on the geometrical aspects of populating spaces with furniture. Complementary to this line of work, Chen et al. [19] created a tool called *Magic Decorator* for automatically assigning materials for objects in indoor scenes. Xu [20], [21] proposed a tool for layout beautification and arrangement. However, none of these approaches considers the stylistic arrangement of wall art items in a scene. We propose a novel approach for modeling common factors such as colors, spatial relationships, and semantic compatibility among a number of art items to generate desirable gallery walls, which could complement the existing automated interior design approaches.

We also note that recently CAD software companies such as Autodesk are applying generative design for automating layout synthesis [22]. Along with this line of work, we believe our generative design tool for semi-automating gallery wall design will find good practical uses, especially for enabling intuitive and interactive design in mixed reality.

C. Interactive Scene Modeling

Typically, users want the ability to control, modify, and visualize their designs during the design process so that they can infuse their personal stylistic preferences in their designs. As such, interactive modeling tools play an important role in the design process. There is a large body of work on interactive modeling tools. We review recently proposed scene modeling interfaces.

Along the direction of suggestive interfaces for scene modeling, Yu et al. [23] proposed a suggestive interface called *ClutterPalette* that uses object distribution statistics learned from real-world scenes for suggesting appropriate furniture items to add to a scene. Matthew et al. [24] devised a context-based search engine for 3D furniture models to add to a scene.

Another line of work focuses on providing users with easy controls for creating objects while modeling a scene. As humans are accustomed to drawing sketches, a promising approach is to devise sketch-based interfaces for modeling scenes. Xu et al. [25] proposed a sketch-based interface for retrieving and placing 3D models in scenes. Li et al. [26] devised an interface called *SweepCanvas* that allows users to perform 3D prototyping on an RGB-D image of a partial scene by sketching. Users can conveniently create virtual objects overlaid on top of the point cloud of a real scene.

D. Augmented Reality for Art Curation

Augmented reality also helps people curate and learn about art [27], [28]. Edmonds et al. [29] delivered interactive art experiences in augmented reality. Dieck et al. [30], [31] proposed the Generic Learning Outcomes (GLO) framework. When visiting an art gallery, users wear Google Glasses which pop up information to help them understand paintings. Coulton et al. [32], [33] created a mobile augmented reality application by which designers can curate art exhibitions at home.

Compared to the existing interfaces, we propose a novel mixed reality-based interface for designing gallery wall layouts *in situ*, with the user seeing the design overlaid on top of the target wall. Seeing how the generated design fits into the real world, the user can easily modify the design by a few intuitive operations. We demonstrate in our evaluation experiments that our mixed reality tool can allow not only designers but also novice users to quickly generate desirable and artistic gallery wall designs.

III. INTERVIEW WITH DESIGNERS ON WORKFLOW

To devise a computational approach and a practical tool for designing gallery walls, we interviewed 4 professional designers from a large furnishings and décor company to better understand the way professional designers create gallery walls under current practice. The detail of the interview is shown in the supplementary material.

We devise our computational approach based on the above observations to automate the conventional design process. Typically, in creating a gallery wall for a common scene like a living room, it takes about 20 minutes for the designer to

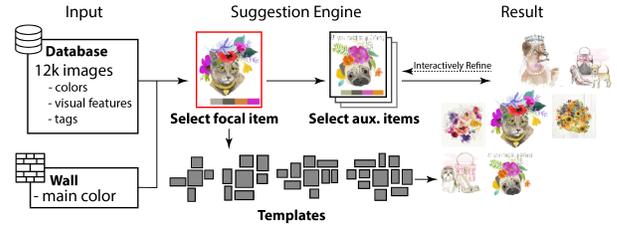


Fig. 2: System overview. Taking a database of art items and a wall as input, our tool suggests a focal art item as the starting point for designing a gallery wall. The user may apply a template to generate an initial design, and then interactively refines the design with the help of our suggestion engine.

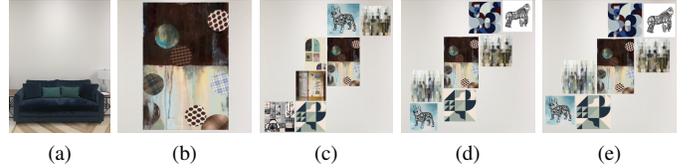


Fig. 3: Designing a gallery wall with our tool. (a) The input empty wall. (b) The user selects a focal item. (c) The user applies a template for initializing the design. (d) The user modifies the design interactively. (e) The finished gallery wall design.

select compatible wall art items and another 20 minutes to arrange the items into a good layout.

IV. SYSTEM OVERVIEW

Our tool is realized as an application that runs on the Magic Leap One mixed reality headset. Figure 1 shows a user designing a gallery wall using our tool. The user wears the headset when using our application to create a gallery wall.

The display on the mixed reality headset shows a user interface as well as the virtual gallery wall design overlaid on the real wall, such that the user can instantly preview how the gallery wall design will look on the real wall while designing. The user interacts with the user interface via a handheld controller (e.g., a Magic Leap One's Control) to perform operations such as selecting and dragging.

Figure 2 shows an overview of our tool. It consists of two major components: a suggestion engine for generating design suggestions and a user interface for modifying the gallery wall design. The tool is connected to a large database of wall art items (pictures) from which suitable art items are automatically retrieved and suggested to the user.

A. Workflow

Figure 3 illustrates the typical user workflow while using our tool for designing a gallery wall. The input is a wall color. The output, which can be reached with as few as three user decisions, is a gallery wall design that goes well with the wall color. Our tool achieves color compatibility by using the wall color to retrieve from the database a number of candidate focal



Fig. 4: Samples of wall art items in our database, which contains more than 12,000 art items.

art items whose colors are compatible with the wall color. A user selects a focal art item from these suggestions, picks a focal item size from among offered art sizes, and either chooses a gallery wall template to launch the synthesis of a gallery wall design or browses recommended auxiliary items. The user can interactively modify the gallery wall design during the design session. The design session ends when the user is satisfied with the gallery wall.

B. Art Items Data

Database. Our tool is built upon a suggestion engine which helps users browse and make selections from a large database of wall art items while designing their gallery walls. The database, which contains 12,000 wall art items created by professional artists, belongs to the company specialized in furniture and interior design business mentioned in Section III. Figure 4 depicts some example art items. The art items are mostly paintings and stylized photographs. Each art item comes with 1 to 5 realistic dimensions (from small to large) whose real replicate can be ordered and put on a real wall.

Annotations. To compute the compatibility between different art items for making suggestions, each item is annotated with: 5 *Dominant Colors*, 256 *Visual Features*, and *Tags*.

- 1) **5 Dominant Colors.** The 5 most dominant colors of the art item are extracted by the k-means clustering algorithm and stored.
- 2) **Visual Features.** 256 visual feature values are computed by a convolutional neural network, which encode the visual style of the art item. We trained a Siamese Network to perform such feature extraction using the aforementioned database containing many art items. The network takes image pairs as input, where a positive pair consists of images of items from the same category and a negative pair consists of images of items from different categories. Each input image is processed by a modified Inception-ResNet [34], whose last layer is changed to a fully connected layer, resulting in a 256-dimensional vector as the output. The network was trained to minimize the contrastive loss [35] of the vectors (embeddings) of the input image pairs.

- 3) **Tags.** Each item also carries tags manually specified by designers of the company. We show the tags in the supplementary material.

Similarity between two art items is computed based on the L2 distance between their annotation vectors: the smaller the distance, the more similar the two items are. We find these annotations very useful in devising our suggestion engine as they characterize the art items and are also common criteria used by designers for comparing art items. By formulating our scoring functions using these three types of annotations, we are able to devise an interface that allows the user to flexibly apply filters using a subset or all of the three types of annotation to retrieve relevant art item suggestions, making it easy for the user to browse through the large database of art items.

V. TECHNICAL APPROACH

We provide details for our design suggestion engine. There are two major components: art items suggestion and templates. By using these components, the user can quickly browse through the database of art items and select items that fit with the wall, as well as obtaining a decent spatial arrangement of the items as an initialization of their design.

Akin to the conventional workflow for designing a gallery wall, our approach starts with considering the wall color. The user wears the Magic Leap One headset and faces the target wall to be decorated. The wall plane is detected and extracted based on the headset’s built-in functionality. The user can manually specify the wall color via a color picker in the user interface, or by using the headset’s camera to take a picture of the wall whose average color is taken as the wall color.

A. Wall Plane and Color

Based on the wall color, a neighbor color within $\pm(60^\circ$ to $90^\circ)$ of the wall color is randomly selected from the HSV circular color space. The complementary color of the neighbor color (180° from the neighbor color in the HSV circular color space [36]) is also selected.

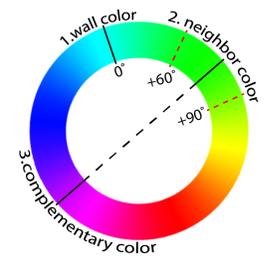


Figure 5 shows an example. The wall color is used as the basis for Fig. 5: Wall’s color retrieving other colors for suggest- palette example. ing relevant art items, akin to the conventional gallery wall design process (Section III).

B. Suggested Focal Items

Our goal in this step is to suggest a list of art items from the database as candidate focal art items for the user. Figure 6 shows some suggested focal art items based on different wall colors.

To achieve this, the wall color, neighbor color and complementary color then form a 3-color palette based on which compatible focal art items are selected. The wall compatibility score S_{foc} of a candidate focal art item ϕ is defined as:

$$S_{\text{foc}}(\phi) = 1 - \frac{1}{9} \sum_{\mathbf{c}_w \in C_w} \min\{d(\mathbf{c}_w, \mathbf{c}_\phi) \mid \mathbf{c}_\phi \in C_\phi\}, \quad (1)$$



Fig. 6: Examples of focal items retrieved by the suggestion engine based on different wall color schemes.

where C_w is a set containing the wall color, neighbor color and the complementary color in the HSV space; C_ϕ is a set containing the 5 dominant colors of the candidate focal art item ϕ in the HSV space. This scoring function evaluates how close the candidate focal art item ϕ 's color palette is with respect to the wall's color palette. The closer they are, the higher the wall compatibility score.

Note that $\mathbf{c}_w, \mathbf{c}_\phi \in \mathbb{R}^3$ are colors in the HSV space. $d(\cdot)$ is a distance metric function to project the two colors \mathbf{c}_w and \mathbf{c}_ϕ into the HSV cone and to compute the squared distance between them [37]:

$$d(\mathbf{c}_w, \mathbf{c}_\phi) = (\sin(H_w)S_wV_w - \sin(H_\phi)S_\phi V_\phi)^2 + (\cos(H_w)S_wV_w - \cos(H_\phi)S_\phi V_\phi)^2 + (V_w - V_\phi)^2,$$

where $H \in [0, 2\pi)$, $S \in [0, 1]$, and $V \in [0, 1]$ are the HSV channel values. The range of $d(\mathbf{c}_w, \mathbf{c}_\phi)$ is $[0, 3]$. Equation (1) sums up the differences of 3 pairs of colors, hence a normalization of 9 is used.

Our approach computes the compatibility scores for all the art items in the database. The top-20 art items are retrieved and displayed in order of the compatibility scores with the highest-scoring item shown first. The user is supposed to select a focal art item from the list of suggested art items. However, if needed, the user can also explore the database to select any other art item as the focal art item using the Item Panel which we describe in a later section.

C. Suggested Auxiliary Art Items

Akin to the conventional gallery wall design approach, the selected focal art item serves as a reference for the suggestion engine to suggest other compatible, auxiliary art items to add to the gallery wall design. To retrieve auxiliary art items from the database as suggestions, an overall compatibility score S_{aux} is computed for each candidate auxiliary art item, which evaluates the style and color compatibility between the auxiliary art item and the selected focal art item:

$$S_{\text{aux}}(\phi) = w_c S_{\text{aux}}^c(\phi) + w_s S_{\text{aux}}^s(\phi), \quad (2)$$

where w_c is the weight of the color compatibility score S_{aux}^c and w_s is the weight of the style compatibility score S_{aux}^s .

Color Compatibility Score: A candidate auxiliary art item ϕ has a high color compatibility score S_{aux}^c if its colors are

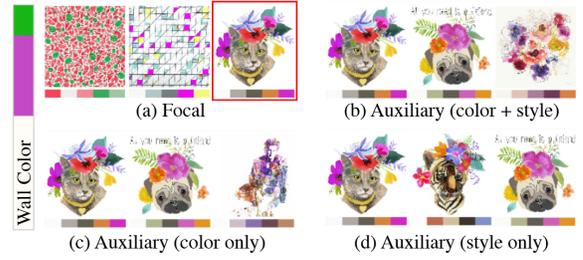


Fig. 7: Art item suggestions. Based on the wall's color palette shown on the left, (a) a number of focal art items compatible in color are suggested. Based on a selected focal art item (highlighted in red), (b) a number of auxiliary art items compatible in color and style are suggested. Auxiliary art items suggested by considering (c) only color or (d) only style.

close to the dominant colors of the selected focal art item. Specifically, the color compatibility score of auxiliary art item ϕ is defined as follows:

$$S_{\text{aux}}^c(\phi) = 1 - \frac{1}{15} \sum_{\mathbf{c}_f \in C_f} \min\{d(\mathbf{c}_f, \mathbf{c}_\phi) \mid \mathbf{c}_\phi \in C_\phi\}, \quad (3)$$

where C_f and C_ϕ are respectively sets containing the 5 dominant colors of the selected focal art item and of the auxiliary art item ϕ . $\mathbf{c}_f, \mathbf{c}_\phi \in \mathbb{R}^3$ are colors in the HSV space. As Equation (3) sums up the differences of 5 pairs of colors, and each difference has a range of $[0, 3]$, a normalization of 15 is used. This scoring function evaluates how close the auxiliary art item ϕ 's color palette is with respect to the selected focal art item's color palette. The closer they are, the higher the score is.

Style Compatibility Score: A candidate auxiliary art item ϕ has a high style compatibility score S_{aux}^s if its visual feature vector is close to the selected focal art item's visual feature vector. The style compatibility score of auxiliary art item ϕ is defined as follows:

$$S_{\text{aux}}^s(\phi) = 1 - \frac{1}{\sqrt{n}} \|\mathbf{v}_f - \mathbf{v}_\phi\|, \quad (4)$$

where $\mathbf{v}_f, \mathbf{v}_\phi \in \mathbb{R}^n$ are the n -dimensional visual feature vectors of the focal item and auxiliary art item ϕ computed by the convolutional neural networks. The size of the dimension in our database is 256.

Overall, the compatibility score is computed for each art item in the database. The user interface displays the top-20 art items sorted in descending order of their compatibility scores as auxiliary art item suggestions. To provide flexibility in retrieving suggestions, our user interface allows the user to turn on and off the consideration of the color or the style compatibility score, which correspond to setting w_c or w_s as 1 or 0. Figure 7 shows an illustration. The user can also select which of the 5 dominant colors C_f of the selected focal art item to consider in computing the color compatibility score.

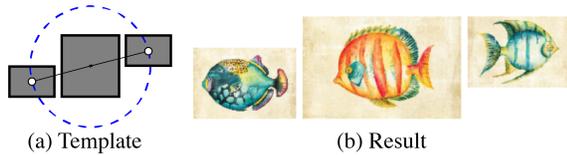


Fig. 8: Applying a template to generate a gallery wall design. (a) A template and its tree structure. (b) A gallery wall created by applying this template.

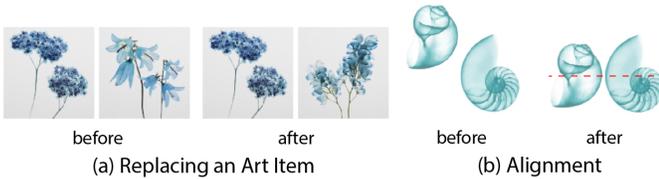


Fig. 9: The user can interactively modify a design in mixed reality using the functionalities of our user interface.

D. Templates

Templates encode spatial relationships of art items that are commonly applied by gallery wall designers. A template arranges groups of auxiliary art items symmetric about and around the focal art item, akin to the gallery wall designs created by the conventional design workflow.

Initializing a Gallery Wall Design. Figure 8 illustrates how to apply a template to generate a gallery wall design. According to the layout of the chosen template, starting from the root (the focal art item), our approach inserts auxiliary art items which are compatible with the focal art item, group by group. More specifically, the items are added on the circumference of a circle centered at the focal item with a random radius $r \in [0.5d, 3d]$, where d is the diagonal length of the focal item. A pair of two auxiliary items would be added at the opposite sides of the circle. A group of three items would be added at the vertexes of a randomly-oriented equilateral triangle circumscribed by the circle. The gallery wall design generation finishes as all groups of auxiliary art items have been placed. The generated design is taken as an initial design based on which the user can modify interactively.

Spatial Refinement. Our approach refines the spatial relationships between the items after the initialization and after every user interaction with the gallery wall design such as adding an item, removing an item, and moving an item.

Snapping: To keep the gallery wall design compact, by default, all auxiliary items steer towards the focal item at the center while they maintain a certain minimum space between each other to avoid overlapping.

Alignment: To keep the gallery wall design neat and uncluttered, by default, our approach aligns neighboring art items either horizontally or vertically by their edges so long as the alignment does not cause overlapping. (see Figure 9(a))

We include more examples of interactively modifying actions in the supplementary materials and video.

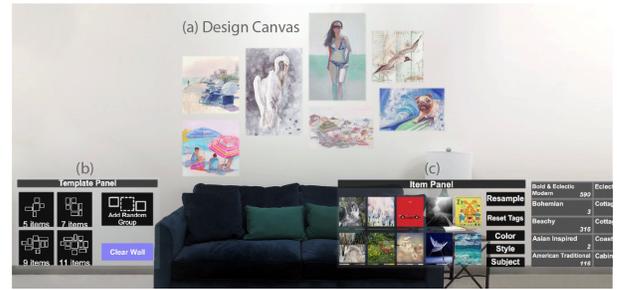


Fig. 10: User interface of our tool via which a user wearing a mixed reality headset visualizes and designs a gallery wall. It consists of three components: (a) Design Canvas; (b) Template Panel; and (c) Item Panel.

VI. USER INTERACTION

Figure 10 shows the user interface of our tool which is displayed in mixed reality. It consists of three components: a) *Design Canvas*, b) *Template Panel*, and c) *Item Panel*. Each of the components provides a number of functionalities that allow the user to refine a gallery wall design conveniently and desirably. We describe them in the following.

Figure 10 shows the user interface of our tool which is displayed in mixed reality. It consists of three components: a) the *Design Canvas* where the user can interactively modify the current gallery wall design visualized on the real wall; b) the *Template Panel* where the user can select and apply a preset template for synthesizing an initial gallery wall design; and c) the *Item Panel* where the user can retrieve art items from the database by specifying different criteria. Each of the components provides a number of functionalities that allow the user to refine a gallery wall design conveniently and desirably. We describe them in the following.

A. Design Canvas

The design canvas visualizes and overlays the current gallery wall design on the real wall via the mixed reality headset’s display. It also provides a number of functionalities for interactively adjusting the art items and the gallery wall layout: add, replace, move, resize, and remove an item. The detailed functionalities are listed in the supplementary material.

- **Add.** The user selects an art item in the current design and retrieves a number of art items from the database that are compatible in terms of dominant colors, visual features and tags, which he can add to the current design.
- **Replace.** The user replaces an art item with another compatible art item from the database. (see Figure 9(a))
- **Move.** The user moves an art item by dragging it.
- **Resize.** The user chooses another size for an art item.
- **Remove.** The user removes an art item.

B. Template Panel

The Template Panel allows the user to quickly generate an initial gallery wall design with items decently placed. It

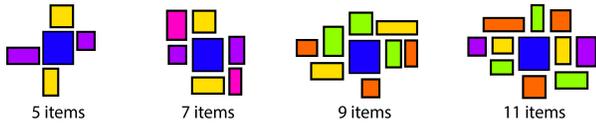


Fig. 11: The templates we use in our system. The blue item represents the focal item. The colored items represent the groups of 2 or 3 auxiliary items.

provides a number of preset templates that the user can apply to synthesize a gallery wall design based on a selected focal art item. It also provides other functionalities to enable automatic refinement of the spatial layout of the current design. A list of functionalities supported: apply a template, add random group, align all, and snap. The detailed functionalities are listed in the supplementary material.

- **Apply a Template.** Based on a placed focal art item, the user applies a template to synthesize a gallery wall design.
- **Add Random Group.** Our tool automatically adds a group of 2 or 3 auxiliary art items, which are compatible with the focal art item, to the current design. The group of auxiliary items are symmetric about the focal item.
- **Align All.** The user triggers our tool to align all the art items with respect to each other. The alignment is done along the horizontal (left, right, or center) or vertical (top, bottom, or center) direction. This functionality comes in handy because it could be tiring and difficult for users to make a lot of precise adjustment in the 3D space [38] using a handheld controller.
- **Snap.** If enabled, an art item is snapped to its neighbor item as the user drags the art item around, i.e., it will steer towards the center of its neighbor item until a minimum spacing between the two items is reached. (see Figure 12)
- **Clear Wall.** All the art items are removed from the design.

C. Item Panel

The Item Panel is connected to the suggestion engine and the database of art items. Its primary function is to display a relevant list of art items that the user can add to the gallery wall design as a focal art item or auxiliary art items. The panel shows a number of buttons that the user can click to set the criteria for retrieving relevant art items from the database. For example, the user can select whether to use color, or style, or both as criteria for determining compatibility between items. The user can also select which color(s) out of the 5 dominant colors of the focal art item to use for determining color compatibility. A list of functionalities supported: *Update Wall's Color Palette*, *Find Focal*, and *Find Auxiliary*.

- **Update Wall's Color Palette.** Our tool re-generates the neighbor and complementary colors of the wall's color palette.
- **Find Focal.** Based on the wall's color palette, our tool retrieves a number of compatible art items as focal art item suggestions (according to equation (1)).



Fig. 12: Snapping example. (a) The user drags an art item, which is (b) snapped towards the center of its neighbor item until a minimum spacing between the two is reached.

- **Find Auxiliary.** The user selects criteria (e.g., colors, visual features, tags) based on which our tool retrieves the compatible art items as auxiliary art item suggestions (our tool retrieves 20 art items in this case). By default, the art items are sorted by their compatibility scores.

VII. USER EVALUATION

We developed our tool using C# on the Unity Game Engine installed with the Magic Leap Lumin SDK. We deployed our tool onto a Magic Leap One headset which we used for our user evaluation experiments.

User Groups. To evaluate our approach, we conducted an IRB-approved user study with 41 participants. We recruited two different groups of users to evaluate our tool.

Group 1: The first group was recruited to evaluate the user experience of designing a gallery wall using our mixed reality interface based on Magic Leap One versus using a 2D interface which mimics a traditional design tool on a laptop. We recruited 17 participants, who are the employees of a company, consisting of 12 males and 5 females, aged from 20 to 45, the average age was 32. All participants did not have experience using the Magic Leap One headset. Each participant designed 2 gallery walls, under *Condition MR* and *Condition 2D* in a random order.

Group 2: The second group was recruited to evaluate the user experience of designing a gallery wall without and without the template functionality. We recruited 24 participants, who are the college students, consisting of 16 males and 8 females, aged from 19 to 24 the average age was 22. Each participant designed 2 gallery walls under *Condition 2D* and *Condition 2DNT* in a random order.

Conditions. We asked the participants to create gallery wall designs under different conditions. The goal of each task was to design a gallery wall that fits with a living room with a pale gray wall and a blue sofa as shown in Figure 10.

- **Condition MR:** The participant used our tool delivered through a mixed reality interface to create a gallery wall.
- **Condition 2D:** The participant used our tool delivered through a 2D interface to create a gallery wall.
- **Condition 2DNT:** The participant used our tool delivered through a 2D interface to create a gallery wall with *no template functionality*. That is, the “Apply a Template” and “Add Random Group” buttons under the Template Panel were disabled (see Figure 10(b)).

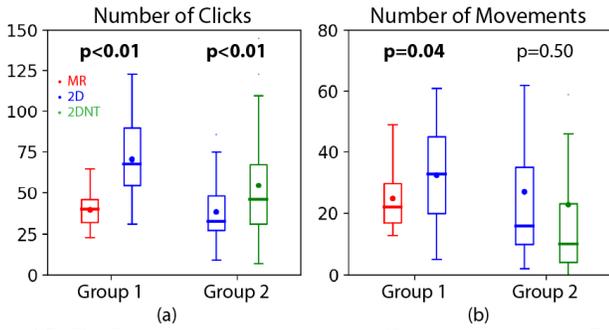


Fig. 13: Performance results under different conditions. Color dots and bars show the means and medians. The p-value of t-test computed between the results of the two conditions in each group is shown. The p-values smaller than 0.05 which reject the null hypothesis are bolded.

Note that, both *Condition MR* and *Condition 2DNT* used a background image showing a pale gray wall and a blue sofa. We include the 2D interface we used for the user evaluation in the supplementary material.

Procedure. Before each task, we briefed and trained the participant in creating a gallery wall design using our tool. We asked the participant to follow a 5-minute tutorial which guided him how to use our tool to create a gallery wall design step by step. Note that we showed the participant the tutorial whether he was tasked with creating a gallery wall using the mixed reality interface (*Condition MR*) or the 2D interface (*Condition 2D*). The participant could ask any question about the user interface to make sure he was familiar with using it.

The participant was then asked to create a gallery wall design that fits with the living room under a given condition. There was no time limit for finishing the designs. The participants stopped when they were satisfied with their designs. Most participants in Group 1 finished their designs within 20 minutes, while most participants in Group 2 finished their designs within 5 minutes.

For *Condition 2D* and *Condition 2DNT*, the participant was presented with a photo of the living room when designing a gallery wall using the 2D interface (please refer to the supplementary material for a screenshot of the interface). Our tool tracked the interaction metrics for later analysis.

Our supplementary material contains the results created by all participants.

VIII. EXPERIMENT RESULTS

We discuss the user evaluation results with regard to performance, usage and user feedback. We use t-tests to evaluate if there is any significant difference between the results obtained under different pairs of conditions by reporting the p-values. We show our results in box plots for easy interpretation. Our supplementary material shows the numeric results and all designs created by participants under different conditions.

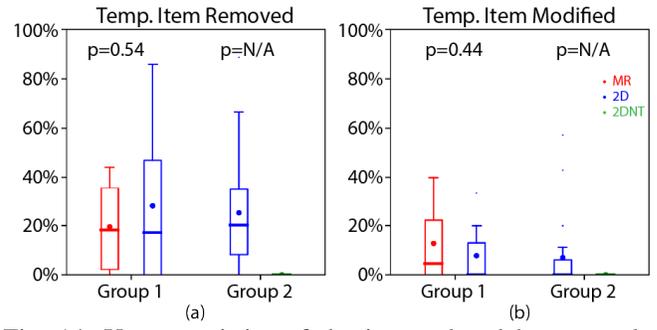


Fig. 14: Usage statistics of the items placed by a template under different conditions.

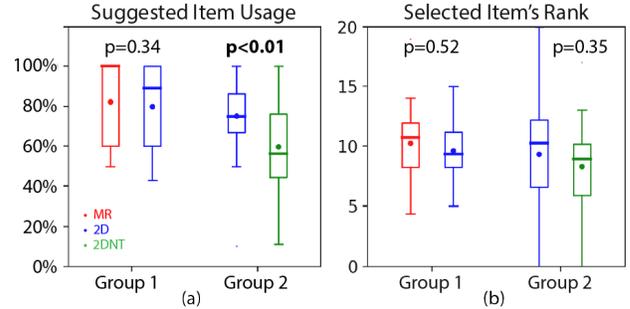


Fig. 15: Usage statistics of the items suggested by the Item Panel under different conditions.

A. Performance

We tracked the performance of retrieving items from the large database we used in our experiments on a desktop computer equipped with a i7-7700K 4.2GHz CPU, 16GB RAM, and an NVIDIA GeForce GTX 1080 8GB graphics card. Retrieving focal items took 3.12 seconds and retrieving auxiliary items took 1.21 seconds. All the other user interface operations ran at an interactive rate.

The performance metrics tracked by our tool are the number of clicks and the number of movements.

Mixed Reality vs. 2D Interface. As the Group 1 results in Figure 13 show, under the MR condition where the participants designed via a Magic Leap headset, they made fewer clicks ($p < 0.01$ in t-test) and movements ($p = 0.04$), compared to the 2D condition where they designed via a 2D interface.

Under the MR condition, the participants could see their design directly visualized on the real wall, which might result in fewer adjustments. In our perceptual study, we find that there is no significant difference between the visual quality of the designs created under the MR and the 2D conditions. In all, the direct visualization brought about by mixed reality allows the participants to create gallery wall designs of similar visual quality (compared to designs created on a 2D screen) with fewer manual adjustments.

Templates vs. No Templates. As the Group 2 results in Figure 13(a) show, under the 2D condition where the participants created designs with the aid of templates they made fewer clicks ($p < 0.01$) compared to the 2DNT condition where they

created designs without the aid of templates. The template functionality helped them create designs more efficiently.

B. Usage

Our tool also tracks the usage of the design suggestions generated by the Template Panel or Item Panel.

Template Items. Figure 14 shows the usage statistics of the items placed by a template under different conditions. As Figure 14(a) shows, the participants removed about 20% and 28% of the items placed by the templates under the MR and 2D conditions respectively. As a template places highly similar items in a design, it seems that the participants tended to replace a few items with other less similar items to introduce some variation or contrast into the overall design. On the other hand, as Figure 14(b) shows, in average only about 13% and 8% of the items were modified under the MR and 2D conditions respectively. Overall, the participants kept a majority of the items placed by a template in their final gallery wall design, suggesting the desirability and compatibility of the items in the template.

Suggested Items. Figure 15 shows the usage statistics of the items suggested by the Item Panel under different conditions. As Figure 15(a) shows, under the MR and 2D conditions, about 82% and 76% of the items used in the final design were chosen from the top-20 suggestions in the Item Panel retrieved according to their specified criteria. The Group 2 results show that there was a significant difference ($p < 0.01$) between suggested item usages under the 2D and 2DNT conditions. Under the 2DNT condition when the participants could not use a template to initialize a gallery wall design, they tended to use items from the database more randomly.

On the other hand, as Figure 15(b) shows, the average rank of the selected items ranges from about 8 to 10 under different conditions. It seems that, in using the Item Panel, the participants tended to choose items that match with their specified criteria but not in a very strict sense.

C. User Feedback

We talked to the participants after the experiments. They generally merited the visualization brought about by our mixed reality tool for showing the design on the real wall, which made the interactive design process more intuitive and interesting, compared to designing on a 2D screen, as they could directly see how their design could fit with the real space. On the downside, some participants reported initial user experience challenges while acclimating themselves to the device's field-of-view and headset fit. We include the participants' comments in our supplementary material.

IX. SUMMARY

We proposed a novel mixed reality-based interactive design tool for gallery walls. By overlaying a virtual gallery wall on a real wall, our tool allows users to directly visualize their design in the real world as an integrated part of the creative process. Our suggestive design interfaces facilitate users in retrieving compatible art items for creating their desired gallery walls.

We used a Magic Leap One headset for experimenting with our tool. The hardware still has limitations for consumer use. For example, the field-of-view is still a bit narrow for the user to see the overall design and the user interface at once; the user has to look around to see different things, which could be inconvenient.

It could be tiring to use the handheld controller for 3D interaction and manipulation in the 3D space for a prolonged period of time. It would be challenging to use a sophisticated design tool that requires tedious and precise user input in mixed reality.

Due to the scope of this project, in this work we only focus on a subset of interior design, namely, gallery wall designs. We show that a mixed reality approach for gallery wall design is feasible. Our approach could be extended to consider and visualize 3D decorations in mixed reality too (Please see the supplementary), though they are less common and we did not consider them in our approach.

It would be helpful to have a large database of gallery wall designs from which our tool could learn the spatial relationships and compatibility between different art items in a gallery wall design, and use such knowledge for synthesizing new designs.

We believe that the visualization of the design in the real space by mixed reality may benefit the users in envisioning and communicating their designs, allowing them to design in the real space directly and intuitively. In the future, it would be interesting to extend our mixed reality design tool to consider more sophisticated context of the 3D scene to facilitate other interior design tasks, such as suggesting furniture placement. Performing advanced scene analysis in real time for enabling interior design applications in mixed reality presents a practical challenge, which could be resolved as the computational power of mixed reality devices continue to increase.

A. Future Work

With the advances of artificial intelligence and natural language processing techniques, we hope that such techniques can be adopted in mixed reality for enabling natural and convenient user interaction experiences in the interior design process. For example, it would be very helpful if a mixed reality interior design software application could understand the scene objects' semantics and the voice commands from the user to decorate a space, minimizing the need for manual user input. With such advances, using a mixed reality software application for interior design will be as natural as talking to an outstanding interior designer who can instantly visualize the design for the user through mixed reality. We believe this is an exciting goal for enabling human-AI collaboration in design.

For commercial purposes, art items could be further annotated with non-aesthetic tags, such as their prices, frequencies of being viewed or selected, and the semantic meanings they carry. Incorporating such additional tags could provide more desirable item recommendations while a user is creating their gallery wall designs.

ACKNOWLEDGMENT

We would like to thank our user study participants. This research project was supported by Wayfair, as well as an NSF CRII award (1565978) and an NSF CAREER award (1942531).

REFERENCES

- [1] E. Zhang, M. F. Cohen, and B. Curless, "Emptying, refurbishing, and relighting indoor spaces," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 6, p. 174, 2016.
- [2] Y.-T. Yue, Y.-L. Yang, G. Ren, and W. Wang, "Scenctrl: Mixed reality enhancement via efficient scene editing," in *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. ACM, 2017, pp. 427–436.
- [3] B. Nuernberger, E. Ofek, H. Benko, and A. D. Wilson, "Snaptoreality: Aligning augmented reality to the real world," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, 2016, pp. 1233–1244.
- [4] H. J. Chae, J.-i. Hwang, and J. Seo, "Wall-based space manipulation technique for efficient placement of distant objects in augmented reality," in *The 31st Annual ACM Symposium on User Interface Software and Technology*. ACM, 2018, pp. 45–52.
- [5] D. Lindlbauer and A. D. Wilson, "Remixed reality: Manipulating space and time in augmented reality," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, 2018, p. 129.
- [6] P. O'Donovan, A. Agarwala, and A. Hertzmann, "Designscape: Design with interactive layout suggestions," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2015.
- [7] —, "Learning layouts for single-pagegraphic designs," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 8, pp. 1200–1213, 2014.
- [8] Y. Qiang, Y. Fu, Y. Guo, Z.-H. Zhou, and L. Sigal, "Learning to generate posters of scientific papers," in *AAAI*, 2016, pp. 51–57.
- [9] X. Pang, Y. Cao, R. W. Lau, and A. B. Chan, "Directing user attention via visual flow on web designs," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 6, p. 240, 2016.
- [10] A. Jahanian, J. Liu, Q. Lin, D. Tretter, E. O'Brien-Strain, S. C. Lee, N. Lyons, and J. Allebach, "Recommendation system for automatic design of magazine covers," in *Proceedings of the International Conference on Intelligent User Interfaces*. ACM, 2013, pp. 95–106.
- [11] C. Rother, L. Bordeaux, Y. Hamadi, and A. Blake, "Autocollage," in *ACM Transactions on Graphics (TOG)*, vol. 25, no. 3. ACM, 2006, pp. 847–852.
- [12] J. Xiao, X. Zhang, P. Cheattle, Y. Gao, and C. B. Atkins, "Mixed-initiative photo collage authoring," in *Proceedings of the 16th ACM International Conference on Multimedia*. ACM, 2008, pp. 509–518.
- [13] P. Merrell, E. Schkufza, Z. Li, M. Agrawala, and V. Koltun, "Interactive furniture layout using interior design guidelines," in *ACM Transactions on Graphics (TOG)*, vol. 30, no. 4. ACM, 2011, p. 87.
- [14] L.-F. Yu, S. K. Yeung, C.-K. Tang, D. Terzopoulos, T. F. Chan, and S. Osher, "Make it home: automatic optimization of furniture arrangement," *ACM Transactions on Graphics (TOG)*, vol. 30, no. 4, p. 86, 2011.
- [15] M. Fisher, M. Savva, and P. Hanrahan, "Characterizing structural relationships in scenes using graph kernels," in *ACM Transactions on Graphics (TOG)*, vol. 30, no. 4. ACM, 2011, p. 34.
- [16] M. Fisher, D. Ritchie, M. Savva, T. Funkhouser, and P. Hanrahan, "Example-based synthesis of 3d object arrangements," *ACM Transactions on Graphics (TOG)*, vol. 31, no. 6, pp. 1–11, 2012.
- [17] K. Wang, M. Savva, A. X. Chang, and D. Ritchie, "Deep convolutional priors for indoor scene synthesis," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, p. 70, 2018.
- [18] T. Weiss, A. Litteneker, N. Duncan, M. Nakada, C. Jiang, L.-F. Yu, and D. Terzopoulos, "Fast and scalable position-based layout synthesis," *IEEE Transactions on Visualization and Computer Graphics*, 2018.
- [19] K. Chen, K. Xu, Y. Yu, T.-Y. Wang, and S.-M. Hu, "Magic decorator: automatic material suggestion for indoor digital scenes," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 6, p. 232, 2015.
- [20] P. Xu, H. Fu, T. Igarashi, and C.-L. Tai, "Global beautification of layouts with interactive ambiguity resolution," in *Proceedings of the 27th Annual ACM Symposium on User interface software and technology*. ACM, 2014, pp. 243–252.
- [21] P. Xu, H. Fu, C.-L. Tai, and T. Igarashi, "Gaca: Group-aware command-based arrangement of graphic elements," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2015, pp. 2787–2795.
- [22] D. Nagy, D. Lau, J. Locke, J. Stoddart, L. Villaggi, R. Wang, D. Zhao, and D. Benjamin, "Project discover: An application of generative design for architectural space planning," in *SimAUD 2017 Conference proceedings: Symposium on Simulation for Architecture and Urban Design*, 2017.
- [23] L.-F. Yu, S.-K. Yeung, and D. Terzopoulos, "The clutterpalette: An interactive tool for detailing indoor scenes," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 2, pp. 1138–1148, 2016.
- [24] M. Fisher and P. Hanrahan, "Context-based search for 3d models," in *ACM Transactions on Graphics (TOG)*, vol. 29, no. 6. ACM, 2010, p. 182.
- [25] K. Xu, K. Chen, H. Fu, W.-L. Sun, and S.-M. Hu, "Sketch2scene: sketch-based co-retrieval and co-placement of 3d models," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 4, p. 123, 2013.
- [26] Y. Li, X. Luo, Y. Zheng, P. Xu, and H. Fu, "Sweepcanvas: Sketch-based 3d prototyping on an rgb-d image," in *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. ACM, 2017, pp. 387–399.
- [27] P. Coulton, E. Murphy, K. Č. Pucihar, R. Smith, and M. Lochrie, "User curated augmented reality art exhibitions," in *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*, 2014, pp. 907–910.
- [28] M. Kljun, K. Č. Pucihar, and P. Coulton, "User engagement continuum: art engagement and exploration with augmented reality," in *Augmented reality art*. Springer, 2018, pp. 329–342.
- [29] E. Edmonds, D. Hills, Y. Ji, and X. Tong, "H space: Interactive augmented reality art," in *Proceedings of the Fourteenth International Conference on Tangible, Embedded, and Embodied Interaction*, ser. TEI '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 683–688. [Online]. Available: <https://doi.org/10.1145/3374920.3375291>
- [30] M. C. Tom Dieck, T. Jung, and D. tom Dieck, *Google Glass Augmented Reality: Generic Learning Outcomes for Art Galleries*, 02 2015, pp. 463–476.
- [31] —, "Enhancing art gallery visitors' learning experience using wearable augmented reality: Generic learning outcomes perspective," *Current Issues in Tourism*, vol. 21, pp. 2014–2034, 05 2018.
- [32] P. Coulton, E. Murphy, K. v. Pucihar, R. Smith, and M. Lochrie, "User curated augmented reality art exhibitions," in *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*, ser. NordiCHI '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 907–910. [Online]. Available: <https://doi.org/10.1145/2639189.2675291>
- [33] P. Coulton, R. Smith, E. Murphy, K. v. Pucihar, and M. Lochrie, "Designing mobile augmented reality art applications: Addressing the views of the galleries and the artists," in *Proceedings of the 18th International Academic MindTrek Conference: Media Business, Management, Content & Services*, ser. AcademicMindTrek '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 177–182. [Online]. Available: <https://doi.org/10.1145/2676467.2676490>
- [34] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [35] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2. IEEE, 2006, pp. 1735–1742.
- [36] J. W. v. Goethe and C. L. Eastlake, *Theory of colours*. Dover, 2006.
- [37] D. Androustos, K. Plataniotiss, and A. N. Venetsanopoulos, "Distance measures for color image retrieval," in *Proceedings of the International Conference on Image Processing (ICIP)*, vol. 2. IEEE, 1998, pp. 770–774.
- [38] I. S. Mackenzie, "Fitts' law as a performance model in human-computer interaction," Ph.D. dissertation, University of Toronto, 1992.