Juxtaform: interactive visual summarization for exploratory shape design

KARRAN PANDEY, University of Toronto, Canada
FANNY CHEVALIER, University of Toronto, Canada
KARAN SINGH, University of Toronto, Canada

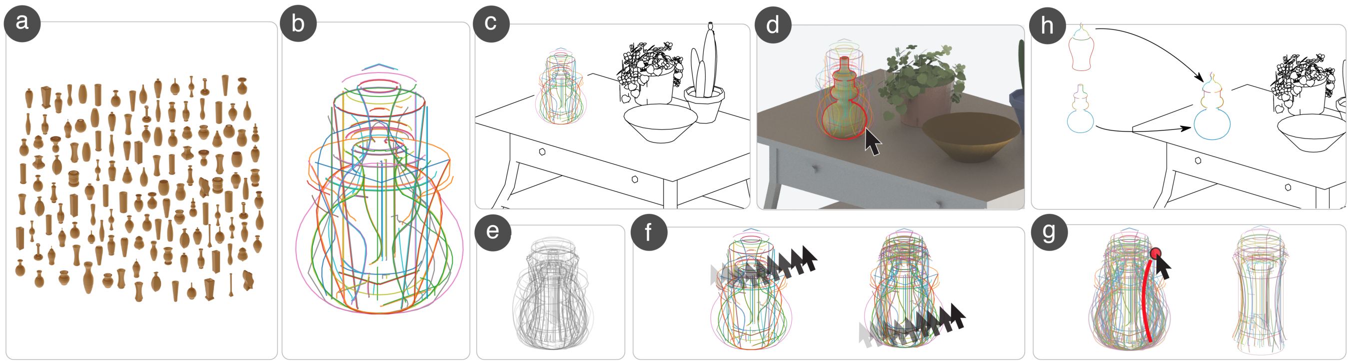

Fig. 1. *Juxtaform* addresses the visual exploration of large shape collections (a) for creative design. Stroke-based shape abstractions (b) are presented juxtaposed in-situ (c,d) to optimally visualize shape corpora with minimal visual clutter. A novel stroke filtering algorithm provides an automatic sketch-like rendering designed to highlight both the most common structures and diverse shape parts in a shape corpus (e). Users can browse, suppress or select currently displayed shapes (d), or reveal unseen parts of the corpus either spatially (f) or by sketching (g), providing a compelling integrated workflow for early sketch ideation (h).

We present *juxtaform*, a novel approach to the interactive summarization of large shape collections for conceptual shape design. We conduct a formative study to ascertain design goals for creative shape exploration tools. Motivated by a mathematical formulation of these design goals, juxtaform integrates the exploration, analysis, selection, and refinement of large shape collections to support an interactive divergence-convergence shape design workflow. We exploit sparse, segmented sketch-stroke visual abstractions of shape and a novel visual summarization algorithm to balance the needs of shape understanding, *in-situ* shape juxtaposition, and visual clutter. Our evaluation is three-fold: we show that existing shape and stroke clustering algorithms do not address our design goals compared to our proposed shape corpus summarization algorithm; we compare juxtaform against a structured image gallery interface for various shape design and analysis tasks; and we present multiple compelling 2D/3D applications using juxtaform.

CCS Concepts: • **Computing methodologies → Graphics systems and interfaces**; **Shape modeling**; **Non-photorealistic rendering**.

Additional Key Words and Phrases: shape exploration, shape design, sketch-based interaction

Authors' addresses: Karran Pandey, University of Toronto, Toronto, Canada, karran@cs.toronto.edu; Fanny Chevalier, University of Toronto, Toronto, Canada, fanny@cs.toronto.edu; Karan Singh, University of Toronto, Toronto, Canada, karan@dgp.toronto.edu.

0730-0301/2023/8-ART $15.00
https://doi.org/10.1145/3592436

ACM Reference Format:
Karran Pandey, Fanny Chevalier, and Karan Singh. 2023. Juxtaform: interactive visual summarization for exploratory shape design. *ACM Trans. Graph.* 42, 4 (August 2023), 14 pages. https://doi.org/10.1145/3592436

1 INTRODUCTION

"The designer does not, as a rule, begin with some preconceived idea. Rather, the idea is the result of careful study and observation, and the design a product of that idea." —Paul Rand

Sketching as an interactive exercise [Schon and Wiggins 1992] and sketch strokes as a visual abstraction [Prats 2007] are both critical in early stage design towards the development of ideas in *divergence-convergence* cycles [Liu et al. 2003] of design iteration [Arias-Rosales 2022]. We present a principled inquiry of such shape design and address the resulting design goals with a novel interactive exploratory shape summarization system called *juxtaform*.

Shape collections and existing designs fuel a rich mix of inspiration, templates, and constraints into this creative design process [Holinaty et al. 2021]. Inspirational shapes often need to be juxtaposed (used synonymously with superimposed) for comparison relative to each other (Figure 2), as well as evaluated *in-situ* within their design context [Arias-Rosales 2022]. Such shape corpora are abundant online and are increasingly produced by generative AI models [Cohen-Or and Zhang 2016; Gao et al. 2022; Rombach et al. 2021], providing both an enormous wealth of design data and a challenge to effectively exploit it. The ability to interactively interleave exploration (divergence) with filtering and refinement (convergence) of such shape corpora presented *in-situ* as juxtaposed, sketchy ensemble summaries thus has the potential to disruptively streamline early stage design [Arias-Rosales 2022]. Yet, design literature does not provide a comprehensive set of design goals that

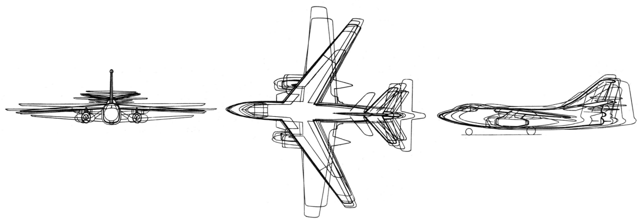

Fig. 2. Examples of superimposition in an airplane design [Nichols 1954]. Designers often use superimposition to explore variations around a common form or visualize similarities and differences between design variations.

can be mathematically optimized as a sketch summary for shape ensemble exploration.

We thus conducted formative interviews with 5 design professionals to formally understand the role of shape collections and generative AI in design applications. Findings consistent with prior work [Arias-Rosales 2022] underline the importance of *rapid* and *ideative* exploration of shape *diversity* in spatial *context* to understand and harness the *creative* potential of a shape collection (§3).

Understanding shape *diversity* and *context* necessitates reasoning of parts of shapes in relation to each other, addressed by a large body of research on shape registration and correspondence [Castellani and Bartoli 2020]. Even given a segmentation and part correspondence across shapes, rapidly conveying a summary view of a large collection in context entails finding an effective visual representation of shape, and then, an optimal balance between overall spatial clutter, showing complete representative shapes, as well as unique shape segments that capture the shape diversity of the collection.

A sketchy stroke-based representation of shape is well-suited to the above design goals: it provides a sparse visual abstraction of partial or complete shapes with minimal clutter suitable for juxtaposition; it aligns with ideative sketching [Arora et al. 2017], where shapes can be imagined as mental compositions of disparate and incomplete strokes; stroke attributes such as color or transparency can convey auxiliary information relevant to their shapes; and the strokes serve as selection and manipulation handles for interactive shape exploration and modeling [Tsang et al. 2004].

While there is a rich body of inter-disciplinary research on sketch-based design (§2), the sketch-based exploration and refinement of juxtaposed shape collections is relatively uncharted. We formulate the design goals distilled from design literature and our formative study (§3) as mathematical objectives to develop a novel algorithm to visually summarize a shape collection (§5). *Juxtaform* (§4) is a system that exploits this visual summary to interactively visualize, explore, filter, and refine large shape corpora using in-situ juxtaposed sketch strokes (Figures 1, 5).

Overview: A review of prior art (§2) is followed by a summary of findings of our formative study on exploratory shape design (§3), with study details in the supplemental. The study provides us with design objectives and a concrete problem statement, several motivating applications, and a framework to position our system *juxtaform* relative to prior art. We then present our interface *juxtaform*, with details on user interaction (§4) and our algorithm for visual summarization of a shape corpus (§5). We also show that compared to our approach, existing stroke and shape aggregation techniques designed for different design objectives are unsurprisingly ill-adapted to visually summarizing a large juxtaposed shape corpus (§6). §7 presents the summary of a formal user study comparing *juxtaform* to a baseline structured image gallery interface (see supplemental). We also show several creative 2D and 3D applications suggested by our formative study and illustrated using *juxtaform* (§8), and conclude with a discussion of limitations and future work (§9).

Our **contributions** are three-fold. We:
1. provide a formative analysis of sketch-based shape corpus exploration whose design goals inform *juxtaform* and future systems.
2. formulate a novel stroke-filtering (summarization) algorithm to balance juxtaposed shape exploration goals and visual clutter.
3. present *juxtaform*, a compelling system for exploratory shape design, with a formal evaluation, and many creative applications.

2 RELATED WORK

Interactive sketch-based exploration of juxtaposed shape collections for early stage design touches areas of shape ideation and exploration [Biasotti et al. 2016a], geometric aggregation [Zhu et al. 2014], and sketch-based shape interaction [Olsen et al. 2009]. We provide an overview of work in these areas before further positioning them relative to design objectives from our formative study in §3.

2.1 Early-Stage Shape Ideation

Exploratory ideation is an essential and challenging part of the creative process. Prior work has focused on both understanding 'ideation' and developing creativity support tools (CSTs) [Shneiderman 1999] to assist people with creative tasks. Despite a variety of multi-modal CSTs, sketching rough shapes around a basic form remains a preferred method for shape ideation [Prats 2007; Schon and Wiggins 1992; Smithers 2001]. In particular, designers typically sketch a basic idea of their vision [Hua 2019; Prats and Earl 2006] and go through cycles of divergence, by exploring variations around this vision, and convergence, by selecting and refining desirable variations, to conceptualize their eventual design [Liu et al. 2003].

Juxtaform provides a novel combination of three main characteristics — the familiar control and environment of sketch-based ideation, exploration and filtering features explicitly aligned with the goal of divergence-convergence cycles around a specific concept, and the computational capability to support the rapid exploration of shape diversity in large shape corpora.

2.2 Shape Exploration

A large body of literature deals with the rich space of shape exploration, and the challenging problem of browsing large collections of graphical objects (Figure 3). From the difficulty of searching [Biasotti et al. 2019], retrieving [Gao et al. 2014; Liu et al. 2013; Schulz et al. 2017; Tangelder and Veltkamp 2008], and exploring [Averkiou et al. 2014; Kleiman et al. 2013] desired shapes to that of imagining new shape ideas from existing collections [Averkiou et al. 2014; Jain et al. 2012; Lee et al. 2011; Xu et al. 2012], shape exploration systems have addressed a diverse set of challenges and user goals pertaining to the exploration process. The majority of research in this area — which is focused on images [Zhu et al. 2014], 3D models,

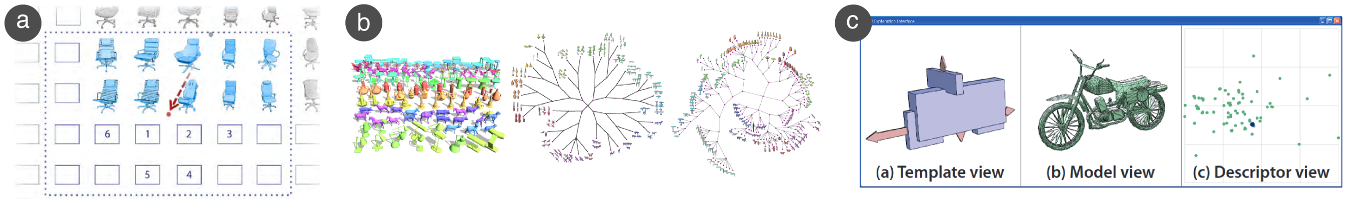

Fig. 3. Three shape exploration systems which illustrate different design choices with respect to visual representation, spatial arrangement, and interaction. (a) A gallery layout with scroll-based interactions [Kleiman et al. 2013], (b) A hierarchical layout with click-based interactions [Huang et al. 2013], and (c) A hybrid layout with edit-based interactions on a structural shape proxy [Ovsjanikov et al. 2011].

and sometimes both [Hueting et al. 2015] — can be characterized by the different application scenarios and exploration goals of the designed systems, as well as the design choices made when building the system.

Our approach is focused on the juxtaposed visual summarization, interactive exploration, and selection of shapes and parts of shapes from large shape collections which are spatially aligned or for which reasonable alignments can be computed [Kim et al. 2012]. Much research is, however, still complementary to ours, and given the multiple ways in which prior art relates to *juxtaform*, we have further positioned this research relative to ours in the context of our problem scope and design goals in the supplementary material.

2.3 Geometric Aggregation

Various algorithms for aggregating collections of graphical objects have been proposed for images [Agarwala et al. 2004; Zhu et al. 2014], vectorized sketches [Liu et al. 2018; Van Mossel et al. 2021], 3D models, and scenes [Fish et al. 2014; Huang et al. 2019; Xu et al. 2014b]. Most relevant to us are techniques that focus on aggregating geometrically similar shapes and sketch strokes.

A notion of similarity between shapes [Van Kaick et al. 2011; Veltkamp 2001] has been used to present organized hierarchical layouts [Huang et al. 2013], discover common structures [Pauly et al. 2008], or support intuitive feature and text-based queries [Biasotti et al. 2016b; Gao et al. 2014]. Stroke aggregation approaches [Grabli et al. 2004; Liu et al. 2018; Van Mossel et al. 2021; Yan et al. 2020] typically assume a single shape and focus on constructing geometrically accurate aggregates for sketch-based applications [Gryaditskaya et al. 2020; Hähnlein et al. 2022; Xu et al. 2014a].

In contrast, our focus is not on averaged strokes or clustered shapes, but on the visual composition of representative strokes from a collection to maximize understanding of the shape collection with minimal clutter.

2.4 Sketch-Based Shape Interaction

The bulk of research in sketch-based modeling pertains to the use of sketch-strokes for the interactive creation of shapes [Olsen et al. 2009]. Our research, focused on the stroke-based exploration of large collections of existing shapes, complements these approaches and easily integrates to provide initial shapes or shape suggestions [Tsang et al. 2004] for further sketch-based refinement. Notably relevant are also techniques like ours, which recognize the value of *in-situ* visualization in design (albeit without juxtaposing multiple shapes) [Arora et al. 2018; Lee et al. 2011; Paczkowski et al. 2011; Ye et al. 2021] and creative ideation [Arora et al. 2017; Hennessey et al. 2016; Holinaty et al. 2021; Orbay et al. 2012]. In the context of designing effective visualization and interaction techniques to maximize the understanding of a visual composition of shapes, our work also takes inspiration from research on composite visualization systems [Javed and Elmqvist 2012], interaction techniques to navigate multiple stacked views [Cockburn et al. 2009; Lam and Munzner 2010], and the use of sketch-based queries for juxtaposed objects [Hassoumi et al. 2019; Matejka et al. 2018].

3 FORMATIVE STUDY: CREATIVE EXPLORATION OF SHAPE COLLECTIONS

We summarize here the findings of a formative study aimed at concretizing design objectives for a system to support the exploration of large shape collections in early stage shape design.

To consolidate our understanding of shape collections exploration in practice, we interviewed 5 participants (2 industrial and 1 product designers, a professional 3D modeler, and an independent artist). The semi-structured interviews included a discussion of creative application(s), workflow(s), example project(s), and the role and challenges of shape exploration within it; and a critique of examples of recent research in shape exploration (ShapeSynth [Averkiou et al. 2014] and SketchSoup [Arora et al. 2017]). We outlined four specific knowledge goals during our interviews: the reasons and goals for exploring a shape collection; criteria for selecting shapes or parts of shapes from the collection; the intended application of the selections; and perceived benefits/challenges with creative exploration of shape. We now report on design insights and objectives for an interactive shape exploration system and refer the reader to the supplementary for further study details.

3.1 Design Insights

We found three typical motivators for shape exploration: *inspiration*, *reference*, and *embellishment*. *Inspirational* search entails a fast but extensive exploration of a shape corpus to quickly gather a wide range of diverse shapes to seed the ideation of a new artefact. Exploration for a *reference* shape, or for shape *embellishment*, involves sifting through one or more shape corpora to find (parts of) shapes that meet a more targeted set of design criteria.

The interviewed creators described a fairly uniform shape exploration workflow they use in practice: they would first narrow their exploration to a tractable set of hundreds of shapes via shape

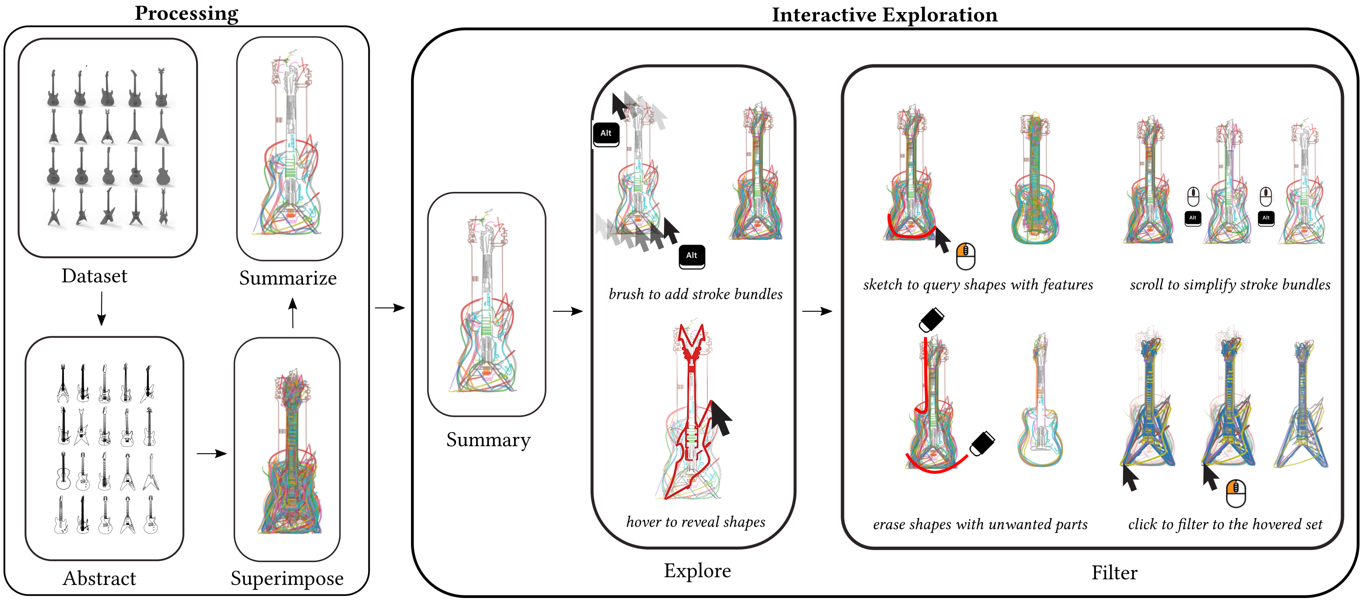

Fig. 4. *Juxtaform* pipeline: shapes in the corpus are first abstracted as a sparse set of sketchy juxtaposed (superposed) strokes. A subset of these strokes acts as both a visual summary of the shape corpus and interaction handles for shape and part exploration, filtering, and refinement.

category queries, resulting in shapes with strong structural and spatial commonalities. These are viewed as a catalogue of thumbnails, and a few (typically 5 or less) shapes or their parts are chosen for further evaluation or processing in-situ within the creator's larger shape design task. Often, multiple iterations of explore, select, refine, and evaluate in-situ are required to accomplish the design task; but comparing disparately presented thumbnails for similarity or diversity was described as both difficult and tedious. Participants were excited about the engaging creative potential of tools like ShapeSynth [Averkiou et al. 2014] (shape extrapolation) and SketchSoup [Arora et al. 2017] (playful and aesthetic visual design).

In summary, we gained **five key insights** from these interviews: early stage design typically entails *imprecise shape search* (**I1**)as shapes are rarely used as-is, or in contexts where precise shape details are not critical; shape exploration is often tied to task deadlines, or perceived as a quick initial step, requiring *high throughput* (**I2**), making effective visual summaries of shape collectives important; creators seek a diverse set of shapes, common and unique, with *rich variation* (**I3**) both for inspiration and for shape or part refinement; shape *understanding in context* (**I4**) is crucial to evaluate juxtaposed shapes relative to each other and in-situ of their design context; and shape exploration, like other aspects of a creative tasks, should be *fun and engaging* (**I5**).

3.2 Design Goals

The findings of our formative study are in strong agreement with general research in creativity support tools [Shneiderman 1999] and the role of sketching, juxtaposition, and shape collections in early stage design [Arias-Rosales 2022]. We distill these collective insights into 4 design goals for creative exploration of shape corpora.

- *Rapid Exploration* (**D1**): Quickly convey the overall essence of shapes in large collections using sketch abstraction (**I1**, **I2**), with fluid interactive tools to select (parts of) shapes and browse variations (**I3**, **I5**).
- *Diversity Exploration* (**D2**): Provide an interactive understanding of shape/part diversity in shape corpora (**I3**).
- *Contextual Exploration* (**D3**): Allow juxtaposed, in-situ presentation of shapes and their parts to enable general shape comparison (**I1**), greater visual throughput (**I2**), and shape understanding in a design context (**I4**).
- *Ideative Exploration* (**D4**): Interactively present shapes and their parts in a manner that aids imagining and creating novel shapes beyond those in the given collection.

3.3 Synthesis

While prior work has explored such design goals in isolation, *juxtaform* is a unique contribution towards the combination of these goals. Inspired by sketchy (NPR) and other perceptual shape abstractions [Lin et al. 2018; Todd 2004], as well as spatial arrangements that juxtapose common structure [Matejka et al. 2018] to aid understanding and comparison in shape collections [Zhu et al. 2014], *juxtaform* employs a juxtaposed sketch-based visual abstraction to define an interactive stroke-based visual summary of a shape corpus. The sparse stroke-rendering of shape parts facilitates legible juxtaposition, which in turn enables easy integration into the spatial context of a design application (Figure 11) and aids designers in mentally imagining shapes beyond the collection (Figure 15).

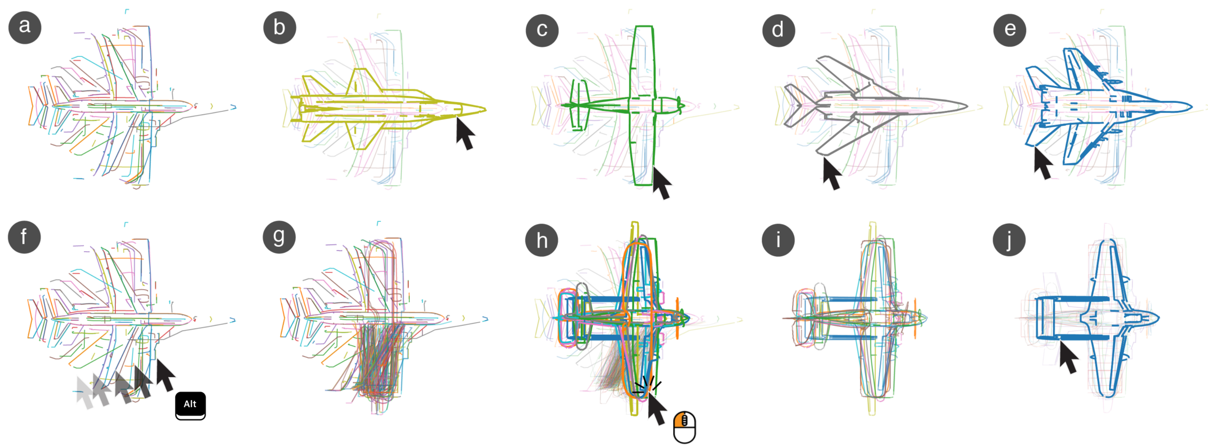

Fig. 5. A typical interaction workflow for *juxtaform* in the context of exploring a collection of airplanes. The user is first shown a visual summary of strokes (a) which they can then explore by hovering over different features of interest like a pointed nose (b), a flat wing (c), a very swept-back wing (d), or a wide tail (e). To explore deeper into the collection, users can brush over regions of interest, such as the wing (f) , revealing a collection of new stroke bundles (g). When hovering over revealed stroke bundles (h) the users can click to 'filter in' sub-collections of interest for further exploration (i). The eventual shape of interest can be selected by clicking on it (j).

4 JUXTAFORM

Guided by these design goals, we designed *juxtaform* as an interactive sketch stroke-based system for the exploration of shape corpora. Critical to the design of *juxtaform* is our choice of *visual representation*, *spatial layout*, and *interaction handles* for shapes in the collection.

Visual representation: We abstract all shapes in the corpus into a set of perceptually salient strokes. Strokes for the entire corpus are then analysed to extract stroke bundles that summarize representative shape and part structure across the corpus (**D1**).

Spatial layout: Our visual summary for a shape corpus is an *in-situ* juxtaposition of a sparse set of stroke bundles that optimize screen space, balancing spatial shape diversity with minimal visual clutter (**D1**, **D2**, **D3**). Note that the displayed strokes capture the entire corpus and may not depict some shapes completely, or at all.

Interaction Handles: The displayed stroke bundles themselves form interaction handles. Hovering over a stroke bundle allows a user to browse and cycle through complete shapes that contribute to the bundle (**D3**). An explore brush enables visualizing more/less of the shape variation in the corpus in a spatial region (**D2**). Visible stroke bundles can also be used to filter shapes in the corpus spatially or using sketch input (**D1**).

4.1 System Pipeline

The *juxtaform* pipeline (Figure 4) comprises an initial processing stage where the shape corpus is analysed to extract a sparse set of strokes filtered to visually summarize the corpus, and an interaction stage where the designer can directly use the summary strokes to further explore or refine parts of the shape corpus.

4.1.1 Processing Stage. The initial processing stage creates a visual summary by (i) abstracting, (ii) superimposing, and (iii) summarizing shapes in the corpus. The input shape corpus is first abstracted to a perceptually salient stroke-based NPR representation for each shape (**D1**). This sketch-stroke abstraction can be computed for 2D/3D shape automatically as shape features like ridges/valleys and view-dependent contour lines from 2D images [Chan and Vese 2001] or 3D models [Bénard and Hertzmann 2019]. The extracted strokes are then segmented into perceptual parts, for example using a corner detection technique [Wolin et al. 2008]. We assume shapes are already rigidly aligned based on these shape features (various auto-registration approaches exist [Van Kaick et al. 2011]). The extracted sketch strokes are then superimposed onto each other (**D3**, **D4**). The superimposed strokes are summarized using our novel stroke summarization algorithm (detailed in §5.1) which balances diversity and visual clutter of depicted strokes from the corpus.

4.1.2 Exploration Stage. The summarized strokes are presented to the user as handles for interactive exploration. We provide a set of interactive features to support three main tasks informed by our formative analysis (§3): exploration (discover / reveal new shapes), search (query and filter shape parts based on stroke features), and analysis (study the distribution of shapes in the corpus).

The user interface (UI) is simple and usable with a mouse or stylus in one of select, explore or draw modes (see supplementary video). Hovering and clicking, as typical, is used to search, highlight, and select sketch strokes which in turn visualizes and selects shapes, or parts of shapes. The modifier 'alt' is used to enter a brush-based exploration mode where mouse motion reveals more/less sketch strokes that comprise the shape corpus. Hotkey 'D' enters draw mode, where mouse strokes are used as sketch input. The mouse wheel can be used to cycle the visualization of searched shapes or to control the clutter parameters for stroke summarization. The hotkey 'C' cycles through multiple color visualizations to aid in understanding the sketch corpus.

Discovery / Exploration: Each displayed stroke is a handle for exploration, allowing a user to discover the shape(s) it belongs to by hovering over it with the mouse cursor. This highlights the complete shape(s) that contain the hovered location, muting unrelated strokes. Users can discover more strokes in regions of interest using the explore brush activated by holding the 'alt' modifier key. As users use the explore brush, they can interactively summarize the added strokes to the desired resolution with the mouse wheel, retaining only novel variations.

Filter / Search: Once users see ideas they like (or do not like) while hovering over a region, they can focus their exploration by clicking to filter in shapes contained in the region, or conversely, remove shapes using an eraser (activated by holding the 'ctrl' modifier). Users can also sketch desired stroke features to query shapes with such strokes in the 'draw' mode, toggled using the 'D' key.

Corpus Analysis: *Juxtaform* uses three color maps (cycled using the 'C' key), to help users study the distribution of shapes and their parts in the corpus at a desired level of abstraction, each of which targets a particular analysis goal. The *object-based color map* assigns one color per object, facilitating visual identification of individual shapes in the collection. The *stroke-based color map* maps each segmented stroke to a different color to help users visualize the part structure of shape(s) and the available exploration handles. The *diversity-based color map* utilizes the VIRIDIS color scheme [Nuñez et al. 2018] to indicate stroke density: violet in regions with a high stroke density in the overall collection (common strokes), to a lighter yellowish color in regions of low stroke density (unique strokes). While the first two color maps help users discern different juxtaposed shapes and features, the diversity score-based color map (see Figure 8) conveys the distribution of shapes in the collection (**D1**, **D2**, **D3**).

5 STROKE FILTERING AND SUMMARIZATION

We now present the two main stroke processing algorithms driving *juxtaform*. (i) Summarizing Stroke Collections: balancing visual clutter and the maximal display of shape diversity; and (ii) Interactive Stroke Filtering: interactively navigating large stroke corpora using spatial filtering.

5.1 Visually Summarizing Large Stroke Collections

The problem of presenting a large shape collection as a set of superimposed strokes can be perceived as optimizing three contending objectives: a) the inclusion of a maximal set of strokes that sufficiently represent the diversity of features in the collection, b) including as-complete-as-possible shapes to provide a strong sense of the overall objects in the collection, and c) minimizing visual clutter to ensure that individual strokes can be perceived and form meaningful handles for selection and exploration. Mathematically, this objective can be expressed as finding a stroke collection S which maximizes an objective E as follows:

$$E(S) = E_d(S) + E_c(S) + E_v(S) \tag{1}$$

where $E_d(S)$ is a measure of diversity captured by the stroke collection, $E_c(S)$ estimates the completeness of shapes in the stroke

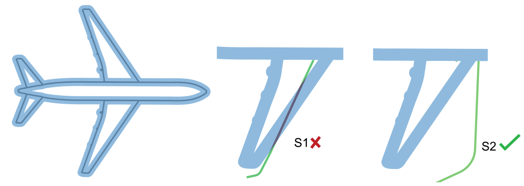

Fig. 6. A δ-neighborhood (in blue) for a set of strokes (in black) and the intersection of new candidate strokes (green) with this neighbourhood (in red). A large fraction of $S1$ intersects with the shape neighbourhood and it is therefore not added to the summary, whereas $S2$ represents a unique feature with almost no intersection with the current neighborhood and is therefore added to the summary.

collection, and $E_v(S)$ is a measure of the visual clarity of the stroke collection. We now describe each of these terms.

$E_d(S)$: We define diversity measure $E_d(S)$ to capture both common and unique features in the collection. We first define a δ-neighborhood for a shape or stroke as the region containing all points within δ distance from the object. The uniqueness of a stroke s, $U(s) \in [0, 1]$, is then defined as the average fraction of s not contained within the respective neighborhoods of other shapes in the collection (Figure 6). Commonness, $1-U(s)$, is the reverse of uniqueness (see Figure 7 (left) for highly common strokes). To capture both common and unique strokes, we define E_d as follows:

$$E_d(S) = \sum_{\forall s \in S} |U(s) - \frac{1}{2}| \tag{2}$$

$E_c(S)$: To estimate the completeness of shapes in the stroke collection, we define a measure called stroke coherence C which captures whether a pair of strokes belong to the same shape. Specifically, we set $C(s_i, s_j) = 1$ for each s_i, s_j that belong to the same shape, and 0 otherwise. $E_c(S)$ is then defined as:

$$E_c(S) = \sum_{\forall s_i, s_j \in S} C(s_i, s_j) \tag{3}$$

Optimizing for E_c also has a desirable side effect of decreasing the likelihood of selecting partial or short floating stroke segments that may be artifacts of a sub-optimal stroke segmentation.

$E_v(S)$: To estimate visual clarity $E_v(S)$, we consider each stroke as the owner of its δ-neighborhood. Given that a collection of strokes belonging to the same shape form a coherent whole, cluttered areas are then defined as regions of the screen whose ownership is contested by strokes belonging to multiple shapes. Visual clutter can thus be estimated as the sum of the areas of such regions, with each area weighted by the number of contesting shapes. E_v, or the lack of visual clutter, is the reciprocal of this sum. E_v thus also implicitly offsets the tendency to select multiple overlapping common strokes due to E_d.

Finding S that maximizes $E(S)$ is an expensive combinatorial selection problem, similar to NP-hard graph problems like Maximum Independent Set. We note, however, that $E(S)$ is based on perceptual heuristics and an efficient/interactive but approximate solution is

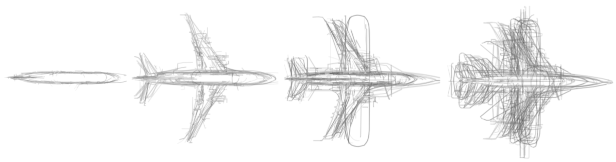

Fig. 7. (From left to right) Dense summaries generated using stroke subsets with increasing uniqueness $U(s)$ rendered using a sketch style that highlights shapes with high average E_d. The ordering of our greedy selection has the quality of reverse-engineering a human sketching session that begins with constructing a scaffolding of the most common parts before adding unique details with high feature diversity.

critically preferable to a slow optimal algorithm. We thus opt for a fast greedy algorithm based on locally optimal stroke objectives.

We now describe our greedy algorithm for stroke summarization. Given a shape collection, we now wish to compute an optimal subset of strokes S that maximizes our objective function E. The underlying idea is to greedily select increasingly diverse, non-conflicting strokes (thus promoting visual clarity E_v and diversity E_d) from the collection on a shape-by-shape basis (thus promoting completeness E_c). We therefore iterate through shapes ordered by the average uniqueness of their constituent strokes (least unique first). For a given shape, we add all of its constituent strokes which have a high uniqueness with respect to the combined δ-neighborhood of shapes that have already been processed ordered by length (also promotes completeness E_c). This is because we can fairly assume that in case of low uniqueness with respect to processed shapes, the feature represented by the stroke is already captured by strokes from previous shapes in the summary view (see Figure 6). A parameter α is used to define the uniqueness threshold for stroke selection.

The order of our greedy selection is designed to include as much diversity as possible, ensuring that the distribution of shapes in the collection is well represented and complete shapes are included as far as possible. The ordering of the shapes helps the summary view stay faithful to the distribution of shapes in the collection. Finally, ordering by length within a shape helps us ensure that we include as-complete-as-possible shapes. Our greedy approach therefore approximates a local optimum for the optimization.

The level of detail in our constructed stroke summary can also be interactively adjusted by varying α and δ (which determine if a stroke is sufficiently represented by already processed shapes) to provide coarse-to-fine summaries of the stroke collection. Higher α (and lower δ) admits more spatially proximal strokes creating more detailed summaries (illustrated in supplementary material).

5.2 Spatial Stroke Filtering

Spatial filtering allows users to select and filter shapes via point-based spatial queries. In particular, if a user specifies a point with a mouse click, all the displayed strokes which intersect the neighborhood of the point are identified, and their owner shapes are selected and filtered in while removing remaining shapes. Conversely, such spatial queries can be used to erase or remove the selected shapes.

Our *sketch-based filter* also works in a similar manner, filtering out any shapes whose δ neighborhoods do not contain the sketched curve. Finally, our stroke summarization algorithm's coarse-to-fine α, δ parameters allow users to control visual clarity adaptively, and globally or in local regions during exploration (illustrated in section 2 of the supplementary material).

6 TECHNICAL EVALUATION

We evaluate our summarization algorithm technically by showing the impact of its design parameters and by comparing it with prior art on geometric aggregation, and evaluate its performance on diverse corpora. A comprehensive evaluation of parameter choices and comparisons with prior art can be found in section 2 of the supplementary material.

6.1 Parameters

Our algorithm relies on two parameters: the δ neighbourhood selected for dilation and the α value which defines the allowed percentage intersection threshold for stroke selection. While both δ and α can be used to control the amount of visual clarity seen in the computed summaries, we fix δ based on the size of the superimposed collection's bounding box (default 1% of the box diameter). Informed by our practical experience with different corpora, we set the default $\alpha = 0.6$ (at least 40% of a stroke's length should convey new visual information).

6.2 Comparison with Geometric Aggregation

Our algorithm is specifically tailored to summarizing a shape corpus subject to our design goals. It is however worth evaluating whether existing geometric aggregation techniques can be trivially repurposed to solve our problem.

To construct a comparable summary with shape aggregation methods, we clustered the shape corpus using out-of-the-box image clustering techniques [Pedregosa et al. 2011; Taskesen 2021] and superimposed the cluster representatives. The number of clusters were chosen to match either the number of strokes or shapes in our summary. As we can see in Figure 9, with comparable stroke density, our summary captures a greater feature diversity and represents a larger portion of shapes in the collection with fewer strokes. On the other hand, with comparable shape representation, shape aggregation-based summaries cannot be visually parsed. Overall, while shape aggregation methods do ensure complete shapes in their summaries, our ability to include incomplete shapes in a coherent part-aware manner helps us capture more feature diversity with less visual clutter.

6.3 Performance

We evaluate the performance of our method on a diverse set of shape corpora consisting of 2D images, 2D fonts, and 3D models with varying size and stroke density. In particular, we measure the amount of sparsification performed by our method with respect to strokes and shapes, and the time taken for pre-processing and summarization. This evaluation, summarized in Table 1, was conducted on an ASUS laptop with 16GB RAM and an NVIDIA RTX 3060 laptop GPU.

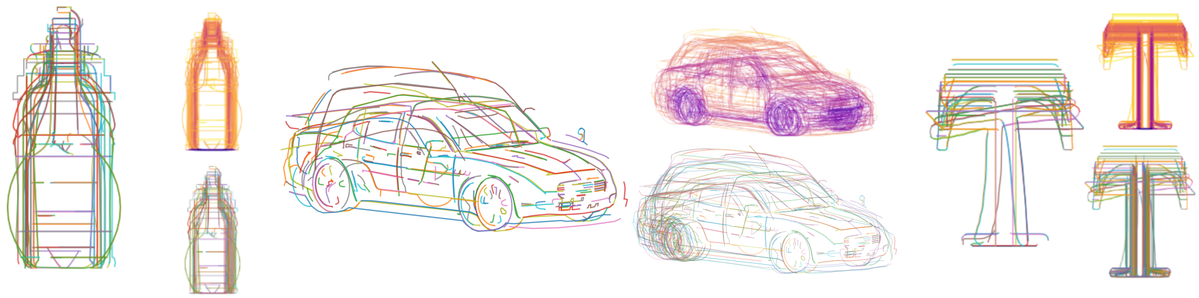

Fig. 8. A gallery of views generated while exploring different shape collections with *juxtaform*. We show two collections of 3D models (bottles, cars) and an image collection of a capital T in different fonts. For each collection, the default visual summary is illustrated, along with a visualization of the strokes in the collection colored by their uniqueness (violet (common) to yellow (unique)) and a view with the diversity of a local region explored by brushing.

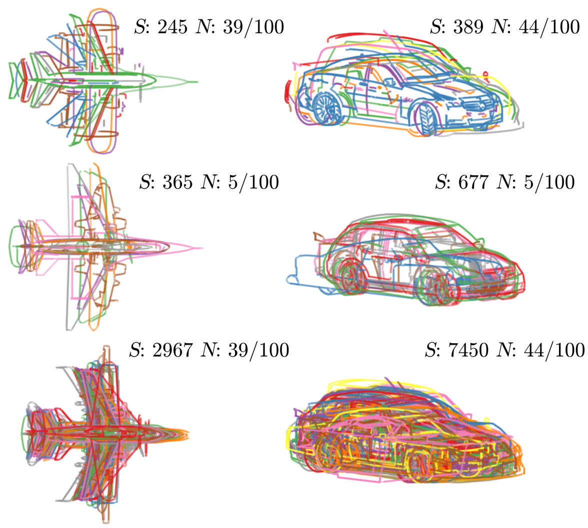

Fig. 9. Our summarization results can represent features from a greater number of shapes N with fewer strokes S and minimal visual clutter (top) compared to shape aggregation-based summaries with a comparable number of strokes (middle) and shape representation (bottom).

Our summarization algorithm robustly handles shape collections of different sizes and stroke densities, ranging from single shape overdrawn sketches to 10K-sized shape repositories. The summaries typically converge on a few hundred strokes, significantly thinning out the input, and feature parts from about 40-100 shapes in larger datasets. Beyond this point, we presumably either hit the visual clarity threshold or there are no further shapes with unique features, making remaining strokes redundant.

With respect to the time taken, the shape-stroke distance calculation is the most expensive computation due to the many shape-stroke intersection computations within it. However, this computation can be perceived as a one-time pre-processing cost for a dataset, as this distance matrix remains constant throughout interactive exploration iterations. Further, in comparison to traditional stroke clustering techniques that require inter-stroke distances our shape-stroke distance computation consists of fewer calculations and is therefore significantly more cost-effective for a large stroke count. The greedy summarization algorithm, on the other hand, is fast and steadily converges in under a second, suitable for interactive use.

Table 1. Performance statistics for our summarization algorithm. S and N are respectively the number of shapes and strokes in the input. S' and N' are corresponding numbers for the summarized output. The time is split into the shape-stroke distance computation (Dist.) and summarization (Summ.).

Dataset	N	S	N'	S'	$\frac{N'}{N}$	$\frac{S'}{S}$	Time (in secs)	
							Dist.	Summ.
Plane2	10K	725K	85	495	0.85	0.07	2835	0.92
Vase	470	15K	236	986	50.22	6.4	625	0.86
Plane1	448	38K	67	399	14.96	1.05	1027	0.8
Guitar	288	12K	47	313	16.32	2.54	313	0.25
Car	225	28K	54	504	24	1.8	960	0.68
Fonts	100	1526	28	169	28	11.08	26	0.01
Bottle	38	773	29	179	76.32	23.16	2.5	0.01
Koala	1	146	1	85	100	58.22	0.3	0.004
Lion	1	362	1	198	100	54.7	0.7	0.01

7 USER STUDY

To evaluate *juxtaform* as a shape exploration system compared against the status-quo traditional gallery layouts, we conducted a within-subject user study with 15 participants of varied vocational backgrounds and experience. Our study included tasks designed to simulate realistic browsing scenarios pertaining to three exploration goals—search, comparison, and understanding tasks; to complete for three different shape collections—guitars (200 3D models [Wu et al. 2015]), airplanes (450 3D models [Wu et al. 2015]), and fonts (Capital T in 100 fonts [Ge et al. 2021]). For each shape collection, participants were asked to use *juxtaform* only, a folder system only (Figure 10b), or either/both as they see fit. A detailed report of the protocol, participants, and results is in the supplementary material. Below, we summarize key take aways from the study.

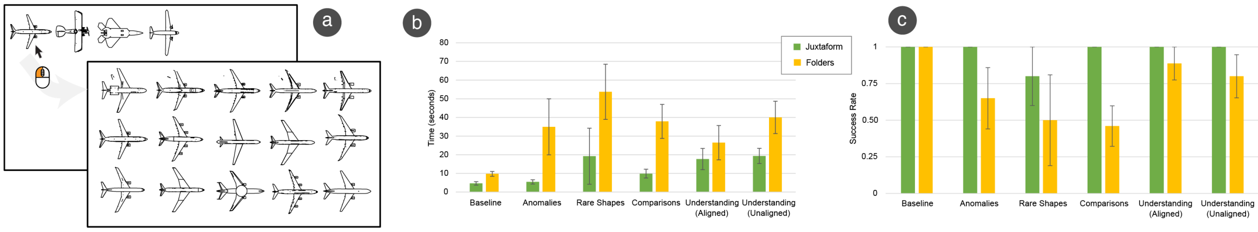

Fig. 10. (a) The folder system used in the user study. Users could view representative icons for each folder (top left) and on double clicking the icon they were shown its contents in a gallery layout (bottom). (b-c) Participants' performance by task: completion time (b) and success rate (c) defined by the average correctness of participant answers, for *juxtaform* (green) and the folder system (yellow). Error bars show 95% confidence intervals.

Performance: Figure 10-b shows a summary of participants' mean completion time and correctness of answers for each task using *juxtaform* or the folder system only.

Search tasks involved finding a particular shape based on a description of a part, geometric feature or semantic style (**D1**), and can be divided into baseline (find generic shapes with many examples in the collection), rare (find uncommon shapes with a rare characteristic feature), and anomalies (find geometric outliers). As we expected, the folder system and *juxtaform* were comparable for baseline search; and *juxtaform* resulted in more efficiency at helping users quickly identify outliers, anomalies, and rare targets in the collections. Participants particularly struggled with the folder system when outlier and rare shapes differed in small local region, compared to a cluster of similar shapes. This can be attributed to the benefits that juxtaposed layout offer over side-by-side layout for quick identification of unique features compared to the general trend, as supported by participants comments: *"It was hit and miss with the folders, because it can be quite overwhelming to scan through a large set of similar shapes looking for a particular feature."* (P9).

Comparison tasks focused on comparative analysis, asking users to identify relative extremes and trends in the context of other shapes in the collection (**D3**); e.g, are shapes of type Y always taller than that of type X?) Arguably, a gallery is *not* designed for comparative analysis; in contrast, *juxtaform* should facilitate such tasks *by design*. As anticipated, results show that participants were able to answer such tasks quickly and with perfect accuracy with *juxtaform*.

Understanding tasks focused on the statistical distribution of shapes in the collection, asking users to identify common features and exemplify the diversity in the collection via examples (**D2**). Because different clusterings in the folder system make tasks easier or harder, we divide understanding tasks based on their alignment with the organization of the folders. As we expected, the folder system performed better on tasks which were aligned with its organization compared to those that were not (e.g. cluster planes by tail similarity makes it easier to process tail features, than it supports understanding of wings). *Juxtaform*'s cluster-agnostic representation performed well irrespective of the task, and users often used the diversity score color to quickly perform such tasks.

Usage Patterns: Analysis of tool usage and interview scripts allowed us to characterize notable usage patterns.

In *juxtaform*, participants balanced the use of the summary view and explore brush in different ways. While most relied on the summary and brushed only when needed, some preferred to see as much of the collection as possible and extensively used the brush, while leveraging different color maps. This suggests that our approach does not prescribe a unique worfklow to solve a task, and therefore effectively accommodates different user styles and strategies.

To be clear, we are not arguing for *juxtaform* to replace galleries: both approaches are complementary, and this was evidenced in our study. For the third shape collection, we gave participants the freedom to use either or both systems as they wished to perform exploration tasks. While most (n=12/15) chose to use *juxtaform* only, complementary usage was also observed, e.g. starting with *juxtaform* to get a sense of the collections then swtiching to folders to visualize details in isolation to make final selections.

Perceptual Challenges: Two prominent and consistent themes that emerged from our interviews analysis were the load on visual memory and spatial attention.

The most pressing challenge participants faced while exploring large collections of shapes with folders was the need to construct and retain a mental map of the different shapes in the collection. One of the most consistent comments was that *juxtaform*'s visualization design allows them to navigate and explore the collection without any mental map at all, e.g. P4 said *"With the folders, the limit really is my memory, it's all about trying to remember where things are. But with* juxtaform *it's just all there in front of me all the time."* (**D1**, **D2**).

The second challenge that participants faced was paying attention to many different spatial objects at once. In the folder system, relating, comparing, and aggregating features across distant objects was overwhelming. With *juxtaform*, the challenge was more about parsing different shapes when multiple got displayed at once, e.g. P5 said *"I would inherently like to see [individual shapes] in isolation. Imagine seeing two colorful shirts on top of each other, it would not give me any sense of what they might look like on their own."* Now, participants felt that the ability to select and filter shapes was essential and effectively mitigated this challenge: P14 said *"juxtaform's eraser and selection tools help so much when there's too many shapes in front of me. Not only does it feel like I am drawing or creating my own shapes when filtering with the sketch and eraser, but within a few clicks or brushes, I can always make sure that I'm looking at a very relevant and easy to visualize set of shapes."* (**D1**, **D2**, **D3**).

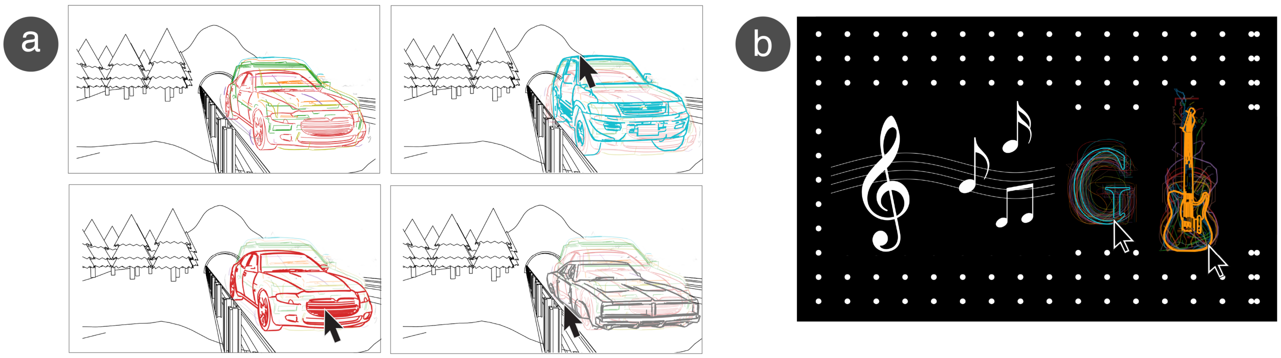

Fig. 11. *Juxtaform*'s potential for in-situ shape exploration is illustrated with two applications: a) selecting assets (a car) for 3D scenes and b) picking fonts and visual elements for a poster design.

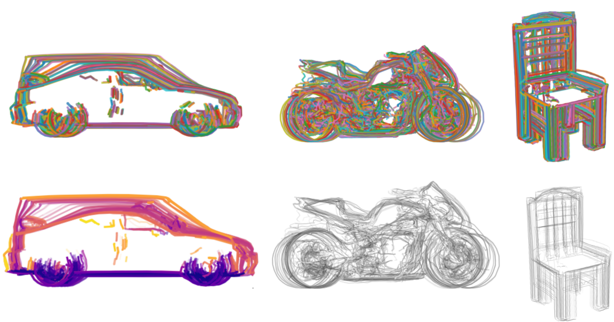

Fig. 12. A shape collection generated by exploring the latent space of GET3D [Gao et al. 2022] is superimposed (top) and the distribution of shape features is visualized using *juxtaform*'s uniqueness score-based rendering (bottom).

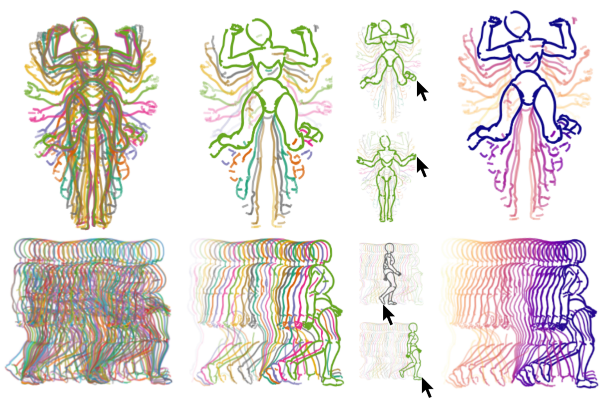

Fig. 13. Shapes constituting an animation cycle are superimposed (left) and summarized by *juxtaform* to retain unique portions of each frame (middle). The trail-like visualization, generated by mapping color and opacity to the frame count, allows users to browse interesting frames (middle) and visualize the animation cycle (right).

8 CREATIVE APPLICATIONS

Juxtaform's visualization and interaction design enables unique creative applications for a shape exploration framework. Inspired by two of our initial design goals, i.e visualizing explored shapes in-situ or in context of a users eventual application scenario (**D3**); and the ability to explore new ideas that go beyond the reference collection (**D4**), we describe four creative applications illustrating the potential of *juxtaform* for ideative exploration (**D4**) that we envisioned, and report on professional artists' impressions on the approach and their ideas on where *juxtaform* could integrate in their creative practice.

8.1 Creative Usage Scenarios

8.1.1 In-situ exploration. Exploring shapes *in-situ* in a given creative context eliminates the need for switching between shape exploration and creative ideation. For example, Figure 11-a illustrates the ability to explore a collection of cars within a 3D scene, and Figure 11-b shows a user exploring different 2D fonts and guitar shapes in the context of a poster design.

8.1.2 Mix-and-Match Creation. Juxtaform's workflow also implicitly supports creative ideation over a novel space of shapes defined by the combination of stroke-based features in the collection. We allow users to explicitly ideate over this rich space of shapes by selecting compatible strokes from different shapes to construct novel designs. As illustrated in Figure 15, users can mix and match shape parts via this workflow to construct novel design variations.

8.1.3 Browsing Personalized Assets. Another creative application of *juxtaform* is the rapid exploration of the different versions, iterations, or articulations of a creative artifact. Artist workflows can produce hundreds of iterations and incremental and varied versions of creative designs and articulated characters, resulting in a rich space of variations of a single shape. As shown in Figure 13, visualizing this space of designs with *juxtaform* instantly highlights regions of similarity and variation allowing artists to conveniently pick out a desired part, version, or articulation of a shape for future use. Access to parts and committed versions in a shape's construction

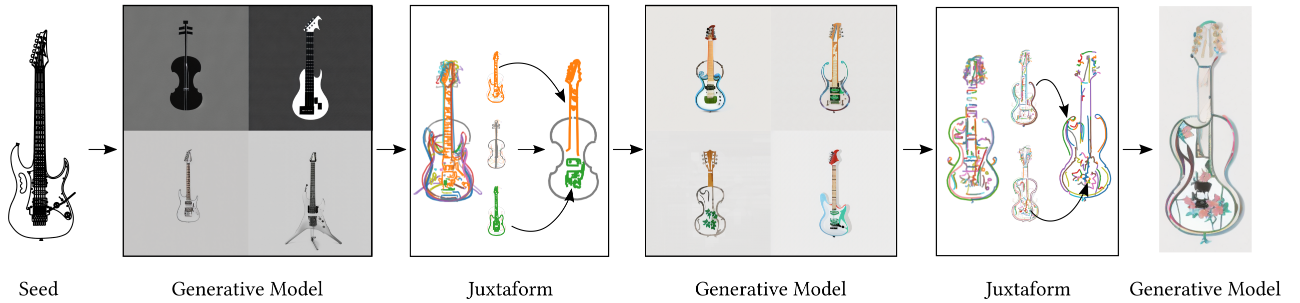

Fig. 14. *Juxtaform* acts as an intuitive interface for artists to explore design ideas using modern diffusion models. Large variation sets generated by the model (here, 50 shapes per iteration) can be loaded into *juxtaform* to browse and construct concept designs that mix and match interesting ideas from the variation set. Constructed designs can be input back into the model to complete an iterative feedback loop which allows users to visually control the generation process.

history is useful since artists often re-use parts and even construction scaffolds. Stroke coloring and transparency in *juxtaform* can further convey the chronology of the shapes, making it easy to visually parse a set of models with a meaningful time ordering. For instance, as we can see in Figure 13, shapes that constitute frames of an animation can be summarized and visualized in the order of the frame count instead of feature commonality to construct compelling summaries with the flavour of an onion-skinning visualization. Such summaries can not only enable artists to instantly access a desired pose from hundreds of unique frames in a dense animation [Assa et al. 2005] but can also act as informative thumbnails to efficiently browse a collection of animation cycles.

8.1.4 Exploratory Design with Generative AI. The emergence of powerful machine learning-based creative tools has made it important for the modern artist to interface with AI-based generative models for early-stage design and ideation. *Juxtaform* has the potential to act as a front-end for modern generative models, enabling artists to quickly incorporate modern machine learning in their creative design workflows. In Figure 14, we illustrate a convergence-divergence design cycle using *juxtaform*'s mix-and-match and stability AI's stable diffusion 2.1 image generator [Rombach et al. 2021]. The image generator is proficient at creating novel variations and polished designs from input text and image guidance. *Juxtaform* can be used to efficiently explore AI generated variations and construct novel concept designs from a combination of desired parts in the variation set. These designs can then be fed back into the image generator as more directed guidance for further exploration. This rapid iterative workflow allows both artist and novice users an intuitive geometric control over the generative model's output, thus making it easier to converge on desirable designs.

In the context of 3D generative models, a common difficulty is to exhaustively explore and visualize a dense latent space of variations corresponding to a particular object class. In Figure 12, we see how *juxtaform*'s diversity score can be used to visualize the distribution of shapes and features in the latent space. In particular, a collection of shapes generated from NVIDIA's GET3D [Gao et al. 2022] are sampled and superimposed in *juxtaform*. Shapes that, on average, contain the most common features, or conversely the most unique, are rendered with greater opacity and deeper colors. Other shapes are assumed to lie within this scaffolding and are gradually hidden away using the same parameters, thus creating a concept sketch-like render of the shape and feature distribution. This visualization gives designers a sense of the shape and part distribution that the generative model is capable of producing, which can then be used to control the model or inspire their own designs.

8.2 Impressions from Artists

We further assessed the creative potential of *juxtaform* in realistic artistic workflows by demonstrating *juxtaform* to 4 creative professionals including a professional illustrator and graphic designer (A1), a product design graduate student (A2), a professional 3D artist (A3), and an experienced developer of creative design tools (A4); with professional experience ranging from 2-3 years (A2, A3) to 5 years (A4) and more (A1).

The creators were generally excited about the novelty of our idea. They felt that it would be a great addition to their day-to-day workflow but also suggested ideas for other creative applications that *juxtaform* could enable.

Creatively Explore Reference Material. A1 said *"The part where the artist is able to sketch and erase to find shapes and forms that match their intended artistic goal really stood out to me as something that I would find beneficial in my day to day as an artist. I imagine it would be a very handy way to ensure I'm following my reference material correctly and find inspiration for new forms to draw."*

They further went on to describe a concrete scenario where they would be excited to use *juxtaform* *"Imagine a production artist tasked to create a set of similar-ish image assets, perhaps a suite of plants or rocks for use in a game or background art. If I was that artist and I had this tool, I would love to feed a collection of my own illustrated assets I've drawn so far, and rely on this tool to help me find new forms and shapes similar to the existing set, while at the same time inspire me to try out new forms and combinations. In a way, I could use this tool to ensure that my output isn't too similar to the existing set."*.

(Collaboratively) Explore Design Options. A2 and A3 alluded to the benefits of *juxtaform*'s ability to integrate shape exploration into different contexts. A2 said *"I think it would be great to churn ideas for quick design proposals. I can also imagine having the interactive output of this tool within my proposal so that the client can explore different options within the context of my proposal."* A3 added: *"I think*

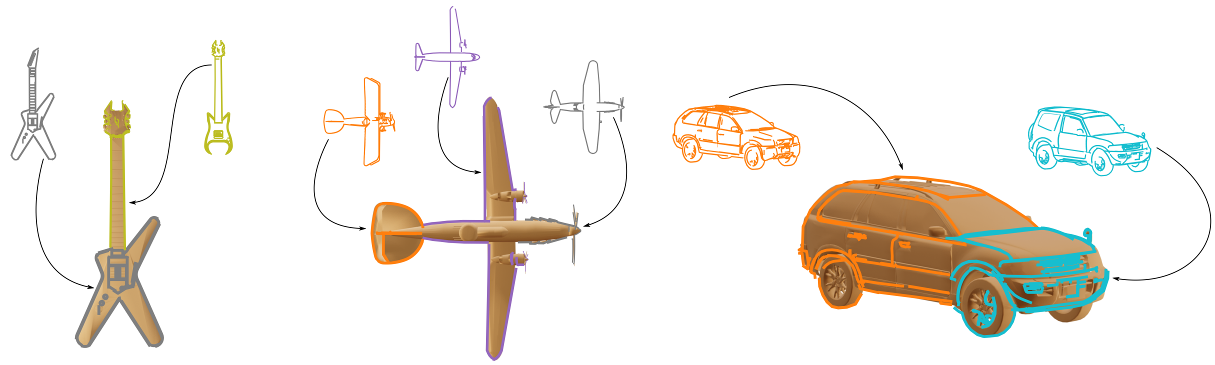

Fig. 15. 3D models made by combining parts from a 3D car repository with MeshMixer [Schmidt and Singh 2010] based on designs created via *Juxtaform*'s mix-and-match-based exploration of the source repository. Design references (highlighted strokes around 3D models) are made using *juxtaform*, based on which parts are queried from the source repository and seamlessly fused together in MeshMixer.

it could be great for concept art and storyboarding to show variations in context.".

Shape Attributes. The possibility of attaching various attributes to the shapes in the tool was an interesting usage mentioned by both A2 and A4. A4 said *"Industrial designers find such exploratory tools useful during ideation process. If you also labeled your shapes with attributes, such as year it was made, etc. – then designers could see the progression of designs and try to imagine where they should take the next year's design."* In a similar vein, and also integrating ideas from above, A2 mentioned *"I can think of it being very useful as a collaborative visual note taker* [...] *designers iterate over designs with clients and I make a lot of rough sketches to visually note what clients like and want* [...] *our design team could color code parts of the design based on the clients comments, and we could then superimpose all our 'visual notes' to explore different design ideas, I would imagine that many good designs would pop out naturally using this tool."*

Composition. Finally, similar to one of our creative usage scenarios, A1 was also interested in using the tool to explore composition, and said *"sometimes I'm looking for compositions instead of just particular details of a subject. For example, if I gave the tool a set containing some of my favourite posters for movies and music that were broken down into sketch lines, I would be curious to see how I would use such a tool to block in new compositions."*.

9 DISCUSSION, LIMITATIONS AND CONCLUSION

Juxtaform addresses the ambitious inter-disciplinary (graphics/HCI/visualization/design) space of shape corpus exploration and refinement for ideative shape design. Given the generality of the creative application and the dimension and domain of shapes, we performed a formal formative study (§3) to formulate clear design objectives to be implemented by an interactive shape exploration system (§4), built around stroke-based shape corpus summarization. We show that existing geometric aggregation approaches are not suited to our problem (§5). Our novel shape corpus summarization (§6) is efficient, robust, and directly tied to our design goals. A user study (§7) and creative applications (§8) showcase the promise of *juxtaform* for shape exploration in interactive shape design for diverse sets of shapes and creative scenarios. Our proposed system is not without limitations, but has been used by designers and study participants with much excitement and positive feedback (see more artist impressions in supplementary material).

Juxtaform is built on the core assumption that input shapes are spatially pre-aligned, both in terms of orientation and scale. When shapes are not aligned (e.g. animation frames in Figure 13), the spatial disparity should be considered as shape diversity/variation. Shapes that are not spatially aligned can result in summary views that are non-intuitive and messy. However, advanced geometric techniques for automatic shape registration, correspondence, and alignment [Van Kaick et al. 2011] offer effective solutions to this limitation, making it easy to pre-deform shapes to be spatially aligned prior to exploration, and this is subject to future work.

Another limitation of our method is the simple Euclidean metric used for spatial dilation and shape similarity. It is possible for our approach to ignore diverse but spatially proximal variations of shape as might be deemed as similar or identical by our distance metric. This is particularly true for unique variations like fine detail and bumps which are perceptually distinct but of small spatial amplitude. However, since our stroke filtering algorithm is agnostic to the shape affinity metric used, in the future we can investigate alternate metrics for shape similarity.

Our approach also assumes shape and application scenarios that are visually well-represented using sparse sketch strokes. Shape filled with color and texture, such as emojis or clothes might need additional visual elements and transparency to be adapted to our framework. Even after that, users might prefer to visualize them in isolation for targeted search tasks. Further, application scenarios like quickly picking out emojis for texting constitute situations where the benefits of *in-situ* exploration may not be as important.

Finally, the size of the shape collection is an important consideration. While our system has clear benefits for a sweet spot of shapes in the hundreds, exploration for smaller shape collections might potentially be as good with simple gallery layouts. Further, larger collections with thousands and millions of shapes may need to be pruned by keywords and other criteria to a tractable size before they can be effectively explored in *juxtaform*. The unique advantages and applications of our method pertaining to comparative analysis and *in-situ* shape design, however, remain relevant contributions that are independent of the size of the dataset.

In conclusion, our work brings together literature and insights from graphics, HCI, and design practice, making a contribution to ongoing research pertaining to ideative shape design. Our formative study design goals can guide future systems. Our stroke summarization algorithm balances spatial variation and clutter, and can be applied to other interfaces looking to achieve compact visual representations of geometric data with rich diversity. Our user study design provides ideas for evaluating creative exploration systems that need to support a broad range of design tasks. In summary, we propose and evaluate a novel shape exploration system that is fast (**D1**), supports the targeted exploration of diversity **(D2)** in context (**D3**), and acts as an easy-to-use enjoyable interface for creative ideation (**D4**).

ACKNOWLEDGMENTS

Our research is funded in part by NSERC. We would also like to thank Disha Tripathy for early discussions, feedback, and help with figures, Zhecheng Wang and Selena Ling for their help with the video submission and figures, the participants of our user studies for their time and input, and the anonymous reviewers for their valuable feedback.

REFERENCES

Aseem Agarwala, Mira Dontcheva, Maneesh Agrawala, Steven Drucker, Alex Colburn, Brian Curless, David Salesin, and Michael Cohen. 2004. Interactive digital photomontage. In *ACM SIGGRAPH 2004 Papers*. 294–302.

Andrés Arias-Rosales. 2022. The perceived value of human-AI collaboration in early shape exploration: An exploratory assessment. *PloS one* 17, 9 (2022), e0274496.

Rahul Arora, Ishan Darolia, Vinay P Namboodiri, Karan Singh, and Adrien Bousseau. 2017. Sketchsoup: Exploratory ideation using design sketches. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 302–312.

Rahul Arora, Rubaiat Habib Kazi, Tovi Grossman, George Fitzmaurice, and Karan Singh. 2018. Symbiosissketch: Combining 2d & 3d sketching for designing detailed 3d objects in situ. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–15.

Jackie Assa, Yaron Caspi, and Daniel Cohen-Or. 2005. Action synopsis: pose selection and illustration. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 667–676.

Melinos Averkiou, Vladimir G Kim, Youyi Zheng, and Niloy J Mitra. 2014. Shapesynth: Parameterizing model collections for coupled shape exploration and synthesis. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 125–134.

Pierre Bénard and Aaron Hertzmann. 2019. Line Drawings from 3D Models: A Tutorial. *Found. Trends. Comput. Graph. Vis.* 11, 1–2 (sep 2019), 1–159. https://doi.org/10.1561/0600000075

Silvia Biasotti, Andrea Cerri, Alex Bronstein, and Michael Bronstein. 2016a. Recent trends, applications, and perspectives in 3d shape similarity assessment. In *Computer graphics forum*, Vol. 35. Wiley Online Library, 87–119.

Silvia Biasotti, Bianca Falcidieno, Daniela Giorgi, and Michela Spagnuolo. 2016b. 3D Objects Exploration: Guidelines for Future Research.. In *3DOR@ Eurographics*.

Silvia Biasotti, Elia Moscoso Thompson, and Michela Spagnuolo. 2019. Context-adaptive navigation of 3D model collections. *Computers & Graphics* 79 (2019), 1–13.

Umberto Castellani and Adrien Bartoli. 2020. 3d shape registration. In *3D Imaging, Analysis and Applications*. Springer, 353–411.

Tony F Chan and Luminita A Vese. 2001. Active contours without edges. *IEEE Transactions on image processing* 10, 2 (2001), 266–277.

Andy Cockburn, Amy Karlson, and Benjamin B Bederson. 2009. A review of overview+ detail, zooming, and focus+ context interfaces. *ACM Computing Surveys (CSUR)* 41, 1 (2009), 1–31.

Daniel Cohen-Or and Hao Zhang. 2016. From Inspired Modeling to Creative Modeling. *Vis. Comput.* 32, 1 (jan 2016), 7–14. https://doi.org/10.1007/s00371-015-1193-9

Noa Fish, Melinos Averkiou, Oliver Van Kaick, Olga Sorkine-Hornung, Daniel Cohen-Or, and Niloy J Mitra. 2014. Meta-representation of shape families. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.

Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. 2022. GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images. In *Advances In Neural Information Processing Systems*.

Lin Gao, Yan-Pei Cao, Yu-Kun Lai, Hao-Zhi Huang, Leif Kobbelt, and Shi-Min Hu. 2014. Active exploration of large 3d model repositories. *IEEE transactions on visualization and computer graphics* 21, 12 (2014), 1390–1402.

Yunhao Ge, Sami Abu-El-Haija, Gan Xin, and Laurent Itti. 2021. Zero-shot Synthesis with Group-Supervised Learning. In *International Conference on Learning Representations*. https://openreview.net/forum?id=8wqCDnBmnrT

Stéphane Grabli, Frédo Durand, and Francois X Sillion. 2004. Density measure for line-drawing simplification. In *12th Pacific Conference on Computer Graphics and Applications, 2004. PG 2004. Proceedings.* IEEE, 309–318.

Yulia Gryaditskaya, Felix Hähnlein, Chenxi Liu, Alla Sheffer, and Adrien Bousseau. 2020. Lifting freehand concept sketches into 3D. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.

Felix Hähnlein, Changjian Li, Niloy J Mitra, and Adrien Bousseau. 2022. CAD2Sketch: Generating Concept Sketches from CAD Sequences. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–18.

Almoctar Hassoumi, María Jesús Lobo, Gabriel Jarry, Vsevolod Peysakhovich, and Christophe Hurter. 2019. Interactive shape based brushing technique for trail sets. (2019).

James W Hennessey, Han Liu, Holger Winnemöller, Mira Dontcheva, and Niloy J Mitra. 2016. How2Sketch: generating easy-to-follow tutorials for sketching 3D objects. *arXiv preprint arXiv:1607.07980* (2016).

Josh Holinaty, Alec Jacobson, and Fanny Chevalier. 2021. Supporting Reference Imagery for Digital Drawing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2434–2442.

Min Hua. 2019. The roles of sketching in supporting creative design. *The Design Journal* 22, 6 (2019), 895–904.

Ruqi Huang, Panos Achlioptas, Leonidas Guibas, and Maks Ovsjanikov. 2019. Limit Shapes – A Tool for Understanding Shape Differences and Variability in 3D Model Collections. *SGP* (2019).

Shi-Sheng Huang, Ariel Shamir, Chao-Hui Shen, Hao Zhang, Alla Sheffer, Shi-Min Hu, and Daniel Cohen-Or. 2013. Qualitative organization of collections of shapes via quartet analysis. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–10.

Moos Hueting, Maks Ovsjanikov, and Niloy J Mitra. 2015. CrossLink: joint understanding of image and 3D model collections through shape and camera pose variations. *ACM Transactions on Graphics (ToG)* 34, 6 (2015), 1–13.

Arjun Jain, Thorsten Thormählen, Tobias Ritschel, and Hans-Peter Seidel. 2012. Exploring shape variations by 3d-model decomposition and part-based recombination. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library, 631–640.

Waqas Javed and Niklas Elmqvist. 2012. Exploring the design space of composite visualization. In *2012 ieee pacific visualization symposium*. IEEE, 1–8.

Vladimir G Kim, Wilmot Li, Niloy J Mitra, Stephen DiVerdi, and Thomas Funkhouser. 2012. Exploring collections of 3d models using fuzzy correspondences. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–11.

Yanir Kleiman, Noa Fish, Joel Lanir, and Daniel Cohen-Or. 2013. Dynamic maps for exploring and browsing shapes. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 187–196.

Heidi Lam and Tamara Munzner. 2010. A guide to visual multi-level interface design from synthesis of empirical study evidence. *Synthesis Lectures on Visualization* 1, 1 (2010), 1–117.

Yong Jae Lee, C Lawrence Zitnick, and Michael F Cohen. 2011. Shadowdraw: real-time user guidance for freehand drawing. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 1–10.

You-En Lin, Yong-Liang Yang, and Hung-Kuo Chu. 2018. Scale-aware black-and-white abstraction of 3D shapes. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–11.

Chenxi Liu, Enrique Rosales, and Alla Sheffer. 2018. Strokeaggregator: Consolidating raw sketches into artist-intended curve drawings. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–15.

Y-C Liu, Amaresh Chakrabarti, and T Bligh. 2003. Towards an 'ideal'approach for concept generation. *Design studies* 24, 4 (2003), 341–355.

Zhen-Bao Liu, Shu-Hui Bu, Kun Zhou, Shu-Ming Gao, Jun-Wei Han, and Jun Wu. 2013. A survey on partial retrieval of 3D shapes. *Journal of Computer Science and Technology* 28, 5 (2013), 836–851.

Justin Matejka, Michael Glueck, Erin Bradner, Ali Hashemi, Tovi Grossman, and George Fitzmaurice. 2018. Dream Lens: Exploration and Visualization of Large-Scale Generative Design Datasets. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3173574.3173943

Harry A Nichols. 1954. *DESIGN FEATURES OF THE DOUGLAS A3D SKYWARRIOR.* Technical Report. SAE Technical Paper.

Jamie R. Nuñez, Christopher R. Anderton, and Ryan S. Renslow. 2018. Optimizing colormaps with consideration for color vision deficiency to enable accurate interpretation of scientific data. *PLOS ONE* 13, 7 (08 2018), 1–14.

Luke Olsen, Faramarz F Samavati, Mario Costa Sousa, and Joaquim A Jorge. 2009. Sketch-based modeling: A survey. *Computers & Graphics* 33, 1 (2009), 85–103.

Günay Orbay, Mehmet Ersin Yümer, and Levent Burak Kara. 2012. Sketch-based aesthetic product form exploration from existing images using piecewise clothoid curves. *Journal of Visual Languages & Computing* 23, 6 (2012), 327–339.
Maks Ovsjanikov, Wilmot Li, Leonidas Guibas, and Niloy J Mitra. 2011. Exploration of continuous variability in collections of 3d shapes. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 1–10.
Patrick Paczkowski, Min H Kim, Yann Morvan, Julie Dorsey, Holly E Rushmeier, and Carol O'Sullivan. 2011. Insitu: sketching architectural designs in context. *ACM Trans. Graph.* 30, 6 (2011), 182.
Mark Pauly, Niloy J Mitra, Johannes Wallner, Helmut Pottmann, and Leonidas J Guibas. 2008. Discovering structural regularity in 3D geometry. In *ACM SIGGRAPH 2008 papers*. 1–11.
Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
Miquel Prats. 2007. *Shape exploration in product design: assisting transformation in pictorial representations*. Open University (United Kingdom).
Miquel Prats and Chris F Earl. 2006. Exploration through drawings in the conceptual stage of product design. In *Design Computing and Cognition'06*. Springer, 83–102.
Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. arXiv:2112.10752 [cs.CV]
Ryan M. Schmidt and Karan Singh. 2010. meshmixer: an interface for rapid mesh composition.. In *SIGGRAPH Talks*, James L. Mohler (Ed.). ACM. http://dblp.uni-trier.de/db/conf/siggraph/siggraph2010talks.html#SchmidtS10
Donald A Schon and Glenn Wiggins. 1992. Kinds of seeing and their functions in designing. *Design studies* 13, 2 (1992), 135–156.
Adriana Schulz, Ariel Shamir, Ilya Baran, David IW Levin, Pitchaya Sitthi-Amorn, and Wojciech Matusik. 2017. Retrieval on parametric shape collections. *ACM Transactions on Graphics (TOG)* 36, 1 (2017), 1–14.
Ben Shneiderman. 1999. User interfaces for creativity support tools. In *Proceedings of the 3rd conference on Creativity & cognition*. 15–22.
Tim Smithers. 2001. Is sketching an aid to memory or a kind of thinking? (2001).
Johan WH Tangelder and Remco C Veltkamp. 2008. A survey of content based 3D shape retrieval methods. *Multimedia tools and applications* 39, 3 (2008), 441–471.
Erdogan Taskesen. 2021. *Python package clustimage is for unsupervised clustering of images*. https://erdogant.github.io/clustimage
James T Todd. 2004. The visual perception of 3D shape. *Trends in cognitive sciences* 8, 3 (2004), 115–121.
Steve Tsang, Ravin Balakrishnan, Karan Singh, and Abhishek Ranjan. 2004. A Suggestive Interface for Image Guided 3D Sketching. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vienna, Austria) *(CHI '04)*. Association for Computing Machinery, New York, NY, USA, 591–598. https://doi.org/10.1145/985692.985767
Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. 2011. A survey on shape correspondence. In *Computer graphics forum*, Vol. 30. Wiley Online Library, 1681–1707.
Dave Pagurek Van Mossel, Chenxi Liu, Nicholas Vining, Mikhail Bessmeltsev, and Alla Sheffer. 2021. StrokeStrip: joint parameterization and fitting of stroke clusters. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–18.
Remco C Veltkamp. 2001. Shape matching: Similarity measures and algorithms. In *Proceedings International Conference on Shape Modeling and Applications*. IEEE, 188–197.
Aaron Wolin, Brian Eoff, and Tracy Hammond. 2008. ShortStraw: A Simple and Effective Corner Finder for Polylines. *SBIM* 8 (2008), 33–40.
Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 2015. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1912–1920.
Baoxuan Xu, William Chang, Alla Sheffer, Adrien Bousseau, James McCrae, and Karan Singh. 2014a. True2Form: 3D curve networks from 2D sketches via selective regularization. *ACM Transactions on Graphics* 33, 4 (2014).
Kai Xu, Rui Ma, Hao Zhang, Chenyang Zhu, Ariel Shamir, Daniel Cohen-Or, and Hui Huang. 2014b. Organizing heterogeneous scene collections through contextual focal points. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–12.
Kai Xu, Hao Zhang, Daniel Cohen-Or, and Baoquan Chen. 2012. Fit and diverse: Set evolution for inspiring 3d shape galleries. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–10.
Chuan Yan, David Vanderhaeghe, and Yotam Gingold. 2020. A benchmark for rough sketch cleanup. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–14.
Hui Ye, Kin Chung Kwan, and Hongbo Fu. 2021. 3D curve creation on and around physical objects with mobile AR. *IEEE transactions on visualization and computer graphics* 28, 8 (2021), 2809–2821.
Jun-Yan Zhu, Yong Jae Lee, and Alexei A Efros. 2014. Averageexplorer: Interactive exploration and alignment of visual data collections. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.