Juxtaform: User Studies and Technical Evaluations

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

In this document we include complete details for the user studies and technical evaluations conducted in our paper "Juxtaform: interactive visual summarization for exploratory shape design". Specifically, we describe the set up, procedure and insights from both the formative user study (section 3.1 in the main manuscript) and the evaluation user study (section 7 in the main manuscript) in full. We further describe impressions from aritsts about the creative potential of *juxtaform* in realistic artistic workflows and an in-depth technical evaluation of our algorithm which includes the effect of parameters, performance and comparisons with relevant prior art.

ACM Reference Format:
Karran Pandey, Fanny Chevalier, and Karan Singh. 2023. Juxtaform: User Studies and Technical Evaluations. 1, 1 (July 2023), 8 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 USER STUDIES

We characterize the problem space of exploring shape collections from the perspective of creative end-users. We highlight insights from formative interviews with artists and designers, who creatively work with shapes on a daily basis. The role and challenges of shape exploration in their workflow inspire a set of design goals for creative shape exploration systems.

1.1 Formative Interviews

We recruited 5 participants comprising of industrial (P1, P2) and product designers (P3), professional 3D modelers (P5)and independent artists (P4) with experience ranging between 2-3 years (P1, P3) to 5 years (P2, P4) and more (P5), through convenience sampling; and interviewed them about their applications, workflow and personal experience exploring shape collections on prior professional projects.

1.1.1 Interview Format, Structure, and Analysis. Each interview was an hour-long semi-structured discussion conducted via videoconference. The interview began with a discussion of the participant's creative application(s), workflow(s) and the role of shape exploration within it. Participants were encouraged to walk us through example projects, and highlight challenges faced when exploring shapes. At the end of the interview, we asked the participant to critique examples of recent research in shape exploration.

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

XXXX-XXXX/2023/7-ART $15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

We outlined four specific knowledge goals for follow-up questions during our interviews:

- *Exploratory Goals*: In the context of their projects, what were creators looking for? What form of shape abstraction did their search terms have? What were the high-level goals for the exploration?
- *Selection Criteria*: How do creators select from a set of candidate shapes and what criteria inform their choice?
- *Applications*: How do creators use the shapes they find?
- *Challenges*: What challenges do creators face in shape exploration? Which aspects of the process are fun/tedious?

Each interview session was video recorded and transcribed. Through thematic analysis, we comparatively analysed transcripts to form a holistic understanding of shape exploration workflows and applications. Quotes from each participant were grouped with respect to the four overarching themes, and further clustered into insights.

1.1.2 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, in current practice had a fairly uniform shape exploration workflow: they would first narrow their exploration to a tractable set of hundreds of shapes via shape category queries to an online search engine, or directly to tagged and labeled shape repositories. These query results, typically comprising shapes with strong structural and spatial commonalities, were then visualized as a catalogue of thumbnails, usually within a web browser or file explorer. Sometimes a few (typically 5 or less) shapes are chosen from the thumbnails, for further processing, such as evaluating these shapes in-situ within a larger design, or using parts of these shapes to refine or embellish an evolving design; such processing is iterative and requires repeated switching between the selected shapes and their application context.

Creators reported that comparing disparately presented shapes in such catalogues, for similarity or diversity, is both difficult and tedious. P2 mentioned: *"If I'm looking for a bottle to add to my scene for example, I can quickly search the term 'bottle' on the internet and I have a huge collection of shape ideas in front of me, that's easy. The difficult and boring part is really to look through this list and try to find what I want. Seeing similar bottles again and again makes me mentally switch off while looking for options and I don't really register differences between them. Its also impossible to see everything in the set of [queried] results so I just pick the best I can see after a point and make the most of it."*. In a similar vein, P5 said *"It gets quite annoying to go through multiple catalogues of shapes looking for new ideas* *after a point. The same ideas repeat themselves in different ways and its hard to [summarily] evaluate whether a collection has something interesting without diving into it and spending a lot of time, which is something we usually don't have time for in a deadline crunch."*

At the end of the interview, we showed participants examples of unique creativity support tools such as ShapeSynth [Averkiou et al. 2014] and SketchSoup [Arora et al. 2017] which combine creative ideation with shape exploration. The participants were excited about the aesthetic visual design of SketchSoup and ShapeSynth's unique ability to generate shapes beyond the initial collection. In particular, they mentioned that such creative elements have the potential to turn the usually tedious task of clicking through shapes into an engaging and creative exploration of ideas.

In summary, we gained five key insights from these interviews.

- *Imprecise Shape Search* (**I1**): Creators often have only a broad sense of desired shape(s), since the chosen shapes in some tasks are rarely used as-is, or are an accessory, whose precise details are not critical to the design task.
- *High throughput* (**I2**): Shape exploration is often accompanied by an approaching deadline or perceived as a quick precursor to impending design tasks. An overarching issue with existing shape collection browsers was the time it could take, even to determine if the shape collection was appropriate for the creative task.
- *Rich Variations* (**I3**): Creators sought a diverse set of shape variations during exploration, both for inspiration and for shape refinement. In particular, there was a desire to understand the overall distribution of shapes in a collection: the common and unique parts and features of shapes and their relationship to each other.
- *Understanding in Context* (**I4**): It is important for creators to be able to pre-view and understand explored shapes, relative to each other and their given design context. It is significantly easier for example, to evaluate and select between candidate 3D shapes, when juxtaposed in-situ, in a 3D scene with spatial constraints.
- *Creative Engagement* (**I5**): Shape exploration, like the creative process itself, should be fun and engaging. Creators we interviewed found ideation interfaces that are playful and aesthetically pleasing (eg. SketchSoup [Arora et al. 2017] for shape, and ColorBuilder [Shugrina et al. 2019] for color), appealing to dabble with, as a source of inspiration and ideas.

1.2 Design Goals

We distilled these collective insights into 4 design goals for creative exploration of shape corpora.

- *Rapid Exploration* (**D1**): Quickly convey using sketch abstraction, the overall essence of shapes in large collections (**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.

1.3 Related Work Relative to Design Goals

*1.3.1 Rapid Exploration (**D1**).* Rapid exploration is important for both objective and subjective tasks.

A large body of research is focused on objective tasks such as the targeted search and retrieval of shape, based on different input queries, such as rough sketches, images or text [Biasotti et al. 2016; Gao et al. 2014; Rehman et al. 2012; Tangelder and Veltkamp 2008]. Relative to *juxtaform*, these approaches would be used either to narrow exploration to still large class of shapes, or downstream for precision in later stage shape design.

Work has also been done on subjective tasks such as free-form browsing, overall understanding of the shape corpus, and ideative exploration of shape diversity [Averkiou et al. 2014; Huang et al. 2013; Ovsjanikov et al. 2011], shape correspondence and structure [Huang et al. 2013; Kim et al. 2012; Xu et al. 2014] to provide an overview of the shape corpus. Low-dimensional feature embeddings or common can also streamline interactive exploration [Arora et al. 2017; Averkiou et al. 2014].

Sketchy (NPR) and other perceptual shape abstractions [Averkiou et al. 2014; Lin et al. 2018; Ovsjanikov et al. 2011; Todd 2004], as well as spatial arrangements and juxtapositions of common structure [Huang et al. 2013; Kleiman et al. 2013; Matejka et al. 2018], aid understanding and comparison in shape collections [Zhu et al. 2014].

Inspired by these themes, *juxtaform* employs a juxtaposed sketch-based visual abstraction, to define an interactive, stroke-based visual summary of a shape corpus that addresses all four design goals.

*1.3.2 Diversity exploration (**D2**).* Many shape exploration systems recognize the artistic need to explore both local and global diversity. Proxies such as a common abstracted part structure [Ovsjanikov et al. 2011], shape grammars [Dang et al. 2015], shape statistics [Matejka et al. 2018], or a low-dimensional feature space [Averkiou et al. 2014], offer interactive modalities to explore local and global shape variations. Critically, these approaches sometimes filter away uniquely interesting and extreme shapes and their parts as outliers, and the presence of explicit intermediate representations can seem foreign, break user flow, and make variation control difficult [Arias-Rosales 2022]. Closer to our goals are sketchy systems that present suggestions from a shape collection based on partial input [Lee et al. 2011; Orbay et al. 2012], or probe the shape collection using regions on individual or averaged shapes [Kim et al. 2012; Zhu et al. 2014].

In contrast to proxy approaches, *juxtaform* interaction remains largely within the domain of pure ideation sketching. It also does not average or alter the input shapes, rather choosing to present an entire shape collection as a juxtaposed sketch stroke-based visual summary, balancing the needs of showing entire representative shapes and common/unique parts of shapes, while minimizing visual clutter.

*1.3.3 Contextual Exploration (**D3**).* While the importance of comparative and in-situ visualization within an application context is well established [Arias-Rosales 2022; Lao et al. 2021; Shireen et al. 2019], most ideative shape exploration systems display shapes one-at-a-time, or in gallery-like spatial layouts [Averkiou et al. 2014; Huang et al. 2013; Kleiman et al. 2013; Ovsjanikov et al. 2011]. Such layouts require dedicated screen space, precluding them from supporting explicit comparisons, and flexible integration into a focused application context. Outside of exploration of shape corpora, various tools have effectively used in-situ visualization for design space exploration in immersive environments [Lao et al. 2021] and sketch-based ideation [Arora et al. 2018; Lee et al. 2011].

Exploration in context is inherent in *juxtaform*, where a sparse stroke-rendering of shape parts frees up valuable screen space to facilitate legible juxtaposition (superposition), which in turn enables convenient comparison of shapes and easy integration into the spatial context of a design application.

*1.3.4 Ideative Exploration (**D4**).* Numerous approaches such as generative models, shape grammars, genetic algorithms [Xu et al. 2012] and latent shape spaces [Averkiou et al. 2014] address ideative shape exploration. Many of these approaches introduce randomness to generate novelty in the shape generation process [Cohen-Or and Zhang 2016], which can be inspiring but can also inhibit control over the shapes, suggesting the need to control any stochasticity in shape variations [Arias-Rosales 2022].

The sparse part-based stroke-rendered visual shape summaries in *juxtaform* are designed to aid designers in mentally imagining shapes beyond the collective (Fig. 15 main paper). Further, *juxtaform* can be a compelling front-end to exploring the ideation space of generative models (Fig. 14 main paper).

1.4 Evaluating Juxtaform

We conducted a user study to evaluate *juxtaform* as a shape exploration system and examine its value in comparison, and for complementary use with traditional gallery-based layouts as identified in our formative interview participants (§1). We set out to learn about: (i) *juxtaform*'s visualization system and interaction workflow for a diverse set of realistic exploration tasks and application scenarios; (ii) how *juxtaform* fares compared to the competition in terms of performance and user experience; (iii) usage patterns of *juxtaform* which complements that of gallery-based exploration system; and (iv) perceptual challenges associated with spatial arrangement in large-scale shape exploration tasks.

We designed a multi-faceted comparative evaluation of *juxtaform* against a baseline exploration system, on tasks involving different browsing scenarios.

1.4.1 Protocol. Our study took place in-person, in a dedicated room in our laboratory. Upon providing consent and filling out a short demographics questionnaire, participants were provided with an overview of the study, where the facilitator explained that they will be tasked to explore three different shape collections (planes, guitars, fonts) to identify characteristics about the collection, or particular shapes of interest, using two exploration systems: *juxtaform* and a Folder system akin to typical gallery layout views as used by professionals for such tasks (§1).

Before completing the tasks relevant to a shape collection with the designated system, participants were given a brief walkthrough of the system. They were also allowed to free-form explore the collection before being shown the tasks, so as to gain a basic sense of the collection (which would be consistent with a real-world scenario) and familiarize themselves with the interface (therefore minimizing time to appropriate basic functionality).

For each shape collection, participants were asked a series of questions which they had to answer through exploration of the shape collection in order, before moving to the next question. Once they answered all questions, participants were invited to take a break, before starting exploration of the next shape collection.

The exploration systems to work with for the first two collections were fixed (i.e *juxtaform* first, then folder system or vice versa). For the third collection, participants were given the freedom to use one or both of the systems as they wished. This allowed us to study potential hybrid usage patterns for exploration tasks, and capture preferences of use in the context where *juxtaform* complements other existing approaches. To reduce potential ordering effects or bias associated with certain shape collections, we balanced systems presentation order and collections across participants. For instance, if P1 used *juxtaform* first to explore the planes collection, then folders for guitars, and finally had a choice for fonts; P2 could be assigned folders for guitars, then *juxtaform* for fonts, and would have a choice for airplanes.

At the end of the study, we performed an unstructured closing interview to discuss the participant's user experience, usage patterns observed during the third task where the participant had the choice of the tool, challenges faced during exploration and a discussion asking them to contrast and compare their experience with both systems. A subset of participants were also given the opportunity to try the additional features of *juxtaform*, such as its different spatial and sketch-based filters, which were developed based on feedback from early participants in the user study.

1.4.2 Tasks and Measures. We curated three shape collections, covering diverse object types, characteristics, and collection size: 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]). Each collection had a set of exploration tasks associated with it, such as identifying items possessing a particular feature, or characterizing global distributions, which we designed based on workflows described by our formative interview participants (§1.1).

Our tasks were designed to simulate realistic browsing scenarios pertaining to three exploration goals; search, comparison and understanding. *Search tasks* involve finding a particular shape based on a description of a part, geometric feature or semantic style (**D1**); and can be divided into three types based on the target; baseline (generic shapes with many examples in the collection), rare (uncommon shapes with a rare characteristic feature) and anomalies (geometric outliers). *Understanding tasks* focus 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**). *Comparison tasks* focused on comparative analysis, asking users to identify relative extremes and trends in the context of other shapes in the collection (**D3**). For e.g, are shapes of type Y always taller than type X? For each collection, we designed a set of 4 search tasks, 3 understanding tasks and 3 comparative tasks.

We recorded completion time for each task, correctness of the answers on a scale of 0 (wrong) to 1 (correct), as well as observation notes whilst participants completed the tasks, along with qualitative feedback collected during the closing interview.

1.4.3 Systems. We used two exploration systems in the study. The first system was a gallery-based layout of stroke-based NPR renders of the shape corpus presented in a folder system (Figure 1-a), which mimics the current solution we found professionals use. The second was a limited version of *juxtaform*. We chose to only encompass *juxtaform* features for which an equivalent capability exists in the baseline to allow for a more fair comparison between the two approaches, i.e. our study did not include spatial or sketch-based filters. Further, we also grouped similar shapes into folders, each with a representative thumbnail, to imbue the traditional gallery with summarization abilities similar to *juxtaform*, allowing us to keep the focus on the juxtapose vs. gallery presentation method: shapes were grouped into folders using a clustering algorithm [Jarman 2020] based on a dilation-based shape similarity score. Given that the shapes were pre-registered, the similarity score for a pair of shapes was defined as the minimum dilation needed for one shape neighbourhood to contain the other shape. As a sanity check, we also ran the study on one participant with a flat unclustered folder structure and noticed that the timing and correctness was at least twice as bad compared to the clustered folders. We therefore decided against including a comparison with a fully unclustered layout in our user study.

1.4.4 Participants and Apparatus. We recruited 15 participants, aged between 23 to 35, from a diverse set of ethnic and vocational backgrounds via convenience sampling. Three participants had extensive prior experience with shape exploration via creative or technical exposure to digital graphics software (P3, P6, P13), while others were limited to an occasional use of text-based searches for online shopping or presentation design. The studies were conducted on an ASUS laptop with 16 GB RAM, an NVIDIA RTX 3060 graphics card and an external monitor at a resolution of 1920 x 1080 pixels, equipped with a mouse and keyboard, with a capture of audio and screen recordings.

1.5 Results

We discuss our main findings with regard to performance, patterns of usage which we synthesize from observing participants completing the tasks, and insights gathered from a thematic analysis of post-study interview transcripts.

1.5.1 Performance. Figure 1-b shows a summary of participants' mean completion time and correctness of answers. We were primarily interested in re-creating in-the-wild exploration scenarios; and while participants were instructed to take a reasonable amount of time to complete each task, no time performance, nor duration limit were set. This allows us to capture completion time that would be closer to reality, i.e. factoring in the time one may take to get more familiar about a dataset in earlier tasks, and a primary goal of correctness as opposed to speed. Completion times should be interpreted in this context. Further, we did not find performance differences across shape collections, suggesting that they were all of equivalent complexity. Below, we discuss our findings per type of tasks.

Baseline Search: As we expected, both the folder system and *juxtaform* performed well for such tasks, with juxtaform having a slight edge with timing which could be attributed to the instant access to diverse shapes via the summary view (**D1**). This finding is however negligible given the fact time efficiency was not instructed as a goal for participants.

Anomaly Search: *Juxtaform*'s efficacy at helping users quickly identify outliers and anomalies in the dataset was evidenced by the notable difference in both timing and answer quality for such tasks (**D1**, **D2**). Participants particularly struggled with the folder system when outliers differed only in small local region (e.g, a guitar with a slanted neck), and got grouped within a larger cluster.

Rare Target Search: Participants effectively used *juxtaform*'s summary view and brushing to find rare shapes in the collection. However, they also mentioned that *juxtaform* filters would have been particularly useful for such tasks. P11 said *"When restricted to the brush and hover, there were times when I could see something I wanted, but it was hard to isolate it, particularly in regions with many strokes."*, providing motivation for the relevance and necessity of the spatial-based and shape-based filters which were not included in our evaluation (**D1**, **D2**). With the folder system, participants missed rare shapes when they belonged to large clusters of similar shapes. P9 said *"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. I found rare shapes only when I chanced upon them at random. I could have sworn that the target shape was not in the collection until I just happened to notice it while browsing the same folder of shapes the fourth time, right before I gave up."*

Comparison: *Juxtaform* had a performance advantage compared to the gallery when it came to comparison tasks. Arguably, a gallery is *not* designed for comparative analysis and therefore participants using the folder system gave low confidence answers and explicit guesses. In contrast, *juxtaform* should facilitate such tasks *by design.* As anticipated, our results show that participants were able to answer such tasks quickly and with perfect accuracy (**D1**, **D3**).

Understanding: Understanding tasks displayed different trends based on the organization of the folder system in relation to the task. For instance, if the folders grouped airplanes by wing types, participants found it easy to perform tasks which required them to characterize different wing types. However, characterizing types of tails with this folder structure could be significantly harder. As a result, we divided understanding tasks based on their alignment with the organization of the folders. As we expected, the folder system performs better on tasks which were aligned with its organization compared to those that were not. *Juxtaform*'s cluster-agnostic representation performed well irrespective of the task. Users often used the diversity score color map to quickly perform such tasks. P1 said *"The diversity score color map is great to get a quick overview of the different common and unique parts in the collection. I found it very*

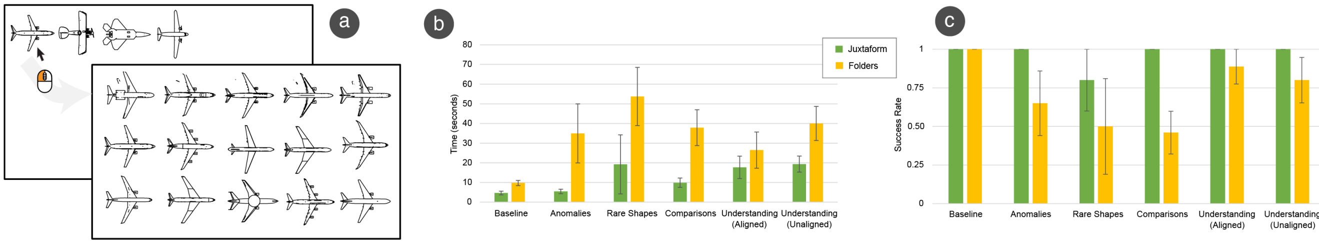

Fig. 1. (a) The folder system used in the user study. Users could view representative icons for each folder (top left) and on double clicking the icon they were shown its contents in a gallery layout (bottom). (b-c) Participants' performance by task: completion time (b) and success rate (c) defined by the average correctness of participant answers, for *juxtaform* (green) and the folder system (yellow). Error bars show 95% confidence intervals.

easy to answer questions about the distribution of shapes using this color map." (**D1**, **D2**).

1.5.2 Usage Patterns. We performed a thorough analysis of tool usage from screen recordings of the study sessions. This allowed us to characterize notable usage patterns, which we further contextualize with insights from our observation notes and relevant quotes from interview transcripts.

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 dataset as possible and extensively used the brush. One such participant, P3, stated that *"While the summary was great, and I love how different shapes pop out when I hover over them; I really enjoyed using the explore brush to reveal more stroke bundles and see everything. I think seeing everything gives me a comfort that I'm not missing anything in the collection. I also think the color maps made it really fun to use the explore brush and discover the different strokes in the dataset"* (**D2**). Like P3, other participants also enjoyed using the different *color maps juxtaform* had to offer. P7 said *"The object-based color map really helped me associate different colors with different objects. It's almost like an additional encoding that also helps you remember the shapes you see. Imagine if everything was black, it would be much harder to remember the shapes. It of course also helps with discerning different shapes when there are multiple shapes displayed at once"* (**D2**). Meanwhile, P2 said *"I loved starting out with the diversity score color map. Brushing in different areas gave me a great sense of the distribution of shapes in the collection like more common or unique features. It also creates very aesthetic cool looking visualizations that I enjoyed exploring with!"* (**D1**, **D2**).

The two approaches are complementary. We gave participants the freedom to use either or both systems as they wished to perform exploration tasks for the third and final collection. This allowed us to investigate combined usage patterns for both exploration frameworks and the potential for their complementary use. All but three of our fifteen participants chose to use *juxtaform* for all the tasks in the third collection. P3 and P8 preferred to start with the folders but switched over to *juxtaform* after one or two tasks, and P6 employed a hybrid use pattern switching between systems based on the different tasks. In our closing interviews, participants noted that *juxtaform* was their preferred starting point to get a high-level overview of the diversity and distribution of shapes in the collection. P13 stated that *"*Juxtaform *is a great starting point for exploration. I think seeing everything together really helps me get a sense of what's really in the collection, what its extents are. Once I know that its easier for me to decide where to explore and how much. I wish all of the shape collections I explored came with this view on the side!"* (**D1**, **D2**). However, they would sometimes prefer to switch to folders to visualize details in isolation to make final selections. P6, the participant with a hybrid use pattern attempted to use the folders for this purpose, but felt that it would be more helpful if there was an explicit correspondence between the two systems. P6 stated *"While* juxtaform*'s combined view is great, I sometimes want to just be able to lay out the curves in front of me. Doing this for 1000s of curves makes no sense, but when I've brought it down to the 5 or 10 that I like, maybe I'd want to just lay them out in front of me and choose. Its hard though because I can't lay out say just the guitar heads that I'm looking at in* juxtaform *now. I have to go back to the folder system completely and then its difficult to even find what I am interested in again."*

1.5.3 Perceptual Challenges. We also discussed the perceptual challenges faced by participants. Two prominent and consistent themes that emerged from our interviews analysis were the load on visual memory and spatial attention.

Visual memory. 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. 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. Whenever I want to pick out a shape or go back to something I found earlier, I just have to brush or hover over a region and there it is!"* (**D1**, **D2**).

Spatial attention. The second challenge that participants faced was paying attention to many different spatial objects on the screen at once. As mentioned earlier (in P9's quote), this could be particularly overwhelming with the folders and lead to missed shapes during exploration. With *juxtaform*, on the other hand, the challenge was more about parsing different shapes when multiple got displayed at once. As P11 quoted earlier, participants felt that the ability to select and filter shapes was essential and effectively mitigated this challenge. When asked to compare, P14 said *"juxtaform's eraser and selection tools help so much when there's too many shapes infront 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. I can imagine it being perfect for picking out that one bag I want from a set of options when online shopping. With the folders, however, I am stuck to just the one way of categorization given to me. I can always try to reorganize everything myself, but I can't expect to do that every time I'm exploring with a different goal!"* (**D1**, **D2**, **D3**).

Within the same context of online shopping, P5 had a different opinion on *juxtaform*'s potential and said *"While* juxtaform *seems great for geometric shapes that have some sort of constant rigid structure, I can't imagine using it for clothes, for example, where there can be so many different orientations, views, colors and patterns. Plus its also something I would inherently like to see 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!"*

2 TECHNICAL EVALUATION AND COMPARISONS

We evaluate our summarization algorithm technically, by showing the impact of its design parameters, by comparing it with prior art on geometric aggregation/simplification, and evaluate its performance on diverse corpora.

2.1 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 clutter 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 be conveys new visual information). In Figure 2 and Figure 3, we see the effect of different choices of α and δ respectively on the computed summaries.

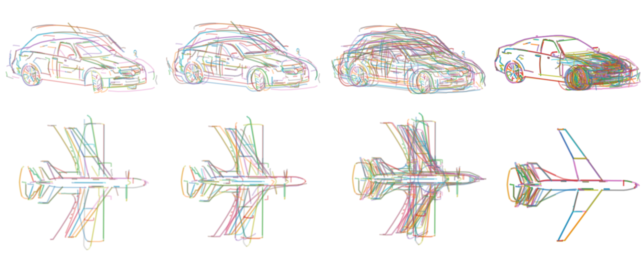

Fig. 3. (Left to Right) Summarization results with three decreasing stroke neighborhoods δ and an adaptive user selected δ in different semantic regions of the shape.

2.2 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. In particular, we consider standard sketch-stroke clustering methods and global shape-aggregation techniques.

Stroke clustering methods assume that they are operating on a single shape, and thus struggle to aggregate a large number of disparate strokes from different shapes within the same spatial context. As shown in Figure 4, we compared our results with a naive spectral k-means clustering approach [Pedregosa et al. 2011] based on the Hausdorff distance and a recent approach for clustering stroke data in overdrawn sketches: StrokeAggregator [Liu et al. 2018] on an input collection of guitars consisting of approximately 8000 strokes. In the case of StrokeAggregator, the input stroke collection required heavy pre-processing and thresholding as the executable struggled to run out-of-the-box on the large number of input strokes. On the pre-processed subset, it returned an overly simplified and aggregated output, assuming that the input strokes are a single overdrawn sketch. The spectral k-means approach, on the other hand, selects a set of representative strokes for each cluster with the number of clusters selected based on the number of strokes in our summary output. Here, we noticed that the selected strokes lack shape and part coherence (i.e. fail to represent whole parts in shapes), but also have regions of heavy visual clutter. Overall, we believe that the overwhelming density of strokes in our use case (thousands of strokes) makes it difficult to capture meaningful information with only inter-stroke distances.

Conversely, while not designed for sketch simplification, we find that our summarization approach, can produce reasonable results, for sketch simplification, treating every stroke in the sketch as a shape in a corpus (Figure 5).

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 6, 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 a complete shape part coherency 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.

2.3 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

Fig. 2. The detail and amount of diversity packed into our summarizations can be interactively adjusted from a coarse overview to a fine collection of detailed geometric features using the α parameter.

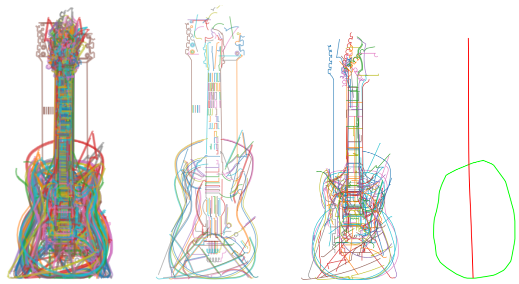

Fig. 4. Our method outperforms traditional stroke clustering techniques on multi-shape summarization both w.r.t visual clutter and coherent part diversity. From left to right, we see the input strokes (8K), our summary (282), a hausdorff distance-based summary (282) and StrokeAggregator output.

Fig. 5. Our summarization algorithm can be used for a fast selection-based sketch simplification, illustrated here on two overdrawn sketches.

summarization. This evaluation, summarized in Table 1, was conducted on an ASUS laptop with 16GB RAM and an NVIDIA RTX 3060 laptop GPU.

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

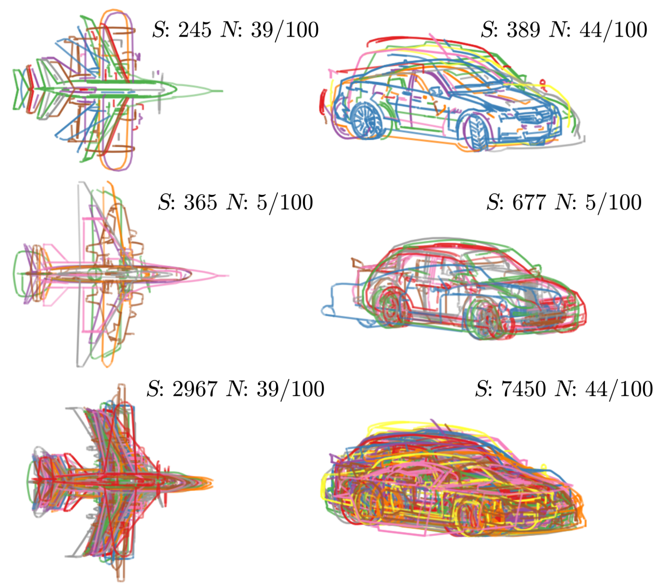

Fig. 6. Our summarization results can represent features from a greater number of shapes N with fewer strokes S and minimal visual clutter (top) compared to shape aggregation-based summaries with a comparable number of strokes (middle) and shape representation (bottom).

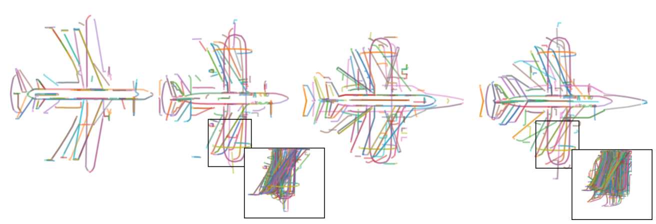

Fig. 7. We evaluate the effect of the dataset size on our summarization results. (Left to right) The dataset increases from 10 to 50, 500 and 10K shapes. Our summarization converges based on the maximum allowed clutter. The inset reveals the density of strokes in the input dataset.

out the input, and feature parts from about 40-100 shapes in larger datasets. Beyond this point, we presumably either hit the visual clutter 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

REFERENCES

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.

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.

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

Daniel Cohen-Or and Hao Zhang. 2016. From inspired modeling to creative modeling. *The Visual Computer* 32, 1 (2016), 7–14.

Minh Dang, Stefan Lienhard, Duygu Ceylan, Boris Neubert, Peter Wonka, and Mark Pauly. 2015. Interactive design of probability density functions for shape grammars. *ACM Transactions on Graphics* 34, CONF (2015), 206.

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

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.

Angur Mahmud Jarman. 2020. Hierarchical cluster analysis: Comparison of single linkage, complete linkage, average linkage and centroid linkage method. *Georgia Southern University* (2020).

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.

Cheryl Lao, Haijun Xia, Daniel Wigdor, and Fanny Chevalier. 2021. Attribute Spaces: Supporting Design Space Exploration in Virtual Reality. In *Symposium on Spatial User Interaction*. 1–11.

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.

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

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.

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.

Mehwish Rehman, Muhammad Iqbal, Muhammad Sharif, and Mudassar Raza. 2012. Content based image retrieval: survey. *World Applied Sciences Journal* 19, 3 (2012), 404–412.

Naghmi Shireen, Halil Erhan, and Robert Woodbury. 2019. Encoding Design Process Using Interactive Data Visualization. In *International Conference on Computer-Aided Architectural Design Futures*. Springer, 231–246.

Maria Shugrina, Wenjia Zhang, Fanny Chevalier, Sanja Fidler, and Karan Singh. 2019. Color Builder: A direct manipulation interface for versatile color theme authoring. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.

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.

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.

Kai Xu, Rui Ma, Hao Zhang, Chenyang Zhu, Ariel Shamir, Daniel Cohen-Or, and Hui Huang. 2014. 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.

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.