Rapid Design of Articulated Objects

JOON HYUB LEE, Department of Industrial Design, KAIST, Republic of Korea
HANBIT KIM, Department of Industrial Design, KAIST, Republic of Korea
SEOK-HYUNG BAE, Department of Industrial Design, KAIST, Republic of Korea

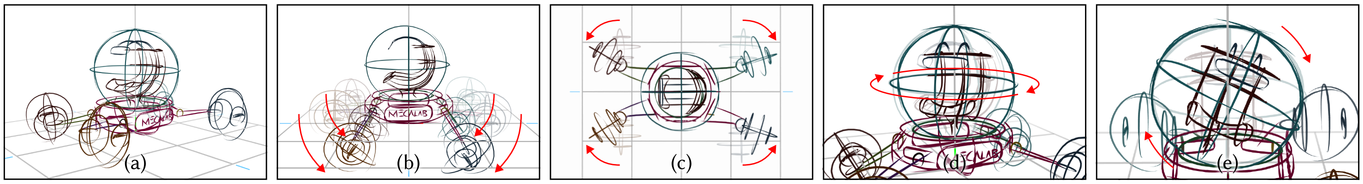

Fig. 1. (a) A lunar rover concept. It features (b, c) 4 omni wheels that swivel for high maneuverability, and (d, e) a pilot seat that rotates inside a spherical cockpit. The concept's designer, who had 11 years of work experience, estimated that it would take 7-8 hours to express and communicate such a simple articulated object with existing tools. With our system, it took only 45 minutes.

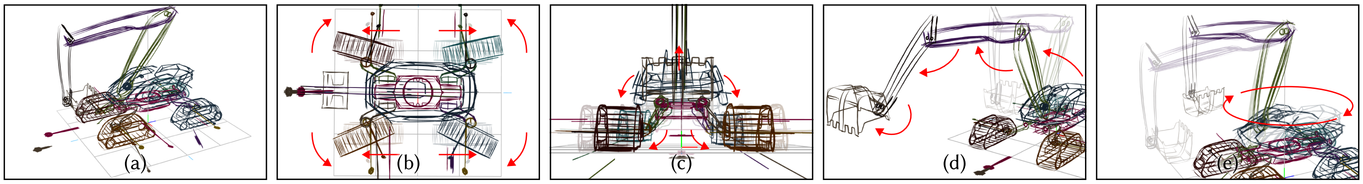

Fig. 2. (a) An autonomous excavator concept. It features (b, c) 4 caterpillars that swivel for high maneuverability, (d) an extendable boom and a bucket connected by multiple links, and (e) a rotating platform. The concept's designer, who had 8 years of work experience, estimated that it would take 1-2 weeks to express and communicate such a complex articulated object with existing tools. With our system, it took only 2 hours and 52 minutes.

Designing articulated objects is challenging because, unlike with static objects, it requires complex decisions to be made regarding the form, parts, rig, poses, and motion. We present a novel 3D sketching system for rapidly authoring concepts of articulated objects for the early stages of design, when designers make such decisions. Compared to existing CAD software, which focuses on slowly but elaborately producing models consisting of precise surfaces and volumes, our system focuses on quickly but roughly producing models consisting of key curves through a small set of coherent pen and multi-touch gestures. We found that professional designers could easily learn and use our system and author compelling concepts in a short time, showing that 3D sketching can be extended to designing articulated objects and is generally applicable in film, animation, game, and product design.

CCS Concepts: • **Human-centered computing** → **Interaction techniques**.

Additional Key Words and Phrases: 3D sketching, segmenting, rigging, posing, filming, bimanual interactions, pen, multi-touch

ACM Reference Format:
Joon Hyub Lee, Hanbit Kim, and Seok-Hyung Bae. 2022. Rapid Design of Articulated Objects. *ACM Trans. Graph.* 41, 4, Article 89 (July 2022), 8 pages. https://doi.org/10.1145/3528223.3530092

1 INTRODUCTION

In recent movies and games, articulated objects often steal the spotlight: a foldable drone that a hero in danger takes out of his pocket, an agile robot animal that roams over rough terrain, a futuristic car that rearranges its interior and exterior parts depending on the driving mode, and a multirole spaceship that transforms when engaging in battle or docking at a station. Many of these objects may even come into existence, thanks to impressive progress in technology.

However, it is highly difficult to design such articulated objects because they must look good and function well not only in one state but also in all the others, and the transitions between the various states should be appealing and coordinated. Meeting these goals requires complex decisions regarding the form, parts, rig, poses, and motions, each of which affects all the others.

Designers make most of these decisions during the early stages of design through repeated trial and error. However, existing CAD software is not the right tool for rapidly generating and discarding rough concepts because it requires them to elaborately construct articulated models with surfaces and volumes. Moreover, they must do so with a user interface that is best suited for precise fine-tuning at the later stages. Therefore, considerable time and effort is expended on initially unnecessary work.

In this paper, we present a lightweight alternative that relies on a small set of coherent pen and multi-touch gestures that resemble marking a few quick and loose pen strokes on paper and physically grabbing an object and moving it back and forth a few times in midair. With a proof-of-concept implementation using off-the-shelf components, we show that articulated concepts for early design decisions can be authored much more quickly (Fig. 1, 2).

Authors' addresses: Joon Hyub Lee, joonhyub.lee@kaist.ac.kr, Department of Industrial Design, KAIST, Republic of Korea; Hanbit Kim, hanbit.kim@kaist.ac.kr, Department of Industrial Design, KAIST, Republic of Korea; Seok-Hyung Bae, seokhyung.bae@kaist.ac.kr, Department of Industrial Design, KAIST, Republic of Korea.

0730-0301/2022/7-ART89 $15.00
https://doi.org/10.1145/3528223.3530092

2 RELATED WORK

We introduce previous works on curve-based 3D sketching, constraint-based multi-touch manipulation, segmentation-based rigging, keyframe-based animation, and systems with similar goals.

Zeleznik et al. [1996] proposed a system that incorporates the speed and expressiveness of sketch-based gestures to 3D modeling. Later, Igarashi et al. [1999], Igarashi and Hughes [2003], and Nealen et al. [2007] proposed systems for novices to author inflated 3D volumes using 2D curves. Conversely, Bae et al. [2008; 2009] proposed a system for authoring 3D NURBS curves through procedural interactions motivated by perspective drawing skills of professional designers, who are also our system's target users.

Subsequent studies found that letting designers define a 3D plane on which 2D pen strokes are projected leads to effective plane-based 3D sketching systems [Bae et al. 2009; Kim et al. 2018; Kim and Bae 2016]. In this study, we extend the 3D plane to propose a sketch plane widget that serves the dual role as a projection surface for creating 3D curves and a handle for moving them in space.

Pen and multi-touch complement each other well, particularly when the pen is used for authoring and the multi-touch is used for manipulation [Hinckley et al. 2010]. Although there are many approaches to mapping 2D multi-touch gestures to an object's 3D translation and rotation [Mendes et al. 2019], explicitly specifying the geometric constraint under which to operate has been found to enhance speed and precision [Au et al. 2012; Cohé et al. 2011] and is the approach we take in our system.

In contrast, VR lets designers view, author, and manipulate 3D curves directly in space through hand gestures in midair. However, the lack of a physical surface makes it difficult to do so precisely and causes fatigue [Arora et al. 2017]. Although many have tried to remedy this problem [Arora and Singh 2021; Hayatpur et al. 2019; Yu et al. 2021a,b], the screen-based input on which we rely may see continued use, especially in professional practice where precision and ergonomics are prioritized.

In making a sketch move, soft-body approaches in 2D [Igarashi et al. 2005] and 3D [Borosán et al. 2012; Dvorožňák et al. 2020; Jin et al. 2015] are suitable for organic forms. In contrast, rigid-body approaches [Davis et al. 2008; Kazi et al. 2014] are suitable for mechanical forms, but they require the sketch to be segmented and rigged and have not been attempted in 3D as in this study.

Although machine learning shows promise for automatically segmenting 2D sketches [Yang et al. 2021] or segmenting and rigging complete 3D scans of known objects [Li et al. 2020], we let designers manually segment and rig to give them fine-grained control over the 3D sketches that are often novel and incomplete.

With traditional keyframe-based animations [Baecker 1969], posing the model and timing the keyframes can be difficult and laborious to get right. Instead, an object's movement as it appears on the screen as a trajectory can be directly manipulated through sketch-based gestures [Choi et al. 2016; Guay et al. 2015; Thorne et al. 2004]. Similarly, in our system, the user interacts with the screen-space trajectory for intuitive temporal navigation [Dragicevic et al. 2008; Karrer et al. 2008] and pacing [Walther-Franks et al. 2012].

Finally, Shao et al. [2013] proposed a system for assisting in the reconstruction of articulated 3D models from multiple 2D sketches, but the system supported only simple 3D primitives. On the other hand, in a preliminary study [Lee et al. 2020], we surveyed 1,669 concepts from a renowned design contest, demonstrating that 38% of them had at least a hinge joint, linear slider, curved slider, or ball joint, and proposed a precursor system that showed promises in a pilot study. However, the precursor allowed only hinge joints, lacked support for animation, and required disparate interactions for various functions.

3 SYSTEM

We propose a system in which professional designers can rapidly switch between sketching, segmenting, rigging, posing, and filming (Fig. 3) to design articulated objects comprising hinge joints, linear and curved sliders, and ball joints. For high learnability and usability, it features a simple but powerful vocabulary of pen and multi-touch gestures motivated by bimanual interactions with physical objects.

The resulting articulated concepts can be exported in a video file format (e.g., MP4) for reviewing without a special application or in a model file format (e.g., FBX) to be consumed by existing CAD software as a reference in later stages of design. (See supplement for examples of exported sketches and follow-up surface modeling.)

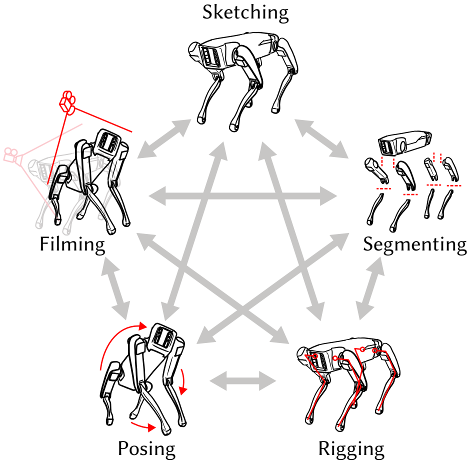

Fig. 3. System overview. The designer transitions back and forth between 5 types of operations: sketching, segmenting, rigging, posing, and filming.

3.1 Interactions & Interface

Interactions. We designed interactions around the way people intuitively handle bimanual tasks for them to be as quick and easy to learn and use as possible: While the nondominant hand (e.g., the left hand) defines the context, the dominant hand (e.g., the right hand) performs the action [Guiard 1987]. For example, while the left hand presses a ruler down or grabs a soda bottle (context), the right hand draws a straight line or rotates the cap (action). Similarly, in our system, while the left hand holds something down, the right hand drags something with fingers or marks it with a pen.

Interface. Our system consists of a small number of on-screen elements (Fig. 4) and functional buttons (b_e: erase, b_a: add to part, b_r: remove from part, b_c: capture keyframe, b_f: follow trajectory, b_{mod}: modifier). Most of its on-screen elements are either types of curve (l) or components that make up a sketch plane widget (π).

3.2 Sketch Plane Widget

Significantly simplifying our pen and multi-touch gesture set is the introduction of a novel sketch plane widget (π) that serves the dual role as a 3D projection surface for creating sketch curves (l_s) and a handle for spatially maneuvering those curves (Fig. 10-14).

The user can set up a sketch plane widget by defining position anchors (π_p), as in static 3D sketching systems [Kim et al. 2018; Kim and Bae 2016] (Fig. 5). However, unlike those systems, our system allows the user to bimanually interact with the widget to quickly and intuitively move it to the desired 3D position and orientation (Fig. 6): The user first selects either one of the axes (π_a), bezel (π_b), or center (π_c) with the left hand (L) and then drags the face (π_f) with 1 or 2 fingers of the right hand (R). For example:

While the left hand holds a sketch plane axis (π_a) down (Fig. 6a), the right hand drags the sketch plane face (π_f) with 1 finger to rotate the sketch plane widget (π) about the axis (π_a) (Fig. 6b) or with 2 fingers to translate the sketch plane widget (π) along the axis (π_a) (Fig. 6c). This description can be shortened as follows:

- L holds $\pi_a \cdots$ R drags π_f:
  - 1 finger $\longrightarrow$ rotate π about π_a;
  - 2 fingers $\longrightarrow$ translate π along π_a.

Similarly, the user can move the sketch plane widget with respect to the bezel (π_b) or center (π_c) through bimanual interactions:

- L holds $\pi_b \cdots$ R drags π_f:
  - 1 finger $\longrightarrow$ translate π on plane (Fig. 6d, e);
  - 2 fingers $\longrightarrow$ translate & rotate π on plane (Fig. 6d, f).
- L holds $\pi_c \cdots$ R drags π_f:
  - 1 finger $\longrightarrow$ orbit π about π_c (Fig. 6g, h);
  - 2 fingers $\longrightarrow$ orbit & spin π about π_c (Fig. 6g, i).

3.3 Sketching

The user can draw a 3D sketch curve by creating a sketch plane widget (Fig. 5), adjusting it (Fig. 6), and then marking a pen stroke that is projected onto it (Fig. 7a). The user can then erase entire or partial sketch curves, sketch plane position anchors, or the entire sketch plane widget through bimanual interactions:

- L holds $b_e \cdots$ R marks $l_s \longrightarrow$ erase l_s (Fig. 7b).
- L holds $b_e \cap b_{mod} \cdots$ R marks $l_s \longrightarrow$ partially erase l_s.
- L holds $b_e \cdots$ R marks $\pi_p \longrightarrow$ erase π_p.
- L holds $b_e \cdots$ R marks $\pi_a \longrightarrow$ erase π (Fig. 7c).

3.4 Segmenting

The user can segment entire or partial sketch curves into separate groups called parts that serve as links in the kinematic chain through bimanual interactions:

- L holds $b_a \cdots$ R marks $l_s \longrightarrow$ add l_s to part (Fig. 8a).
- L holds $b_a \cap b_{mod} \cdots$ R marks $l_s \longrightarrow$ partially add l_s to part.
- L holds $b_r \cdots$ R marks $l_s \longrightarrow$ remove l_s from part (Fig. 8b).
- L holds $b_r \cap b_{mod} \cdots$ R marks $l_s \longrightarrow$ partially remove l_s from part.

Performing the above without specifying a part creates a new one to which the marked curves are added. The user can also select an existing part to which all newly drawn curves are automatically added after entering the exploded view through a bimanual interaction:

- L holds $l_{s,\,part} \cdots$ R drags o:
  - 3 fingers $\longrightarrow$ enter exploded view centered on part (Fig. 9).

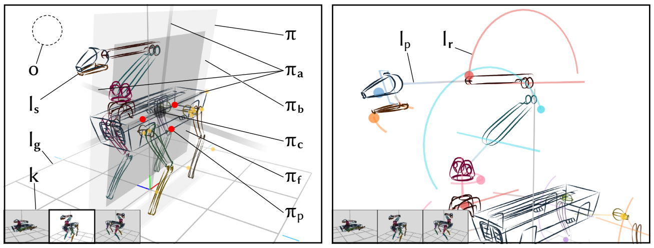

Fig. 4. On-screen elements. o: blank screen, π: sketch plane widget, π_a: sketch plane axis, π_b: sketch plane bezel, π_c: sketch plane center, π_f: sketch plane face, π_p: sketch plane position anchor, l_s: sketch curve, l_g: grid line, l_p: kinematic pair, l_r: joint range of motion, k: keyframe thumbnail.

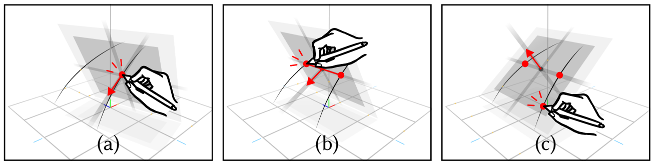

Fig. 5. Creating a sketch plane widget. Mark short pen strokes that intersect existing sketch curves or grid lines to specify 1-3 position anchors in 3D space that the plane should contain. The normal of the plane is determined as (a) the tangent of the sketch curve at 1 point, (b) the average of the tangents at 2 points, or (c) the uniquely possible normal for 3 points.

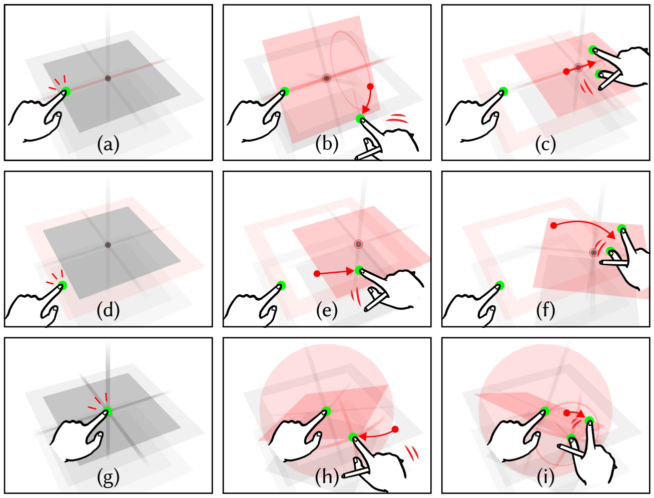

Fig. 6. Manipulating a sketch plane widget. (a) Hold the axis and (b) rotate the sketch plane widget about it with 1 finger or (c) translate along it with 2 fingers. (d) Hold the bezel and (e) translate the sketch plane widget on its plane with 1 finger or (f) translate and rotate on it with 2 fingers. (g) Hold the center and (h) orbit the sketch plane widget about it with 1 finger or (i) orbit and spin about it with 2 fingers. (See Supplement for details.)

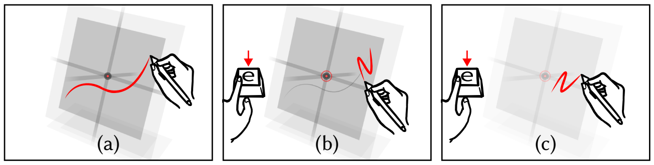

Fig. 7. Drawing and erasing. (a) Mark a pen stroke that is projected onto the sketch plane widget. Hold the erase button and mark the (b) sketch curve or (c) sketch plane widget to be erased. The pictogram on the left side indicates the button being held down.

3.5 Rigging

The user can create a joint by repeatedly demonstrating the desired motion of a part. Doing so leaves behind a trail, from which the system infers a joint's type, position, orientation, and range of motion. After strategically positioning and orienting a sketch plane widget, the user can perform bimanual interactions utilizing it:

- L holds $\pi_a \cdots$ R drags $\pi_f \cap l_{s,\text{ part}}$ repeatedly:
  - ├ 1 finger → create hinge joint (Fig. 10);
  - └ 2 fingers → create linear slider (Fig. 11a-c).
- L holds $\pi_b \cdots$ R drags $\pi_f \cap l_{s,\text{ part}}$ repeatedly:
  - ├ 1 finger → create linear slider (Fig. 11d-f);
  - └ 2 fingers → create curved slider (Fig. 12).
- L holds $\pi_c \cdots$ R drags $\pi_f \cap l_{s,\text{ part}}$ repeatedly:
  - └ 1 finger → create ball joint (Fig. 13).

The user can establish a kinematic pair so that a child part can move with respect to a parent part through a bimanual interaction:

- L holds $l_{s,\text{ parent}} \cdots$ R drags $l_{s,\text{ child}}$:
  - └ 1 finger → connect kinematic pair.

After rigging, the user can alter the kinematic chains in the exploded view through bimanual interactions:

- L holds $b_e \cdots$ R marks l_p → disconnect kinematic pair.
- L holds $b_e \cdots$ R marks l_r → erase joint.

3.6 Posing

The user can create desired poses by actuating joints between an ancestor part and a descendant part with forward kinematics (FK) or inverse kinematics (IK) through bimanual interactions:

- L holds $l_{s,\text{ ancestor}} \cdots$ R drags $l_{s,\text{ descendant}}$:
  - ├ 1 finger → actuate 1 joint with FK (Fig. 14a, b);
  - └ 2 fingers → actuate all connecting joints with IK (Fig. 14d).

The user can actuate many joints at the same time with FK by dragging many sketch curves that each belong to a different part (Fig. 14c). The user can perform IK that satisfies geometric constraints through bimanual interactions utilizing the sketch plane widget:

- L holds $\pi_a \cdots$ R drags $\pi_f \cap l_{s,\text{ part}}$:
  - ├ 1 finger → IK such that $l_{s,\text{ part}}$ rotates about π_a;
  - └ 2 fingers → IK such that $l_{s,\text{ part}}$ translates along π_a.
- L holds $\pi_b \cdots$ R drags $\pi_f \cap l_{s,\text{ part}}$:
  - ├ 1 finger → IK such that $l_{s,\text{ part}}$ translates on plane (Fig. 14e, f);
  - └ 2 fingers → IK such that $l_{s,\text{ part}}$ translates & rotates on plane.
- L holds $\pi_c \cdots$ R drags $\pi_f \cap l_{s,\text{ part}}$:
  - ├ 1 finger → IK such that $l_{s,\text{ part}}$ orbits about π_c;
  - └ 2 fingers → IK such that $l_{s,\text{ part}}$ orbits & spins about π_c.

For IK, the system calculates the pose solution that minimizes the screen-space distances between the touch points and contact points on the object [Reisman et al. 2009].

3.7 Filming

To showcase an articulated object, the user can compose a short animation from a sequence of keyframes. The user can create a keyframe that defines a desired pose seen from a desired viewpoint (Fig. 15a-c). The user can rearrange the order of a keyframe in a sequence, shown as a row of thumbnails, by dragging it to the left or the right or erase the keyframe by dragging it upward.

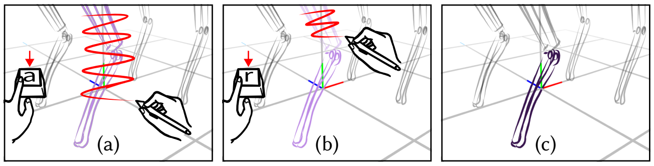

Fig. 8. Segmenting. (a) Hold the add button and mark sketch curves to be added to a part, or (b) hold the remove button and mark curves to be removed from a part (c) to create a separate part.

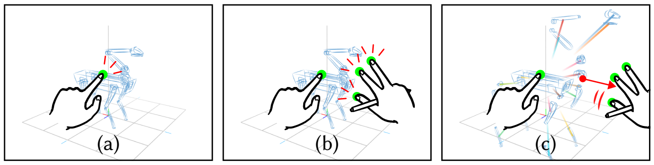

Fig. 9. Entering the exploded view. (a) Hold a part to keep it still and (b-c) pull other parts apart across the screen with 3 fingers.

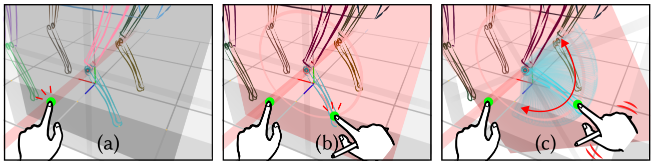

Fig. 10. Creating a hinge joint. (a) Hold an axis and (b) rotate a part about it with 1 finger (c) repeatedly, leaving behind a trail from which a corresponding joint is inferred.

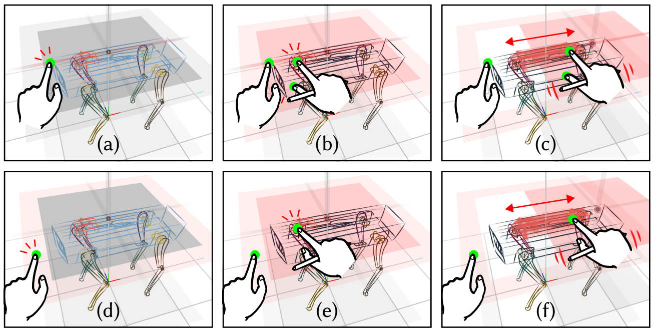

Fig. 11. Creating a linear slider. (a) Hold an axis and (b) translate a part along it with 2 fingers (c) repeatedly or (d) hold a bezel and (e) translate a part on its plane with 1 finger (f) repeatedly, leaving behind trails from which corresponding joints are inferred.

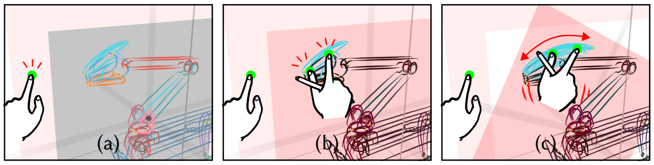

Fig. 12. Creating a curved slider. (a) Hold a bezel and (b) translate and rotate a part on its plane with 2 fingers (c) repeatedly, leaving a trail behind from which a corresponding joint is inferred.

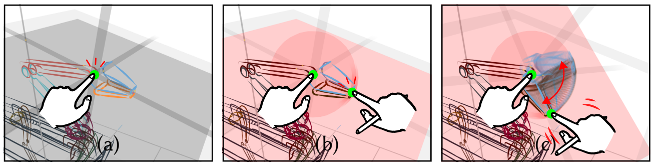

Fig. 13. Creating a ball joint. (a) Hold a center and (b) orbit a part about it with 1 finger (c) repeatedly, leaving a trail behind from which a corresponding joint is inferred.

The user can recall a particular pose and viewpoint by selecting a keyframe by touching on the thumbnail. When a keyframe is recalled, the user can create a motion connecting the sequence of keyframes through a bimanual interaction:

- L holds b_f ··· R drags l_s:
  └ 1 finger → create motion.

Doing so reveals the touched point's screen-space trajectory. The varying rate at which the user drags his or her finger along the trajectory determines the animation's playback speed [Dragicevic et al. 2008; Karrer et al. 2008; Walther-Franks et al. 2012] (Fig. 15d-f).

4 PROOF OF CONCEPT

To determine whether the intended users could learn to use the system in a short time and quickly create compelling results, we implemented our system using Unity 3D and BioIK [Starke et al. 2017] and recruited 5 expert designers (P1-5).

The designers had training in product design (P3, 5) or transportation design (P1, 2, 4) and an average of 6.2 years (min: 2 years; max: 11 years) of work experience. They had used Autodesk Alias (P1, 2, 4), Blender (P3, 4, 5), or Rhinoceros 3D (P3, 5) to design articulated objects in previous projects related to cars (P1, 2, 4), electronics (P3), transforming robot toys (P5), concept art (P1-5), or films and animations (P3, 5). All were experts in their respective fields, having won renowned international design awards (P1, 2, 4) or played prominent roles in multinational teams (P3, 5).

After an hour-long step-by-step tutorial, the designers sketched, segmented, rigged, posed, and filmed any concepts they imagined. To allow them to push the system to its limits, we did not impose any restrictions on the subject matter or time. Finally, they filled out a survey, and we interviewed them individually for approximately 1.5 hours. (See supplement for full interviews.)

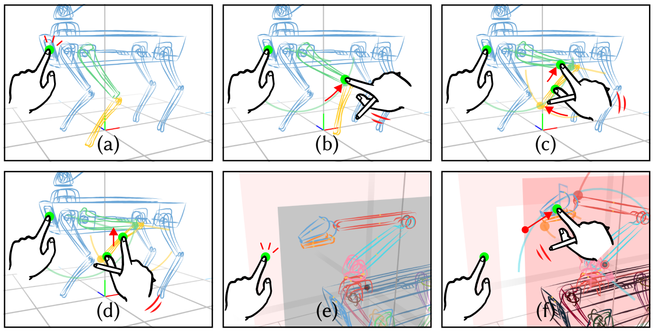

Fig. 14. Posing. For FK, (a) hold an ancestor part and move (b) single or (c) multiple descendant parts with 1 finger on each part. (d) For IK, hold an ancestor part and move a descendant part with 2 fingers on the same part. For IK with a planar constraint, (e) hold a bezel and (f) translate a part on its plane with 1 finger.

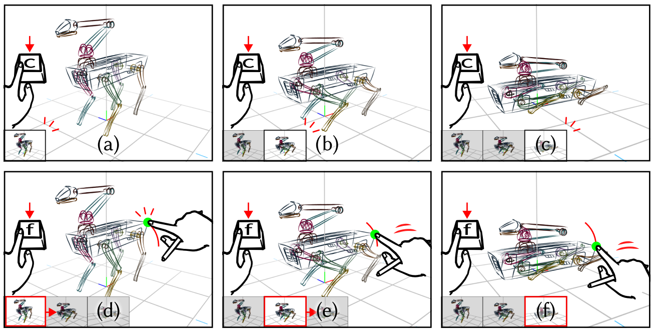

Fig. 15. Filming. (a-c) Add keyframes of desired poses and viewpoints by pressing the capture button. (d-f) To create a motion that connects the keyframes, recall a keyframe, hold the follow button, touch a sketch curve, and drag along the point's screen-space trajectory with 1 finger.

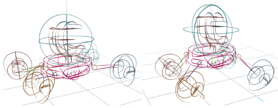

P1-C (0:45:13)
Lunar rover

Curve count		268
Part count		17
Joint count	Hinge	12
	Linear	0
	Curved	0
	Ball	0
	Total	12
Keyframe count		10

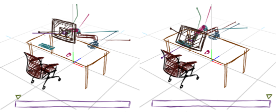

P2-B (1:07:07)
Designer workspace

Curve count		559
Part count		13
Joint count	Hinge	6
	Linear	2
	Curved	0
	Ball	0
	Total	8
Keyframe count		4

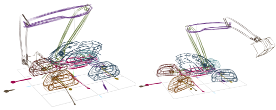

P2-C (2:52:53)
Autonomous excavator

Curve count		973
Part count		23
Joint count	Hinge	17
	Linear	5
	Curved	0
	Ball	0
	Total	22
Keyframe count		10

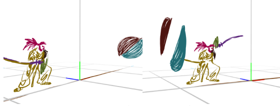

P3-B (0:27:56)
Samurai

Curve count		48
Part count		8
Joint count	Hinge	3
	Linear	2
	Curved	0
	Ball	0
	Total	5
Keyframe count		4

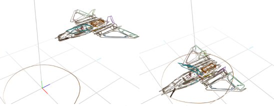

P3-C (2:10:09)
Spacecraft

Curve count		336
Part count		13
Joint count	Hinge	9
	Linear	3
	Curved	0
	Ball	0
	Total	12
Keyframe count		5

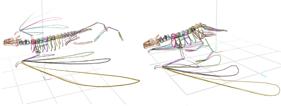

P3-E (0:55:28)
Skeleton dragon

Curve count		148
Part count		43
Joint count	Hinge	36
	Linear	0
	Curved	0
	Ball	0
	Total	36
Keyframe count		3

P4-A (1:09:07)
Assassin cyborg

Curve count		239
Part count		25
Joint count	Hinge	21
	Linear	2
	Curved	0
	Ball	0
	Total	23
Keyframe count		4

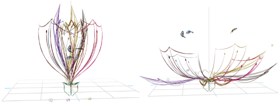

P4-B (1:01:00)
Flower and butterfly

Curve count		296
Part count		19
Joint count	Hinge	15
	Linear	0
	Curved	0
	Ball	0
	Total	15
Keyframe count		6

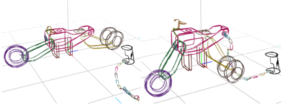

P4-F (1:05:20)
Electric tricycle

Curve count		191
Part count		25
Joint count	Hinge	5
	Linear	3
	Curved	0
	Ball	13
	Total	21
Keyframe count		4

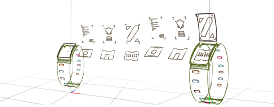

P4-G (0:23:12)
Holographic smartwatch

Curve count		167
Part count		16
Joint count	Hinge	11
	Linear	1
	Curved	0
	Ball	0
	Total	12
Keyframe count		6

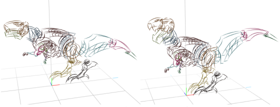

P5-A (0:59:11)
Mecha T-Rex

Curve count		444
Part count		17
Joint count	Hinge	14
	Linear	0
	Curved	0
	Ball	2
	Total	16
Keyframe count		2

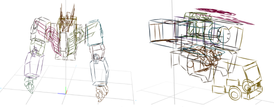

P5-B (1:11:09)
Transformer robot

Curve count		708
Part count		16
Joint count	Hinge	10
	Linear	5
	Curved	0
	Ball	0
	Total	15
Keyframe count		7

Fig. 16. Concepts of articulated objects (12 out of 20) that 5 professional designers (P1-5) authored using our system (time taken in parentheses).

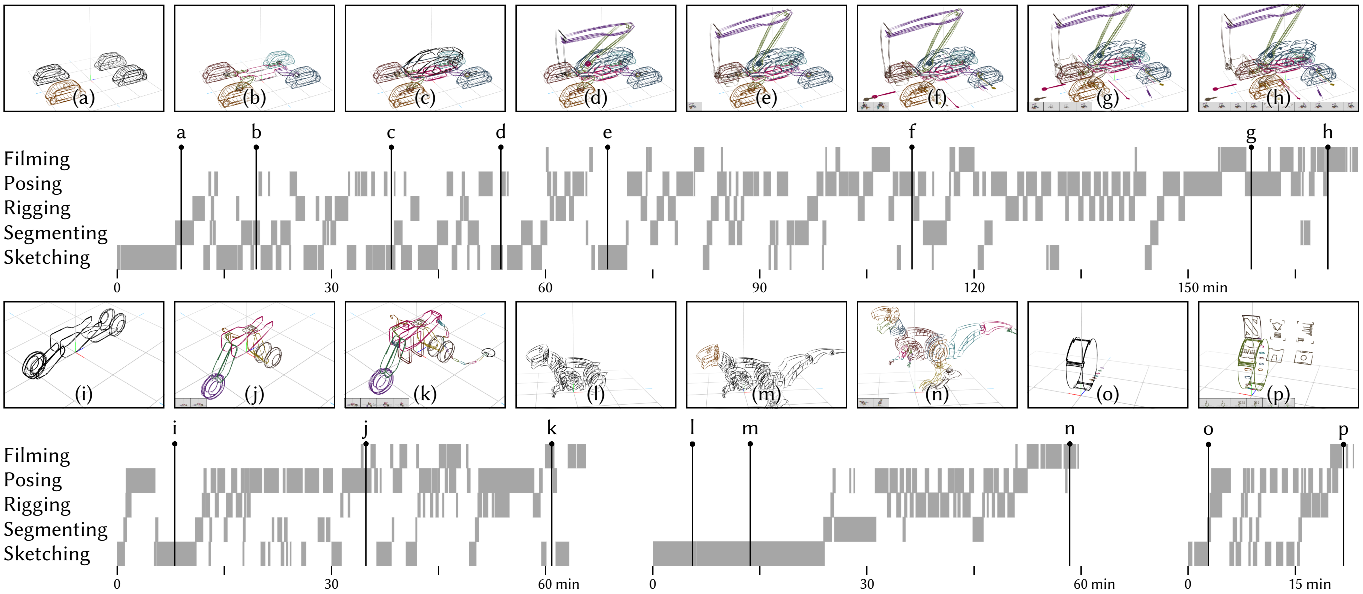

Fig. 17. Timelines of performed operations. The 5 bars represent the sketching, segmenting, rigging, posing, and filming operations in ascending order. The insets show the progress at the corresponding moments. Frequent switches between operation types call for an integrated system. (a-h) Autonomous excavator (P2-C). (i-k) Electric tricycle (P4-F). (l-n) Mecha T-Rex (P5-A). (o, p) Holographic smartwatch (P4-G). (See supplement for time-lapse footage of system usage.)

Collectively, they used our system for 23 hours and 53 minutes to author 20 articulated concepts (Fig. 16). (See supplement for all concepts.) On average, excluding the time taken for the tutorial, each concept took 1 hour and 12 minutes and comprised 333.1 curves (48-973), 18.6 parts (8-43), 15.4 joints (5-36), and 5.7 keyframes (2-12). On average, 12.6 hinge joints (81.5%), 1.9 linear sliders (12.0%), 0.05 curved sliders (0.3%), and 1.0 ball joints (6.2%) were used.

Easy to learn. The designers learned and masterfully used all the system's functions after only a 1-hour tutorial to produce concepts that were rough but creative, functional, and aesthetically pleasing, such as a wearable accessory, a flower, furniture, a samurai, a transformer robot, a vehicle, a spaceship, and a robot dinosaur.

In particular, they produced motions with dramatic rhythms simply by dynamically adjusting the speed at which they dragged their fingers along the screen-space trajectory, for example in sword wielding (P3-B) and a heroic landing (P4-A) (Fig. 16).

Such results represent a sharp reduction in learning time compared to existing modeling and animating tools for professional use, which typically require 1-6 months to learn the basics and 1.5-2 years to master, according to all designers (P1-5).

Rapid. As can be seen in the case where a holographic smartwatch concept (P4-G) was authored in 23 minutes with 16 parts, 12 joints, and impressive motion (Fig. 16, 17o, p), a fully articulated concept can be produced using our system as rapidly as when using existing systems for static concepts, which take anywhere from 36 [Kim et al. 2018] to 63 minutes [Bae et al. 2008] on average for a concept.

The designers noted that it would take much longer to convey the equivalent amount of information needed for design decision making with existing tools and processes, which typically involve multiple 2D sketches, 3D surface models, and presentation slides of key poses. They estimated that a simple object, such as the lunar rover (P1-C) (Fig. 1, 16), which took only 45 minutes using our system, would take 7-8 hours (P1), and that a complex object, such as the autonomous excavator (P2-C) (Fig. 2, 16, 17a-h), which took only about 3 hours using our system, would take 1-2 weeks (P2).

Iterative. The usage pattern revealed the considerable benefits of an integrated system (Fig. 17). On average, the designers switched 104.9 times (46-209) between sketching, segmenting, rigging, posing, and filming per concept and stayed in each for 41 seconds.

Notably, they kept revising the forms. In 14 cases (70%), they returned to sketching after the midpoint in time, and in 11 cases (55%), even during the last quarter. On average, they returned to sketching 11.2 times, most often immediately after they had finished making a new pose that necessitated a revision (5.9 times) (Fig. 18).

In existing workflows, performing different operations in different software can slow down the iteration pace. Reducing this friction is especially important during the early stages of design, in which designers spontaneously try many novel combinations of form, parts, rig, poses, and motion.

Satisfying. The interactions' average score (Fig. 19) was higher than neutral (Q1-32, 3.79). By type, the sketching (Q1-7, 3.77), segmenting (Q8-13, 4.07), rigging (Q14-24, 3.53), posing (Q25-28, 4.15), and filming (Q29-32, 4.85) scores were all higher than neutral.

The average score of the system's usefulness was also higher than neutral (Q33-38, 4.00). The designers liked that it requires only drawing skills and a few touch gestures (P2), believed that visualizing 3D forms and motions early in the process would help them prevent unexpected problems later (P2) especially in large-scale projects (P1), and looked forward to using our system at work (P4) and expanding their capacities as designers (P3).

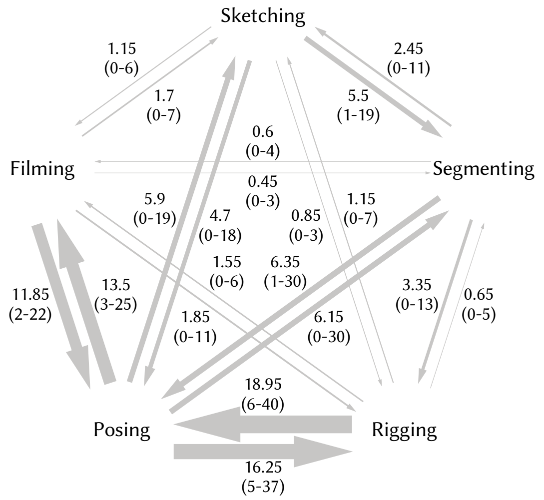

Fig. 18. Operation switch counts averaged over 20 concepts. The arrow's thickness is proportional to the count (min-max in parentheses). Switches to and from posing occurred most frequently, indicating that designers often try out different poses immediately before and after making changes to the form, parts, rig, and motion.

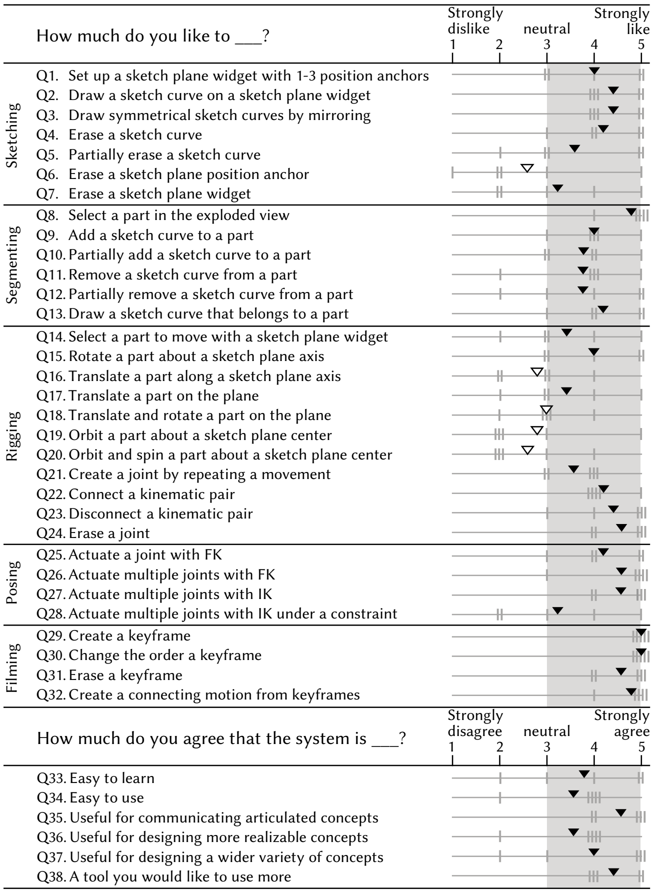

Fig. 19. The 5-point Likert scale survey on the system's usability and usefulness. Each tick corresponds to a designer's response. An average score of 3.94 (±2 SE: 0.94) indicates a satisfying overall experience.

5 LIMITATION & FUTURE WORK

We identify 3 limitations of this study and suggest future work to overcome them: the types of geometry the system supports, the direct manipulation interactions, and the duration and extent of the user study.

First, planar 3D curves are useful for expressing forms quickly and easily. However, by supporting surfaces and volumes, the system could assist in design decisions concerning physical factors, such as interference, collision, and dynamics [Liu and Popović 2002; Zhang et al. 2018], earlier, and be used to design more advanced apparatuses, such as legged robots, flying cars, and folding space telescopes.

Second, although directly touching and moving on-screen objects is intuitive and effective, it is sometimes challenging to do so when the screen is crowded with too many of them or when the hands block the screen or interfere with each other. The interactions that were more susceptible to these issues (Q6, 16, 18-20) received lower scores (Fig. 19) than others and hinted at opportunities for improvements.

Third, briefly working with professional designers showed us that our system could potentially alter the way many designers work in the future. However, long-term real-world studies in which we apply our system to actual film, animation, game, and product design are needed to evaluate the positive difference the system can make in the context of the entire production pipeline.

6 CONCLUSION

Designing articulated objects is difficult to the extent that something having "many moving parts" is synonymous with something being "too complicated and bound to go wrong." This is especially true in the absence of appropriate tools to try out various possibilities and quickly make decisions with far-reaching consequences during the early stages of design.

In this study, we proposed a lightweight system that facilitates frictionless transitions between 5 types of operations (sketching, segmenting, rigging, posing, and filming) to rapidly author concepts of articulated objects consisting of 4 types of joints (hinge joint, linear slider, curved slider, and ball joint), with only a small set of coherent pen and multi-touch gestures.

For proof of concept, we recruited 5 professional designers with 6.2 years of work experience on average, and we collected 24 hours of system usage, 20 concepts of articulated objects, scores for 38 system-related questions, and 7.5 hours of interviews. We found that our system can be learned and mastered in roughly 1 hour, facilitate a satisfying back and forth workflow, and help designers rapidly author concepts suitable for early-stage design decisions in much less time than existing tools and processes.

ACKNOWLEDGMENTS

This research was supported by the National Research Foundation of Korea (NRF) (2020X1A3A1098434). We thank Jeongche Yoon, Junghwan Kwak, Sangyun Kim, Hyewon Seo, Youngbin Jeon, Jaeho Sung, Donghyeok Ma, Sang-Hyun Lee, Junwoo Yoon, Taegyu Jin, Seung-Jun Lee, and Haena Cho for their help. We also thank Klio Design, Sketchsoft Inc., and God Brave Studio for hosting us.

REFERENCES

Rahul Arora, Rubaiat Habib Kazi, Fraser Anderson, Tovi Grossman, Karan Singh, and George Fitzmaurice. 2017. Experimental evaluation of sketching on surfaces in VR. In *Proc. CHI '17*. 5643–5654. https://doi.org/10.1145/3025453.3025474

Rahul Arora and Karan Singh. 2021. Mid-air drawing of curves on 3D surfaces in virtual reality. *ACM Trans. Graph.* 40, 3, Article 33 (Jul 2021), 17 pages. https://doi.org/10.1145/3459090

Oscar Kin-Chung Au, Chiew-Lan Tai, and Hongbo Fu. 2012. Multitouch gestures for constrained transformation of 3D objects. *Comput. Graph. Forum* 31, 2 (2012), 651–660. https://doi.org/10.1111/j.1467-8659.2012.03044.x

Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2008. ILoveSketch: as-natural-as-possible sketching system for creating 3D curve models. In *Proc. UIST '08*. 151–160. https://doi.org/10.1145/1449715.1449740

Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2009. EverybodyLovesSketch: 3D sketching for a broader audience. In *Proc. UIST '09*. 59–68. https://doi.org/10.1145/1622176.1622189

Ronald M. Baecker. 1969. Picture-driven animation. In *Proc. AFIPS '69 (Spring)*. 273–288. https://doi.org/10.1145/1476793.1476838

Péter Borosán, Ming Jin, Doug DeCarlo, Yotam Gingold, and Andrew Nealen. 2012. RigMesh: automatic rigging for part-based shape modeling and deformation. *ACM Trans. Graph.* 31, 6, Article 198 (Nov 2012), 9 pages. https://doi.org/10.1145/2366145.2366217

Byungkuk Choi, Roger Blanco i Ribera, J. P. Lewis, Yeongho Seol, Seokpyo Hong, Haegwang Eom, Sunjin Jung, and Junyong Noh. 2016. SketchiMo: sketch-based motion editing for articulated characters. *ACM Trans. Graph.* 35, 4, Article 146 (Jul 2016), 12 pages. https://doi.org/10.1145/2897824.2925970

Aurélie Cohé, Fabrice Dècle, and Martin Hachet. 2011. tBox: a 3D transformation widget designed for touch-screens. In *Proc. CHI '11*. 3005–3008. https://doi.org/10.1145/1978942.1979387

Richard C. Davis, Brien Colwell, and James A. Landay. 2008. K-Sketch: a "kinetic" sketch pad for novice animators. In *Proc. CHI '08*. 413–422. https://doi.org/10.1145/1357054.1357122

Pierre Dragicevic, Gonzalo Ramos, Jacobo Bibliowitcz, Derek Nowrouzezahrai, Ravin Balakrishnan, and Karan Singh. 2008. Video browsing by direct manipulation. In *Proc. CHI '08*. 237–246. https://doi.org/10.1145/1357054.1357096

Marek Dvorožňák, Daniel Sýkora, Cassidy Curtis, Brian Curless, Olga Sorkine-Hornung, and David Salesin. 2020. Monster mash: a single-view approach to casual 3D modeling and animation. *ACM Trans. Graph.* 39, 6, Article 214 (Nov 2020), 12 pages. https://doi.org/10.1145/3414685.3417805

Martin Guay, Rémi Ronfard, Michael Gleicher, and Marie-Paule Cani. 2015. Space-time sketching of character animation. *ACM Trans. Graph.* 34, 4, Article 118 (Jul 2015), 10 pages. https://doi.org/10.1145/2766893

Yves Guiard. 1987. Asymmetric division of labor in human skilled bimanual action. *J. Mot. Behav.* 19, 4 (1987), 486–517. https://doi.org/10.1080/00222895.1987.10735426

Devamardeep Hayatpur, Seongkook Heo, Haijun Xia, Wolfgang Stuerzlinger, and Daniel Wigdor. 2019. Plane, ray, and point: enabling precise spatial manipulations with shape constraints. In *Proc. UIST '19*. 1185–1195. https://doi.org/10.1145/3332165.3347916

Ken Hinckley, Koji Yatani, Michel Pahud, Nicole Coddington, Jenny Rodenhouse, Andy Wilson, Hrvoje Benko, and Bill Buxton. 2010. Pen + touch = new tools. In *Proc. UIST '10*. 27–36. https://doi.org/10.1145/1866029.1866036

Takeo Igarashi and John F. Hughes. 2003. Smooth meshes for sketch-based freeform modeling. In *Proc. I3D '03*. 139–142. https://doi.org/10.1145/641480.641507

Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. 1999. Teddy: a sketching interface for 3D freeform design. In *Proc. SIGGRAPH '99*. 409–416. https://doi.org/10.1145/311535.311602

Takeo Igarashi, Tomer Moscovich, and John F. Hughes. 2005. As-rigid-as-possible shape manipulation. *ACM Trans. Graph.* 24, 3 (Jul 2005), 1134–1141. https://doi.org/10.1145/1073204.1073323

Ming Jin, Dan Gopstein, Yotam Gingold, and Andrew Nealen. 2015. AniMesh: interleaved animation, modeling, and editing. *ACM Trans. Graph.* 34, 6, Article 207 (Oct 2015), 8 pages. https://doi.org/10.1145/2816795.2818114

Thorsten Karrer, Malte Weiss, Eric Lee, and Jan Borchers. 2008. DRAGON: a direct manipulation interface for frame-accurate in-scene video navigation. In *Proc. CHI '08*. 247–250. https://doi.org/10.1145/1357054.1357097

Rubaiat Habib Kazi, Fanny Chevalier, Tovi Grossman, and George Fitzmaurice. 2014. Kitty: sketching dynamic and interactive illustrations. In *Proc. UIST '14*. 395–405. https://doi.org/10.1145/2642918.2647375

Yongkwan Kim, Sang-Gyun An, Joon Hyub Lee, and Seok-Hyung Bae. 2018. Agile 3D sketching with air scaffolding. In *Proc. CHI '18*. Article 238, 12 pages. https://doi.org/10.1145/3173574.3173812

Yongkwan Kim and Seok-Hyung Bae. 2016. SketchingWithHands: 3D sketching handheld products with first-person hand posture. In *Proc. UIST '16*. 797–808. https://doi.org/10.1145/2984511.2984567

Joon Hyub Lee, Hyung-Gi Ham, and Seok-Hyung Bae. 2020. 3D sketching for multi-pose products. In *CHI EA '20*. Article 261, 8 pages. https://doi.org/10.1145/3334480.3382870

Xiaolong Li, He Wang, Li Yi, Leonidas J. Guibas, A. Lynn Abbott, and Shuran Song. 2020. Category-level articulated object pose estimation. In *Proc. CVPR '20*. 3706–3715. https://doi.org/10.1109/CVPR42600.2020.00376

C. Karen Liu and Zoran Popović. 2002. Synthesis of complex dynamic character motion from simple animations. *ACM Trans. Graph.* 21, 3 (Jul 2002), 408–416. https://doi.org/10.1145/566654.566596

Daniel Mendes, Fabio Marco Caputo, Andrea Giachetti, Alfredo Ferreira, and Joaquim Jorge. 2019. A survey on 3D virtual object manipulation: from the desktop to immersive virtual environments. *Comput. Graph. Forum* 38, 1 (2019), 21–45. https://doi.org/10.1111/cgf.13390

Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. 2007. FiberMesh: designing freeform surfaces with 3D curves. In *Proc. SIGGRAPH '07*. 41–es. https://doi.org/10.1145/1275808.1276429

Jason L. Reisman, Philip L. Davidson, and Jefferson Y. Han. 2009. A screen-space formulation for 2D and 3D direct manipulation. In *Proc. UIST '09*. 69–78. https://doi.org/10.1145/1622176.1622190

Tianjia Shao, Wilmot Li, Kun Zhou, Weiwei Xu, Baining Guo, and Niloy J. Mitra. 2013. Interpreting concept sketches. *ACM Trans. Graph.* 32, 4, Article 56 (Jul 2013), 10 pages. https://doi.org/10.1145/2461912.2462003

Sebastian Starke, Norman Hendrich, Dennis Krupke, and Jianwei Zhang. 2017. Evolutionary multi-objective inverse kinematics on highly articulated and humanoid robots. In *Proc. IROS '17*. 6959–6966. https://doi.org/10.1109/IROS.2017.8206620

Matthew Thorne, David Burke, and Michiel van de Panne. 2004. Motion doodles: an interface for sketching character motion. *ACM Trans. Graph.* 23, 3 (Aug 2004), 424–431. https://doi.org/10.1145/1015706.1015740

Benjamin Walther-Franks, Marc Herrlich, Thorsten Karrer, Moritz Wittenhagen, Roland Schröder-Kroll, Rainer Malaka, and Jan Borchers. 2012. DRAGIMATION: direct manipulation keyframe timing for performance-based animation. In *Proc. GI '12*. 101–108.

Lumin Yang, Jiajie Zhuang, Hongbo Fu, Xiangzhi Wei, Kun Zhou, and Youyi Zheng. 2021. SketchGNN: semantic sketch segmentation with graph neural networks. *ACM Trans. Graph.* 40, 3, Article 28 (Aug 2021), 13 pages. https://doi.org/10.1145/3450284

Emilie Yu, Rahul Arora, Tibor Stanko, J. Andreas Bærentzen, Karan Singh, and Adrien Bousseau. 2021a. CASSIE: curve and surface sketching in immersive environments. In *Proc. CHI '21*. Article 190, 14 pages. https://doi.org/10.1145/3411764.3445158

Xue Yu, Stephen DiVerdi, Akshay Sharma, and Yotam Gingold. 2021b. ScaffoldSketch: accurate industrial design drawing in VR. In *Proc. UIST '21*. 372–384. https://doi.org/10.1145/3472749.3474756

Robert C. Zeleznik, Kenneth P. Herndon, and John F. Hughes. 1996. SKETCH: an interface for sketching 3D scenes. In *Proc. SIGGRAPH '96*. 163–170. https://doi.org/10.1145/237170.237238

He Zhang, Sebastian Starke, Taku Komura, and Jun Saito. 2018. Mode-adaptive neural networks for quadruped motion control. *ACM Trans. Graph.* 37, 4, Article 145 (Jul 2018), 11 pages. https://doi.org/10.1145/3197517.3201366