

Squidgets: Sketch-based Widget Design and Direct Manipulation of 3D Scenes

JOONHO KIM, University of Toronto, Canada
KARAN SINGH, University of Toronto, Canada




Fig. 1. Humans naturally sketch strokes (cyan) over graphical scenes to communicate desired scene changes (a,b). These strokes when associated with a curve abstraction of the scene (c) can visually quantify the desired change to scene attributes to match the user’s sketch (d). Squidgets or ‘sketched widgets’ are scene curves users draw to directly manipulate a graphical scene (e,f). Squidgets can be inherent to the scene (c) or user created as curve widgets to explicitly control a set of scene attributes (g). User-authored in-situ curve controllers also integrate seamlessly into the Squidget interaction framework (h).

Squidgets or ‘sketch-widgets’ is a novel stroke-based UI framework for direct scene manipulation. Squidgets is motivated by the observation that sketch strokes comprising visual abstractions of scene elements implicitly provide natural handles for the direct manipulation of scene parameters. Configurations of such strokes can further be explicitly drawn by users to author custom widgets associated with scene attributes. Users manipulate a scene by simply drawing strokes: a squidget is selected by partially matching the drawn stroke against both implicit scene contours and explicitly authored curves, and used in-situ to interactively control scene parameters associated with the squidget. We present an implementation of squidgets within the 3D modeling animation system *Maya*, and report on an evaluation of squidget creation and manipulation, by both casual users and professional artists.

1 INTRODUCTION

Sketching, is a well established tool for visual communication. In particular, sketch-based techniques are ubiquitous in interactive design and graphical content creation. User control of scene objects and attributes/parameters in such applications however, is typically a disparate combination of traditional UI components, pre-defined widgets and stroke gestures [SSB08]. While there have been a number of seminal sketch-based systems for content creation spanning over half a century [Sut64, ZHH96, IMT99, BBS08, XCS⁺14], their operational interface has largely relied on pre-defined stroke gestures and traditional UI elements.

Pen, tablet and stroke-based interfaces are also increasingly used to control general computing applications. A vein of research has thus adapted the interaction of components like menus/buttons

[AG04] and 3D widgets [SSB08], from a point-and-click to stroke-friendly design. Inspired by recent research [PCS23], that uses sketchy-renderings of objects for interactive visualization and exploration, we seek to exploit such renderings for general stroke-based scene interaction.

We thus present Squidgets (*sketch-widgets*), as a novel stroke-based UI framework for direct scene manipulation (Figure 1). Our work uses the insight that graphical scenes inherently possess implicit in-situ handles for scene manipulation, that are based on a user imagined visual abstraction of the scene (Figure 1(a-c)). Humans naturally sketch strokes over such imagined scene abstractions to visually indicate desired changes to the scene. We aim to capture this natural interaction computationally, using squidgets.

Conceptually this is an inverse rendering problem: compute a minimal change to the current scene attribute values that will result in a manipulated scene, where the over-sketched stroke matches some part of the scene’s curve abstraction (Figure 1(c,d)).

The benefits of solving such a problem are twofold. First, the over-sketching interaction is natural, direct, and in-situ for 2D or 3D applications. Second, no explicit gestures, widgets or UI components for the scene need to be defined or learned by the user.

There are however, three non-trivial and potentially ambiguous aspects to solving this ambitious problem: inferring the user-imagined curve abstraction of a graphical scene; discerning what part of this curve abstraction to associate with the over-sketched stroke; and computing the set of scene attribute value changes that manifest in a scene whose associated visual abstraction best matches the over-sketched stroke.

The squidgets framework allows each of these sub-problems to be explored, constrained and addressed independently. Inherent stroke-based visual abstractions of scenes have long been studied in the non-photorealistic rendering (NPR) literature [GG01]. Such abstractions can be defined using occluding and other suggestive contours in a scene [Her99] (Section 4), or using image-based approaches to differential rendering [LMLH07]. Given a user stroke and any curve based visual scene abstraction, we define a curve similarity metric (Section 4) that is able to match the stroke to (partial) curves using a mix of spatial and shape proximity. This enables the association of user strokes to a curve segment from a visual scene abstraction. Finally, scene attributes that affect the visual appearance of the associated curve segment are sampled around their present values to edit the scene to visually match the user stroke.

In addition to implicit squidget curves suggested by the visual scene abstraction, squidget curves can be explicitly authored as custom widgets to control scene attributes within our framework (Figure 1(g,h)). Explicitly authored squidget curves address two of the three sub-problems: the curves explicitly define the user imagined visual scene abstraction; and since squidget curve configurations are explicitly co-related to one or more scene attribute values by simple invertible functions, finding attribute values to match an over-sketched stroke is straightforward.

We provide an overview of related work (Section 2), followed by details of our squidget framework (Section 3). Section 4 presents our approach to stroke matching, attribute inference and other implementation details of a squidget interaction prototype built within the modeling and animation system *Maya*. Section 5 presents a number of applications and squidget usage scenarios. We discuss the outcomes of a user study (Section 6) followed by limitations and avenues for future work on squidgets (Section 7).

2 RELATED WORK

Squidgets relate to techniques for direct, interactive scene manipulation. A large body of research spanning 50 years [Sut64], has specifically addressed the transformation (eg. translation/rotation/scaling) of 3D scene objects in scenes. We refer readers to a recent survey paper [MCG⁺19]) tracing this research from desktop to immersive devices. Techniques designed for

constrained tasks like docking [Bie90], have also been compared across various input modalities [BIAI17]. Object manipulation techniques interact with a 3D spatial context that can be physical (such as tangible 3D proxy), or virtual (like mouse controlled 3D widgets) [BKLP04]. Squidget curves provide our spatial context, controlled by 2D (Figure 1) or 3D (Figure 11) stroke input.

We further note, that transforming 3D objects with squidgets (Figure 7) is but one example, of our ability to homogeneously manipulate arbitrary scene object attributes using squidget strokes (Figure 4(c),(d)). In this regard our work is related to the paradigms of *sketch-based gestures*, *scene proxies*, *custom widgets*, and *direct in-situ manipulation*.

2.1 Sketch-based gestures

The recognition and use of sketch or stroke-based gestures can be roughly categorized into gestures that are hard-coded or visually matched[JGHD09]. Hard-coded gestures tend to be expressive, but context specific and difficult to customize. Visual matching, such as the simple \$1 Gesture Recognizer [WWL07], geometrically match a given strokes to a user-given set of gesture templates. These gestures are typically scene agnostic and can be mapped to general directives (eg. invoking undo using a scratch gesture). Similar to visually matched gestures, squidgets match user drawn strokes to the set of squidget curves in a graphical scene.

2.2 Scene proxies and custom widgets

UI widgets are visual 3D elements designed to provide an in-situ interface to manipulating objects and aspects of a virtual scene [Bie86]. Widgets are typically hand-designed to capture the form/function of scene attributes they control, and can range from simple spatial transform widgets, to custom controllers specific to complex objects such as the human face in Figure 1(h). Point-click-and-drag interaction of widgets can be improved using stroke-based techniques [SSB08]. Our squidget framework is able to homogeneously interact using over-sketching with such curve-based widgets Figure 1(h).

Curves have also been used as visual proxies and manipulation handles for deformable objects [SF98, NSACO05]. Interaction with these curves are manifested as a deformation of the object vertices. Squidgets are inspired by such techniques, and generalize their control beyond deformable objects (eg. the mug in Figure 4(c)) to any scene elements whose visual appearance can be controlled by some set of attributes (eg. the cone-angle of a spotlight in Figure 4(c)).

2.3 Direct in-situ manipulation

In-situ visualization and control of object attributes in a scene can greatly streamline a sketch-based workflow [XAGW16]. Squidgets take such a design further, allowing user to directly sketch the visual change they expect as a result of changes in object attribute values.

Such direct visual control requires an inverse mapping from the visual abstraction of objects and other scene elements to their attribute values. Inverse computation for direct control is popular in a number of domains on Computer Graphics such as skeletal kinematics (IK) [ALCS18, GCR13], facial expressions [LA10, KS21], CAD modeling [MB21], and animation [IMH07, CiRL⁺16]. Squidgets are inspired by these solutions from specific graphical contexts to built a general framework for manipulating scene attributes using sketch-strokes.

2.4 Non-photorealistic Rendering (NPR)

Finally, the fundamental insight behind squidgets is that imagined curves in a scene provide natural manipulation handles. Crucial to understanding and algorithmically computing such implicit curves that comprise a visual scene abstraction, is grounded in non-photorealistic rendering research.

There is a rich body of literature in Computer Graphics [GG01] and perceptual psychology [Tod04] on understanding and creating salient stroke-based NPR abstractions of graphical scenes. This sketch-stroke abstraction can be computed for 2D/3D shape automatically as shape features like ridges/valleys and view-dependent contour lines from 2D images [CV01] or 3D models [BH19]. The extracted strokes can be segmented into perceptual parts, for example using a corner detection technique [WEH08]. We rely on this body of work to compute a set of strokes that define a visual sketch abstraction for 2D or 3D graphical scenes.

3 SQUIDGETS




Fig. 2. Implicit Squidgets comprise a set of perceived curves (2D or 3D) that are already scene curves or computed using an NPR contour rendering of scene objects. In the default *control mode*, user drawn strokes are matched against all scene squidget curves (explicit and implicit), scene attributes corresponding to the closest matched squidget are then computed and set so the squidget curve best-fits the user stroke, and further be interactively dragged to refine the associated attributes values in real-time.

Squidgets serve as visual proxies and manipulation handles to scene attributes, facilitating rapid referencing and direct visual adjustment of scene attributes. A **squidget** comprises a curve element C and a corresponding state or values of a set of attributes A . Changes to the values of A result in a visual change to C in the scene. Squidgets explore the idea that scene attributes A , can be naturally manipulated by directly specifying the desired visual change to an associated curve C .

Formally, let a function f map attribute values A to a squidget curve $C = f(A)$. Manipulating a scene with squidgets requires computing the inverse mapping. Given an user stroke S , we seek attribute values A' , such that the edited squidget curve ($C' = f(A')$) $\approx S$. For example in Figure 1(c), drawing the cyan user stroke S , results in re-computing the translation attribute of the red squidget curve C (and associated lamp object), so that the translated squidget curve C' matches S (Figure 1(d)).

3.1 Squidget Classification

Every visually perceived curve in a scene is a potential squidget. As shown in Figure 4, perceived curves include both *real* rendered scene curves, and *virtual* curves that can be visually imagined from the scene context.

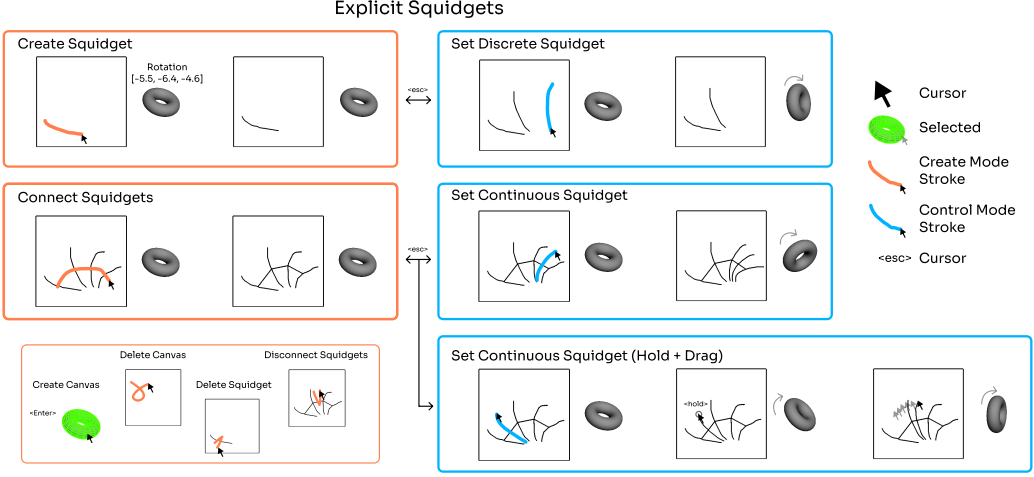


Fig. 3. Explicit squidgets comprise a set of curves (2D or 3D) that are explicitly authored in (*create mode*). During creation, curves are drawn (on planar canvases or directly in 3D) and associated with selected scene object attributes, to form discrete squidget curves, which serve as a bookmark for the configuration of attribute values when created. Multiple discrete squidgets can be connected to define a continuous squidget that interpolates the curves, as well as the union of their bookmarked attribute values. In *control mode*, user drawn strokes on canvases or 3D objects are matched against explicit and implicit squidgets and the closest matching squidget is selected to manipulate its associated attributes.




Fig. 4. All *perceived* scene curves are potential squidgets including: *real* curves (left) that are part of the scene’s construction (green), or explicitly authored squidgets (white); and *virtual* curves (center) imagined by a user (highlighted in red). A user draw stroke (cyan) can manipulate scene attributes (right), such that their associated squidget curve matches the user stroke.

Implicit Squidgets: We refer to curves (both *real* and *virtual*) that are already part of any given scene as *implicit* squidgets. Graphical scenes often include objects constructed using *real* curves, such as the orange mug with an extruded circular profile (green) in Figure 4(a). We also perceive boundaries of high visual contrast in a scene as *virtual* curves [BH19]. Such curves including boundaries, ridge and valley curves, silhouettes and other suggestive contours on objects (red curves on mug in Figure 4(b)) are well-known in NPR literature and easy to automatically compute from a given scene. Curves such as the elliptical boundary marking the cone angle of a spotlight on the plane in Figure 4(b) can also present a virtual squidget.

Implicit squidgets are thus visually inherent curve handles associated with scene objects and attributes that a user can exploit to manipulate a scene. Note, that while there is no a-priori coherence between virtual curves and scene attributes (i.e. NPR contours can change discontinuously with

small changes in associated scene attributes), as manipulation handles, we can treat virtual squidgets as curves embedded on scene object(s), that deform and change with these objects and attributes.

Explicit Squidgets: We capture the increasingly popular approach to hand-crafting *in-situ* curve controllers (Figure 12), by allowing users to "sketch widgets" as *explicit* squidget curves. A curve C can be explicitly drawn using our framework (Figure 3) and associated with the values of any selected set of attributes A , serving as an *in-situ* scene bookmark for a discrete configuration of values of A . A number of such discrete squidgets can further, be combined in a sequence to define a continuous interpolated squidget. For example, two discrete squidgets C_0, A_0 and C_1, A_1 can be combined into a weight $w \in [0, 1]$ interpolated, continuous squidget curve $\text{interp}(C_0, C_1, w)$ (with $C_0 = \text{interp}(C_0, C_1, 0)$, $C_1 = \text{interp}(C_0, C_1, 1)$), and similarly interpolated attributes $\text{interp}(A_0, A_1, w)$ for a set A' that could be the union or intersection of attribute sets A_0, A_1 .




Fig. 5. Explicitly created squidget curves can be manifest on surfaces (a), canvases in space (b-c), or using VR controls (d).

Any modeling and selection tools can be used to create an explicit *discrete squidget* curve C and associated scene attributes A , for example, curves can be drawn directly projected on a desired scene object (Figure 5(a)), or directly in 3D in VR (Figure 5(d)). We also facilitate the drawing of curves on planar canvases in 3D (Figure 5(b-c)). 3D drawing canvases are popular [DXS⁺07], as they allow constrained and predictable planar curves to be drawn in 3D, and provide a natural visual grouping of collections of squidget curves and scene object attributes that they control (Figure 5(c)).

Users can further connect multiple discrete squidgets in sequence into a *continuous squidget* (Figure 3). Continuous squidgets are visually indicated by a piecewise linear path connecting the center points of the discrete squidget curves.



Fig. 6. Explicitly authored curves can be associated with discrete configurations of face attributes to capture a smile (left) and a laugh (right); and further combined into a continuous squidget, where the curves and related face attributes can be interpolated (center).

As an illustration, 3D faces are often controlled by many (typically over 30) psuedo-muscle attributes [KS21]. An expressive smile and laugh (Figure 6(left and right)) each produced by a configuration of a subset of these attributes, can be discretely associated with explicitly authored squidget curves (white). These curves can be combined into a continuous squidget that captures an animated expression, as the two curves and their attributes are interpolated (Figure 6(center)).

In general, a squidget is thus a family of curve shapes with a corresponding family of scene attribute values, that enable *in-situ* user interaction with scene attributes.

3.2 Squidget Interaction and Control

Users interact with the graphical scene by simply drawing strokes. Strokes perform two functions: selecting a squidget from the set of scene squidgets (based on curve similarity Section 4.4); and manipulating the scene attributes for the selected squidget, so it best matches the user stroke. For example, the cyan stroke in Figure 1(b) is used to first select the implicit red squidget (due to strong curve shape similarity); then the squidget and associated lamp are translated so the red curve best matches the cyan input stroke.

For small or incremental changes to a widget curve a single user stroke Figure 1(b) can suffice to both select and manipulate the squidget. For large changes to squidgets in a scene with many squidgets, it may be necessary to draw two strokes: an approximately replicating stroke to unambiguously select a desired squidget; and a subsequent manipulation stroke. In Figure 4(c) for example, the cyan manipulation stroke is closer in shape to the handle (white curve) of the mug, than to the silhouette (red curve), and would thus require a prior selection stroke to deform the mug silhouette curve. This is easily accomplished using a modifier while drawing a stroke to indicate that it is the first of a pair of strokes, where the stroke being drawn will be used for selection, and will be followed by a manipulation stroke. Unless specifically stated however, we assume the more efficient workflow where a single stroke is used to both select and manipulate a squidget.

The stroke is matched for similarity (Section 4.4) against both implicit and explicitly authored squidgets. The best matching squidget curve is selected. Optionally, nothing can be selected if all matching scores are poor and below a given threshold.

Once a squidget is selected, its associated attributes need to be set to match the user stroke. For explicitly authored squidgets, this is simply the associated attribute configuration for discrete squidgets, and weight interpolated attribute values based on where the user stroke crosses the piecewise linear path for a continuous squidget.




Fig. 7. The user stroke is matched for similarity against all scene squidgets. The cuboid contour best matches the stroke in (a), causing the object and contour to translate to fit the curve (b). Once selected, squidgets can be interactively transformed (c). The torus' contour is selected (d) and translated (e) to match a second stroke.

For implicit squidgets, one needs to find a minimal change to the associated set of attributes such that the change deforms the curve to match the user drawn stroke. For common attributes like deforming points on the squidget curve (Figure 4(c)) the change is direct, and for others like transformations (translate/rotate/scale) (Figure 7), the change in attribute value can be analytically computed. In general, We can optimize for such attribute values using a local search from the

current values of associated attributes. For example, the cyan stroke drawn in Figure 4(d), can edit the spotlight cone-angle associated with the elliptical red squidget in Figure 4(c), by searching for a nearby cone-angle value whose hot-spot best fits the user stroke.

Implicit and continuous squidgets further allow users to interactively refine the controlled attributes by dragging their values in real-time. A user stroke held at its end-point (0.3 seconds in our implementation) remains active allowing the stroke to be dragged around spatially. This behavior interactively manipulates the user stroke along the interpolation path for continuous widgets, or as a transformed stroke for implicit squidgets (Figure 3). Modifiers can also be employed to specify or restrict the subset of attribute values to change. For example, modifiers with implicit squidgets can force the interactive refinement to be restricted to translate/rotate/scale (Figure 3).

4 IMPLEMENTING A SQUIDGET FRAMEWORK

While the squidget framework supports a wide range of implementations for a variety of applications and workflows, in this section we discuss details for all aspects of an implementation with the 3D animation system *Maya* 2024’s C++ API. User input can be provided by a mouse, track-pad, or Wacom Cintiq 24HD Touch tablet.

There are two modes to our squidget framework: *create*, dealing with the creation/deletion of squidgets and their association with scene attributes; and *control* dealing with the interactive stroke-based manipulation of scene attributes.

4.1 Squidget Creation

A squidget is a scene curve C (*implicit* or *explicit*) whose appearance can be controlled by an associated set of scene attributes A .

Implicit squidgets comprise existing curves in the scene, as well imagined curves that can be automatically created using NPR, from silhouettes and other feature curves on scene objects. In both cases, the curves are associated with the attributes that define their respective scene objects. We use an NPR toon-rendering in *Maya*, to create imagined curves from scene objects. Arbitrary object attributes can visually change toon curves, illustrated in our implementation using object transforms. We note that while view-dependent toon curves can change discontinuously upon object transformation, we can treat them 3D curves pasted on their respective objects, for squidget control.

Explicit squidgets are curves authored and associated with an arbitrary set of scene attributes using our framework. While these curves can be created using any curve drawing tools in *Maya*, we streamline their creation and attribute association management using canvases.

A canvas B can serve both, as a planar surface in 3D upon which to project a user stroke, and as a repository for selected object attributes to be associated with strokes drawn on the canvas. Given a screen space user stroke S as a poly-line with points $[x_0, \dots, x_k]$, we project x_i onto a canvas plane B , by casting a ray through the screen to intersect with the canvas as a curve $B_S = [x_0^B, \dots, x_k^B]$. In our implementation, we simply utilize *Maya* planes for canvases, designed with translucency to allow users to view and interact with the scene through the plane. Canvases are associated with a set of scene object attributes $B.A$, recursively computed from selected object hierarchies upon canvas creation, that all squidgets drawn on that canvas control. When an explicit curve is drawn on a canvas B , the current attribute values $B.A$ are associated with a discrete squidget. We experimented with other attribute selection schemes such as automatically selecting recently changed attributes, or re-selecting attributes for each drawn squidget, but found our proposed canvas-centric approach provided a good balance between simplicity and flexibility of selection.

Canvases can also be used to constrain the best-match squidget search for a given user stroke. We found that strokes that project entirely onto and within a first visible canvas, were typically

associated with squidgets on that canvas. This allows us to explicitly constrain the processing of such strokes, to a single canvas and its squidgets, making our system more robust to ambiguities in stroke matching, specially in complex scenes with many canvases and squidgets (Figure 17(a)).

A stroke in create mode can be responsible for discrete squidget creation/deletion, continuous squidget connection/disconnection, and canvas deletion. Simple gesture processing can be used to distinguish the stroke operation as follows: discrete squidgets or continuous squidget paths are deleted using a cross-out stroke gesture, where the stroke crosses the curve representing a canvas, a squidget, or a path, at least twice; A stroke crossing a sequence of multiple discrete squidget curves exactly once is used to create a continuous squidget; all other strokes are used to create discrete squidgets.

Discrete squidget curves are created from strokes by fitting a NURBS spline to the projected stroke point sequence $[x_1^B, \dots, x_k^B]$, and then re-sampled as a poly-line with a fixed number of points (default 30).

Continuous squidget curves, are created by strokes crossing m discrete squidget curves $[C_1, \dots, C_m]$. We establish an m point poly-line path through the mid-points of each curve in $[C_1, \dots, C_m]$. The path is parameterized to control an interpolation weight $w \in [0, 1]$, that uniformly defines a piece-wise interpolation path between curve pairs C_{i-1}, C_i , for $w \in [(i-1)/m, i/m]$.

4.2 Squidget Control

In Control mode, a stroke is first matched against a list of implicit and explicit squidget curves to *select* the best matched squidget, and then *manipulate* the selected squidget's associated attributes to match the stroke or a subsequent manipulation stroke (signified by a modifier key during the first stroke).

Squidget Selection: When user strokes are completely contained within the extents of a canvas, we restrict our search to squidget curves belonging to that canvas.

For explicit squidgets, we compute a matching distance (similar to [WWL07]) from a stroke S , to a candidate squidget curve C , by first re-sampling the stroke to have the same number of control points as C , and then summing the pairwise distance between corresponding points (i.e. $dist(S, C) = \sum_i (|x_i^S - x_i^C|^2)$). The distance is also computed against the reverse ordering of stroke points. To make the matching spatially agnostic, points can be translated to their respective local origins using the point averages for the stroke and the curve. We also note that for strokes S drawn in 2D the distance computation and matching above is performed in 2D between points on S and points on C projected in screen space.

For continuous squidgets, when the stroke cross the path at some parameter w , the distance is computed against the interpolated squidget curve at parameter w . When the stroke does not cross the path, we can analytically compute a least squares best-fit curve whose shape most closely matches the stroke, or the value w where the $dist(S, C(W))$ is minimized.

Implicit squidgets require the matching distance to be computed against an imagined sub-segment of the NPR toon curves. We do this by first segmenting the NPR toon curves representing implicit squidgets, using corner detection [VEH08]). We can further employ a partial matching between the toon curve and the user stroke. We map our stroke S onto C (in screen space in case of 2D input strokes) to identify a mapping of S to a portion of C . In practice, each stroke point is projected to its closest point on C , to create a partial curve C^S , and a mapping of stroke points to toon curve points.

Spatial vs. Shape distance: We can further decompose the $dist$ function to match a stroke and curve, into a spatial and shape component. In other words, we can compute a best-fit transform A , such that: $\arg \min_A (dist(S, AC^S))$. Matrix (A) can be decomposed into a Rotation and Translation

matrix: $A = \begin{bmatrix} R & \vec{T} \\ 0 & 1 \end{bmatrix}$. We can find these components using polar decomposition and average point translation similar to the method found in [MHTG05]. The matrix (A) can be defined as a local object transform $A_1 = X^{-1}AX$, where X is the local to screen space transform for the object in the scene.

When trying to select from multiple Implicit Squidgets, we select the best A that maximizes the score:

$$\text{score} = (1 - \lambda) * \frac{1}{\text{dist}(S, AC^S)} + \lambda * \frac{1}{\text{dev}(A_1^0, A_1)}$$

where A_1^0 is the original local object transformation, $\text{dev}(A_1^0, A_1)$ is the squared vector norm of the matrix translation components: $|A_1^0.t - A_1.t|^2$, and $\lambda = 0.7$.

Squidget Manipulation: Once a squidget is selected, its associated attribute values need to be set to cause the squidget curve to best-fit match the user manipulation stroke.

These attribute values are directly obtained based on the selection computation for explicit squidgets, and for spatial object transformation using implicit squidgets. For arbitrary attributes like the spotlight cone-angle in Figure 4(d), a local numerical search for attribute values that minimize the distance between the resulting implicit squidget and the user stroke, can be employed.

Finally, for further interactive control of squidgets, users can employ a hold-and-drag technique to interactively refine the attribute values like a virtual slider: for example by moving along the path of a continuous squidget to refine the weight w and its consequent attribute values; or by interactively moving the object transform of an implicit squidget. We further show that modifiers can be used to selectively control the changed attributes, for example to perform constrained translation/rotation (Figure 15(a)).

5 APPLICATIONS

In this section, we present many examples of applications incorporating squidgets.

5.1 Spatial Scenes

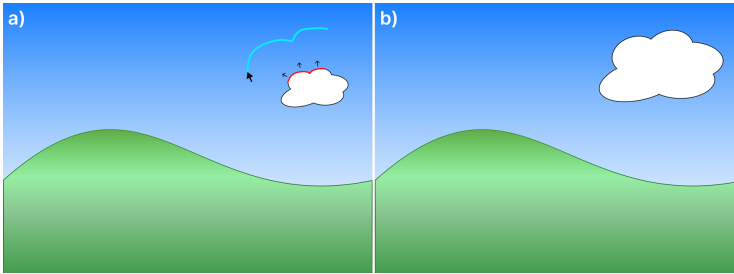


Fig. 8. (a) An implicit stroke is drawn to match the contour of the cloud. (b) The cloud resizes to fit along the stroke shape.

5.1.1 2D Applications. (Figure 8) In this example, we present how a user could change the size of a cloud to fit the stroke shape. 2D applications, which often naturally support drawing gestures, can seamlessly incorporate squidgets to allow quick editing of scene parameters through matching object contours with a user stroke.




Fig. 9. A self-authored, continuous squidget with four keyed attribute states that translates and rotates a torus.

5.1.2 Shape Transformation. (Figure 9) In this example, we present a transformation with four keyed parameter states combined into a continuous squidget. This continuous squidget exhibits visual guidance through two mediums: 1) the torus’s motion path is conveyed through the squidget’s track, and 2) the torus’s spatial orientation matches the key elements along the path. The user can use the draw and drag method to interpolate torus transform states.




Fig. 10. Multiple squidgets control different parameters of a light cone on one canvas: (a) light visibility, (b) light intensity, and (c-d) spot-light position through altering rotation.

5.1.3 Scene Lighting. (Figure 10) In this example, we present several hand-authored squidgets for scene lighting. The flexibility of curve shapes allows users to dictate squidget utility through visual form. An on and off switch is comprised from two discrete squidgets with opposite orientation while a linear slider shape determines light intensity. These squidgets present familiar indicators we observe from real world light switches. We also replace the local nature of performing rotations with a global control by specifying light location rather than the local light cone rotation.

5.1.4 VR. (Figure 11) In this example, we present a method of using squidgets in VR environments. Squidget curves can be formed in VR environments not bound to an object surface or canvas to key and interpolate parameters outside of 2D viewports. Users can pose beyond limitations of 2D visual mediums to more accurately pose objects that may otherwise be limited from depth perception.

5.2 Characters and Faces

5.2.1 Face Control. (Figure 12) In this example, we present several hand-authored squidgets for facial posing. Face rigs often have many handles for each part of the face, but posing often requires manipulating multiple handles together. Squidgets allow users to generate higher-level abstraction handles (such as smile and frown handles) to simplify facial animation by controlling these self-authored handles without losing local control.




Fig. 11. Within a VR environment, curves can be drawn to control an wood-chopper (a-b). 3D curves drawn in VR (c) can be used to match body contours at a specified camera view (d) to correct the body tilt (e-f).




Fig. 12. Squidgets can pose combinations of individual FACs control handles to drive complex facial pose interpolations: (b) smile, neutral, frown, jaw open, (c) eye-lid open and close, (d) eye brow raise. The authored squidgets can be used synchronously to create complex facial expressions (e-f).




Fig. 13. Squidgets can control character rigs to key multiple poses (a-c) for a jumping motion (d-g).

5.2.2 Character Animation. (Figure 13) In this example, we present the motion sequence for a jump connected together through a continuous squidget. Squidgets can be used to visualize motion sequences in real time by keyframing poses and inbetweening using sketch gestures.




Fig. 14. Using previously hand-authored squidgets to move the boat and moon, we can create a parent squidget to control previous continuous squidgets (a) by keying squidget values (b-d). By controlling this parent squidget (e), we can drive multiple squidgets together (f-h)

5.2.3 Squidgets for Squidgets. (Figure 14) In this example, we present a squidget controlling other squidgets. Because self-authored squidgets themselves contain scene parameters, these squidget parameters can be driven by other squidgets. Nested squidgets provided a multi-timeline animation workflow where individual motion sequences are encapsulated into larger sequence containers. Thus characteristic motion sequences keyed by squidgets can be re-used by other squidgets to modularize complex actions.

6 EVALUATION




Fig. 15. Study scenarios with various tasks: arranging object using implicit squidgets (a), creating squidgets for a boat scene (b), and using both squidgets for character control (c).

We evaluated squidgets two ways: with an informal 7 participant user study; and by a free-form interactive session with 4 animation professionals.

6.1 User Study

We evaluated our framework with 5 casual users C_1, \dots, C_5 and 2 experienced artists E_1, E_2 against a series of tasks. Our study was broken down into three scenarios: In scene 0, users arranged shapes using implicit squidgets and *Maya*'s internal transformation tool (move and rotate). In scene 1, users followed a tutorial on how to author explicit squidgets to complete a series of tasks. Tasks involved creating squidgets to control different sets of parameters (transform and rendering attributes). In scene 2, users used both implicit squidgets and explicit squidgets to control a character rig.

For each of the three scenarios, users were prompted through a text window with directives such as: "Create a simple squidget that raises her arm above her head and waves." or "Use our Tool to draw strokes on the scene to overlay the yellow objects over the blue objects." Before the study, users viewed a tutorial video explaining how to use squidgets and later filled out a post-study form with questions pertaining to their experience using the tool. Participants completed the study using either a mouse or pen-interactive display.

Results

For scene 0, we randomly assigned the tool usage order (Squidgets or *Maya*) with participants and measured the time spent using each tool and accuracy of shape placement. We found that on average participants took around 40% more time to arrange the shapes using Squidgets than using *Maya*'s internal tools despite having similar placement accuracy. Results from our study revealed a balanced preference among participants, with 3 of the participants preferring *Maya*'s tools and 4 preferring squidgets (including the 2 experienced artists). Though many participants found our tool more responsive, noting how translations were quick and easy to control, difficulty with out of view-plane rotations and inadvertent snapping of objects due to an incorrectly inferred squidget, were frustrating. Although we provided an undo feature, we found that participants preferred prioritizing proximity to the mouse location over stroke shape when selecting a shape. This can easily be fixed with parameter tuning. We recognize that there is a trade-off with fine-tune controls that *Maya* provides and natural quickness of squidgets, and imagine a workflow involving quick parameter tuning with Squidgets and precise adjustments with a traditional interface or constraint snapping.

The explicit widgets in scene1 and scene2, received widespread acclaim from all participants. Universally, participants found squidgets to be user-friendly and intuitive. Multiple participants praised the simplification of key-framing motions, highlighting how squidgets made this process




Fig. 16. Post-study survey questions and results from six participants.

significantly easier and quicker (C_1, C_4, E_1). They also appreciated how squidget curves provide visual cues and previews for animations (C_1, C_4). We observed how many participants used squidgets to imitate temporal parametric control rather than spatial controls to visualize animations (Figure 17). Furthermore, squidgets enabled users to efficiently control multiple parameters simultaneously, especially when manipulating multiple facial handles to create complex facial expressions (E_1).

Some participants provided feedback requesting additional features like the ability to refine a continuous squidget by inserting more discrete squidget curves, or controlling the timing along the interpolation path, indicating a desire for more nuanced control over animations, including non-linear animation curves. All these features can be easily incorporated with a more comprehensive implementation of our framework.

Remarkably, participants expressed genuine enjoyment while using the tool, with some even describing it as fun. At the conclusion of scenes 1 and 2, participants were asked whether they wished to continue using our tool. All but one chose to extend their interaction, even beyond the scheduled study duration to create intricate animation sequences (Figure 17), underscoring the tool's engaging and user-friendly nature for even those who unfamiliar *Maya*. The study duration was estimated to be 1 hour long, but some users spent over 2 hours playing with the tool.

Overall, the study while somewhat informal, highlighted the playful, versatile, and naturally engaging nature of the squidget framework as a comprehensive, stroke-based interface for creative applications.

6.2 Interactive sessions with experts

We also performed an in-depth free-form session (30 mins. to 1.5 hours long) with 4 animation professionals P_1, \dots, P_4 , interactively demonstrating squidgets and having qualitative discussions about our idea. Here we distill the pros and cons of both the concept and our current prototype from the sessions:

Positives:

- Explicit squidgets are "fun", "very intuitive", and "promising".

- Explicit squidgets can be used to modularize key-frame animation by creating animation sequences for smaller modules and controlling each squidget module with a meta squidget.
- Explicit Squidgets could be created in-situ for non-tangible scene parameters (such as lighting) to provide artists more accessibility towards difficult to reach parameters.
- Implicit squidgets representing skeletons could be a powerful way to roughly pose characters by allowing users to draw character outlines of the pose [GCR13]. This would "save artists hundreds of hours" by allowing squidgets to pose the character and then the artist can perform edits and tweaks to polish the pose. Additionally, stroke color could enforce constraints on different body parts for more accurate parameter adjustments.

Negatives:

- Implicit squidgets may suffer from ambiguity if selected from too many near-by squidgets or too-similarly shaped curves (multiple scene objects with similar shapes).
- Non-planar rotations with implicit squidgets can be challenging.
- Explicit Squidgets have the potential for visual clutter if too many squidgets are created, obstructing the viewport.
- Explicit Squidgets may have problems if the same attribute is coupled to multiple different squidgets. Changing one squidget may thus affect the states of the others.

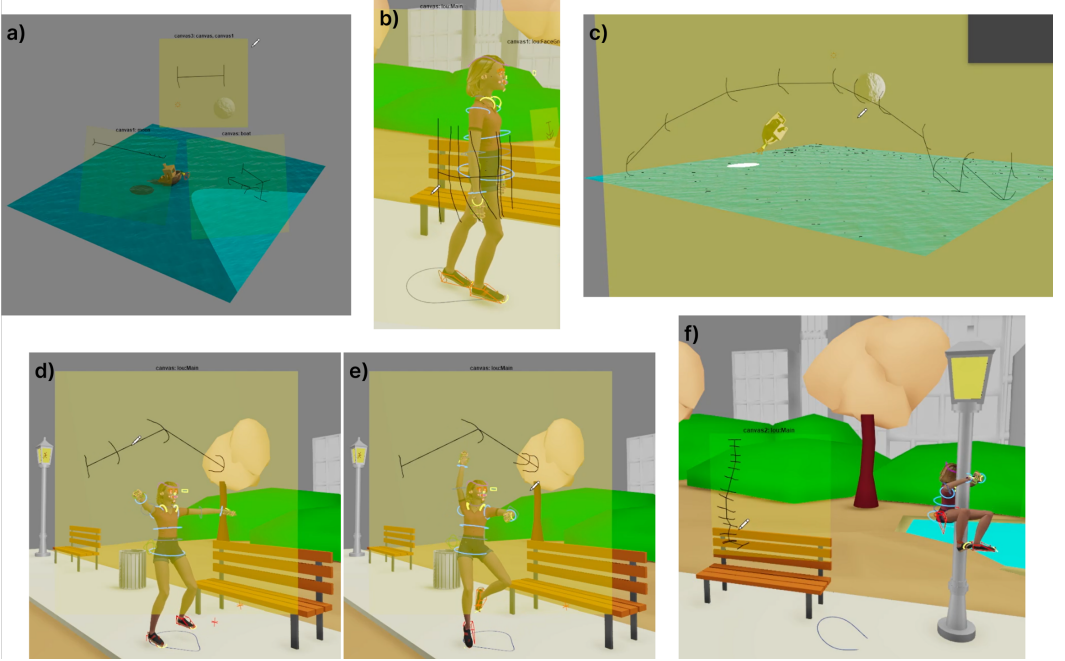


Fig. 17. Examples of self-authored widgets from user-study participants: squidget controlling other squidgets (a), walking animation (b), bouncing moon (c), ballerina pirouette (d-e), and climbing a lamppost (f).

7 CONCLUSION

Inspired by hand-crafted custom widgets, and the observation that users naturally sketch desired changes over graphical scenes, we propose squidgets, UI interaction framework, where user strokes indicate visual change that is manifest by directly manipulating scene attributes. We present and

evaluate an example implementation of the squidget concept, within the 3D modeling and animation system *Maya*.

Limitations and Future work: Our work has many limitations. Conceptually, squidgets capture an ill-posed and ambiguous inverse control problem. Despite this, we found that by constraining the number of squidget curves in a scene and their mapping to scene attributes, interacting with squidgets can be predictable and satisfying. Our example implementation makes a number of assumptions to further restrict the problem. For example, we only show control of transform attributes using implicit widgets, leaving general attribute control to future work. There may also be inherent redundancy in scene attributes. For example, a stroke indicating a larger sphere in *Maya*, can be realised by a change to the sphere's radius, its scale transform, or an infinite combination of the two. Such redundancy can be mitigated by authoring explicit squidgets that clarify the attributes to controlled by the squidget.

In summary, the overall concept of squidget based interaction was received with much enthusiasm in our user study, and we hope that this work will fuel further work in direct, in-situ, sketch-based scene interaction.

Acknowledgements: Acknowledgements are anonymized during this review process.

REFERENCES

- [AG04] Georg Apitz and François Guimbreti re. Crossy: a crossing-based drawing application. In Steven Feiner and James A. Landay, editors, *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology*, Santa Fe, NM, USA, October 24-27, 2004, pages 3–12. ACM, 2004.
- [ALCS18] Andreas Aristidou, Joan Lasenby, Yiorgos Chrysanthou, and Ariel Shamir. Inverse kinematics techniques in computer graphics: A survey. *Computer Graphics Forum*, 37:35–58, 09 2018.
- [BBS08] S.H. Bae, Ravin Balakrishnan, and Karan Singh. ILoveSketch: as-natural-as-possible sketching system for creating 3d curve models. In *Proc. User Interface Software and Technology*, 2008.
- [BH19] Pierre B nard and Aaron Hertzmann. Line drawings from 3d models: A tutorial. *Found. Trends. Comput. Graph. Vis.*, 11(1–2):1–159, sep 2019.
- [BLAI17] Lonni Besan on, Paul Issartel, Mehdi Ammi, and Tobias Isenberg. Mouse, tactile, and tangible input for 3d manipulation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, page 4727–4740, New York, NY, USA, 2017. Association for Computing Machinery.
- [Bie86] E. A. Bier. Skitters and jacks: Interactive 3d positioning tools. In *Proc. I3D 86*, pages 183–196, 1986.
- [Bie90] E. A. Bier. Snap-dragging in three dimensions. In *Proc. I3D 1990*, pages 193–204, 1990.
- [BKLP04] Doug A. Bowman, Ernst Kruijff, Joseph J. LaViola, and Ivan Poupyrev. *3D User Interfaces: Theory and Practice*. Addison Wesley Longman Publishing Co., Inc., USA, 2004.
- [CIRL⁺16] Byungkuk Choi, Roger Blanco i Ribera, J. P. Lewis, Yeongho Seol, Seokpyo Hong, Haegwang Eom, Sunjin Jung, and Junyong Noh. Sketchimo: Sketch-based motion editing for articulated characters. *ACM Trans. Graph.*, 35(4), jul 2016.
- [CV01] Tony F Chan and Luminita A Vese. Active contours without edges. *IEEE Transactions on image processing*, 10(2):266–277, 2001.
- [DXS⁺07] Julie Dorsey, Songhua Xu, Gabe Smedresman, Holly Rushmeier, and Leonard McMillan. The mental canvas: A tool for conceptual architectural design and analysis. In *Proceedings of the 15th Pacific Conference on Computer Graphics and Applications*, PG '07, page 201–210, USA, 2007. IEEE Computer Society.
- [GCR13] Martin Guay, Marie-Paule Cani, and R mi Ronfard. The line of action: An intuitive interface for expressive character posing. *ACM Trans. Graph.*, 32(6), nov 2013.
- [GG01] Bruce Gooch and Amy Gooch. *Non-photorealistic rendering*. CRC Press, 2001.
- [Her99] Aaron Hertzmann. Introduction to 3d non-photorealistic rendering: Silhouettes and outlines. *Non-Photorealistic Rendering. SIGGRAPH*, 99(1), 1999.
- [IMH07] T. Igarashi, T. Moscovich, and J. F. Hughes. Spatial keyframing for performance-driven animation. In *ACM SIGGRAPH 2007 Courses*, SIGGRAPH '07, page 25–es, New York, NY, USA, 2007. Association for Computing Machinery.
- [IMT99] Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. Teddy: A sketching interface for 3d freeform design. *SIGGRAPH '99*, pages 409–416, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.
- [JGHD09] Gabe Johnson, Mark Gross, Jason Hong, and Ellen Do. Computational support for sketching in design: A review. *Foundations and Trends in Human-Computer Interaction*, 2:1–93, 01 2009.

- [KS21] Joonho Kim and Karan Singh. Optimizing ui layouts for deformable face-rig manipulation. *ACM Trans. Graph.*, 40(4), jul 2021.
- [LA10] John P Lewis and Ken-ichi Anjyo. Direct manipulation blendshapes. *IEEE Computer Graphics and Applications*, 30(4):42–50, 2010.
- [LMLH07] Yunjin Lee, Lee Markosian, Seungyong Lee, and John F. Hughes. Line drawings via abstracted shading. *ACM Trans. Graph.*, 26(3):18, 2007.
- [MB21] Elie Michel and Tamy Boubekeur. Dag amendment for inverse control of parametric shapes. *ACM Transactions on Graphics*, 40(4):173:1–173:14, 2021.
- [MCG⁺19] D. Mendes, F. M. Caputo, A. Giachetti, A. Ferreira, and J. Jorge. A survey on 3d virtual object manipulation: From the desktop to immersive virtual environments. *Computer Graphics Forum*, 38(1):21–45, 2019.
- [MHTG05] Matthias Müller, Bruno Heidelberger, Matthias Teschner, and Markus Gross. Meshless deformations based on shape matching. *ACM transactions on graphics (TOG)*, 24(3):471–478, 2005.
- [NSACO05] Andrew Nealen, Olga Sorkine, Marc Alexa, and Daniel Cohen-Or. A sketch-based interface for detail-preserving mesh editing. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, 24(3):1142–1147, 2005.
- [PCS23] Karran Pandey, Fanny Chevalier, and Karan Singh. Juxtaform: Interactive visual summarization for exploratory shape design. *ACM Trans. Graph.*, 42(4), jul 2023.
- [SF98] Karan Singh and Eugene Fiume. Wires: A geometric deformation technique. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '98*, 1998.
- [SSB08] Ryan Schmidt, Karan Singh, and Ravin Balakrishnan. Sketching and Composing Widgets for 3D Manipulation. *Computer Graphics Forum*, 27(2):301–310, 2008.
- [Sut64] Ivan E. Sutherland. Sketch pad a man-machine graphical communication system. In *Proceedings of the SHARE Design Automation Workshop, DAC '64*, pages 6.329–6.346, New York, NY, USA, 1964. ACM.
- [Tod04] James T Todd. The visual perception of 3d shape. *Trends in cognitive sciences*, 8(3):115–121, 2004.
- [WEH08] Aaron Wolin, Brian Eoff, and Tracy Hammond. Shortstraw: A simple and effective corner finder for polylines. *SBIM*, 8:33–40, 2008.
- [WWL07] Jacob O. Wobbrock, Andrew D. Wilson, and Yang Li. Gestures without libraries, toolkits or training: A \$1 recognizer for user interface prototypes. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology, UIST '07*, page 159–168, New York, NY, USA, 2007. Association for Computing Machinery.
- [XAGW16] Haijun Xia, Bruno Araujo, Tovi Grossman, and Daniel Wigdor. Object-oriented drawing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16*, page 4610–4621, New York, NY, USA, 2016. Association for Computing Machinery.
- [XCS⁺14] Baoxuan Xu, William Chang, Alla Sheffer, Adrien Bousseau, James McCrae, and Karan Singh. True2form: 3d curve networks from 2d sketches via selective regularization. *ACM Transactions on Graphics*, 33(4), 2014.
- [ZHH96] R. C. Zeleznik, K. P. Herndon, and J. F. Hughes. Sketch: An interface for sketching 3d scenes. In *Proc. SIGGRAPH '96*, pages 163–170, 1996.