國立臺灣大學電機資訊學院資訊網路與多媒體研究所

碩士論文

Graduate Institute of Networking and Multimedia

College of Electrical Engineering and Computer Science

National Taiwan University

Master's Thesis

視覺 Fidget 互動：範例、設計框架與混合實境中的創作工具

Visual Fidget Interactions: Examples, Design Frameworks, and Authoring Tools in Mixed Reality

梁舜勛

Shun-Hsun Liang

指導教授: 鄭龍磻 博士

Advisor: Lung-Pan Cheng Ph.D.

中華民國 115 年 2 月

February, 2026

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

摘要

Fidgeting 泛指人類為了自我調節而進行的低強度、重複性的動作。雖然傳統上這種行為根植於觸覺體驗（如擠壓或輕敲），但近期的研究開始探索透過視覺介面來豐富這類行為，並擴展其表達潛力。然而，現有的方法鮮少將「Fidget 動作」與「回饋」之間的映射關係正式化，且往往依賴於特定的工具，這與日常自然的動作有所脫節。為了彌補這一缺口，我們開發了一套混合實境創作系統，旨在以日常的 Fidgeting 為基礎，進行視覺化的設計。我們展示了該介面的工作流程，並詳細介紹了從前置研究中提取出的基礎可供性映射。研究結果展示了 10 款由使用者創作的視覺設計及其相應的映射策略。

關鍵字：混合實境、Fidgeting、具身互動、創作工具

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

Abstract

Fidgeting involves low-effort, often repetitive hand movements used for self-regulation. While traditionally grounded in haptic experiences such as squeezing or tapping, recent research explores visual medium to enrich fidgeting and extend its expressive potential. However, these approaches rarely formalize gesture–feedback mappings and often rely on specialized tools that diverge from natural, everyday gestures. To address this gap, we developed a mixed reality (MR) authoring system that enables visual fidget design grounded in everyday fidgeting behaviors. We demonstrate the workflow of our interface and walk through the underlying affordance mapping extracted from our formative study. Our results show 10 user-created visual fidget designs with mapping strategies.

Keywords: Mixed Reality, Fidgeting, Embodied Interaction, Authoring Tool

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

Contents

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

List of Figures

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

Chapter 1 Introduction

Fidgeting is a common human behavior involving low-effort, often patterned hand movements—performed consciously or unconsciously—for self-regulation [7, 12, 21]. Traditionally, fidgeting has been associated with haptic experiences such as squeezing stress balls or tapping on tables. However, recent work has begun to explore programmable effects as a means to visually enrich the fidgeting experience beyond what traditional fidget tools afford, which we call ”Visual Fidget.” Examples include those with screen-based virtual fidgets [24, 31], 2D AR overlays [10], custom tangible devices [11], and programmable swarm robots [13]. Evidence suggests that such visual enrichment enables strategies to promote emotional well-being beyond traditional fidget approaches [10, 11, 31].

Across these systems, designers map fidget gestures to responsive animations, highlighting the potential for visual rather than purely haptic fidgeting. However, few studies formalize these mappings into reusable design principles, and many rely on specialized input devices that diverge from natural, everyday fidgeting. These gaps point to key needs: visual fidget examples grounded in familiar fidgeting behaviors and a design framework for gesture-animation mapping.

In this paper, we propose MR as a flexible medium for exploring visual fidget design.

doi:10.6342/NTU202600777

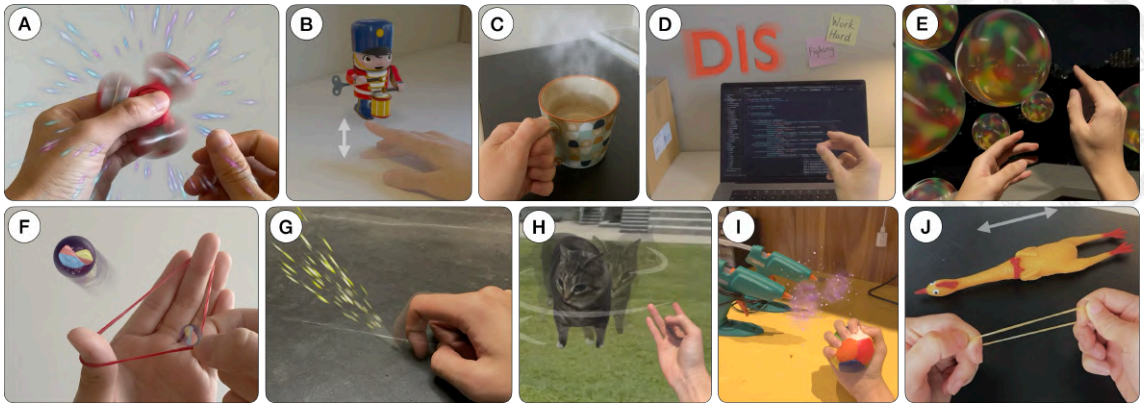

Figure 1.1: Visual fidget designs created by participants. (A) Spinning a fidget spinner generates outward radial shockwaves, amplifying the perceived kinetic energy of the spin. (B) Finger-drumming animates a miniature toy drummer marching in place. (C) Rubbing the mug handle releases rising wisps of steam, evoking warmth and enhancing the tactile quality of the gesture. (D) Rubbing on the desk surface triggers a vibrating motivational reminder, offering an uplifting prompt. (E) Swinging both hands while softly humming causes bubbles to vibrate in sync, forming an orchestral-like visual rhythm. (F) Plucking a rubber band launches marbles outward, resembling a slingshot. (G) Scratching across the surface emits sparks, simulating friction-induced energy. (H) A quick flick triggers a spinning cat-meme animation. (I) Squeezing a stress ball activates dual heat-gun nozzles that eject smoky particles. (J) Stretching a rubber band elongates a rubber chicken in a playful deformation effect.

Unlike physical tools, MR enables interaction unconstrained by object affordances or surrounding space. Prior work has shown that MR systems can support visual programming without explicit coding [1, 25, 28], enabling fast prototyping and exploration. Building on this, our system captures real-time fidget gestures—bypassing the need for preregistration—and its contextual affordance by using a large language model (LLM) to suggest animations, streamlining the design exploration process. To guide the development, we conducted a formative study to establish design frameworks, developed an authoring system based on the derived insights, and followed by a validation study where participants used this interface to create their own visual fidget interactions on top of their everyday fidgeting practices.

Our contributions are: (1) Introduced the concept of visual fidget and underlying affordance mappings that support personal fidget transformation, supported by 10 user-

doi:10.6342/NTU202600777

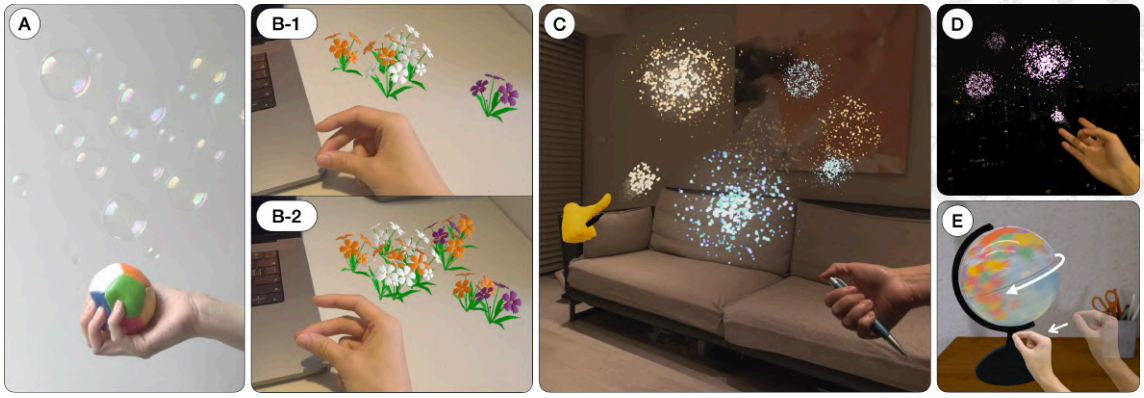

Figure 1.2: Visual fidget prototypes used in the formative study. (A) MR Stress Ball releases floating bubbles when squeezing a stress ball, externalizing released tension through calming visual. (B) Meditative Garden visualizes blooming of flowers incrementally when snapping, fostering self-awareness and reflecting on the progression of their snapping habits over time. (C) Multi-User Fireworks turns annoying pen-clicking from an annoyance into a playful, social interaction, triggering shared firework bursts between remote users. (D) Fingertip Fireworks reinterprets flicking as a trigger to launch a burst of fireworks, amplifying momentary motion into expressive visual trajectories. (E) Spinning Globe enables continuous scroll of virtual globe model from a distance, translating hand motion into spatial navigation of visual inertia.

generated (Figure 1.1) and 4 study-validated examples (Figure 1.2). (2) Developed visual fidget, a mixed reality–based authoring system that enables rapid prototyping of visual fidgets.

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

Chapter 2 Related Work

This section provides background on fidgeting as a meaningful interaction modality, its design trend of visual enrichment, and how MR support design exploration.

2.1 Fidgeting

Fidgeting involves low-effort and patterned behaviors that occur with or without one's conscious awareness [7, 21] using everyday objects [12]. Although what constitutes fidgeting can be highly personal and subjective, fidgeting generally serves a common function: provide stimulation that sustains engagement and emotional balance by countering sensory monotony. [3, 8, 23].

From a concentration perspective, too little stimulation can leave individuals under-aroused, and fidget tools help by providing sensory input that supports sustained attention and task completion [14]. Prior studies support this benefit: stress balls have been shown to improve learning and concentration [27], and fidget spinners have increased on-task time for students with ADHD [2]. On the other hand, releasing pent-up energy through fidgeting also serves an emotional function. For example, squeezing helps release anger and regulate emotions [5, 31]. Similarly, research has revealed how fidget behaviors like pen-clicking and surface-scratching offer emotional satisfaction [12, 23].

doi:10.6342/NTU202600777

Despite existing evidence highlighting the functional outcomes of fidgeting, the underlying mechanisms that bridge these endpoints remain insufficiently understood, posing challenges for design exploration. Recent work, however, provides theoretical grounding. Perrykkad et al. [21] suggest that fidgeting may act as a form of self-confirmation, wherein individuals seek "sense of control" by sensory feedback that aligns with their expectations of causality. While our objective is not to directly test this hypothesis, it informs our approach to developing an affordance mapping of fidget behaviors. In doing so, we offer empirical support for this perspective and articulate design guidelines that translate such theoretical insights into actionable principles. This contribution opens up new directions for future exploration and validation in fidget tool design.

2.2 Enrich Fidget Experience through Visual Effects

While fidgeting has traditionally been grounded in haptic engagement, an increasing body of research explores how visual effects can enrich the fidgeting experience. Conventional fidgeting is often constrained by the physical environment and the objects immediately available, which can render the experience repetitive and monotonous over time. By leveraging the programmability of digital media, recent work enhances fidgeting through dynamic visual feedback that sustains engagement and extends its regulatory potential.

For instance, AR Fidget demonstrates how swiping or tapping on AR glasses can be coupled with immersive visual metaphors—such as tapping the side to pop a virtual bubble or sliding to extinguish virtual fire—creating a sense of contextual causality [10]. Similarly, Zhou et al. map squeezing gestures onto soothing on-screen particle effects, highlighting the motional affordance of squeezing for emotional regulation and stress re-

doi:10.6342/NTU202600777

lief [31]. Kim's SwarmFidget introduces swarm robotic motion as a programmable fidget outlet, enabling users to interact with dynamic responses that reinforce physical causality [13]. Ross maps cursor movement to the virtual fluid swirls to help users sustain attention during on-screen tasks [24]. On the other hand, Karlesky uses the bubble-wrap-like textures on screen to prompt uses the correlation of pressing input and vibration feedback [11].

2.3 Design Exploration through Mixed Reality

Prior research on authoring tools has highlighted the potential of mixed reality in facilitating rich and expressive design exploration. Systems such as GesturAR have demonstrated how gestures can be visually programmed to trigger predefined manipulations, enabling rapid prototyping through gesture-based interactions [28]. Similarly, MagicalHands explores how gestural input can modulate properties of animated effects, focusing on the response layer of interaction design [1]. However, both systems primarily rely on predefined gesture classification and are optimized for operational manipulation tasks. In contrast, inspired by programming-by-demonstration paradigm in GestureCanvas [25], gestures in our work are user-recorded initially and are used to trigger personalized visual effects users map to. By omitting pre-defined gestures, we are able to cater to diverse nature of fidget behaviors.

To expand the semantic scope of interaction design, we further build on the concept of Augmented Object Intelligence as proposed in XR-Objects [6]. XR-Objects bypasses pre-registration by fusing real-time vision inputs (e.g., object bounding boxes, scene context) with language-based reasoning via a multimodal large language model. Inspired by

doi:10.6342/NTU202600777

this architecture, our system adopts a similar inference pipeline tailored to the domain of fidgeting. We combine input modalities collected from recording, such as image sequences, motion data (e.g., fingertip velocity), and user metadata (e.g., handedness) into a single prompt for semantic reasoning. This enables the system to infer gesture intent, retrieve high-level descriptions of fidget behavior, and suggest context-appropriate visual responses based on the fidget affordance.

doi:10.6342/NTU202600777

Chapter 3 Formative Study

The goal of this formative study was to explore, in an open-ended manner, how different visual design might shape the perception and experience of fidgeting. We began with an informal brainstorming session among seven HCI researchers, reflecting on personal fidget habits and sketching possibilities for visualizing fidgeting beyond the constraints of physical objects. These early explorations suggested that, although programmable visuals offer substantial expressive flexibility, animations must retain an intuitive sense of causality in order for gestures to feel low-effort and satisfying. For instance, launching fireworks from a closed fist felt mismatched when triggered by a snapping gesture as it does not inherently imply an outward, projectile-like motion, whereas a pen click provided a more intuitive match.

Guided by these early observations, we selected a set of everyday fidget gestures for further exploration: two object-based gestures (squeezing a stress ball and clicking a pen) and two bare-hand gestures (snapping and flicking). We also included a pinch-and-drag gesture—not typically viewed as a fidget, but often used playfully in virtual environments where small, low-effort manipulations can feel inherently "fidget-like." The resulting prototype set is illustrated in Figure 1.2.

We evaluated these prototypes in both primary-task and secondary-task contexts to

doi:10.6342/NTU202600777

reflect the dual nature of fidgeting—sometimes intentional and foregrounded, other times incidental and backgrounded. Following prior work [13], we incorporated a standardized set of rating scales to provide baseline points of comparison across conditions. These measures were not used for formal benchmarking; instead, they served as reference markers that helped contextualize participants' qualitative reactions and surface early design considerations for visual fidget interactions.

3.1 Participants

We recruited 20 participants (10 females and 10 males), with ages ranging from 19 to 57 years ($\mu = 25.4, \sigma = 8.0$), following approval from the local Institutional Review Board. Participants were recruited through the social media platforms to include individuals with diverse fidgeting habits and preferences. To participate in this study, participants needed to be aged 18 or older, have basic proficiency in using MR-related devices, and possess adequate vision, hearing, verbal communication, and social interaction abilities required for daily life. Participants were excluded if they wore glasses that impaired their ability to see the MR environment clearly, had recently experienced visual or auditory impairments or discomfort in MR environments, or had conditions such as epilepsy, dizziness, or other illnesses unsuitable for prolonged MR device use. Before participation, all participants signed an informed consent form detailing the study's objectives, procedures, and data protection measures.

doi:10.6342/NTU202600777

Figure 3.1: The user is wearing the Apple Vision Pro device and is squeezing the ball to trigger the fidgeting interactions.

3.2 Apparatus

The study involved the five visual fidget prototypes shown in Figure 1.2, running on an Apple Vision Pro. Participants were instructed to interact with these interactions while seated comfortably in front of a desk (see Figure 3.1). For the MR Stress Ball interaction, participants were provided with a physical stress ball; for the Multi-User Fireworks interaction, participants were provided with a pen. Additionally, a physical globe was placed on the desk approximately 80 cm from the participants' seating location for use in the Spinning Globe interaction. Audio feedback was delivered through the speakers of the Apple Vision Pro. A laptop was placed in front of the participants to play the videos for testing scenarios where fidgeting is a secondary task.

3.3 Task

Participants engaged with the interactions under two conditions: fidget (1) as a primary task and (2) as a secondary task. In the primary task condition, participants focused

doi:10.6342/NTU202600777

solely on interacting with a prototype. This condition aimed to simulate scenarios where people fidget without engaging in another task. In the secondary task condition, a separate main task was introduced to simulate situations where people fidget while concentrating on another task, such as listening to a lecture or watching a tutorial video.

For the secondary task condition, participants watched a series of video tutorials from the Coursera course "Learning How to Learn." A total of five recordings were used, each lasting an average of 3 minutes. The content focused on learning science, specifically addressing topics related to procrastination and memory. These videos were selected to appeal to a broad range of participants. The goal is to draw participants' main attention away from their fidgeting actions and the visual feedback provided by the prototypes. To further emulate real-world scenarios, participants had the freedom to choose whether to fidget, as well as the temporality and duration of their fidgeting while watching the videos.

3.4 Procedure and Measures

The study was divided into the following phases:

Introduction and Familiarization. Participants were briefed on the concept of fidgeting and its relevance to the study, followed by an introduction to traditional fidget tools (pen, stress toy, globe, strings, chair, and zigzag) to establish a baseline understanding. A semi-structured interview was then conducted to explore their typical fidgeting experiences, triggers, and perceptions of their own and others' fidgeting behaviors. This phase aimed to help participants recall their past fidgeting habits and familiarize themselves with the traditional fidgets.

Interaction and Evaluation. After being fitted with the Apple Vision Pro headset

doi:10.6342/NTU202600777

and completing calibration, participants experienced each of the five prototypess in random order. For each interaction, they engaged with it first as a primary task for approximately one minute and then as a secondary task while watching a tutorial video. Following each interaction, participants were evaluated with a quiz of two multiple-choice questions per recording, with each question worth 1 point to evaluate their attention to the primary task. Upon completion, participants filled out a questionnaire that included a self-report on whether they considered the interaction to be fidgeting, as well as 7-point Likert scale ratings of ease of use, pleasantness, intuitiveness, usefulness, and likelihood of future use [13]. They also provided qualitative feedback comparing the mixed reality interactions to traditional fidget tools and suggesting potential improvements.

Post-Study Reflection. Upon completing all interactions, participants participated in a second semi-structured interview to discuss their overall experiences. They reflected on the differences between mixed reality and traditional fidgeting, how those differences influenced their fidgeting behaviors, and their ideas for designing custom mixed reality fidget interactions.

Throughout the study, participants' feedback and perspectives within the MR environment were recorded using the device's built-in video capture feature. In addition, we gathered data on the fidgeting behaviors, such as temporality and positioning, throughout the experiment. Both qualitative and quantitative responses were collected.

The total study duration, including onboarding, interaction time, and interviews, was approximately 70 minutes per participant. To ensure consistency, all participants were tested individually under the same experimental setup.

doi:10.6342/NTU202600777

3.5 Analysis

We conducted a thematic analysis to interpret the qualitative feedback. This approach allowed for a structured comparison of participants' experiences across predefined and emerged categories identified during analysis, reflecting broader insights and unexpected patterns. Detailed codes, definitions, examples, and emergent themes can be found in Appendix A.

Quantitative measures, including ratings, post-video quizzes, and fidget behavior data from the prototypes experienced by participants, complemented the qualitative analysis, were also collected. Although the primary goal was not to achieve statistically significant results, these numerical evaluations provided high-level insights. Additionally, we used a paired t-test to assess differences in fidgeting frequencies across tasks.

3.6 Results

In this section, we summarize the study findings, including qualitative insights obtained from post-experiment interviews and quantitative assessments of the tested fidgeting experiences.

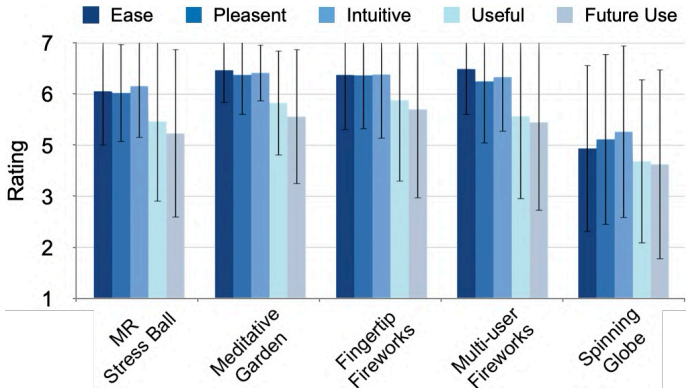

Figure 3.2: The participants' ratings across different MR interactions.

doi:10.6342/NTU202600777

We begin by presenting the participants' overall perceptions. In particular, we analyze their ratings of these interactions, identify which ones are considered appropriate for fidgeting, and explore the factors that influence users' evaluations of what qualifies as a fidget interaction.

Figure 3.2 illustrates the ratings of the five different interactions on Ease, Pleasantness, Intuitiveness, Usefulness, and Future Usage. Overall, most interactions received high scores (avg > 5) across all dimensions. The Spinning Globe scored lower in Ease and Pleasantness, with higher variability (SD = 1.93 and 1.98, respectively), but remained above 4. Furthermore, the data reveals that usefulness and future usage tend to have larger variations across all interactions. For instance, future usage in Fireworks and Multi-user Fireworks had standard deviations of 2.06 and 2.05, respectively, indicating that these dimensions elicited more mixed responses. This trend suggests that while users generally agreed on ease and pleasantness for most interactions, there was less consensus on the practical usefulness and likelihood of future usage.

We asked our participants to categorize which visual fidget prototypes they perceived as fidgeting. As shown in Figure 3.3, Meditative Garden was unanimously recognized as a fidget interaction, with 100% of participants agreeing. MR Stress Ball was also widely regarded as fidgeting, with 90% of participants categorizing it as such. Fingertip Fireworks and Multi-user Fireworks received moderate recognition, with 80% of participants identifying them as fidgeting interactions. In contrast, Spinning Globe was categorized as fidgeting by only 50% of participants, indicating significant variability in perception. These results highlight that while some interactions, such as Meditative Garden and MR Stress Ball, achieved near-unanimous agreement, others showed more diverse participant interpretations.

doi:10.6342/NTU202600777

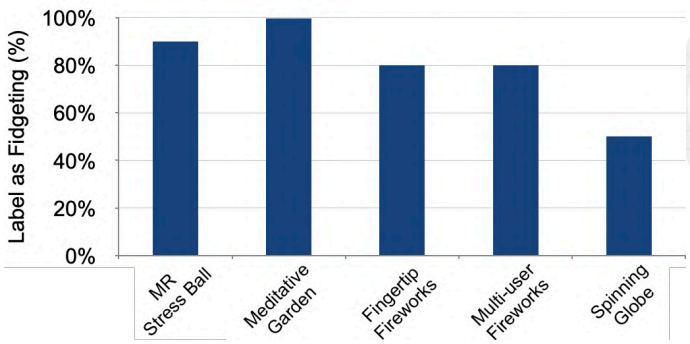

Figure 3.3: Fidget Labeling. The bar chart illustrates the percentage of participants who identified various MR interactions as fidgeting. The MR Stress Ball and Meditative Garden were most frequently labeled as fidgeting, while the Spinning Globe had the lowest percentage.

Based on the system log data (see Figure 3.4), participants fidgeted more frequently when fidgeting was their primary task (mean = 32.52 counts/min) compared to when it was a secondary task (mean = 19.12 counts/min). A statistically significant difference between fidgeting frequencies when fidgeting was used as a primary task versus a secondary task (T-statistic: 5.76, $p < 0.05$). Among the five interactions, Meditative Garden exhibited the highest fidget frequency as a primary task (mean = 39.27 counts/min, SD = 15.26) and a secondary task (mean = 18.27 counts/min, SD = 18.34). Multi-user Fireworks was the only interaction where fidgeting was slightly more frequent during the secondary task (mean = 36.84 counts/min, SD = 28.53) than the primary task (mean = 33.51 counts/min, SD = 14.58). In contrast, Spinning Globe recorded the lowest fidget frequencies for both task types (primary task: mean = 29.46 counts/min, SD = 13.00; secondary task: mean = 9.78 counts/min, SD = 12.20). The overall average quiz score for the video tutorials during the secondary task condition was 86.3 out of 100, indicating that participants maintained a high level of attentiveness to the primary task even while fidgeting.

doi:10.6342/NTU202600777

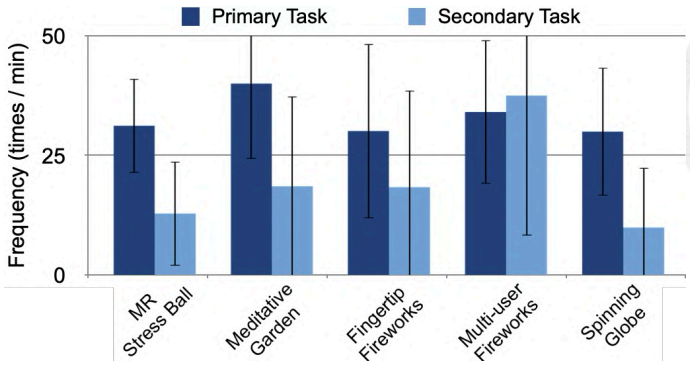

Figure 3.4: The participants' fidgeting frequencies across different MR interactions as the primary and the secondary tasks.

3.7 Findings

Across the five interactions, participants developed a nuanced sense of how fidget gestures should relate to the resulting animations. Their feedback emphasized that enjoyable visual fidgeting depends on maintaining a clear sense of agency and requiring minimal effort—conditions supported when the animation's movement and timing feel naturally tied to the gesture that drives it. Two recurring themes surfaced across the study.

Motional Affordance. Participants often reasoned about what kind of animation "made sense" for a given gesture by drawing on familiar physical or tactile analogies. In MR Stress Ball, they expected squeezing to produce something soft, gradual, and volume-releasing—reflected in comments about bubbles that "flow out" with pressure (P2, P7, P10, P11). In Fingertip Fireworks, flicking felt sharp and directional, leading participants to imagine trajectories that shoot outward from the fingertip (P2, P7, P11). Similarly, in Meditative Garden, snapping suggested a small but decisive impulse, so participants intuitively accepted a flower blooming in response to each snap (P3, P5, P11). These reflections show that users naturally interpreted gestures as having implied motion patterns or directional launch—and expected animations to follow these implied dynamics.

doi:10.6342/NTU202600777

Causal Temporality. Participants also attended closely to how the timing of a gesture aligned with the timing of the animation. They intuitively distinguished between discrete gestures—short, punctuated actions with clear onset and offset—and continuous gestures—sustained motions whose intensity varies over time. In MR Stress Ball, participants emphasized that squeezing is continuous and should yield a correspondingly continuous release of bubbles; a single burst felt mismatched compared to a gradual flow (P1, P3, P5, P7, P10, P11). In Multi-User Fireworks, pen-clicking was seen as inherently discrete, making a crisp, toggle-like firework burst feel fitting (P1, P2, P3, P6, P7, P8). By contrast, participants found the Spinning Globe less intuitive when the rotation animation triggered only after the gesture ended, noting that the continuous swiping motion should have produced continuous rotational feedback (P3, P4, P6, P8, P9). These comparisons echo observations in other interactive systems [18, 28, 30] and underscore the importance of temporal alignment [4] in causally linking gesture and animation in visual fidget interactions.

In reviewing these interactions as well as participants' broader fidgeting habits, interviews revealed that users naturally differentiated between gestures performed with objects and those performed with the bare hand, noting that each evokes different expectations for how animations should respond. While earlier examples in our system focused primarily on bare-hand motions inferred from hand trajectory, participants frequently referenced the behavior of physical objects—such as spinners, sliders, and squeezable materials—when imagining what "should happen" visually. These distinctions, summarized in Figure 4.4 based on participants' lived fidgeting experiences, highlight that object-based and bare-hand gestures prompt different intuitive predictions about gesture–animation correspondence, shaping what users perceive as natural or satisfying in visual fidget interactions.

doi:10.6342/NTU202600777

Chapter 4 Authoring System

Based on our findings in formative studies, we introduce our visual fidget interface by walking through the creation of the *MR Stress Ball* example as shown in Figure 1.2A. The user starts by recording a fidget behavior (Figure 4.1A), and in the background, the system packs up features and recorded video frames to the inference model to gain suggestions of both visual themes and animations for the following setup.

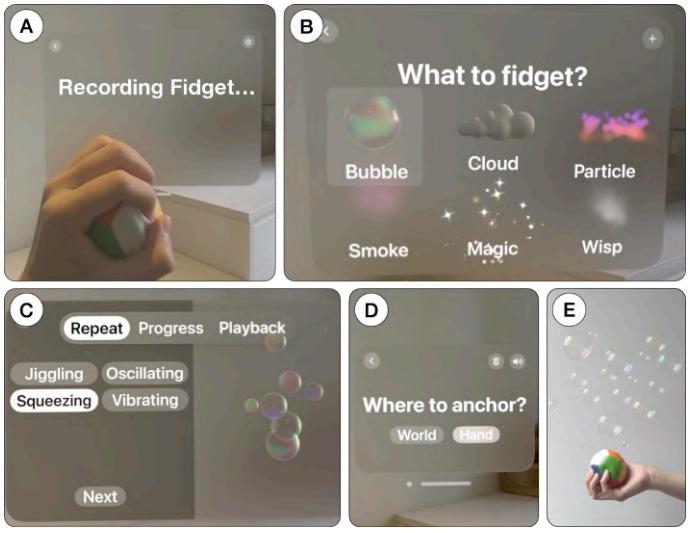

Figure 4.1: (A) The visual fidget interface begins by recording a user-performed fidget gesture at their natural pace. The system then formulates a gesture trigger based on this input. (B) A visual library appears, showing system-suggested visual models tailored to the context of the recorded fidget gesture. (C) Upon selecting a visual theme—here, a bubble—the system further suggests a set of context-relevant animations. These animations are categorized into tabs: repetitive, progressive, and playback, each offering different temporal characteristics. The user selects a ”squeezing out” animation under the repetitive category. (D) The system then prompts the user to anchor the fidget either to their hand or to the world. (E) After choosing to anchor to the hand, the system attaches the animation relative to the detected hand position. The user can now squeeze a ball to release bubbles that visualizes exerted pressure.

doi:10.6342/NTU202600777

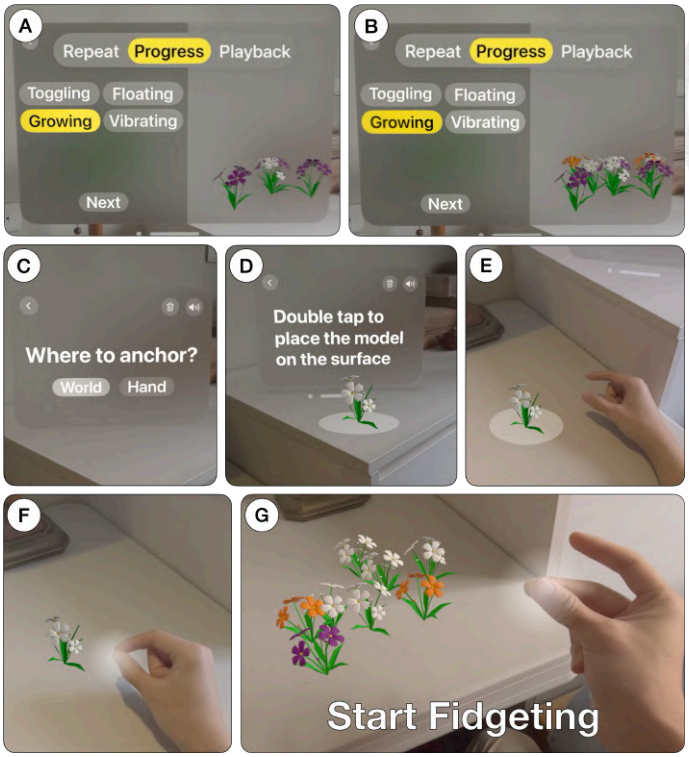

Figure 4.2: (A, B) The user here chooses the progressive animation, growing, and previews the daisy growing. (C, D, E) After the user chooses to anchor to the world, the system places the visual effect at the intersection of the user's head orientation and a detected surface. (F) The user confirms the placement by double tapping fingers. (G) The user starts fidgeting to form garden over time.

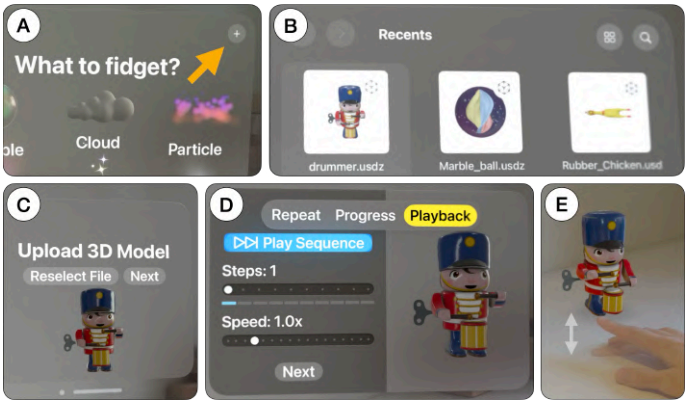

Figure 4.3: The system allows users to import custom models and edit their embedded animation playback, supporting versatile integration of rich 3D assets into personalized fidget interactions.

Next, the user chooses a visual theme from the suggested set of visual library; in this case, a bubble is selected (Figure 4.1B). Optionally, the user can import their own 3D object from the file system by tapping on the '+' button (Figure 4.3). After importing, the user can resize the 3D model by pinching the model in two hands, setting its default

doi:10.6342/NTU202600777

size for later use in the library. Then, the user choose animations from three options: repetitive, progressive, and playback. The first two options are from suggestion set, which are applicable to general visual theme without rigging, and the playback mode refers to the embedded animation from the 3D asset so that the user can utilize its orignal design by trimming, segmenting, or speeding up preset animations (Figure 4.3). Here, the user chooses the squeezing out animation under the repetitive tab (Figure 4.1C).

Finally, the system facilitates the user to anchor the configured fidget in their physical context. Anchoring offered two distinct experiences. Space-anchored visuals allowed users to choose whether to approach feedback, adding spatial anticipation and strengthening connections to contextual affordances (e.g., decorating a far skyline with fireworks). Hand-anchored visuals, by contrast, supported immediate, portable use similar to traditional handheld fidgets yet hand-free, such as spawning fireworks ephemerally. As the user chooses to anchor to the hand (Figure 4.1D), the user can now squeeze a stress ball to release bubbles that visualizes exerted pressure.

Figure 4.2 illustrate another example creation, the *Meditative Garden* (Figure 1.2B). This time, the user setup the snapping gesture, daisy visual, and growing animation under the progressive tab. The system then shows the preview of the daisy growing into a cluster (Figure 4.2A&B). As the user chooses to anchor to the world, the system automatically places the content on one of the detected surfaces (Figure 4.2D). The user can look around to explore alternative surfaces for placement—for example, repositioning the anchor onto a tabletop by double tapping to confirm (Figure 4.2E&F). The user can now generate flowers when snapping.

This workflow grounds a broader design framework that explores how mixed real-

doi:10.6342/NTU202600777

ity enables properties of visually enriched fidgeting that are non-achievable in traditional fidgeting.

4.1 Affordance Mapping

Affordance mapping explains how user gestures are translated into responsive animations by aligning the physical character of an action with the visual character of its response. Based on previous findings (section 3.7), to preserve the strong sense of agency central to fidgeting, both gesture motion and animation behavior must satisfy the same underlying constraints: the type of motion being expressed and the temporal character with which it unfolds. We therefore organize gesture–animation mappings along two coupled dimensions adapted from [19]: (a) the motional affordance provided by hands or objects, and (b) the temporality of the action. Together, these dimensions structure a broad space of fidget interactions (Figure 4.4).

4.1.1 Motional Affordance

Motional affordance describes how a gesture's physical motion—whether constrained by an object or defined by the hand' s trajectory and contact pattern—can be mirrored or abstracted in animation. It links how the gesture moves to how the animation should move, aligning visuals with users' embodied expectations.

Linear. The linear gesture is defined by motion that follows a mostly straight or gently arced path, whether performed with an object or with the bare hand. When using objects, the hand drives a rigid item through a one-degree-of-freedom translation—either reciprocating or stepwise—as in clicking a pen or pushing a slider. When using the bare

doi:10.6342/NTU202600777

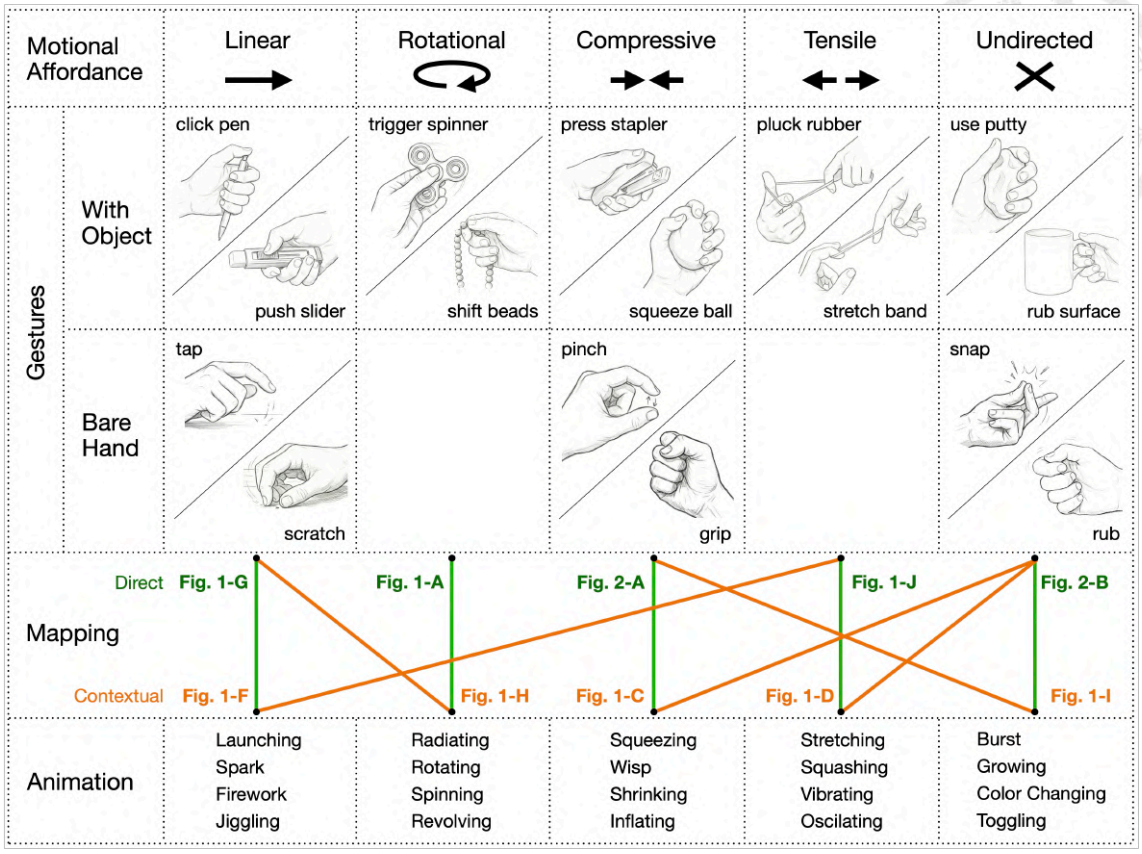

Figure 4.4: Each motional-affordance category (columns) includes object-based and bare-hand gesture examples, with discrete (top) and continuous (bottom) variants illustrating the causal temporality. Animations listed below each column are designed for that motional affordance. Pairing a gesture with an animation from the same category yields a direct mapping (green), while pairing across categories creates a contextual mapping (orange).

hand, the digits similarly move in a consistent direction with minimal shape change, as in scratching or tapping a surface. Across both cases, design can leverage animations that highlight path and directionality—for example, the directional spark triggered by scratching a table in the Figure 1.1G example or the drumming animation triggered by sequential tapping in the Figure 1.1B example.

Rotational. The rotational gesture involves 1-DOF rotation of a rigid object around an axis—continuous or oscillatory—as in spinning a fidget toy or shifting beads. This motion cannot be produced with the bare hand, as finger anatomy does not permit true axial rotation or the clicking, fidget-like feedback such objects provide. Design can lever-

doi:10.6342/NTU202600777

age rotating or centrifugal-style animations, such as the radiating animation triggered by spinning motion in the Figure 1.1A example.

Compression. The compressive gesture applies inward normal force to reduce the thickness or volume of a deformable target, whether through an object or the bare hand. With objects, compression occurs through pulsed or discrete force—such as squeezing a ball or pressing a stapler—with elastic restoration on release. With bare-hand actions, multiple digits converge toward a common focus (finger, thumb, or palm), increasing pressure and collapsing the distance between contact points, as in pinching or gripping. In both cases, design can use animations that highlight inward force and release—for example, the bubbles-squeezed-out animation triggered by squeezing in the Figure 1.2A example.

Tensile. The tensile gesture involves applying a pulling force to increase the distance between contact points on a deformable object, producing elongation along the pulling axis and elastic recoil upon release—as in plucking a rubber band or stretching a band. This gesture is inherently object-dependent, as the bare hand cannot produce tensile deformation or elastic rebound on its own. Design can leverage animations that elongate or tether to convey this elastic tension, such as the shooting animation triggered by plucking in the Figure 1.1F example.

Non-directional. Non-directional gestures lack a stable path or explicit focal direction, instead producing local or impulsive movements characterized by micro-shear or irregular deforming—such as snapping or rubbing fingers, or object-based actions like manipulating putty or rubbing a surface. Design can leverage in-place animations that reflect this localized motion, such as the growing effect triggered by snapping in the Figure 1.2B example.

doi:10.6342/NTU202600777

4.1.2 Causal Temporality

Another dimension of maintaining causal mapping is temporality, as gestures within the same motional affordance category can differ in being discrete or continuous, requiring animations that reflect these temporal profiles. For example, in Figure 4.4, pen-clicking and slider pushing are both linear, with-object gestures, yet clicking is discrete while pushing is continuous—demanding distinct visual responses to preserve intuitive causality. Our system supports both modes by analyzing the gesture's velocity curve during registration. Discrete gestures are event-based, with clear start/stop points, thresholded force release, and impulsive motion; animations for them benefit from punctuating effects like bursts, sparks, shots, or pops. Continuous gestures are flow-based, with smooth, sustained motion that can evolve or reverse at any point; animations here emphasize accumulation and resonance, such as gradual build-up, morphing, or lingering trails.

4.2 Animation Procedure

Varied sensory stimulation is fundamental to fidgeting, sustaining interest, regulating arousal, and providing an outlet for releasing pent-up energy. Our system extends this variety beyond traditional repetitive fidgeting by introducing two modes of feedback:

Progressive mode. Each input builds on the previous state, amplifying or modifying feedback over time. For example, the Meditative Garden scene gradually populated a user's desk with blooming flowers in response to continued fidgeting, which participants described as calming, rewarding, and fostering a sense of achievement.

Playback mode. To support more intricate designs, our system leverages the abun-

doi:10.6342/NTU202600777

dance of 3D assets and crafted animations. Users can segment or trim animations, as well as adjust playback speed, to explore diverse fidget interactions. For instance, drumming gestures triggered segmented animations of a drummer model, creating a sequential performance synchronized with user actions (Figure 1.1B).

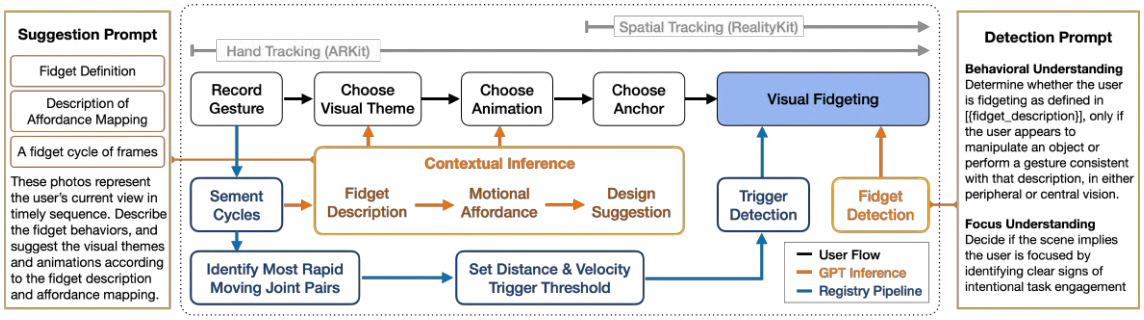

Figure 4.5: System workflow of visual fidget on Apple Vision Pro. Users sequentially configure fidget gesture input, visual theme, animation, and anchor placement. The system leverages GPT-based contextual inference to suggest visual effects and enhance trigger awareness. Hand and spatial tracking are enabled through ARKit and RealityKit.

4.3 Development

We implemented our system (Figure 4.5) on Apple Vision Pro using Swift, building on ARKit and RealityKit. The system consists of three tightly integrated components—a gesture registration pipeline and an inference module with a runtime adaptation layer—that together support capturing fidget gestures, interpreting their visual affordances, and rendering responsive MR animations.

4.3.1 Gesture Registration Pipeline

To support flexible gesture definition, the system begins with a brief prerecording phase that determines the user's handedness—right, left, or bilateral—based on wrist-joint placement and relative motion. Over a 2-second sampling window, the system monitors

doi:10.6342/NTU202600777

wrist positions: if both wrists remain closely aligned, bilateral use is inferred; otherwise, the higher wrist is selected as the active hand. This handedness estimation is used to anchor subsequent motion analysis.

The system then initiates a 5-second gesture-recording session using ARKit's 60 Hz joint-tracking pipeline. During this interval, video frames and 3D joint positions are continuously captured and normalized to the active wrist. Because many fidget gestures exhibit cyclic or repeated patterns, the captured signal is segmented by identifying intervals of high joint-speed summation. From these segments, the system computes the average movement speed for each joint pair and selects the three pairs with the greatest relative motion. These high-motion pairs are used to automatically derive distance and velocity vector thresholds that later drive real-time gesture detection. As our registration method is optimized for flexibility rather than fine-grained accuracy, this lightweight representation enables the system to accommodate a wide range of natural fidget behaviors. After recording, all video frames and motion features are passed to the Inference Module for interpretation.

4.3.2 Inference Module

Although the design space for MR visual effects is effectively open-ended, animations that align with a gesture's implied affordances tend to preserve a stronger sense of agency and thus serve fidgeting more effectively. Inspired by XR-Objects [6], the system leverages a large language model (GPT-5) to interpret the recorded gesture and generate visual candidates. The model receives the affordance-mapping prompt along with sampled video frames and outputs two ranked suggestion lists—one for visual themes and one for animations—filtered through the system's existing visual library. This inferred

doi:10.6342/NTU202600777

description is then used as the reference for subsequent fidget-behavior detection.

4.3.3 Runtime Adaptation

To support fidgeting without becoming disruptive [20, 24], the system integrates real-time context analysis inspired by Lindlbauer et al. [17]. Every 0.05 s, a first-person screenshot from Apple Vision Pro is processed by GPT-4o, chosen for its strong visual-scene understanding and low latency. GPT-4o performs two parallel inferences:

Fidget-behavior detection. Compare the live scene against the gesture description captured during registration (e.g., "spinning a fidget spinner with the right hand") to confirm whether the user is currently fidgeting.

Focus detection. Inspect cues such as hand posture (typing, writing), tool engagement, and gaze direction toward screens or work surfaces to estimate the user's attentional state.

Based on these signals, the system dynamically modulates the opacity of visual feedback—for instance, fading animations during deep focus or amplifying them during idle moments. This scene-understanding approach was considered as a potentially more adaptable alternative to proximity-based heuristics in complex environments. Although GPT-4o's inference time averages 1.5 s, this latency is well-matched to the intention of detecting sustained fidgeting and sustained attention states rather than momentary fluctuations, enabling smooth and non-intrusive visual modulation.

doi:10.6342/NTU202600777

Chapter 5 Validation Study

To validate whether visual fidget allows users to create visually enriched fidgeting, we conducted a study to observe how users use our interface described in chapter 4 and collect their feedback.

5.1 Participants

We recruited 9 participants (5 female, 4 male), aged between 20 and 59 years ($\mu = 31.4$, $\sigma = 14.4$). Using the same inclusion and exclusion criteria as in Study 2, participants were recruited from various on-campus locations to diversify testing scenarios and potential interaction outcomes.

5.2 Task and Procedure

The study consisted of three main phases conducted over a 40-minute session. It began with an introduction and familiarization period, during which participants were introduced to the concept of fidgeting and its relevance to the study. After fitting the Apple Vision Pro headset and completing calibration, participants engaged in a guided trial to become familiar with the visual fidget interface and interaction workflow. They

doi:10.6342/NTU202600777

were then encouraged to independently explore and create their own fidget interactions, experimenting with different combinations of gestures, visuals, animations, and spatial placements. At the end, a semi-structured interview was conducted to gather participants' reflections on their favorite creations and to assess the usability of the system. Throughout the session, in-situ experiences were recorded using the headset's built-in video capture feature.

5.3 Findings

Overall, participants found the system intuitive and effective at transforming their everyday fidget habits into visually enriched, often magical experiences (Figure 1.1). As they iterated on their designs, they reflected on how they chose and shaped animations, revealing several recurring mapping strategies that preserved an intuitive sense of motion–effect causality. These relationships are summarized in Figure 4.4.

- **Direct Mapping**. Many participants selected animations that closely mirrored the kinetic dynamics of their gestures. P21 drew from fire performance when designing a fidget spinner that emits radial shockwaves as it spins (Figure 1.1A), emphasizing the centrifugal dynamics. P22 mapped the frictional motion of scratching to sparking animation (Figure 1.1G), describing the momentary sensation as "like having Superman's strength." P23 built a squeezing interaction similar to Figure 1.2A but enhanced it with a translucent particle outburst to evoke a sci-fi feel.

- **Contextual Mapping**. Other participants chose animations that were not direct kinetic matches but became meaningful through imagined mechanics or tactile associations. P21 connected plucking a rubber band to a slingshot, creating a marble-

doi:10.6342/NTU202600777

launching animation that conveyed tension and release (Figure 1.1F). P24 mapped a linear flick to a rotational spin by treating the flick as a trigger for a virtual rotor. P28 used a squeezing gesture to imply pressure driving material out of melt-gun nozzles, paired with an in-place bursting animation (Figure 1.1I). Tactile associations also shaped decisions: P29 linked rubbing a mug handle to the generation of warmth, releasing rising steam (Figure 1.1C), while P26 amplified the micro-tactile feel of rubbing with noticeable vibration animation (Figure 1.1D).

Beyond these gesture-to-animation mappings, participants also explored environmental visual concepts, treating the surrounding space as a creative fidget canvas. P25 assembled vibrating bubbles into an orchestra-like sculpture (Figure 1.1E), evoking the sensation of conducting the *1812 Overture*. The same participant decorated a city skyline with fireworks, describing this as "environmental fidgeting." P26 used a laptop screen as a site for motivational text and vibration triggered by rubbing during stressful moments, while P27 scattered jiggling flowers throughout a lecture room to self-entertain when facing unresponsive students. These examples show how participants extended their designs into the environment when imagining richer or more atmospheric experiences.

Participants also praised the system's support for custom asset import, which enabled more personalized and imaginative outcomes. P24, for example, animated a spinning "OIIA OIIA" cat meme triggered by a flick gesture (Figure 1.1H). At the same time, several participants (P21, P23, P24, P25) suggested providing more descriptive animation previews to clarify the behavior of suggested or imported effects before finalizing their designs.

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

Chapter 6 Discussion

The results of our study offer evidence supporting the integration of visual effects in fidgeting. In this section, we discuss the findings based on the insights and feedback gathered from our participants, including emergent themes derived from questionnaire responses and post-study semi-structured interviews. Additionally, we identify potential avenues for future improvements based on these findings.

6.1 Visual Effect as a Catalyst for Fidgeting

P2 captured the shared enthusiasm of our participants by stating, "*I can do this whole day*" (P1, P2, P3, P5, P7, P8, P9, P10, P11, P15, P19). The incorporation of visual effect into fidgeting seems to foster a stronger motivation to engage in fidgeting, even among those who typically do not fidget. P12 admitted, "*I don't usually fidget, but I'd love to start if I could take your device with me.*"; or as P15 and P16 noted, "I never thought I'd enjoy playing with a stress ball this much—turns out the fancy visuals make all the difference," demonstrating visual fidget's potential to form new behaviors and habits. P3, P5, and P18 described how actions not typically linked to fidgeting, such as flicking without targeting an object, became redefined as fidgeting when it binds to launching virtual fireworks. P2 reflected, "After using MR Fidget, I wanted to develop a fidgeting habit." Participants

doi:10.6342/NTU202600777

also praised the amplified stimulation compared to traditional fidgets. As P5 put it, "The fireworks are so cool…more extra, more crazy, and more fun." P2, P10, P11, and P19 noted that even small gestures, such as clicking a pen to trigger particles, became more engaging and playful. These observations underscore how visual effect can expand the vocabulary of fidgeting.

While visual fidget is grounded in mixed reality, its underlying design framework offers broader applicability to the study of visually enriched fidgeting and interactive media. For example, our affordance mapping approach could be extended to smart environments: inspired by EchoSight's look-and-control paradigm [15], a system might infer functional affordances of IoT devices—such as a Philips Hue lightbulb—and allow users to modulate color or rhythm through continuous fidget gestures. Similarly, drawing from Kim et al.'s swarm robotics [13], a system could map the popping motion of a virtual bubble wrap to invisible force dispersions across a matrix of robots, generating emergent movement patterns based on gesture dynamics. These speculative extensions demonstrate how the logic of programmable fidgeting can inform future applications across ubiquitous computing and embodied interaction design.

6.2 Enhancing Self-Awareness and Reducing Disruptive Habits

Fidgeting is often perceived as socially inappropriate by many, as evidenced by prior research on misokinesia [9]. Participants frequently expressed concerns about their own disruptive behaviors and the possibility of bothering others. This highlights a dual need: fostering self-awareness to regulate one's own fidgeting and reducing the negative social

doi:10.6342/NTU202600777

impact on bystanders.

Incorporating progressiveness into visual fidget tools enhances self-awareness and improves behavior regulation more effectively than the repetitive feedback typically provided by traditional fidget tools. Progressive feedback, like snapping fingers to grow flowers, helped users become more aware of their fidgeting. P3 observed, *"the progressive feedback offered by the Meditative Garden makes me aware of my fidgeting frequency"*. P7 and P19 noted that progressiveness turned unconscious fidgeting into a conscious action. This progressive feedback helps users stay aware of and control their fidgeting. For example, P1 and P9 mentioned that the color of the flowers motivates them to slow down, as they prefer the calmness represented by white flowers, prompting them to fidget more mindfully. This suggests that this type of ambient display not only reduces the barrier to self-awareness but also presents opportunities for guiding users' fidgeting habits.

6.3 Emotion Regulation

visual fidget shows potentials in facilitating emotion regulation by aiding users in shifting focus and managing emotions across various contexts. For instance, P1 described visual fidget as a means to "escape to a personal bubble", which effectively redirected their attention and alleviated anxiety at the beginning of the user study; P14 described the firework interactions as "a good way to release some feelings or thoughts by looking at the fireworks". P2 indicated a desire to use visual fidget for attention redirection during high-pressure situations, such as performance preparation. P10 appreciated using the MR Stress Ball to generate bubbles symbolizing stress, finding comfort in watching them drift away. Additionally, P1, P3, P5, P9, P11, P15, P16, and P19 reported that visual fidget helped

doi:10.6342/NTU202600777

alleviate boredom while watching the video tutorials, which they found tedious. P10 even suggested that visual fidget could facilitate "micro-meditation," offering a quicker, more accessible way to achieve a meditative state compared to traditional practices, making mental regulation more integrated and effortless in everyday tasks. These findings are promising and suggest that future research should investigate the use of visual fidget for emotion regulation.

doi:10.6342/NTU202600777

Chapter 7 Limitations & Future Work

Overall, though participants responded positively to visual fidget, feedback highlighted opportunities to improve its long-term engagement. To address these limitations, we outline several directions for future work:

7.1 System Constrains

While our system supports a wide range of fidget behaviors, certain types remain unsupported due to limitations in system design. For example, because the system relies on camera-based hand tracking from a headset, fidget behaviors that are out-of-sight or occluded cannot be captured, such as face-scratching or fabric-rubbing, which could benefit from non-optical tracking methods like gloves [29]. Furthermore, as our system is designed for hand-involved fidget behaviors, whole-body movements like shaking legs or rocking a chair) are not supported. These gaps echo user feedback—for example, participants P2, P10, and P18—who expressed a desire for broader support beyond hand-based gestures.

Additionally, while our inference model is primarily designed to support the creative process within the authoring tool, its capabilities remain underutilized. Currently, contextual inference relies only on the name and metadata of the visual themes and anima-

doi:10.6342/NTU202600777

tions, limiting its usefulness in recognizing nuanced affordances. A promising direction for future work is enabling the model to analyze the geometry of 3D assets to give out suggestions [26]. This would not only empower more expressive design creation but also facilitate more intelligent suggestions within the authoring interface. Furthermore, technical evaluation will be necessary to rigorously assess the model's ability to generalize across diverse fidget scenarios and to unlock its full potential beyond contextual streamlining.

7.2 Experiment Design

This study focused on exploring fidget design insights, but future work should address potential novelty effects and include baseline comparisons to strengthen the findings.

Novelty Effect. While the short interaction period in our study was intentionally designed to familiarize participants with the MR fidget interactions and prevent fatigue, it has been an effective approach for generating fidget design insights, as evidenced in similar research [13, 16]. However, the potential influence of the Novelty Effect remains underexplored. For example, several participants (P7, P8, P11, P14) experienced excessive excitement during the study, as it was their first time using an MR headset. This initial excitement may have influenced their evaluations of the interactions.

To address this issue, future research could adopt a longitudinal, field study approach, as suggested by [11], to better understand fidgeting behaviors over time. Longitudinal methods, often used in psychology research [22, 27], have primarily focused on measuring performance outcomes influenced by fidgeting behaviors. However, applying such methods to fidget design studies may introduce challenges such as recall bias, contextual

doi:10.6342/NTU202600777

variability, and difficulty in capturing serendipitous insights, all of which merit further investigation.

Baseline Comparisons with Traditional Tools. While traditional fidget tools were introduced to help participants recall their fidgeting experiences and familiarize themselves with the counterparts of our MR fidget interactions, we did not measure the usage of these traditional tools as a baseline for comparison.

Some evidence suggests that MR redefined actions not typically associated with fidgeting—such as flicking in air to launch virtual fireworks (P3, P5, P18)—as fidgeting behaviors. Furthermore, MR feedback was reported to enhance enjoyment and engagement with interactions that participants (P15, P16) had not previously considered part of their fidgeting habits. However, incorporating quantitative analysis of both traditional tools and their MR counterparts could provide a stronger foundation for these claims.

While the goal of this paper is to explore fidget design insights beyond soley haptic experience, we urge future studies to integrate baseline measures for both traditional and MR fidget tools. Such comparisons would enhance the rigor and generalizability of findings regarding the impact of MR mediation on fidgeting behaviors.

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

Chapter 8 Conclusion

In this work, we introduced visual fidgets and a gesture–animation affordance mapping grounded in everyday fidgeting. We built an MR authoring system that captures real-time gestures, leverages LLM-driven animation suggestions, and enables rapid exploration. Our studies show how users can transform familiar fidget behaviors into expressive MR interactions that support playful, personal well-being.

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

References

[1] R. Arora, R. H. Kazi, D. M. Kaufman, W. Li, and K. Singh. Magicalhands: Mid-air hand gestures for animating in vr. In Proceedings of the 32nd annual ACM symposium on user interface software and technology, pages 463–477, 2019.

[2] K. B. Aspiranti and D. M. Hulac. Using fidget spinners to improve on-task classroom behavior for students with adhd. Behavior analysis in practice, 15(2):454–465, 2022.

[3] L. Biel. Fidget toys or focus tools. Autism File, 74:12–13, 2017.

[4] E. D. Chase, T. Gerstenberg, and S. Follmer. Realism of visual, auditory, and haptic cues in phenomenal causality. In 2023 IEEE World Haptics Conference (WHC), pages 306–312. IEEE, 2023.

[5] S. B. da Câmara, R. Agrawal, and K. Isbister. Identifying children's fidget object preferences: toward exploring the impacts of fidgeting and fidget-friendly tangibles. In Proceedings of the 2018 Designing Interactive Systems Conference, pages 301–311, 2018.

[6] M. D. Dogan, E. J. Gonzalez, K. Ahuja, R. Du, A. Colaço, J. Lee, M. Gonzalez-Franco, and D. Kim. Augmented object intelligence with xr-objects. In Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology, pages 1–15, 2024.

doi:10.6342/NTU202600777

[7] P. J. Drew, A. T. Winder, and Q. Zhang. Twitches, blinks, and fidgets: important generators of ongoing neural activity. The Neuroscientist, 25(4):298–313, 2019.

[8] K. Isbister. Fidget toys aren't just hype. The Conversation, 2017.

[9] S. M. Jaswal, A. K. De Bleser, and T. C. Handy. Misokinesia is a sensitivity to seeing others fidget that is prevalent in the general population. Scientific reports, 11(1):17204, 2021.

[10] C. Ji and K. Isbister. Ar fidget: augmented reality experiences that support emotion regulation through fidgeting. In CHI Conference on Human Factors in Computing Systems Extended Abstracts, pages 1–4, 2022.

[11] M. Karlesky and K. Isbister. Designing for the physical margins of digital workspaces: fidget widgets in support of productivity and creativity. In Proceedings of the 8th international conference on tangible, embedded and embodied interaction, pages 13–20, 2014.

[12] M. Karlesky and K. Isbister. Understanding fidget widgets: Exploring the design space of embodied self-regulation. In Proceedings of the 9th Nordic Conference on Human-Computer Interaction, pages 1–10, 2016.

[13] L. H. Kim, V. Domova, Y. Yao, and P. Rajabi. Swarmfidget: Exploring programmable actuated fidgeting with swarm robots. In Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, pages 1–15, 2023.

[14] S. L. Kriescher, D. M. Hulac, A. M. Ryan, and B. L. King. Evaluating the evidence for fidget toys in the classroom. Intervention in School and Clinic, 59(1):66–69, 2023.

doi:10.6342/NTU202600777

[15] J. Li, Q. Yang, K. Xu, Y. Zhang, and C. Xu. Echosight: Streamlining bidirectional virtual-physical interaction with in-situ optical tethering. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, pages 1–18, 2025.

[16] R.-H. Liang, B. Yu, M. Xue, J. Hu, and L. M. Feijs. Biofidget: Biofeedback for respiration training using an augmented fidget spinner. In Proceedings of the 2018 CHI conference on human factors in computing systems, pages 1–12, 2018.

[17] D. Lindlbauer, A. M. Feit, and O. Hilliges. Context-aware online adaptation of mixed reality interfaces. In Proceedings of the 32nd annual ACM symposium on user interface software and technology, pages 147–160, 2019.

[18] K. Monteiro, R. Vatsal, N. Chulpongsatorn, A. Parnami, and R. Suzuki. Teachable reality: Prototyping tangible augmented reality with everyday objects by leveraging interactive machine teaching. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, pages 1–15, 2023.

[19] R. Nguyen, C. Gouin-Vallerand, and M. Amiri. Hand interaction designs in mixed and augmented reality head mounted display: a scoping review and classification. Frontiers in Virtual Reality, 4:1171230, 2023.

[20] R. Nyqvist. Fidgeting for creativity. 2016.

[21] K. Perrykkad and J. Hohwy. Fidgeting as self-evidencing: A predictive processing account of non-goal-directed action. New Ideas in Psychology, 56:100750, 2020.

[22] J. Persia. Examining the impacts of subtle fidget jewelry on anxiety, stress, and attention. 2023.

[23] O. Ricciardi, P. Maggi, and F. D. Nocera. Boredom makes me'nervous': fidgeting as

doi:10.6342/NTU202600777

a strategy for contrasting the lack of variety. International journal of human factors and ergonomics, 6(3):195–207, 2019.

[24] S. H. Ross, N. Sullivan, and J. A. Yoon. Virtual fidgets: Opportunities and design principles for bringing fidgeting to online learning. In Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems, pages 1–6, 2023.

[25] A. Sayara, E. L. Chen, C. Nguyen, R. Xiao, and D. Yoon. Gesturecanvas: A programming by demonstration system for prototyping compound freehand interaction in vr. In Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, pages 1–17, 2023.

[26] J. Shi, Z. Guo, S. Cheng, Y. Liu, M. Zhang, and Z. Xiao. 3d object recognition based on point cloud geometry construction and embeddable attention. In Image and Graphics : 12th International Conference, ICIG 2023, Nanjing, China, September 22–24, 2023, Proceedings, Part III, page 235 – 246, Berlin, Heidelberg, 2023. Springer-Verlag.

[27] S. Stalvey and H. Brasell. Using stress balls to focus the attention of sixth-grade learners. Journal of At-Risk Issues, 12(2):7–16, 2006.

[28] T. Wang, X. Qian, F. He, X. Hu, Y. Cao, and K. Ramani. Gesturar: An authoring system for creating freehand interactive augmented reality applications. In The 34th Annual ACM Symposium on User Interface Software and Technology, pages 552–567, 2021.

[29] D. Weimer and S. K. Ganapathy. A synthetic visual environment with hand gesturing and voice input. In Proceedings of the SIGCHI Conference on Human Factors

doi:10.6342/NTU202600777

in Computing Systems, CHI '89, page 235–240, New York, NY, USA, 1989. Association for Computing Machinery.

[30] J. O. Wobbrock, M. R. Morris, and A. D. Wilson. User-defined gestures for surface computing. In Proceedings of the SIGCHI conference on human factors in computing systems, pages 1083–1092, 2009.

[31] N. Zhou, Y. Sun, S. Devleminck, and L. Geurts. Squeezable interface for emotion regulation in work environments. In Proceedings of the Eighteenth International Conference on Tangible, Embedded, and Embodied Interaction, pages 1–7, 2024.

doi:10.6342/NTU202600777

doi:10.6342/NTU202600777

Appendix A — Codebook

Code Name	Definition	Example Excerpt	Application Notes
Motional Mapping	The perceived match or mismatch between a user's motion and the resulting visual or animated feedback.	"I enjoy how squeezing the ball makes bubbles gently float out—it feels just like pressing a sponge." (P16, MR Stress Ball)	Use this code when participants note how well or poorly their gestures align with visual or animated feedback.
Animation Procedures	Repetitive feedback replicates traditional fidget tools, while progressive feedback builds cumulatively.	"The progressive feedback offered by the Meditative Garden makes me aware of my fidgeting frequency" (P3, Meditative Garden)	Use for feedback that is either consistent and repetitive or accumulates over time, such as growing flowers.
Adjustable Feedback	Modifications to visual or auditory feedback, including muting, swapping, or augmenting feedback.	"Being able to customize the speed, stimulation, and colors of the fireworks was a nice touch for personalization" (P17, Multi-user Fireworks)	Includes changes to sounds, visuals, and pseudo-haptics to enhance fidgeting experiences.
Beyond Physical Constraints	Unique MR-specific interaction methods that go beyond physical constraints.	"I can move it indirectly with my resting hand… It doesn't require much effort to move it." (P2, Spinning Globe)	Use for gestures, indirect interactions, or novel MR-enabled input methods that are not possible with traditional tools.
Multi-user	Social dynamics and interactions involving multiple users in a shared MR fidgeting experience.	"Shooting fireworks with others feels like co-creation or bonding to me." (P16, Multi-user Fireworks)	Applies to collaborative or competitive multi-user scenarios, whether co-located or remote.
Context Awareness	Adaptation of feedback or interaction based on spatial, social, or environmental context.	"Fidgeting in MR, such as growing a flower on a table, fosters a deeper connection to the environment." (P15, Meditative Garden)	Includes spatial or social awareness that tailors the interaction to the user's surroundings.
Playfulness	User enjoyment and engagement derived from creative, exploratory, or humorous aspects of MR interactions.	"The fireworks are so cool, and there are so many colors and lights… it's just more extra, and more crazy, and more fun" (P5, Fingertip Fireworks)	Emergent theme reflecting the fun and playful nature of interactions.
Habit Formation	Potential for MR interactions to foster new or redefined fidgeting behaviors.	"After using it, I wanted to develop a fidgeting habit." (P2)	Includes comments about the motivation or ability to form consistent fidgeting habits.
Customization	Flexibility to modify interactions or use personal objects as fidget tools.	"I could see myself turning a desk figurine into a fidget tool by programming it to respond to my movements." (P10)	Includes both user-directed customization and adaptability of the MR system.
Self-Awareness	Users' recognition of their own fidgeting behaviors.	"It's like a life record. I look forward to seeing how it will grow the next time after working sessions." (P13, Meditative Garden)	Applies to moments of reflection or realization of fidget behaviors.
Peripheral Use	Fidgeting interactions designed to remain in the user's peripheral vision, supporting secondary tasks.	"It doesn't occlude my visual" (P10, Spinning Globe)	Includes comments about interactions that do not obstruct primary tasks.
Emotional Regulation	Use of MR interactions to manage emotions or redirect focus in stressful situations.	"Bubbles are reminiscent of stress, and releasing them makes me feel alleviated." (P10, MR Stress Ball)	Applies to interactions that aid in stress relief, anxiety management, or emotional balance.
Concentration	Feedback on the role of MR interactions in affecting focus, either positively or negatively	"Sudden firework explosions helped me refocus during tutorial videos" (P11, Fingertip Fireworks)	Use this code when participants discuss how MR interactions impact focus, including supportive or critical feedback.
Fidget Without Hands	Exploration of fidgeting methods beyond hand-based interactions	"I would like to use MR to generate ripples when doing foot-tapping." (P18)	Emergent theme exploring alternative modes of engagement beyond hand-based fidgeting.

Figure A.1: The codebook used for analyzing the qualitative data. It contains code names, definitions, example excerpts, and application notes for each code.

doi:10.6342/NTU202600777