

Grounding Pedagogical Intents in Embodied AI Teachers: A Human-Centered Framework for Designing and Evaluating Instructional Gestures

Authors

Lai Wei[†]

Brain, Language, and Computation Lab

Department of Language Science and Technology
The Hong Kong Polytechnic University
Hong Kong, China

lai.wei@polyu.edu.hk

0000-0002-5476-5450

Sark Pangrui Xing[†]

School of Design

The Hong Kong Polytechnic University
Hong Kong, China

sark.xing@connect.polyu.hk

0000-0002-2273-4772

Kenny K.N. Chow

School of Communication

Hong Kong Baptist University

Hong Kong, China

knchow@hkbu.edu.hk

0000-0002-8368-0157

Stephen Jia Wang^{*}

School of Design

The Hong Kong Polytechnic University

Hong Kong, China

stephen.j.wang@polyu.edu.hk

0000-0001-9835-9932

[†]These authors contributed equally to this work.

^{*}Corresponding author.

Abstract

The proliferation of generative AI in HCI offers new possibilities for creating embodied agents, yet a significant gap persists between high-level principles and concrete design practices. This gap is particularly acute for co-speech gestures of pedagogical agents (PAs), where fully automated text-to-gesture generation often fails to capture the pedagogical nuance that effective instruction demands. To bridge this gap, we adopted a Research-through-Design approach to develop a human-centered design framework that systematically translates educators' pedagogical intent into implementable gesture specifications for embodied AI teachers. The framework structures a collaborative and iterative process across four stages: (1) *Preparation*: multimodal corpus analysis of authentic teaching to identify gesture patterns and their instructional functions; (2) *Human PA Acting*: performance-based rehearsal by the educator and designer to externalize tacit pedagogical knowledge; (3) *Embodied PA Acting*: human-to-agent motion transfer via video-based pose estimation and manual joint correction; and (4) *EPA-assisted Course Delivery*: student-centered evaluation through qualitative interviews and thematic content analysis. Applying this process yielded a 14-minute EPA-assisted course module, which was delivered to and evaluated by 38 university students. Qualitative analysis revealed that students perceived the PA as professional and approachable, describing enriched engagement through cohesive gestures that smoothed conceptual transitions and beat gestures that reinforced instructional rhythm. This work advances existing knowledge in three ways: first, it operationalizes HCAI principles through a replicable four-stage framework that bridges generative synthesis and pedagogical intent; second, it empirically links specific gesture design choices (i.e., word-synchronized cohesive and beat gestures) to student perceptions of instructional quality; third, exploratory findings suggest that structured prompts with temporal phases and linguistic anchors enable more coherent pedagogical outcomes in generative AI tools.

Keywords: Human-Centered AI; Design Methodology; Pedagogical Agent; Gesture; Generative AI

1 Introduction

Generative AI (GenAI) has become a novel tool in Human-Computer Interaction (HCI) and design research, transforming the ways in which we conduct HCI research and design practices [1]. Notably, scholars have employed GenAI to generate Embodied Agents (EAs) with co-speech gestures, and it is progressively becoming capable of performing conversational tasks [2, 3]. Recently, joint efforts have advanced to leverage GenAI for developing gesture synthesis systems that generate semantically relevant gestures aligned with speech rhythm and content for Embodied Conversational Agents (ECAs) [4, 5]. Meanwhile, in the landscape of digital learning, the application of Embodied Pedagogical Agents (EPAs)¹ has emerged as a transformative field for enhancing educational engagement and learning performance [6, 7, 8]. The development and design of these agents encompass several dimensions, including EPAs’ appearance [9, 10], acoustic attributes [11], interaction [12], personality [13, 14, 15], and voices [16, 17, 18, 19]. Additionally, particular attention has been given to nonverbal cues [20, 21], which have been extensively studied across various disciplines, including education, psychology, and computing [22, 23, 24, 25].

EPAs in digital learning environments have been further explored in areas such as educational virtual reality (VR) [26, 27] and intelligent EPAs [28]. Among these applications, anthropomorphic characteristics, such as co-speech gestures, have become a primary focus [29, 30, 31, 32]. These gestural cues direct learner attention, clarify concepts, support retention, and guide transitions between knowledge points—making their pedagogically grounded design central to effective learning outcomes [33, 34, 35]. Regarding existing motion synthesis for generating PA, there are two approaches. One reconstructs motion through pose estimation and retargeting, inferring full-body movement from sparse data and refining articulation through kinematic modeling, as seen in DeepMotion [36]—this approach is time-intensive and closely dependent on the motion-tracked performer’s behavior [37, 38]. The other generates motion through video synthesis, predicting coherent spatiotemporal sequences for fluid, context-aware movement, exemplified by large vision models

¹Throughout this paper, we use PA (Pedagogical Agent) as the broader category that includes both non-embodied and embodied instructional agents, and EPA (Embodied Pedagogical Agent) to refer specifically to agents with a three dimensional physical or virtual form capable of producing body movements. When describing the visual style of our specific design artifact, a minimalist and faceless 3D avatar, we use 3D animated avatar as an appearance descriptor rather than a category designation.

such as KlingAI and Sora [39]. This paradigm operates through prompt-driven generation (text or image to video), producing outputs that vary considerably across runs, with limited user control over gesture semantics, timing, or pedagogical alignment [40, 41]. Leveraging these technologies, commercial platforms such as HeyGen ², Synthesia ³, and iFlyTek ⁴ offer AI-generated instructors featuring anthropomorphic characteristics—including facial expressions, lip synchronization, and natural body movements—through text-to-speech (TTS) and automated animation.

While a substantial body of literature affirms that gestures performed by EPAs positively enhance student learning outcomes [42, 43, 44, 6], recent findings suggest that these benefits are often inconsistent. This variability indicates that the pedagogical value of gestures is highly contingent upon specific gesture categories [33], their interaction with other modalities (e.g., agent appearance [30]), and the context of use [32]. Such complexity underscores an urgent need for foundational design frameworks that remain largely unestablished. Although recent advancements have integrated computer vision with biomechanical skeletal modeling to achieve high-fidelity movement reconstruction [45], significantly enhancing the alignment between visual pose estimations and skeletal rigging, this technical precision has yet to be guided by the pedagogical design standards necessary for effective instructional behavior.

Specifically, while existing research has made significant technical advances in gesture synthesis (pose estimation [45], semantic retrieval [4]), and documented gesture effectiveness in education [22, 46], a critical gap remains between technical precision and pedagogical authenticity: no systematic framework translates pedagogical intent into implementable gesture specifications grounded in authentic teaching practices. Current research has yet to address whether the proportional distribution of gesture types—as observed in naturalistic teaching—is essential for EPA realism and efficacy, or how specific gesture designs influence students’ perceptions of instructional quality. Without such ecological behavioral benchmarks, Pedagogical Agent (PA) gesture design remains predominantly intuitive and fragmented. Consequently, current generative technologies have not resolved the fundamental challenge of creating PAs that are both behaviorally realistic and pedagogically effective. We identify three critical, interrelated problems:

²HeyGen. Available at: <https://www.heygen.com/>

³Synthesia. Available at: <https://www.synthesia.io/>

⁴iFlyTek. Available at: <https://iflytek.com>

- (1) **Lack of a Design Framework:** There is no structured, replicable framework to guide the design of instructional gestures. This absence impedes collaboration between educators and designers, making it difficult to systematically translate specific **instructional needs** into meaningful, co-speech gestures that are aligned with pedagogical goals.
- (2) **The Connectivity Gap Between Human Behavioral Data and PA Design:** A significant gap persists in translating empirical research on human instructional gestures into PA design. Specifically, few studies correlate authentic teacher behaviors (e.g., cohesive or beat gestures) with measurable student perceptions in the learning process. Without understanding how a PA's gestures influence key HCI metrics, such as learner perceptions of credibility and professionalism, evidence-based design choices remain difficult. Consequently, the lack of behavioral benchmarks leaves even advanced generative models disconnected from the nuanced pedagogical functions of human movement.
- (3) **Unexplored Potential and Pitfalls of Generative AI:** Current text-to-video generation tools for instructional content lack systematic evaluation. It remains unclear how to structure prompts to capture pedagogical intent and temporal coherence. A critical assessment of generation workflows is necessary to establish effective design strategies for pedagogically grounded gesture creation.

Research has demonstrated that instructional gestures for embodied agents can foster stronger social connections, which in turn enhance student learning performance [22, 46]. By enriching the social agency of virtual instructors, these gestural behaviors directly cater to the burgeoning trend toward AI-driven personalized and immersive learning environments [47, 48]. Given the increasing demand for such high-fidelity interaction, this study seeks to establish formalized design frameworks to support the systematic development of gestures in AI teachers. As such, we investigate the collaborative instructional gesture design process between PA designers and educators, explore how generative technologies can be leveraged within this workflow, and examine student perceptions of the resulting multimodal instructional information.

This study is grounded in Research-through-Design (RtD) [49], where an educator and a PA designer collaboratively explore the design of embodied PAs with instructional meaning, using state-of-the-art technologies to fine-tune PAs as desired. Initially, we observed thirteen design-related lectures to analyze the frequency of different gesture types used by educators. This was followed by the formation of a team consisting of a PA designer and educator, who worked together through collaborative acting and rehearsal. Throughout the gesture design process, the team aligned their objectives to ensure the gestures met educators’ instructional needs and improved the clarity of the PA’s expressions. Using the RtD framework, the team went through four stages of the design process: Preparation, Human PA Acting, Embodied PA Acting, and EPA-Assisted Course Delivery. During each stage, both team members engaged in iterative practices, delivering the prepared presentation, making gestures intuitively, and emulating real-life teaching scenarios. The PA designer observed the educator’s gestures, and both reflected on how to better express co-speech gestures to convey instructional meaning. Building on insights from previous research on word-synchronized gestures [50, 51, 52] and the alignment of gestures with parts of speech (PoS) [25, 53]⁵, we explored the use of DeepMotion, Sora, and KlingAI to evaluate the quality of animated instructional gestures for PAs. The result of this design and acting process was a 14-minute course featuring instructional gesture-based PAs. To evaluate learning efficacy, we tested the course with 38 validated Hong Kong-based university students.

This paper makes the following contributions to the field of HCAI methodologies:

- (1) **A Human-Centered Framework for Gesture Design:** We detail a replicable **human-centered design framework** for creating pedagogically effective gestures. Unlike prior work that focuses on technical naturalness or general gesture generation, our framework operationalizes the translation of educators’ tacit pedagogical knowledge into implementable specifications through structured corpus analysis, collaborative rehearsal, and student-centered validation. This four-stage process is grounded in the analysis of authentic teaching and uses reflective, collaborative practices between educators and designers to translate pedagogical intent into

⁵Parts of Speech refer to the grammatical categories used to describe the function of a word in a sentence, including nouns, verbs, adjectives, and adverbs.

agent embodiment, serving as a practical benchmark for advancing teaching practices via multimodal instructional design and educational technology development.

- (2) **An Empirical Evaluation of the PA and Its Impact:** We conducted a formal **empirical evaluation** of a pedagogical agent built using our framework. The results validate our approach: the agent's use of cohesive and beat gestures led to high student ratings of its professionalism, credibility, and approachability, with students describing enriched engagement and positive learning experiences.
- (3) **Insights into Generative AI and Future Directions:** Through exploratory testing and quantitative evaluation, we demonstrate that structured prompts with temporal phases and linguistic anchors substantially improve instructional coherence in text-to-video generation. This perspective supports a "human-in-control" workflow and reveals practical considerations for developing future instruction-focused AI systems.

2 Background and Related Work

2.1 Human-centered AI and the methodological gap

The notion of Human-Centered AI (HCAI) advocates for a paradigm shift, emphasizing AI systems that amplify and augment human capabilities rather than merely replace them [1, 54]. Key principles of HCAI include ensuring human control, transparency, fairness, and promoting well-being [55]. These principles are further substantiated by evidence showing that outcomes of decision-making augmented by algorithmic systems are strongly influenced by how humans perceive, trust, and engage with these systems [56, 57]. However, many organizations and researchers have noted the difficulty in translating these high-level principles into concrete design actions [58]. The core challenge lies in developing methodologies that are both structured enough to be repeatable and flexible enough to handle the complexity of human contexts.

2.2 Generative AI for HCI and design

GenAI refers to a class of artificial intelligence models designed to create new content, such as images, text, and music, by learning patterns from existing data. One of its most prominent applications is AI image synthesis, where techniques such as Generative Adversarial Networks (GANs) and diffusion models play crucial roles. GANs are particularly known for generating high-fidelity images through a competitive learning process between a generator and a discriminator [59]. However, with the advancement of diffusion models and large language models (LLMs), modern Text-to-Image Models, such as DALL·E, Stable Diffusion, and MidJourney, have outperformed GANs especially in image diversity, enabling a broader audience to create high-quality, stylized images from textual descriptions [60].

GenAI has become an integral tool in HCI and design research, shaping design workflows [61, 62], decision-making [1, 63], and user experience studies [64]. It is widely applied in design research and practice across various fields, both within and beyond the HCI community [65]. While widely applied in ideation [66] and artifact generation [67, 68], traditional text-to-x models often struggle to capture the fluid and multifaceted nature of design intent. Consequently, multimodal AI systems have emerged as a significant trend, offering integrated frameworks that align diverse data streams—such as visual, textual, and spatial inputs—to support more nuanced explorations [69, 70, 71]. By fusing these modalities, these systems provide designers with higher degrees of creative agency and more intuitive ways to translate complex, non-verbal intentions into tangible artifacts [72, 73, 74].

Despite these advances, GenAI poses significant challenges. One major issue is *control*, designers often struggle to steer AI outputs due to the opacity of AI decision-making, leading to results that require iterative refinement as stated in Gao et al. [75]. Additionally, while *multi-modal* inputs such as audio and images are being explored to facilitate more intuitive designer-AI interactions, text-based prompts remain the dominant input method, potentially limiting creativity and flexibility.

2.3 Co-speech gesture synthesis

Co-speech gesture synthesis refers to the automatic generation of realistic and contextually appropriate hand and body gestures that accompany spoken language in EAs. Recently, GenAI has been

widely adopted to generate co-speech gestures for these agents. There are two long lasting methodologies to create such gestures: hybrid rule-based techniques and data-driven methods [76, 2, 3, 77]. The earliest study [78] on the hybrid rule-based approach employed probabilistic models to generate gestures based on linguistic and contextual cues. In contrast, the data-driven approach utilizes deep learning to train agents' gestures using Mocap or video data from sources like dyadic conversations and TED presentations [79, 3, 80, 81]. Both methods have demonstrated effectiveness in producing naturalistic bodily movements. To develop AI-generated EAs that enhance authenticity and immersion, integrating co-speech gestures with semantic meaning is essential. However, the aforementioned approaches are inherently labor-intensive and face scalability challenges due to technological constraints. Moreover, research on generative methods for **gesture synthesis**, especially on defining gesture rules for AI-generated EAs remains limited.

In recent years, a few co-speech gesture generation models with semantic retrieval have emerged, combining rule-based and deep learning-based systems [82, 4, 5]. Evaluations of these semantics-aware gesture synthesis systems demonstrate that these systems outperform purely deep-learning approaches in generating semantically meaningful and rhythmical gestures, thereby enhancing the human-like expressiveness of AI-generated EAs. For instance, a semantic-aware co-speech gesture synthesis system leverages a GPT-based motion generator and a LLM-driven retrieval framework to ensure high-quality, semantically relevant gestures that align with speech rhythm and content [4, 83]. Inserting retrieved semantic gesture identifiers into rhythmic gestures enhances the perceived expressiveness of AI-generated ECAs in social interactions. Despite the advantages of using Mocap datasets for precise full-body movement tracking, inconsistencies in semantic annotations across datasets pose a challenge [4, 5]. Additionally, while these systems employ semantic motion generation via new text prompts to replicate several classic semantic gestures [84, 4, 5], such as "pointing left", "disagreement", and "a large amount", human-produced semantic gestures are often more nuanced and hierarchical. Consequently, the current prompt-based segmentation of gestures is relatively simplified compared to McNeill's Gesture Phases and functional gesture classifications [43]. However, this simplification may introduce inconsistencies, such as incomplete or unclear gestures, or conflicts in the temporal sequencing of semantic and rhythmic gestures.

Moreover, most current datasets are primarily sourced from conversational gestures [85, 86, 87, 88]. In contrast, generating co-speech instructional gestures necessitates an understanding of educators' instructional intentions to effectively convey educational content and account for students' interpretations of these gestures. For example, educators may deliberately adjust the rhythm, amplitude, and direction of their co-speech gestures to emphasize key concepts, while students may perceive these gestures differently in terms of their communicative intent. This highlights the need for further research on educators' gesture usage and students' perceptions of instructional gestures performed by AI-generated PAs. Past evaluations of user perception have often relied on scales measuring dimensions such as naturalness and human-likeness [3, 89, 5, 4]. However, these metrics may not fully capture how AI-generated PAs are perceived in real-world educational settings. Factors such as professionalism, approachability, and credibility are also crucial in educational contexts. Bridging this gap requires not only technical advances in synthesis but an understanding of educators' instructional intentions and learners' perceptual needs—dimensions that gesture synthesis research, operating largely within computational paradigms, has yet to address.

Though research has investigated data-driven co-speech gesture generation that emphasizes speech rhythm and lexical embedding [50, 51, 52, 81], there remains substantial scope for examining the synchronization of coherent gestures, such as cohesive and beat gestures, with words. Word-synchronized gestures represent minimal units in the alignment of speech and gesture and hold promise for enhancing the design of gesture transitions and connections. A pioneering study conducted by Wei and Chow [53] has begun to show these alignments. They have discovered that cohesive gestures are associated with specific PoS, such as adverbs and logical connectors, whereas beat gestures are associated with pronouns, adjectives, and action-oriented verbs. Cohesive gestures appear to strengthen the logical coherence and comprehension of speech, whereas beat gestures provide rhythmic emphasis and visual cues. The study identifies four patterns of alignment between educators' PoS, lexis, and gestures [53], providing insights into the design of instructional gestures. It also calls for a closer analysis of how human teachers use these gestures in delivering educational knowledge. Building on this, this study explores the design of PAs with instructionally coherent gestures, particularly cohesive and beat gestures, and evaluates their impact on the learning experience. Furthermore, it

informs the intentional integration of word-synchronized gestures in PA design and extracts insights from evaluation findings to optimize semantic motion generation via text prompts.

2.4 Pedagogical agents' gesture design

Educators who employ gestures can enhance both social engagement and learning outcomes among students [90, 91, 92]. These gestures transform cognitive understanding into multimodal information that students can easily process [93]. McNeill categorized gestures into five distinct types — *cohesives*, *beats*, *deictics*, *iconics*, and *metaphorics* — each serving a distinct communicative role [43]. While iconic and metaphoric gestures are often tied to specific visual content, this study focuses on cohesive and beat gestures due to their frequent appearance in discourse and their fundamental role in reinforcing coherence and prosody. Cohesive gestures are movements that connect different parts of a narrative or discourse, creating logical links between concepts. They help to maintain continuity and structure the flow of information, for example, by using a hand to sweep from one idea to the next to show a transition; Beat gestures are small, rhythmic motions, such as a flick of the hand or a tap of the fingers, that are synchronized with the prosody of speech. Expanding on McNeill's framework, research on PA gesture design has aimed to mirror human gestural behavior in educational contexts. Studies suggest that integrating beat, deictic, iconic, and metaphoric gestures enhances attention, retention, and comprehension, with notable advantages for second-language learners[42, 94, 95, 44]. Yet, while students intuitively comprehend iconic and metaphoric gestures that depict semantic information, this is not the case for discipline-specific terminology [96]. The way gestures distribute within speech is also tightly linked to linguistic content; when speakers describe visual imagery, they are more likely to use iconic gestures[97]. Additionally, an individual's fluency in spoken language influences gesture production [98]. The effectiveness of iconic and metaphoric gestures, therefore, hinges on subject-specific terminology, the extent of visual descriptions, and language fluency.

By contrast, cohesive and beat gestures frequently appear in discourse, reinforcing coherence and prosody. They do not carry semantic meaning on their own but function to emphasize or add prominence to specific words or phrases, much like verbal stress. Their fundamental role in communication makes them a promising avenue for further research on gesture design. Cohesive gestures help

maintain continuity across conversational turns and organize spatial references [99, 100, 101], making them a natural tool for structuring discourse. Investigations into instructional gestures in PAs, including spatially communicative gestures [102] and rhythmic hand motions [44], shed light on how fluid transitions between movements contribute to a more natural delivery of speech. Beyond their inherent communicative functions, gestures are also shaped by the instructional setting. Researchers argue that factors such as classroom layout, seating arrangement, and material placement influence both the frequency and type of gestures used by teachers [103]. Additionally, variations in gesture frequency contribute to different learning outcomes [35]. Enhanced gesture frequency significantly improved cued recall and recognition, while average frequency does not, highlighting the role of social cue strength in learning effectiveness. Thus, PA gesture design should consider gesture types, teachers's instructional intentions, frequency, and contextual factors. This study provides insight into the natural integration of instructional gesture transitions in PAs and their influence on student perceptions. Yet, much remains to be explored regarding gestures that encode semantic information, particularly cohesive and beat gestures. Translating these findings into actionable design specifications for AI-generated agents, however, requires a methodological bridge that neither gesture synthesis research nor pedagogical literature has established on its own.

2.5 Instructional materials design

Educational material design encompasses the systematic planning, development, and evaluation of instructional resources to achieve specific learning objectives [104]. Seminal frameworks provided by scholars such as Merrill [105], Gagné [106], and Reigeluth [107] emphasize learner-centered approaches, advocating for the deliberate alignment of content, guided by the educator's instructional intent, with cognitive processes and the specific learning context. In the digital era, the integration of technology has received significant attention, with Mayer's Principles of Multimedia Learning [108, 109] providing a theoretical foundation for utilizing digital media to foster dynamic engagement and support diverse learning styles.

A critical dimension of contemporary instructional design is the duration of digital content. While early empirical studies on massive open online courses (MOOCs) observed a decline in engagement for videos exceeding six minutes [110, 30], more recent pedagogical literature suggests that this

"six-minute rule" may not apply to complex, university-level curricula. Research indicates that durations ranging from 12 to 20 minutes are effective when supported by instructional scaffolding and high-quality production [111, 112]. From a social agency perspective, extended exposure is often necessary to establish a stable perception of an agent's persona and to allow for the cumulative effect of continuous gestural cues [113, 114]. Furthermore, the efficacy of such materials is influenced by the pacing of information delivery. While speaking rates in educational contexts vary widely, typically ranging from 48 to 254 words per minute (wpm), an optimum of approximately 160 wpm is often recommended for clarity [115]. A moderated pace (e.g., 140-150 wpm) is frequently employed in complex technical subjects to manage cognitive load without compromising engagement [116].

Beyond structural parameters, the design of instructional materials integrates cultural responsiveness and accessibility to address diverse learning needs [117]. This focus on inclusivity aligns with the movement toward Open Educational Resources (OER), which aims to increase collaboration and knowledge democratization [118]. Within this context, Human-Computer Interaction (HCI) serves as a particularly effective instructional subject for a broad student population; its interdisciplinary nature not only aligns with foundational curricula [119] but also provides a versatile platform for exploring the intersection of design, technology, and human behavior [120]. By applying Mayer's principles to HCI-oriented materials, instructional designers can create visual and auditory elements tailored to the demographic and cognitive profiles of modern learners. This strategic alignment ensures that content is relevant and engaging, as HCI's inherent user-centricity mirrors the pedagogical goal of creating accessible, culturally resonant resources. Such a dual-layered approach reflects the educator's intent to ensure that instructional material is not only cognitively accessible but also professionally resonant for students across various disciplines.

As AI becomes increasingly integrated into instructional delivery through automated agents, synthetic voices, and generated gestures [6, 28, 56], how educators' instructional intent survives translation into automated systems emerges as a central design challenge. Whether students and educators trust such systems depends in part on whether this intent remains coherent and legible in the agent's behavior [121, 122, 123, 124]. Lee and See [125] address this directly, identifying performance (demonstrated capabilities), process (how the system operates), and purpose (the goal-oriented intent embedded by designers) as the three bases of trustworthy automation. PA research has made

Manuscript submitted to IJHCI

considerable progress on performance and process, including evaluating gesture naturalness, improving synthesis techniques, and refining agent appearance, yet pedagogical intent, as the purpose dimension of PA design, has received comparatively little attention. How an educator's instructional goals are encoded and preserved across content, pacing, and embodied behavior remains underexplored, and constitutes the design challenge this work addresses.

2.6 Design processes and reflective practices

In the field of HCI, the design process is widely recognized as a sequence of decision-making and reflection steps that foster iterative development. Prior work addresses the potential for embedding new knowledge within design artifacts, illustrating the role of design processes and prototypes in advancing research agendas — a concept encapsulated in Research-through-Design (RtD) [49] and articulated through *Strong Concepts* that solidify abstract design principles [126]. Schön's notable concept of *reflective practice* characterizes this iterative design as a continuous cycle of "Reflection-in-action", in which practitioners dynamically evaluate and adapt their methods in-situ [127]. Designers, through reflective engagement in their activities such as sketching, prototyping, or creating representational artifacts - gain insights that inform and refine their design thinking in real time. This approach resonates with Ingold [128]'s idea of *thinking through making*, where design thinking emerges organically from iterative making practices, progressing from broad concepts to intricate details [129]. Moreover, Brown et al. [130] introduces the notion of *indexing*, which emphasizes that learning is inherently contextual and illustrates how deep immersion in design activity strengthens experiential learning, analogous to language acquisition through contextual immersion.

Reflective practices extend into instructional activities, such as the gestures acted by PA instructors and the design of educational materials. Although PA gesture design might seem peripheral to traditional design discourse, multimedia learning research has shown that effective knowledge transmission requires an intentional alignment of verbal and visual elements [109]. Consequently, PA gesture design could benefit from design research methodologies, opening up a less-explored space of design in educational contexts. Additionally, recent studies suggest that iterative, experimental approaches involving collaborative partnerships can help navigate risks and enhance innovation outcomes [131]. Agile design thinking methodologies emphasize prototyping and human-centered perspectives to

tackle such complex educational design challenges [132, 133, 134]. By incorporating a blend of design research methodologies — conceptual framing, experimentation, and user-focused evaluation — the development of multimedia instructional materials and educational practices can achieve improved effectiveness. In the following section, we will explore how design research methodologies can guide the iterative design of educational materials and PA gestures, underscoring an exploratory design approach tailored to educational contexts.

2.7 Theoretical framing

This research synthesizes three bodies of work: Human-Centered AI methodologies that operationalize high-level principles into concrete design practices [55, 58], co-speech gesture synthesis that aligns gestural behavior with linguistic structure [4, 53, 43], and pedagogical research on multi-modal instruction [109, 93]. Each domain, while productive independently, addresses only part of the challenge: HCAI principles lack concrete design processes for artifacts like gestural PAs; gesture synthesis prioritizes technical fluency over instructional grounding; and pedagogical research documents gesture effectiveness without translating findings into implementable agent design specifications. By adopting HCAI as our theoretical lens, we reframe PA gesture design as a collaborative process [49, 127] that preserves educators’ agency while leveraging computational tools. This perspective positions our four-stage methodology as a systematic approach to capture tacit pedagogical knowledge and translate it into implementable specifications validated through student perceptions [90, 92]. Situated at the intersection of these three bodies of work, the framework treats pedagogical intent as the *purpose* dimension of trustworthy PA design [125], positioning it as a traceable design input alongside the technical and evaluative processes that carry it through to the final agent.

3 Design Methodology

Our methodology was developed through a Research-through-Design process. It is structured to ensure that the design of the PA is continuously grounded in the instructional needs of the educators and the perceptual needs of the learners. The process involves a close collaboration between an educator and a PA designer, spanning five key steps from gesture analysis to EPA integration (Fig. 1; panels A–E are labeled in the figure).




FIG. 1. Key Steps in our Human-Centered Design Methodology: **A)** *Gesture analysis*: Understanding teachers' intents by observing a senior teacher's gestures. **B)** *Educator's improvised gestures*: Capturing tacit instructional knowledge through performance. **C)** *PA Designer mirrors educator*: Facilitating interdisciplinary collaboration and refinement. **D)** *Human-to-EPA transformation*: Upper body motions extracted via DeepMotion for 3D animation. **E)** *EPA and teaching materials integration*: 3D animation embedded alongside teaching materials.

3.1 Collaborative team and course preparation

The teaching material preparation was a collaborative effort between the PA designer and the educator, each contributing their domain expertise. *PA designer*, with over ten years of experience in animation production and design education, was responsible for gesture design and animation. The *educator*, trained as a designer and with over five years of teaching experience in design education, delivered the keynote using pre-scripted presentation notes. To ensure alignment with the desired presentation quality, the educator and PA designer engaged in iterative discussions to articulate a shared vision of the instructional delivery. These sessions led to the formulation of preliminary design criteria for performance-based teaching, which informed both the structure of a 14-minute lecture and the design of word-synchronized co-speech instructional gestures.

As part of the effort to eliminate the influence of prior knowledge earned by the participants in the field of HCI, the team performed an exclusive literature inquiry strategy on ACM's digital library. This search was conducted using the following formula [Title: "guide"] OR [Title: "tutorial"] AND [E-Publication Date: Past year], which returned 92 results. Following the removal of duplicates and the application of the "full text research article" filter, 37 entries were selected for the screening repository. the two collaborators independently screened the titles and abstracts of these entries, subsequently reviewing the full texts to include only those papers deemed fit for college-level comprehension, ensuring the inclusivity of the material for all participants regardless of the participants' study backgrounds. Although these efforts yielded four suitable entries (i.e., [135, 136, 137, 138]) deemed appropriate for teaching, one particular case arose to our attention [138] recently published

on ACM International Conference on Tangible, Embedded, and Embodied Interaction (TEI). It utilizes a pictorial format that centralizes visuals and diagrams to describe a novel material-centered design process (i.e., *analysis, synthesis, and detailing*) of crafting interactive materiality in a step-by-step manner. The two collaborators agreed to use [138] as the final teaching material due to its flow of articulation and comprehensibility for college students (Fig. 2).




FIG. 2. Flowchart for literature selection process

3.2 Developing the instructional materials

The overall structure of the prerecorded video closely mirrors the structure of the article. It begins with an introduction to the literature concerning material-centered and interactive materiality (for further details, please refer to the original paper [138]), and then provides a brief overview of the design artifact known as Puffy. Then, it outlines the analysis-synthesis-detailing (A-S-D) design and

Manuscript submitted to IJHCI

implementation process in a total of thirteen methodical steps, which are reflected in the presentation slides. The *educator* also prepared presentation notes for each slide to allow a smooth teaching experience. While developing these materials, a gap was observed as there was no clear guidance in the existing literature to support the design implementation of the PAs at an operational level. Yet, the two collaborators recognized that Puffy’s design process might serve as a valuable reference. It exemplifies an exploratory procedure that encompasses 1) analyzing various shape and material explorations, 2) synthesizing suitable samples, and 3) detailing processes. the two collaborators were inspired by that and started to unfold the PA design in the following section.

3.3 Implementing the Methodology: From Human Action to Embodied Agent




FIG. 3. Visual overview of the entire Pedagogical Agent (PA) gestures design procedures.

The core of our methodology consists of four iterative stages for implementing the PA’s gestures: 1) Preparation, 2) Human PA Acting, 3) EPA Acting, and 4) EPA-assisted course delivery and evaluation (Fig. 3). Each stage builds upon the last, ensuring that human requirements gathered early on are carried through to the final AI artifact.

Stage 1: Preparation (Multimodal Corpus Analysis). The initial iteration aimed to investigate the natural gestures employed in an educational context (Fig. 4).




FIG. 4. Stage 1: Preparation

Preparation. The initial iteration aimed to investigate the natural gestures employed in an educational context, specifically within an undergraduate design course. As motivated by the word-synchronized gesture study [53], this phase was initiated by *a) observing and analyzing* (Fig. 1-A). Eleven teachers (8 male, 3 female), each with more than five years of teaching experience, were recruited via email. They taught design- and art-related courses, with video data collected from 13 one-hour sessions (7 offline, 4 online, and 2 hybrid). Their courses featured detailed case studies, increasing the likelihood of using metaphoric and iconic gestures [139, 43]. During the lectures, full-body movements and presentation slides were recorded using an iPad Pro. Researchers analyzed both gestures and verbal language from the video recordings. Gestures were manually annotated using ELAN [140], following McNeill’s five-category gesture taxonomy (cohesive, beat, deictic, iconic, metaphoric) [43]. Two researchers underwent structured training prior to coding: both studied McNeill’s classification criteria and jointly coded a calibration video, during which definitional boundaries for each category were discussed and formalized into a shared coding protocol. Subsequent annotations were conducted through a collaborative consensus process: where a gesture’s classification was disputed, both coders discussed the case against the established criteria until agreement was reached. This negotiated agreement procedure is recognized as a valid reliability mechanism in qualitative content analysis, as it ensures that category boundaries are systematically applied and that ambiguous cases are resolved through principled deliberation rather than arbitrary assignment [141]. Speech was transcribed through Otter.ai [142] and subsequently reviewed for accuracy. To

Manuscript submitted to IJHCI

investigate word-synchronized gestures, a chronological coding approach was applied, grouping one to three consecutive words into microcosmic units based on the lecture timeline. Using Natural Language Processing (NLP) with Python [143], we systematically coded 71,252 samples of gestures aligned with their respective words and PoS. Results showed that cohesive gestures accounted for 44.34% (33,031 occurrences) of the total, followed by beat gestures at 38.26% (34,532 occurrences). Deictic (8.64%), metaphoric (2.98%), and iconic gestures (2.53%) were less frequently used. The distribution across PoS revealed that nouns were the most common (36.59%), followed by verbs (23.8%), pronouns (16.63%), adverbs (14.91%), adjectives (8.32%), and conjunctions (7.24%). Cohesive gestures are essential for establishing referential connections and enhancing the clarity of lectures. Following our Pearson's Chi-squared correlation analysis, we discovered that the use of nouns, adverbs, particles, and conjunctions in gestures enhances communication coherence, effectively reinforcing the delivery of key educational concepts (Fig. 4a). The researchers concluded that among these, cohesives and beats constituted the most frequently employed yet unobtrusive categories compared to deictics, iconics, and metaphors. As highlighted in the related work section, cohesive and beat gestures can augment the coherence of PAs' bodily movements, thereby enhancing students' learning experiences. These types of gestures were selected to further develop PA with naturalistic instructional gestures.

Subsequent substantial effort was devoted to *b) gesture samples*. The *PA designer* randomly selected one student from seven out of the eleven lectures to recruit a paired group of seven participants (Male=3; Female=4; aged 20-22 years). English is a second language for these students, who aimed to interpret selected cohesive and beat gestures extracted from the lecture videos (Fig. 4b). When students observed the instructors' cohesive and beat gestures without any accompanying audio cues, 12 out of 13 were able to easily infer the contextual meanings. For example, when instructors repeatedly performed upward or downward hand movements, students recognized these gestures as an intentional emphasis on specific concepts—often accompanied by nouns or adverbs indicating a strong attitude toward the definition. Similarly, when instructors linked gestures in sequence, students identified this as a transition between concepts, with the likely presence of conjunctions or similar connecting terms. As informed by these insights, the next phase involved *c) educator gesture mimicking* (Fig. 1-B). The *educator* initially executed a trial featuring a sequence of both intentional

and non-intentional gestures, grounded in the prepared presentation notes (Fig. 4c). The *educator* and the *PA designer* then engaged in a detailed discussion to reach a consensus on criteria for word-gesture alignment, drawing from the trial's gesture sequences. Throughout the process, the *educator* reflected on key moments and logical connections in the course content where key concepts emerged, while also identifying how gestures could reinforce these concepts and facilitate knowledge transitions to advance the course flow. Meanwhile, the *PA designer* documented these reflections. This dialogue directly influenced the design of naturalistic instructional gestures for subsequent phases.

Stage 2: Human PA Acting (Performance-Based Rehearsal). In the second iteration, the Human PA acting phase, the educator acted as a "human prototype" to externalize tacit pedagogical knowledge through performance (Fig. 5).




FIG. 5. Stage 2: Human PA acting

In the second iteration, known as the Human PA acting phase, involving the *educator* as the central figure. During this stage, the educator primarily focused on professionally delivering the prepared presentation, integrating instructional gestures summarized from the previous phase. The educator seamlessly incorporated them into natural, rhythmic hand movements, designing for a smooth flow that closely mirrors everyday teaching interactions (see Fig. 5d). This stage consisted of multiple rounds of *a) presentation rehearsals and acting out* to ensure that gestures, speech rate, and intonation were synchronized timely and coherently with the presentation contents.

Following these rehearsals, the educator formally delivered the lecture in English, adhering to the presentation notes, visual aids, and gesture design criteria developed in the prior stage. The *PA designer* observed the gestures performed by the educator (Fig. 1-C), and both parties reflected on

Manuscript submitted to IJHCI

how to better express word-synchronized gestures to *b) correct and synthesize the human gesture samples* (see Fig. 5e). For instance, when an image of a swelling pufferfish was displayed on the screen, the teacher made a obvious raising-hands gesture to emphasize the new concept. The entire process was recorded using an iPhone with the voice captured simultaneously using AirPods Pro as a separate audio track, ensuring high-quality audio.

Subsequently, the recording was trimmed and edited into a 14-minute presentation, serving as a key reference for the *PA designer*. However, flaws such as inconsistent transitions between gestures appeared in the draft recording. Therefore, the PA designer shadowed the educator's gestures in context, a process that was also being recorded, yielding a smoother and more well-organized *c) presentation video with synthesized educator's gestures* acting out, thereby maintaining consistency in the gestures and transitions between them (see Fig. 5f).

Stage 3: EPA Acting. (Human-to-Agent Motion Transfer). The third phase, *Embodied PA Acting*, emphasized transforming the human performance into an embodied agent form through human-to-agent motion transfer via video-based pose estimation and manual joint correction (Fig. 1-D; Fig. 6g).




FIG. 6. Stage 3: EPA acting.

The third phase emphasized transforming the human PA form into an embodied agent form, specifically concentrating on the post-production aspects of 3D avatar design, mainly executed by the *PA designer*. Moving from human performance video to an EPA form serves a purpose beyond replication: it decouples gestural content from the identity of a specific presenter, enables systematic modification of gesture parameters across iterations, and provides a controlled research instrument for isolating gesture-specific effects on student perceptions. We opted for minimalist, robotic avatars

over humanoid configurations to focus participants’ attention on the instructional gestures—the primary research focus. While prior research indicates that agent appearance significantly influences student perceptions [9, 10], our faceless design choice aimed to isolate gesture-specific effects by minimizing confounding factors related to facial features, physique, and attire. We retained the original human-voiced soundtrack to mitigate the potential bias introduced by synthesized speech quality, which could impact the learning experience[29]. This also ensures the provision of appropriate social cues for student engagement, while maintaining a level of abstraction to prevent unintended associations or biases related to physical appearance (see Fig. 5d). To achieve this transformation, several AI-based motion capture and tracking techniques were explored, where a 3D animated avatar was constructed to mirror the educator’s presentation, including platforms such as Plask Motion [144], MOVE Ai [145], and DeepMotion [36]. We settled on DeepMotion due to its robust pose estimation capabilities, particularly its accessibility and precise hand-tracking. Subsequently, the previously recorded human PA acting video was uploaded to DeepMotion and then transformed to 3D animated avatar video.

The *PAs designer* refined human PA video, excising repetitions and reorganizing the footage into seven 2-minute sequences to optimize readability and transform capability on the DeepMotion platform. Upon completion of the transition from human to 3D PA, minor adjustments were made to the avatar joints to rectify issues, especially concerning the conversion of 2D videos into 3D animations, where the depth information was frequently either imprecise or entirely absent (see Fig. 6h). In such instances, the PA designer manually moved the joint nodes to properly show the gestures, adhering to the principle of texturing the gestures as naturally and human-like as possible. Subsequent to the rendering of the 3D animated PA and its integration into the presentation slides using Adobe After-Effect, we yielded a *c) presentation video with rendered PA gestures* (see Fig. 6i). The half-body size of the PA is positioned on the right-hand side of the video frame to optimize visual clarity. This intentional design configuration facilitates an unobstructed view of both the educational materials and the PA’s gestural expressions, thereby enhancing the students’ capacity for cognitive engagement.

Stage 4: EPA-assisted course delivery. The final stage involved deploying the EPA-assisted presentation (see Fig. 1-E) in an actual learning context to evaluate its effectiveness and gather user feedback (Fig. 7).




FIG. 7. Stage 4: EPA-assisted course delivery

3.3.1 EPA-assisted course delivery. To investigate the efficacy and user experience of a virtual lecture featuring a gestural PA customized to address teachers' multimodal instructional needs. We conducted *a) user study* on campus (see Fig. 7j). The student participants were recruited through on-site poster advertising. English is their second language, and all participants possess an IELTS score of 6 or equivalent, enabling them to understand courses primarily taught in English. Eligibility criteria stipulated that candidates possess limited prior knowledge of the subject matter but exhibit a keen interest in interactive materiality. Upon arrival at a designated empty classroom, each participant was furnished with a laptop and noise-canceling headphones. The self-paced learning module had an average duration of 45 minutes. A researcher was on-site to facilitate the study and address any queries, ensuring an environment devoid of extraneous disruptions. Comprehensive audio-visual recordings, along with on-site notations, were secured for subsequent analysis (see Fig. 7k).

4 User Study

4.1 Method and materials

Our user study aimed to evaluate the learning experiences and perceptions of EPA-assisted courses customized according to the instructional needs of educators. The study was conducted through a structured three-phase evaluation protocol (on-campus setup, see Fig. 7j). The initial phase consisted of a pre-questionnaire encompassing a demographic inquiry and a pre-transfer test to assess participants' baseline knowledge of the subject. The second phase featured a 14-minute video-based instructional session. This duration was selected based on pedagogical recommendations for university-level content [111], suggesting that 12-20 minutes provide sufficient exposure for students to form stable perceptions of an agent's social presence and benefit from cumulative gestural cues [113]. The 14-minute format also enabled comprehensive coverage of complete instructional content while maintaining engagement, aligning with our research focus on perceptual validation rather than long-term learning outcome assessment. The final phase comprised a post-questionnaire that included a post-transfer test, seven subquestions adopted from the Agent Persona Instrument (API) [146], which covered four aspects: engagement, person-like qualities, instructor-like qualities, and credibility, along with two open-ended questions designed to elicit in-depth reflections on the educational experience. In the pre-questionnaire, we gathered data on participants' gender, age, major, and educational level. Additionally, we assessed their English proficiency, familiarity with the theme of the course we designed, and their level of interest in the course. For these assessments, we employed a 7-point scale for each question. The pre- and post-transfer tests consist of the same set of 15 questions. These questions are designed based on an average of one key concept appearing for every two slides in the lecture videos, with a total of 30 slides. The lecture script was designed at CEFR B2 level (Flesch-Kincaid Grade Level = 11.6; Flesch Reading Ease = 42.3) [147, 148], and the narration was delivered at a speaking rate of 142 words per minute (wpm) — a moderated pace suitable for complex technical subjects that balances cognitive load management with engagement [116, 115]. The teacher explains related concepts using various methods, such as presenting relationship

diagrams and providing examples. Transfer test scores are determined as follows: 3 points for integrating video knowledge with personal explanation and examples; 2 points for reciting video content with examples; 1 point for partial recall or somewhat relevant examples; 0 points for unrelated answers. The maximum score per lecture is 45 points. The open-ended questions were formulated as follows: *"Can you provide a detailed description of your experience with today's virtual learning process, offering specific examples where possible?"* and *"What are your impressions of the PA? Please elucidate on aspects such as professionalism, personality traits, and the efficacy of body movements and communicative gestures"*.

4.2 Subjects

The study initially screened candidates to ensure they met specific pre-test criteria, resulting in a final sample of 38 students recruited from a university in Hong Kong. This sample comprised 32 females and 6 males, aged between 18 and 26 years old ($M=21$; $SD=2.05$). The participants come from diverse academic backgrounds including Business ($n=17$), Engineering ($n=9$), Health Science ($n=7$), Computing ($n=2$), Design ($n=2$), and Linguistics ($n=1$). Among them, 27 were undergraduates and 11 were postgraduates. As English is the medium of instruction at Hong Kong universities, all participants were required to demonstrate a minimum English proficiency equivalent to an IELTS score of 6 or higher, ensuring their ability to comprehend the lecture script (CEFR B2 level). While the majority of participants were local or from mainland China and Southeast Asia—speaking English as a second language—they self-rated their proficiency for academic communication as relatively high ($M = 4.89$, $SD = 0.89$). Their familiarity with the subject of our tested course was relatively low ($M = 2.95$, $SD = 1.49$), yet their interest in the course was high ($M = 4.84$, $SD = 1.41$). Following the lecture, participants rated the perceived difficulty of the content at $M = 4.16$ out of 7 ($SD = 1.35$), indicating that the material—designed at CEFR B2 level—provided a sufficient cognitive challenge without exceeding their comprehension levels, which corroborates the alignment between the script's readability metrics and students' actual learning experience.

4.3 Data analysis

Our primary research interest centered on students' qualitative perceptions and experiences of gesture-based PAs. To ensure the validity of these interpretations, we first established that participants engaged meaningfully with the instructional material and formed coherent impressions of the PA. Paired-sample t-tests on pre- and post-transfer scores verified knowledge acquisition during the session, while descriptive statistics for API dimensions (engagement, person-like qualities, instructor-like qualities, and credibility) confirmed that participants developed stable perceptions of the PA's social and instructional presence. These quantitative checks served as essential prerequisites for interpreting the subsequent qualitative data.

The core analysis focused on open-ended interview responses, which underwent thematic content analysis [149]. Two independent coders reviewed the qualitative feedback, identifying recurring themes related to the PA's professionalism, human-likeness, approachability, and gesture effectiveness, guided by a shared coding protocol developed during the Stage 1 calibration phase. Where categories were disputed, both coders discussed until consensus was reached, following established consensus coding procedures [141]. Representative quotes were then organized into a two-dimensional feedback categorization matrix, systematically mapping student perceptions across dimensions of human-likeness versus professionalism and human-likeness versus approachability. This approach enabled us to synthesize patterns in how specific gesture types (cohesive, beat) influenced students' subjective experiences and social impressions of the PA, thereby generating grounded insights to inform future design iterations in gesture-based embodied pedagogy.

4.4 Results

Upon completion of data collection and subsequent data cleaning, we performed a paired-sample t-test, revealing statistically significant disparities between pre- and post-transfer scores (pre-transfer $M = 3.95$, $SD = 2.99$; post-transfer $M = 7.82$, $SD = 4.37$; $t(37) = 6.02$, $p < .001$, Cohen's $d = 0.98$), confirming that participants engaged meaningfully with the instructional material. Additionally, Students' scores for API engagement ($M = 4.46$, $SD = 1.55$), API person-like qualities ($M = 4.39$, $SD = 1.56$), API instructor-like qualities ($M = 4.97$, $SD = 1.16$), and API credibility ($M = 4.68$, $SD = 1.58$) reflect

a generally positive impression of the PA, as students perceived it to be human-like and capable of conveying important learning information, along with a high level of credibility. These results confirm that participants engaged meaningfully with the lecture and formed stable perceptions of the PA. Then we engaged in the *b) synthesized feedback* garnered from student participants. Eighteen out of thirty-eight respondents reported that the instructional materials were understandable and well-organized, while seventeen expressed a desire for opportunities to continue learning related courses. Their feedback described the course as "interesting," "satisfying," and "meaningful." Fourteen respondents noted that the PA resembled a real person, seven commented on the naturalness of its movements, and another seven described the PA as amicable. Nineteen respondents perceived the PA as professional, with remarks such as, "the degree of professionalism is the same as that of real people," along with observations about its organized speech and occasional humor. However, four respondents mentioned that the PA's hand movements were noticeably repetitive and exaggerated, using phrases like "very intelligent use of a lot of body language."

Following the analytic approach described above, the feedback categorization matrix presents our findings across four distinct dimensions of the PA's human-like characteristics: human-like and professional, human-like and amicable, professional yet divergent from human attributes, and amicable yet distinct from human attributes. Students who perceived the PA as both human-like and professional offered feedback such as, "*He is very professional and has a calm personality*" and "*accompanied by rich body movements, it can accurately convey information*". Conversely, certain students expressed that while the PA possessed a professional demeanor, the semblance to a real human was not entirely convincing. One student commented, "*he is not similar to the real teacher. but it's fine, because I can also learn the knowledge from him.*" Meanwhile, a majority of students conveyed that the PA remarkably resembled a genuine human and projected an agreeable personality. They cited instances such as "*occasionally shows a slight sense of humor*" and "*body movements are coherent and can attract people to listen*". However, a subset of students contended that the PA's authenticity was compromised by repetitive gestures, with one student noting, "*the PA's keeps moving, and the action continues to repeat, which is a little hindering me from reading the words on the slides*", and by the absence of interactions with the surrounding content, as voiced by another student, "*there is no engagement between the PA and the visual elements on the slides*".

Based on the matrix analysis, we deduced four distinct directions for further refinement. The initial direction underscores that the integration of cohesive and beat gestures into the PA design can engender a heightened sense of professionalism among students. The synchronization between speech rhythms and beat gestures effectively punctuates crucial information, mimicking the discernible behavior of a genuine teacher. Secondly, cohesive gestures play a pivotal role in facilitating smooth transitions and connections within the PA's gestural repertoire, thereby cultivating a sense of approachability among students. Thirdly, consideration should be accorded to the integration of the PA's bodily movements with instructional materials and the audience. This strategic alignment bolsters students' engagement and engenders a heightened sense of social involvement. Lastly, the tendency for repetitive gestures to enhance visual awareness should be recognized. However, excessive use can negatively affect students' perceptions, leading to feelings of unreality and emotional detachment. Overall, incorporating pedagogically grounded gestures into PAs enhances students' social presence and perceived instructional quality, thereby enriching their learning experience and creating conditions conducive to engagement and knowledge acquisition.

5 Insights for the design of instructional gesture generation

Building on the student feedback patterns identified above, this section synthesizes insights from our four-stage methodology—from multimodal analysis of authentic teaching (Stage 1) through collaborative prototyping (Stage 2), technical implementation (Stage 3), to student evaluation (Stage 4)—into actionable design principles. Our approach offers a replicable process for developing pedagogically grounded gestures: analyzing gesture patterns in authentic instruction, collaboratively externalizing educators' tacit knowledge through rehearsal, navigating technical constraints during implementation, and validating designs through student perceptions. Our focus on cohesive and beat gestures is grounded in both empirical prevalence (these types dominated observed instructional discourse) and theoretical transferability: unlike iconic or metaphoric gestures tied to disciplinary terminology [97, 96], cohesive and beat gestures serve fundamental communicative functions—linking concepts, marking emphasis, regulating rhythm—applicable across domains [43, 100]. What varies across contexts is the linguistic anchors (which words trigger emphasis, where transitions occur) and timing parameters, not the gesture types themselves. The five instructional functions we identify (emphasis,

Manuscript submitted to IJHCI

flow, emotion, rhythm, connection) and the methodology for operationalizing them show potential for application across educational domains.

5.1 Teaching objectives and multimodal instruction design

The corpus analysis of 71,252 gesture samples from 13 design and art lectures (Stage 1) revealed that cohesive (44.34%) and beat gestures (38.26%) dominated instructional discourse. This distribution extends prior findings on word-synchronized gestures [53], providing further evidence that these gesture types frequently accompany specific linguistic structures in teaching contexts. This foundation informed our gesture selection for the PA design case. One of the first steps in the collaboration was to clarify the teaching objectives, as these directly influence the type of gestures required. Educators would articulate their instructional goals, such as explaining and emphasizing key concepts or guiding students through a thought process. From a multimodal instruction design perspective, instructional gestures structure information flow, regulate cognitive load [109], and amplify verbal content [93]. Based on these goals, PA designers would translate them into indexes and labels for instructional intents. The educator's instructional intents can be categorized into five key types. These include:

- (1) **Emphasis on Key Concepts:** This is achieved by varying the pace, size, and amplitude of gestures, which helps draw attention to important knowledge or transitions.
- (2) **Natural Flow of Communication:** Changes in the direction of the gestures' movement contribute to a smooth and natural flow of interaction, facilitating better communication with the students.
- (3) **Emotional Expression:** The contrast between forward and backward gestures, as well as the directional movement of the gestures, helps convey emotional tone and intent, aiding in emotional engagement and emphasis.
- (4) **Adjustment of Course Rhythm:** Pauses or breaks in the gesture flow help regulate the pace of the lesson, providing students with the opportunity to process information.
- (5) **Connection Between Concepts:** Smooth transition gestures are employed to link key points of the lesson, maintaining continuity and coherence in the instructional flow.

For instance, when the educator needed to explain an abstract concept, the instructional gesture could be labeled as an "emphasis gesture" by the LLM-based generative retrieval system, manifested through an increase in amplitude and described as "hands moving vertically." This gesture would signal important knowledge points to the students, highlighting shifts in the educator's communication and helping students identify key concepts or transitions in the material, thus enhancing focus and comprehension. In another example, when the educator's script indicated a semantic shift, such as "Now let's move on to the next step in the design process," the LLM-based generative retrieval system could label the gesture as a "cohesive gesture," guiding students through the transition and helping them follow the change in focus.

5.2 Translating student feedback into design specifications

Based on the feedback categorization matrix, we identified four main needs from students that can guide the refinement of PA gestures. These needs span across gesture variability, engagement with content, professional demeanor, and human-likeness:

- (1) **Gesture Variability:** Some students noted that repetitive gestures hindered their engagement. Therefore, there is a need for more varied gestures to maintain students' focus and to enhance the authenticity of the PA.
- (2) **Engagement with Instructional Content:** Students expressed a desire for the PA to interact more dynamically with the instructional materials, such as slides. This would improve student engagement and create a more interactive learning experience.
- (3) **Professional Demeanor:** Students who perceived the PA as professional yet lacking in human-like qualities indicated the need for gestures that align more closely with human communication patterns, enhancing both the professionalism and relatability of the PA.
- (4) **Human-Likeness:** While many students appreciated the PA's friendly demeanor, they suggested that more social and emotional cues, such as humor or varied gestures, would improve the overall learning experience.

By addressing these needs—enhancing gesture variability, improving interaction with instructional content, and balancing professional and human-like qualities—PA designs can be refined to

Manuscript submitted to IJHCI

better engage students and support their learning process. For example, when students provide feedback such as, "The gesture is unclear or confusing," the LLM-based generative retrieval system can generate prompts based on the semantic context to adjust the direction of the gestures. This ensures that the gestures' aiming direction is clearer and aligns with the visual content, enhancing the overall clarity and effectiveness of the communication.

5.3 Synchronization with speech

Synchronizing gestures with speech was a key aspect of the collaborative design process. This involved aligning gestures not only with the semantic content and instructional intent but also with temporally appropriate moments in speech. Building on research correlating gesture types with parts of speech [53], we developed operational heuristics through Stage 2 rehearsals that guided gesture-speech integration. While the specific timing parameters emerged from English prosody and our case context, the underlying principle—that gestures should synchronize with linguistic structure—applies across languages, though local calibration may be needed for different rhythmic patterns. To facilitate this, the gesture generation system should incorporate gesture timing annotations relative to PoS, enabling precise temporal alignment that allows gestures to integrate seamlessly with verbal communication rather than appear disjointed.

5.4 Iterative refinement

Throughout the process, iterative testing and refinement were central to the collaboration. The PA designer and the educator continuously worked together, conducting mock sessions and gathering feedback from students. This iterative approach allowed for fine-tuning both the gestures and their synchronization with the lecture content. The PA designer's shadowing process (Stage 2c) exemplifies this iterative refinement: while the educator's initial performance featured rich semantic content, some transitions appeared abrupt when isolated for animation. By mirroring the educator's movements and experimenting with interpolated transitions, the PA designer smoothed the gestural flow, ensuring each movement had clear onset, stroke, and retraction phases—a principle rooted in gesture phase theory [43] that transcends specific instructional contexts. Feedback loops helped ensure

that the gestures were not only pedagogically effective but also culturally sensitive and contextually relevant to the students' learning environment.

5.5 Exploration of GenAI tools



FIG. 8. KlingAI generated PA in a sequence of gestures animated by using the prompt, "A male teacher is enthusiastically describing an academic definition. He performs an emphasizing gesture, briefly pauses, and slightly turns his upper body to look toward the right side of the students. The shot only includes the teacher, with the camera fixed at a medium-close range, and the teacher remaining stationary." The first gesture highlights emphasis, and the second demonstrates intent, as the PA slightly turns their upper body toward the right side, as if engaging with students on that side.

We generated instructional gesture-based teaching videos using commercial text-to-video generation tools (i.e., Sora and KlingAI) to explore their pedagogical applicability. In our iterative testing, we observed that Sora produced highly realistic visuals with minimal body deformation and executed simple gestures like directional movements effectively, while KlingAI offered motion trajectory brushes that enabled finer control over gesture sequences. Through testing various educator-intent prompts and researcher discussions, we found that prompt structure plays a critical role in shaping gesture rhythm, transitions, and emphasis. Prompts specifying temporal phases, "*performs an emphasizing gesture, briefly pauses, slightly turns*" - yielded significantly more coherent output than semantic labels alone. For instance, "cohesive gesture" produced ambiguous results, but "a lateral hand sweep from left to right, synchronized with a specific timepoint" generated recognizable transitions. This extends insights from Stages 1-2: linguistic anchoring and temporal structuring improve GenAI output quality. As an example, we produced a 5-second PA video in which the teacher explains an abstract concept using an emphasis gesture followed by a directional body shift, enhancing communicative flow (see Fig.8). Without such explicit structuring, generated videos often lack temporal coherence, reducing communicative clarity. These findings suggest that future text-to-video systems

should incorporate timing, transition, and emphasis cues as explicit prompt parameters to support pedagogically grounded generation.

To quantitatively evaluate these generation approaches and compare the effectiveness of different prompt strategies, we conducted a controlled comparison of four generation methods: (1) trajectory-augmented textual prompts (KlingAI 1.5 with motion paths), (2) structured textual gesture prompts (Sora), (3) structured textual gesture prompts (KlingAI 2.6), and (4) baseline with speech script only (KlingAI 2.6, no gesture instructions). Through online recruitment, we screened respondents based on having at least one year of experience in educational practice and one year of familiarity with AI tools, ultimately selecting 12 qualified experts (aged 29.67 ± 4.50 years; 4 male, 8 female). Each expert rated the four videos on nine rating items using 7-point Likert scales; items were grouped into six dimensions: gesture execution quality, prompt adherence, temporal coherence, visual realism, instructor authenticity, and instructional suitability.

Repeated measures ANOVAs revealed significant main effects for all six dimensions (all $F > 3.99$, $p < .05$, partial $\eta^2 > .27$), with prompt adherence showing the largest effect ($F(3, 33) = 15.26$, $p < .001$, partial $\eta^2 = .58$). Post-hoc pairwise comparisons with Bonferroni correction ($\alpha = .008$) indicated that trajectory-augmented guidance (condition 1) significantly outperformed other methods on multiple dimensions (all $p < .008$, $d = 0.40$ – 2.04). Critically, all gesture-enhanced conditions (1-3) significantly surpassed the baseline (all $p < .008$, $d = 0.66$ – 2.93). These findings are consistent with our exploratory observations, providing quantitative evidence that structured gesture prompts (temporal structuring, linguistic anchors, and motion trajectories) substantially improve instructional quality and enable a human-in-control workflow for pedagogically grounded video generation.

5.6 Methodological contributions and scope

Advancing existing knowledge. First, while prior research documents gesture importance [22, 46], systematic frameworks for translating pedagogical intent into implementable specifications remain limited—our framework operationalizes this through a structured methodology that encompasses corpus analysis, collaborative rehearsal, technical implementation, and student validation. Second, existing synthesis research evaluates naturalness [3, 4], yet few studies empirically link teaching behaviors to student perceptions of pedagogical dimensions—our feedback matrix shows API ratings

and qualitative evidence validate word-synchronized gesture design. Third, our GenAI exploration reveals structured prompts (temporal phases and linguistic anchors) produce more instructionally coherent output than semantic labels alone.

Contextual grounding. This methodology was developed and validated through one instructional case: a 14-minute lecture on material-centered design processes, delivered in English to university students (primarily undergraduate level) in a Hong Kong educational context where English serves as the medium of instruction yet most students are second-language speakers. The educator explained a procedural workflow through step-by-step demonstration, using cohesive gestures to connect sequential phases and beat gestures to emphasize key concepts. This linguistic context shaped both the lecture’s moderate speaking rate (142 wpm) and the design emphasis on multimodal cues beyond verbal content. The case reflects our human-centered approach: we began by observing authentic teaching to identify prevalent patterns, captured an educator’s tacit pedagogical knowledge through collaborative rehearsal, and validated designs through student perceptions of instructional quality. This case demonstrates how the framework operationalizes HCAI principles—preserving educator agency while leveraging computational tools—and the design heuristics derived from this process inform future applications where similar collaborative, reflective practices can translate instructional intentions into embodied agent behavior.

6 Discussion

Our work, centered on designing and evaluating pedagogical agent (PA) gestures, is presented as a methodological case study for the broader field of Human-Centered AI (HCAI). This discussion reflects on our process and findings in relation to the core challenges of HCAI, particularly the need to bridge the gap between high-level principles and the concrete, practical work of designing and implementing intelligent systems. We structure our contributions into three main areas: the human-centered design framework we developed, the empirical results of its application, and a critical assessment of how our process informs the future of Generative AI.

6.1 Novel aspects of the pedagogical agent design process

6.1.1 Collaborative reflective practices. A distinctive feature of our process was the integration of collaborative *reflection-in-action*[127] during the *Human PA Acting* phase. While reflective practice traditionally centers on individual cognition, our approach extended this to a dyadic context where the *educator* and *PA designer* engaged in the iterative refinement of gestures, resembling an ongoing decision-making process. Within this collaborative setting, gestures were continuously explored, evaluated, and adjusted through real-time feedback and shared reflection. Although the *educator* worked from a pre-scripted presentation, multiple rehearsal rounds proved essential for this reflective cycle. This process revealed a key benefit of collaborative reflection: the *educator*'s propensity for subconscious or extraneous gesturing was identified by the *PA designer*, who, serving as an external observer, provided immediate feedback. This intervention was crucial for mitigating unintended movements and suggesting enhancements for clarity. Through this dynamic of observation and adjustment - a process reminiscent of collaborative human-AI interaction — the *PA designer*'s role was instrumental in ensuring that gestures were precisely aligned with pedagogical goals, thereby enhancing their overall communicative efficacy.

6.1.2 Improvisational practices in gesture design. A core characteristic of our design process was the use of structured improvisation for gesture creation. In contrast to conventional teaching where educators' gestures are often spontaneous and may lack consistent alignment with verbal content, our methodology employed iterative rehearsals across the *Human PA Acting* and *EPA Acting* stages. This was done to situate gestures meaningfully within the instructional narrative. During the *Human PA Acting* phase, the rehearsal process afforded the *educator* the flexibility to adapt and refine gestures in real time, fostering an alignment with the teaching content that was both natural and intentional.

This improvisational flexibility ensured that gestures were not only synchronized with speech but also contextually tailored to the educational material, thereby enhancing their pedagogical relevance. This method established a robust foundation for the *EPA Acting* stage, generating a repository of refined human movements that the *PA designer* could translate into a cohesive gestural performance for the agent. While this approach ensures high fidelity between the human-performed gestures and

their EPAs' embodiment, it introduces a practical trade-off: the process is time-intensive, requiring multiple iterations to achieve the desired level of refinement in both gestures and posture.

6.1.3 Re-contextualizing role-playing as a design method. In our PA design process, we re-contextualized the established design research method of role-playing—also known as "acting out" [150]—specifically for the design of PA gestures. Originally employed to give designers an embodied understanding of the user experience, we adapted this method to generate and refine gestural interactions. During the *Human PA Acting* phase (see Fig. 3), the *educator* engaged in an immersive enactment of the script, concentrating on the coherent and timely synchronization of gesture with speech. This performance was augmented by multiple explorative rehearsals aimed at discovering the most intuitive gestural expressions.

Concurrently, the *PA designer* adopted the dual roles of observer and proxy audience. In this capacity, they provided critical feedback to calibrate the *educator's* gestures for optimal precision and expressiveness, preventing them from being either exaggerated or overly subtle. This interaction also engendered a sense of social presence for the *Educator*, reinforcing the awareness that they were addressing students and thereby sharpening their focus on effective gestural communication [93]. Ultimately, this application of role-playing fostered greater empathy with the prospective students' experience and provided invaluable first-hand feedback integral to the design process.

6.2 Toward human-centered AI for gesture generation

A key contribution of this work is the presentation of a structured, four-stage methodology that guides the design of PA with instructional gestures. This framework responds to the challenge identified in previous research [58] regarding the difficulty of translating abstract human-centered AI principles—such as amplifying human capabilities [57, 55]—into concrete and replicable design practices. By grounding the process in educators' instructional intents and fostering iterative collaboration, the approach provides a practical and flexible methodology to operationalize these principles, offering valuable guidance for the design and development of embodied AI systems. In doing so, it contributes to advancing human-centered AI by informing the creation of PAs that have the potential to enhance teaching and learning experiences.

Our methodology is distinguished by two key features. First, it centers human expertise. In contrast to purely data-driven approaches that train models on generic data, our method begins with and iteratively returns to the embodied, tacit knowledge of a human expert. The *Human PA Acting* stage is a critical methodological step for eliciting and capturing nuanced performative requirements that are otherwise difficult to articulate; Second, it provides a clear pathway from requirements to implementation. It establishes a transparent and traceable link from the initial analysis of human gestures (Stage 1) to the final EPA (Stage 3). This ensures that the nuanced requirements identified during the ideation and acting phases are faithfully carried into the downstream stages of design and development.

This four-stage process may offer a structured blueprint for developing future HCAI systems where human expertise actively guides and refines AI-driven outputs. Here, AI serves not as a replacement for the designer but as a powerful collaborative partner. Much of contemporary AI development, even when labeled "human-in-the-loop," remains data-centric, focusing on collecting vast amounts of data where the human role is often reduced to that of a data producer. In contrast, our workflow advocates for an intent-centric perspective. The primary input is not raw data, but the structured, embodied, and articulated pedagogical intent of an expert. In Lee and See's [125] terms, pedagogical intent constitutes the purpose dimension of trustworthy PA design—the educator's goals inherited by the system through deliberate design. The framework makes this intent explicit and traceable: surfacing tacit knowledge (Stage 1) through collaborative rehearsal (Stage 2), embedding it through motion transfer (Stage 3), and validating it through student perceptions (Stage 4). The *Human PA Acting* stage is specifically designed to capture a high-fidelity representation of this intent. For HCAI research, this marks a crucial shift: it provides an epistemological basis for building AI systems that are not merely statistically accurate relative to a dataset, but are semantically aligned with the goals and values of their human collaborators.

The framework is structured as a human-AI collaboration in which educators and designers retain authorship over instructional intent, while AI tools handle motion capture and gesture generation in service of that intent. The human investment across stages is the means by which pedagogical intent is extracted, formalized, and made available as a standard for automated systems to build upon. As intent-aware generation tools mature, individual stages of this process become natural candidates

for progressive automation, guided by the pedagogical standards established through the human-centered process.

The development of gestural PAs raises important questions regarding pedagogical authority and student trust in AI instructors. Our study found positive ratings on credibility and professionalism, yet responsible deployment requires transparency about AI involvement [55], clear communication of system limitations, and institutional frameworks that position PAs as supplementary tools preserving human educators' essential roles in mentorship and adaptive support. Research indicates that trust in algorithmic systems depends critically on users' understanding of system boundaries and institutional context [56], highlighting the need for thoughtful implementation policies.

6.3 Empirical validation and design insights

6.3.1 Impact of gesture types on student perceptions. Our findings affirm the potential impact of cohesive and beat gestures on students' perceptions, suggesting that these gestures can improve professionalism and approachability in PAs. This aligns with previous research [90, 91, 92], which suggests that gestures could enhance social engagement and learning experiences among students in educational settings. The findings also provide empirical evidence that extends the existing understanding of how gestures contribute to enhancing the effectiveness of educational interactions [30, 92, 42]. In particular, the study's identification of cohesive and beat gestures as enhancing students' sense of professionalism and approachability on PA expands upon the previous research's focus on the multifaceted functions of gestures in instructional communication.

6.3.2 Context as the core of human-centered AI gesture design. A key insight from our study is the necessity of designing for context, a dimension frequently overlooked by current generative AI. While many AI models can generate realistic human motion, they often lack an understanding of the situational factors that make a gesture meaningful. Our work demonstrates that the synergy between a presenter's bodily movements, the specific educational materials, and the engagement of the audience is what fosters a vital sense of social involvement. This aligns with previous findings that the educational environment itself shapes gestural communication [103].

Ultimately, both our study and prior work argue that for gestures to enhance the situated learning experience, they cannot be generic; they must be contextually appropriate. This has profound implications for the design of human-centered AI, particularly for PAs in immersive settings like VR [26]. The future goal must be to move beyond mere kinematic mimicry and toward creating agents whose gestures are semantically rich and socially aware, thereby transforming an immersive environment into a truly collaborative and intelligent educational space.

6.3.3 Balancing gesture frequency and variety. The study highlights the potential pitfalls of excessive repetitive gestures and the need for a balanced approach to gesture design. This recognition of balance is reminiscent of studies that have examined the optimal use of gestures to augment learning. For instance, Pi and her colleagues [44] discussed the complexity and frequency of gestures in relation to their impact on student engagement and perception. This contributes to the broader discussion on how gesture frequency and variety intersect with learning experiences. Furthermore, the call for a balanced approach aligns with previous investigations into gesture complexity and variability. Studies have shown that the use of gestures, including both iconic and deictic gestures, can enhance students' engagement with different concepts [151]. Similarly, the negative implications of excessive repetition resonate with concerns raised in previous research regarding the potential cognitive overload caused by constant or monotonous gestures. For further development of gestural PA design, PA designers should consider utilizing various gestures and carefully balancing their frequency to enhance the sense of realism and meaningfulness while also avoiding cognitive overload for students.

6.4 Generative AI and future directions

6.4.1 Fields of Application. The design and application of gestures in PAs have broad implications for various fields within and beyond educational technology. Building upon the foundational insights provided by Petersen et al. [26], regarding the potential of PAs to improve educational experiences within VR environments, we investigated and contributed to the intricacies of gesture design for PAs, specifically focusing on word-synchronized gestures. For online education in particular, gesture-enabled PAs could serve as virtual teacher clones designed to supplement educators

during fatigue or technical difficulties. This opens up the possibility for more engaging and authentic AI-driven presentations and addresses practical problems such as teacher absenteeism. Another field is that of *AI Spokespersons* (e.g., AI Humans [152, 153]) for multimedia content, particularly in virtual platforms where the agent serves as a stand-in for human presenters. Previous research in AI spokespersons has focused primarily on vocal intonation and facial expressions, and gesture design could be an essential addition. Another promising field is *e-commerce*, which engages in product promotions. However, the heavy reliance on human streamers for constant endorsements can be problematic. This dependence creates a significant manpower workload, misleading consumers with human streamers' emotional cues and subtle behaviors. To address these issues, gesture-rich virtual agents could be leveraged as a substitute for the real streamer. Given that behavioral cues can significantly impact online user engagement [154], a gesture-rich agent could substantially improve customer interaction metrics. This alleviates the workload on human presenters and provides a level of consistency and engagement that can be programmed and fine-tuned for maximum impact.

6.4.2 Implications for generative AI teacher design. Building on our exploration of collaborative and reflective practices in PA gesture design, the future of embodied AI teacher development holds promising opportunities for incorporating advanced technologies from HCI, such as Computer Vision (CV), AI-Generated Content (AIGC), and Large Language Models (LLMs). These technologies offer untapped potential to enhance efficiency and flexibility in the PA design process by automating or augmenting aspects of gesture creation and content delivery.

Toward transcript-prompted AI teacher synthesis. Tools that leverage AIGC in creative media, as demonstrated by He et al. [155], where an AI generative model can process the educator's speech transcripts and instructional intents as input, generating multiple options for AI-generated gesture synthesis. Moreover, drawing inspiration from the combination of rule-based and learning-based techniques to select naturalistic instructional gestures [82, 4, 5], this research offers valuable insights into crafting prompts for semantic motion generation. This approach can augment the PA's alignment with teaching intents, ultimately enhancing AI's efficiency in delivering multimedia instructions within a virtual learning environment.

Furthermore, our preliminary findings indicate that cohesive and beat gestures contribute positively to students' perceptions of engagement and enjoyment. This insight makes a strong case for additional quantitative studies. One possible approach would be to use computer vision techniques to analyze the relationship between gesture types and parts of speech within large video datasets such as TED Talks. Insights from such analysis could inform the development of Generative Adversarial Network models that optimize gesture and speech coordination, resulting in more authentic and engaging pedagogical agents. AI-generated content technology also provides flexibility to design highly customizable and diverse virtual agent appearances, surpassing the limitations of traditional 3D animation platforms. While AI-generated content tools can handle complex visual customization, large language models could generate responsive instructional content and support real time interaction, which would create a more immersive and dynamic learning experience. Looking ahead in gesture design, a longer term trajectory involves shifting from pre-scripted content to agents capable of adapting gestures in real time. Future systems could utilize multimodal learner sensing, including gaze patterns, facial expressions, or affective signals, to adjust gesture intensity, pacing, and type according to real time engagement states. This would mark a transition from the human in control workflow featured here to a synchronous and adaptive instructional system in which pedagogical intent is not simply encoded at the outset but is continually negotiated between the agent and learner.

Toward highly controllable and versatile PA gestures. In this study, we primarily focused on cohesive and beat gestures as foundational elements for PAs in interactive environments. However, these gestures represent only the starting point of a broader exploration. Future research could expand this foundation by incorporating a diverse range of gesture types, each providing unique functionality depending on context and application. One promising direction involves deictic gestures, which can be used to precisely identify, point to, or highlight specific elements within the user interface or learning environment. These gestures are particularly useful in educational or interactive settings where users need to focus attention on particular details. Another area of growth is the inclusion of iconic gestures, which visually represent the content being manipulated or discussed. Such gestures could have immense utility in domains that rely on visualization, such as science and engineering, where complex structures and phenomena need to be depicted dynamically through

PA’s multimodal expressions. Additionally, metaphoric gestures can enable abstract and symbolic representations, which would be valuable in fields such as literature, social sciences, or philosophy, where concepts are often intangible and need a more creative and expressive mode of interaction. The future evolution of PA gestures can therefore offer users a high degree of control and versatility, enhancing the richness and expressiveness of human-computer interactions.

Enriching and diversifying training datasets. The use of diffusion-based models for generating co-speech gestures has gained considerable attention. While existing studies predominantly rely on TED talks or interview-style videos, these datasets primarily feature conversational content where gestures are often spontaneous and context-dependent. This focus on informal or conversational settings raises several concerns regarding the applicability and generalizability of the model in instructional or scripted settings. To enhance the model’s performance and applicability, we propose expanding and diversifying the training datasets, particularly by focusing on instructional videos that exhibit a broader range of gesture types. In these instructional settings, gestures are typically more deliberate and serve distinct communicative purposes — whether to emphasize key points, clarify instructions, or convey emotions. Videos of educators, coaches, and instructors should be included in the dataset, as these settings naturally offer a greater variety of gestures intended for clear communication. Moreover, including content that specifically focuses on the role of gestures in conveying meaning could significantly enhance the model’s understanding and generation of context-appropriate gestures.

6.5 Limitations

6.5.1 Experimental design and scope. This study employed a within-subjects design focused on perceptual validation, consistent with Research-through-Design (RtD) methodologies that prioritize understanding user experience and informing design principles [49]. We did not include a control condition (e.g., lecture without gestures or static slides), as our primary objective was to validate whether the collaboratively designed PA elicited coherent and positive student perceptions. The observed pre-post transfer gains served as manipulation checks to confirm meaningful engagement, rather than to isolate gesture-specific learning effects. Future controlled experiments with counterbalanced conditions (e.g., gesture-based PA vs. static PA vs. slides-only) are necessary to isolate

Manuscript submitted to IJHCI

the specific causal contributions of instructional gestures to learning effectiveness. Additionally, the 14-minute instructional duration, while aligned with pedagogical recommendations for establishing stable agent persona perceptions [111, 113], imposed practical limitations. The brief session restricted the complexity and depth of content covered and may have prevented students from fully experiencing sustained instructional interaction typical of authentic classroom settings. Future research should examine student perceptions and learning experiences in authentic classroom settings with extended instructional sessions. Additionally, our evaluation relied on self-reported perceptions and transfer test scores. Complementing these with objective physiological measures, such as gaze-tracking or EEG-based cognitive load assessment—would provide more direct empirical evidence of gestural contributions to learning processes. We recognize this as a meaningful direction for future work.

6.5.2 Sample population. Another limitation concerns the characteristics of the sample population. The sample exhibited gender imbalance, which may influence perceptions of agent characteristics. The study findings, derived from a specific age group and second language speakers in a Hong Kong educational context where English instruction begins early, may not be generalizable to a broader or different demographic. Future research should employ balanced gender sampling and recruit participants from diverse linguistic backgrounds (e.g., native speakers of Japanese, Russian, or other languages) and educational systems. It is crucial to explore how PAs interact with diverse age groups, cultural backgrounds, and learning styles to gain more comprehensive insights.

6.5.3 Instructional material selection. The teaching material was selected for its pictorial format and procedural structure to ensure cross-disciplinary accessibility. However, we did not explicitly analyze how disciplinary background might mediate comprehension across the diverse majors represented in our sample. Future research should systematically assess materials’ disciplinary adaptability and test the framework with content from multiple knowledge domains.

6.5.4 Pedagogical agent appearance design. The design of the PAs used in the study poses a limitation. We employed a minimalist virtual avatar lacking facial features to isolate gesture-specific effects. However, this design choice was not comparatively evaluated against more realistic agent

appearances. Given that agent appearance significantly influences student perceptions [9, 10], the faceless design may limit the external validity and generalizability of findings to real-world applications where more realistic agents are common. Future research should examine how varying levels of agent realism interact with gesture design across different instructional contexts.

6.5.5 Generative AI tool exploration and evaluation. Our exploration focused on two commercial platforms (Sora and KlingAI), representing a subset of available technologies. Other tools and open-source models may exhibit different capabilities in instructional gesture generation. The quantitative evaluation compared specific model versions, providing initial evidence but limited by small sample size and tool-specific focus. Given rapid technological evolution, findings reflect a particular moment. Future research should expand to diverse platforms, larger samples, and longitudinal assessments to establish broader design principles.

6.5.6 Interactivity of pedagogical agents. The limitation on PA interactivity goes beyond technical issues to include methodological challenges in gesture analysis. A key constraint is the limited gesture repository, which restricts PA diversity and user experience. We addressed this by analyzing gestures from experienced teachers in real scenarios, but future work should expand the gesture database and include more subjects.

7 Conclusion

In conclusion, our work addresses the critical gap between high-level principles and the practical design of co-speech gestures for pedagogical agents (PAs). We introduced and validated a human-centered design framework that empowers educators and designers to systematically translate instructional intentions into meaningful agent gestures. Our evaluation, involving a 14-minute course module assessed by 38 university students, demonstrated the framework’s effectiveness. The findings confirmed that an EPA with human-designed gestures was perceived as more professional and approachable, with students describing enriched engagement and positive learning experiences. This research offers three main contributions to the HCI community: 1) We provide a replicable four-stage design framework that enables a reflective and collaborative process for embodying pedagogical

Manuscript submitted to IJHCI

goals in an agent. 2) We offer empirical evidence that well-designed cohesive and beat gestures improve student perceptions and support meaningful learning. 3) Through exploratory testing and quantitative evaluation of text-to-video generation tools, we demonstrate that structured prompts with temporal phases and linguistic anchors improve instructional coherence, providing concrete design principles for future AI-generated instructional content.

Funding

This project is supported by the University Grants Committee (UGC) Funding Scheme (RHCE & G.73.xx.R006) from The Hong Kong Polytechnic University.

Declaration of Interest Statement

The authors report there are no competing interests to declare.

References

- [1] Jiayi Zhou, Renzhong Li, Junxiu Tang, Tan Tang, Haotian Li, Weiwei Cui, and Yingcai Wu. 2024. Understanding nonlinear collaboration between human and ai agents: a co-design framework for creative design. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–16.
- [2] Simbarashe Nyatsanga, Taras Kucherenko, Chaitanya Ahuja, Gustav Eje Henter, and Michael Neff. 2023. A comprehensive review of data-driven co-speech gesture generation. In *Computer Graphics Forum* number 2. Vol. 42. Wiley Online Library, 569–596.
- [3] Pieter Wolfert, Nicole Robinson, and Tony Belpaeme. 2022. A review of evaluation practices of gesture generation in embodied conversational agents. *IEEE Transactions on Human-Machine Systems*, 52, 3, 379–389.
- [4] Zeyi Zhang, Tenglong Ao, Yuyao Zhang, Qingzhe Gao, Chuan Lin, Baoquan Chen, and Libin Liu. 2024. Semantic gesticulator: semantics-aware co-speech gesture synthesis. *ACM Transactions on Graphics (TOG)*, 43, 4, 1–17.
- [5] Yihao Zhi, Xiaodong Cun, Xuelin Chen, Xi Shen, Wen Guo, Shaoli Huang, and Shenghua Gao. 2023. Livelspeaker: towards semantic-aware co-speech gesture generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 20807–20817.
- [6] Wenjing Li, Ziyi Kuang, Xiaoxue Leng, Richard E Mayer, and Fuxing Wang. 2024. Role of gesturing onscreen instructors in video lectures: a set of three-level meta-analyses on the embodiment effect. *Educational Psychology Review*, 36, 3, 67.
- [7] G. Kizilkaya and P. Aşkar. 2008. The effect of an embedded pedagogical agent on the students' science achievement. *Interactive Technology and Smart Education*, 5, 4, (Nov. 2008), 208–216. doi:10.1108/17415650810930893.
- [8] M. Savin-Baden, G. Tombs, and R. Bhakta. 2015. Beyond robotic wastelands of time: abandoned pedagogical agents and new pedalled pedagogies. *E-Learning and Digital Media*, 12, 3-4, (June 2015), 295–314. doi:10.1177/2042753015571835.
- [9] Maik Beege, Felix Kriegelstein, and Caroline Arnold. 2022. How instructors influence learning with instructional videos - the importance of professional appearance and communication. *Computers & Education*, 185, (Aug. 2022), 104531. doi:10.1016/j.compedu.2022.104531.

- [10] Youssef Shiban, Iris Schelhorn, Verena Jobst, Alexander Hörnlein, Frank Puppe, Paul Pauli, and Andreas Mühlberger. 2015. The appearance effect: influences of virtual agent features on performance and motivation. *Computers in Human Behavior*, 49, (Aug. 2015), 5–11. doi:10.1016/j.chb.2015.01.077.
- [11] J. Ceha and E. Law. 2022. Expressive auditory gestures in a voice-based pedagogical agent. In *Conference on Human Factors in Computing Systems - Proceedings*. ACM, (Apr. 2022). doi:10.1145/3491102.3517599.
- [12] Mariana Serras Pereira, Jolanda de Lange, Suleman Shahid, and Marc Swerts. 2018. A perceptual and behavioral analysis of facial cues to deception in interactions between children and a virtual agent. *International Journal of Child-Computer Interaction*, 15, (Mar. 2018), 1–12. doi:10.1016/j.ijcci.2017.10.003.
- [13] S. Castillo, P. Hahn, K. Legde, and D.W. Cunningham. 2018. Personality analysis of embodied conversational agents. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents, IVA 2018*. Vol. 2018-May. ACM, (Nov. 2018), 227–232. doi:10.1145/3267851.3267853.
- [14] P. Hahn, S. Castillo, and D.W. Cunningham. 2018. Look me in the lines: the impact of stylization on the recognition of expressions and perceived personality. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents, IVA 2018*. ACM, (Nov. 2018), 339–340. doi:10.1145/3267851.3267881.
- [15] S. Thomas, Y. Ferstl, R. McDonnell, and C. Ennis. 2022. Investigating how speech and animation realism influence the perceived personality of virtual characters and agents. In *Proceedings - 2022 IEEE Conference on Virtual Reality and 3D User Interfaces, VR 2022*. IEEE, (Mar. 2022), 11–20. doi:10.1109/VR51125.2022.00018.
- [16] Leon Reicherts, Yvonne Rogers, Licia Capra, Ethan Wood, Tu Dinh Duong, and Neil Sebire. 2022. It’s good to talk: a comparison of using voice versus screen-based interactions for agent-assisted tasks. *ACM Transactions on Computer-Human Interaction*, 29, 3, 1–41.
- [17] Michael Bonfert, Nima Zargham, Florian Saade, Robert Porzel, and Rainer Malaka. 2021. An evaluation of visual embodiment for voice assistants on smart displays. In *Proceedings of the 3rd Conference on Conversational User Interfaces*, 1–11.
- [18] Sunok Lee, Minji Cho, and Sangsu Lee. 2020. What if conversational agents became invisible? comparing users’ mental models according to physical entity of ai speaker. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4, 3, 1–24.
- [19] Thomas Mildner, Orla Cooney, Anna-Maria Meck, Marion Bartl, Gian-Luca Savino, Philip R Doyle, Diego Garaialde, Leigh Clark, John Sloan, Nina Wenig, et al. 2024. Listening to the voices: describing ethical caveats of conversational user interfaces according to experts and frequent users. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–18.
- [20] Lijia Lin, Paul Ginns, Tianhui Wang, and Peilin Zhang. 2020. Using a pedagogical agent to deliver conversational style instruction: what benefits can you obtain? *Computers & Education*, 143, (Jan. 2020), 103658. doi:10.1016/j.compedu.2019.103658.
- [21] Amy L. Baylor and Soyoung Kim. 2009. Designing nonverbal communication for pedagogical agents: when less is more. *Computers in Human Behavior*. Including the Special Issue: State of the Art Research into Cognitive Load Theory 25, 2, (Mar. 2009), 450–457. doi:10.1016/j.chb.2008.10.008.
- [22] Wenjing Li, Fuxing Wang, Richard E. Mayer, and Tao Liu. 2022. Animated pedagogical agents enhance learning outcomes and brain activity during learning. *Journal of Computer Assisted Learning*, 38, 3, (June 2022), 621–637. doi:10.1111/jcal.12634.
- [23] Alyssa J Kersey, Cristina Carrazza, Miriam A Novack, Eliza L Congdon, Elizabeth M Wakefield, Naureen Hemani-Lopez, and Susan Goldin-Meadow. 2024. The effects of gesture and action training on the retention of math equivalence. *Frontiers in psychology*, 15, 1386187.

- [24] Alyssa P. Lawson, Richard E. Mayer, Nicoletta Adamo-Villani, Bedrich Benes, Xingyu Lei, and Justin Cheng. 2021. Recognizing the emotional state of human and virtual instructors. *Computers in Human Behavior*, 114, (Jan. 2021), 106554. doi:10.1016/j.chb.2020.106554.
- [25] Youngwoo Yoon, Keunwoo Park, Minsu Jang, Jaehong Kim, and Geehyuk Lee. 2021. Sgtoolkit: an interactive gesture authoring toolkit for embodied conversational agents. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, (Oct. 2021), 826–840. ISBN: 978-1-4503-8635-7. doi:10.1145/3472749.3474789.
- [26] Gustav Bøg Petersen, Aske Mottelson, and Guido Makransky. 2021. Pedagogical agents in educational vr: an in the wild study. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, (May 7, 2021), 1–12. ISBN: 978-1-4503-8096-6. doi:10.1145/3411764.3445760.
- [27] Wan-Lun Tsai, Li-wen Su, Tsai-Yen Ko, Cheng-Ta Yang, and Min-Chun Hu. 2019. Improve the decision-making skill of basketball players by an action-aware vr training system. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, (Mar. 2019), 1193–1194. doi:10.1109/VR.2019.8798309.
- [28] Chih-Pu Dai, Fengfeng Ke, Yanjun Pan, Jewoong Moon, and Zhichun Liu. 2024. Effects of artificial intelligence-powered virtual agents on learning outcomes in computer-based simulations: a meta-analysis. *Educational Psychology Review*, 36, 1, 31.
- [29] Laduona Dai, Merel M Jung, Marie Postma, and Max M Louwerse. 2022. A systematic review of pedagogical agent research: similarities, differences and unexplored aspects. *Computers & Education*, 190, (Dec. 2022), 104607. doi:10.1016/j.compedu.2022.104607.
- [30] Robert. Davis. 2018. The impact of pedagogical agent gesturing in multimedia learning environments: a meta-analysis. *Educational Research Review*, 24, (June 2018), 193–209. doi:10.1016/j.edurev.2018.05.002.
- [31] Isaac Wang and Jaime Ruiz. 2021. Examining the use of nonverbal communication in virtual agents. *International Journal of Human–Computer Interaction*, 37, 17, (Oct. 2021), 1648–1673. doi:10.1080/10447318.2021.1898851.
- [32] Richard E. Mayer and C. Scott DaPra. 2012. An embodiment effect in computer-based learning with animated pedagogical agents. *Journal of Experimental Psychology: Applied*, 18, 3, 239–252. doi:10.1037/a0028616.
- [33] Wenjing Li, Fuxing Wang, Richard E. Mayer, and Huashan Liu. 2019. Getting the point: which kinds of gestures by pedagogical agents improve multimedia learning? *Journal of Educational Psychology*, 111, 8, (Nov. 2019), 1382–1395. doi:10.1037/edu0000352.
- [34] Jewoong Moon and Jeeheon Ryu. 2021. The effects of social and cognitive cues on learning comprehension, eye-gaze pattern, and cognitive load in video instruction. *Journal of Computing in Higher Education*, 33, 1, 39–63.
- [35] Robert O Davis, Joseph Vincent, and Lili Wan. 2021. Does a pedagogical agent’s gesture frequency assist advanced foreign language users with learning declarative knowledge? *International Journal of Educational Technology in Higher Education*, 18, 1–19.
- [36] DeepMotion. [n. d.] Create 3d animations from video using ai. (). Retrieved Aug. 15, 2023 from <https://www.deepmotion.com>.
- [37] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, Chris Landreth, and Karan Singh. 2020. Rignet: neural rigging for articulated characters. *arXiv preprint arXiv:2005.00559*.
- [38] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. 2018. Deepmimic: example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)*, 37, 4, 1–14.
- [39] Yixin Liu, Kai Zhang, Yuan Li, Zhiling Yan, Chujie Gao, Ruoxi Chen, Zhengqing Yuan, Yue Huang, Hanchi Sun, Jianfeng Gao, et al. 2024. Sora: a review on background, technology, limitations, and opportunities of large vision models. *arXiv preprint arXiv:2402.17177*.

- [40] Tianjun Zhang, Yi Zhang, Vibhav Vineet, Neel Joshi, and Xin Wang. 2023. Controllable text-to-image generation with gpt-4. *arXiv preprint arXiv:2305.18583*.
- [41] Zheng Zhu, Xiaofeng Wang, Wangbo Zhao, Chen Min, Bohan Li, Nianchen Deng, Min Dou, Yuqi Wang, Botian Shi, Kai Wang, et al. 2024. Is sora a world simulator? a comprehensive survey on general world models and beyond. *arXiv preprint arXiv:2405.03520*.
- [42] Maik Beege, Manuel Ninaus, Sascha Schneider, Steve Nebel, Julia Schlemmel, Jasmin Weidenmüller, Korbinian Moeller, and Günter Daniel Rey. 2020. Investigating the effects of beat and deictic gestures of a lecturer in educational videos. *Computers & Education*, 156, (Oct. 2020), 103955. doi:10.1016/j.compedu.2020.103955.
- [43] David McNeill. 2011. Hand and mind. In *Hand and Mind*. De Gruyter Mouton, (Dec. 2011), 351–374. ISBN: 978-3-11-087425-9. doi:10.1515/9783110874259.351.
- [44] Zhongling Pi, Fangfang Zhu, Yi Zhang, Louqi Chen, and Jiumin Yang. 2022. Complexity of visual learning material moderates the effects of instructor’s beat gestures and head nods in video lectures. *Learning and Instruction*, 77, (Feb. 2022), 101520. doi:10.1016/j.learninstruc.2021.101520.
- [45] Kenneth Shaw, Shikhar Bahl, Aravind Sivakumar, Aditya Kannan, and Deepak Pathak. 2024. Learning dexterity from human hand motion in internet videos. *The International Journal of Robotics Research*, 43, 4, 513–532.
- [46] Sascha Schneider, Maik Beege, Steve Nebel, Lenka Schnaubert, and Günter Daniel Rey. 2022. The cognitive-affective-social theory of learning in digital environments (castle). *Educational Psychology Review*, 34, 1, 1–38.
- [47] Heng Zhang, Yuhan Liu, Meilin Jiang, Juanjuan Chen, Minhong Wang, and Fred Paas. 2025. Emotional artificial intelligence in education: a systematic review and meta-analysis. *Educational Psychology Review*, 37, 4, 106.
- [48] Shunan Zhang, Yincen Li, Guangji Gan, Sirui Pang, Jang Hyun Kim, et al. 2025. Effects of social cues of artificial intelligence-powered pedagogical agents: a multilevel meta-analysis. *Educational Research Review*, 49, 100746.
- [49] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in hci. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, (Apr. 2007), 493–502. doi:10.1145/1240624.1240704.
- [50] G. Ali, M. Lee, and J.-I. Hwang. 2020. Automatic text-to-gesture rule generation for embodied conversational agents. *Computer Animation and Virtual Worlds*, 31, 4, (July 2020). doi:10.1002/cav.1944.
- [51] Taras Kucherenko, Patrik Jonell, Sanne Van Waveren, Gustav Eje Henter, Simon Alexandersson, Iolanda Leite, and Hedvig Kjellström. 2020. Gesticulator: a framework for semantically-aware speech-driven gesture generation. In *Proceedings of the 2020 international conference on multimodal interaction*. ACM, (Oct. 2020), 242–250. doi:10.1145/3382507.3418815.
- [52] Xian Liu, Qianyi Wu, Hang Zhou, Yuanqi Du, Wayne Wu, Dahua Lin, and Ziwei Liu. 2022. Audio-driven co-speech gesture video generation. In *Advances in Neural Information Processing Systems*. Vol. 35. Curran Associates, Inc., 21386–21399. doi:10.48550/arxiv.2212.02350.
- [53] Lai Wei and Chow Kenny K.N. 2023. When gestures and words synchronize: exploring a human lecturer’s multimodal interaction for the design of embodied pedagogical agents. In *In Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing (CSCW ’23)*. The 2023 Conference on Computer Supported Cooperative Work and Social Computing (CSCW ’23). Association for Computing Machinery, New York, NY, USA.
- [54] Wei Xu. 2019. Toward human-centered ai: a perspective from human-computer interaction. *interactions*, 26, 4, 42–46.
- [55] Ben Shneiderman. 2020. Human-centered artificial intelligence: reliable, safe & trustworthy. *International Journal of Human–Computer Interaction*, 36, 6, 495–504.
- [56] Jason W Burton, Mari-Klara Stein, and Tina Blegind Jensen. 2020. A systematic review of algorithm aversion in augmented decision making. *Journal of behavioral decision making*, 33, 2, 220–239.

- [57] Ozlem Ozmen Garibay, Brent Winslow, Salvatore Andolina, Margherita Antona, Anja Bodenschatz, Constantinos Coursaris, Gregory Falco, Stephen M Fiore, Ivan Garibay, Keri Grieman, et al. 2023. Six human-centered artificial intelligence grand challenges. *International Journal of Human-Computer Interaction*, 39, 3, 391–437.
- [58] Wei Xu and Zaifeng Gao. 2023. Enabling human-centered ai: a methodological perspective. *arXiv preprint arXiv:2311.06703*.
- [59] Andrew Brock. 2018. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*.
- [60] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34, 8780–8794.
- [61] Jakob Tholander and Martin Jonsson. 2023. Design ideation with ai-sketching, thinking and talking with generative machine learning models. In *Proceedings of the 2023 ACM designing interactive systems conference*, 1930–1940.
- [62] Savvas Petridis, Michael Terry, and Carrie J Cai. 2024. Promptinfuser: how tightly coupling ai and ui design impacts designers’ workflows. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, 743–756.
- [63] Gun Woo Park, Payod Panda, Lev Tankelevitch, and Sean Rintel. 2024. The coexplorer technology probe: a generative ai-powered adaptive interface to support intentionality in planning and running video meetings. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, 1638–1657.
- [64] Severi Uusitalo, Antti Salovaara, Tero Jokela, and Marja Salmimaa. 2024. ” clay to play with”: generative ai tools in ux and industrial design practice. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, 1566–1578.
- [65] Justin D Weisz, Jessica He, Michael Muller, Gabriela Hoefler, Rachel Miles, and Werner Geyer. 2024. Design principles for generative ai applications. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–22.
- [66] Adam Hollmén Larsen and Jichen Zhu. 2024. Ideary: facilitating electronic music creation with generative ai. In *Companion Publication of the 2024 ACM Designing Interactive Systems Conference*, 275–278.
- [67] Peter Kun, Matthias Anton Freiburger, Anders Sundnes Løvlie, and Sebastian Risi. 2024. Genframe–embedding generative ai into interactive artifacts. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, 714–727.
- [68] Hongbo Zhang, Pei Chen, Xuelong Xie, Chaoyi Lin, Lianyan Liu, Zhuoshu Li, Weitao You, and Lingyun Sun. 2024. Protodreamer: a mixed-prototype tool combining physical model and generative ai to support conceptual design. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, 1–18.
- [69] Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. 2023. Visual chatgpt: talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*.
- [70] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.
- [71] Shuai Liu, Zhe Luo, and Weina Fu. 2024. Fcdnet: fuzzy cognition-based dynamic fusion network for multimodal sentiment analysis. *IEEE Transactions on Fuzzy Systems*, 33, 1, 3–14.
- [72] Xiaohan Peng, Janin Koch, and Wendy E Mackay. 2024. Designprompt: using multimodal interaction for design exploration with generative ai. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, 804–818.
- [73] Bekzat Tilekbay, Saelyne Yang, Michal Adam Lewkowicz, Alex Suryapranata, and Juho Kim. 2024. Expressedit: video editing with natural language and sketching. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*, 515–536.
- [74] Setareh Aghel Manesh, Tianyi Zhang, Yuki Onishi, Kotaro Hara, Scott Bateman, Jiannan Li, and Anthony Tang. 2024. How people prompt generative ai to create interactive vr scenes. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, 2319–2340.

- [75] Weiyue Gao, Yihan Mei, Henry Duh, and Zhibin Zhou. 2024. Envisioning the incorporation of generative artificial intelligence into future product design education: insights from practitioners, educators, and students. *The Design Journal*, 1–21.
- [76] Bohong Chen, Yumeng Li, Youyi Zheng, Yao-Xiang Ding, and Kun Zhou. 2025. Motion-example-controlled co-speech gesture generation leveraging large language models. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, 1–12.
- [77] Youngwoo Yoon, Pieter Wolfert, Taras Kucherenko, Carla Viegas, Teodor Nikolov, Mihail Tsakov, and Gustav Eje Henter. 2022. The genea challenge 2022: a large evaluation of data-driven co-speech gesture generation. In *Proceedings of the 2022 International Conference on Multimodal Interaction*, 736–747.
- [78] Justine Cassell. 1998. A framework for gesture generation and interpretation. *Computer vision in human-machine interaction*, 191–215.
- [79] Gilwoo Lee, Zhiwei Deng, Shugao Ma, Takaaki Shiratori, Siddhartha S Srinivasa, and Yaser Sheikh. 2019. Talking with hands 16.2 m: a large-scale dataset of synchronized body-finger motion and audio for conversational motion analysis and synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 763–772.
- [80] Youngwoo Yoon, Bok Cha, Joo-Haeng Lee, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. 2020. Speech gesture generation from the trimodal context of text, audio, and speaker identity. *ACM Transactions on Graphics*.
- [81] Youngwoo Yoon, Woo-Ri Ko, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. 2019. Robots learn social skills: end-to-end learning of co-speech gesture generation for humanoid robots. In *2019 International Conference on Robotics and Automation (ICRA)*. 2019 International Conference on Robotics and Automation (ICRA). IEEE, (May 2019), 4303–4309. doi:10.1109/ICRA.2019.8793720.
- [82] Najmeh Sadoughi and Carlos Busso. 2019. Speech-driven animation with meaningful behaviors. *Speech Communication*, 110, 90–100.
- [83] Haozhou Pang, Tianwei Ding, Lanshan He, Ming Tao, Lu Zhang, and Qi Gan. 2025. Llm gesticulator: leveraging large language models for scalable and controllable co-speech gesture synthesis. In *Eighth International Conference on Computer Graphics and Virtuality (ICCGV 2025)*. Vol. 13557. SPIE, 1355702.
- [84] Qingrong Cheng, Xu Li, and Xinghui Fu. 2024. Siggester: generalized co-speech gesture synthesis via semantic injection with large-scale pre-training diffusion models. In *SIGGRAPH Asia 2024 Conference Papers*, 1–11.
- [85] Ikhsanul Habibie, Weipeng Xu, Dushyant Mehta, Lingjie Liu, Hans-Peter Seidel, Gerard Pons-Moll, Mohamed Elgharib, and Christian Theobalt. 2021. Learning speech-driven 3d conversational gestures from video. In *Proceedings of the 21st ACM International Conference on Intelligent Virtual Agents*, 101–108.
- [86] Shiry Ginosar, Amir Bar, Gefen Kohavi, Caroline Chan, Andrew Owens, and Jitendra Malik. 2019. Learning individual styles of conversational gesture. (2019). <https://arxiv.org/abs/1906.04160> arXiv: 1906.04160 [cs.CV].
- [87] Haiyang Liu, Zihao Zhu, Naoya Iwamoto, Yichen Peng, Zhengqing Li, You Zhou, Elif Bozkurt, and Bo Zheng. 2022. Beat: a large-scale semantic and emotional multi-modal dataset for conversational gestures synthesis. In *European conference on computer vision*. Springer, 612–630.
- [88] Hongwei Yi, Hualin Liang, Yifei Liu, Qiong Cao, Yandong Wen, Timo Bolkart, Dacheng Tao, and Michael J Black. 2023. Generating holistic 3d human motion from speech. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 469–480.
- [89] Sicheng Yang, Zhiyong Wu, Minglei Li, Zhensong Zhang, Lei Hao, Weihong Bao, Ming Cheng, and Long Xiao. 2023. Diffusestylegesture: stylized audio-driven co-speech gesture generation with diffusion models. *arXiv preprint arXiv:2305.04919*.
- [90] Eri Nakagawa, Motofumi Sumiya, Takahiko Koike, and Norihiro Sadato. 2021. The neural network underpinning social feedback contingent upon one’s action: an fmri study. *NeuroImage*, 225, (Jan. 2021), 117476. doi:10.1016/j.neuroimage.2020.117476.

- [91] A.M. Sinatra, K.A. Pollard, B.T. Files, A.H. Oiknine, M. Ericson, and P. Khooshabeh. 2021. Social fidelity in virtual agents: impacts on presence and learning. *Computers in Human Behavior*, 114, (Jan. 2021), 106562. doi:10.1016/j.chb.2020.106562.
- [92] Sascha Schneider, Felix Krieglstein, Maik Beege, and Günter Daniel Rey. 2022. The impact of video lecturers' nonverbal communication on learning – an experiment on gestures and facial expressions of pedagogical agents. *Computers & Education*, 176, (Jan. 2022), 104350. doi:10.1016/j.compedu.2021.104350.
- [93] Susan Goldin-Meadow. 2015. Gesture as a window onto communicative abilities: Implications for diagnosis and intervention. *Perspectives on language learning and education*, 22, 2, (Mar. 2015), 50–60. pmid: 26366247. doi:10.1044/lle22.2.50.
- [94] S.D. Craig, J. Twyford, N. Irigoyen, and S.A. Zipp. 2015. A test of spatial contiguity for virtual human's gestures in multimedia learning environments. *Journal of Educational Computing Research*, 53, 1, (Sept. 2015), 3–14. doi:10.1177/0735633115585927.
- [95] R.O. Davis and J. Vincent. 2019. Sometimes more is better: agent gestures, procedural knowledge and the foreign language learner. *British Journal of Educational Technology*, 50, 6, (Nov. 2019), 3252–3263. doi:10.1111/bjet.12732.
- [96] Lai Wei and Kenny K. N. Chow. 2023. How students perceive lecturers' gestures? an exploration in gesture-meaning matching toward embodied pedagogical agent design. In *the tenth Congress of the International Association of Societies of Design Research (IASDR 2023)*.
- [97] Ingrid Masson-Carro, Martijn Goudbeek, and Emiel Krahmer. 2017. How what we see and what we know influence iconic gesture production. *Journal of nonverbal behavior*, 41, 367–394.
- [98] Kawai Chui. 2005. Temporal patterning of speech and iconic gestures in conversational discourse. *Journal of Pragmatics*, 37, 6, 871–887.
- [99] Kazuki Sekine and Sotaro Kita. 2015. Development of multimodal discourse comprehension: cohesive use of space by gestures. *Language, Cognition and Neuroscience*, 30, 10, 1245–1258.
- [100] Hassan Belhiah. 2013. Using the hand to choreograph instruction: on the functional role of gesture in definition talk. *The Modern Language Journal*, 97, 2, 417–434.
- [101] David McNeill and Elena T Levy. 1993. Cohesion and gesture. *Discourse processes*, 16, 4, 363–386.
- [102] Qi Wu, Cheng-Ju Wu, Yixin Zhu, and Jungseock Joo. 2021. Communicative learning with natural gestures for embodied navigation agents with human-in-the-scene. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. IEEE, (Sept. 2021), 4095–4102. doi:10.1109/iro51168.2021.9636208.
- [103] Lai Wei and Kenny K. N. Chow. 2022. Who shapes the network of a pedagogical space? clues from the movements in the physical places. In *Advances in Mobile Computing and Multimedia Intelligence: 20th International Conference, MoMM 2022, Virtual Event, November 28–30, 2022, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 143–153. ISBN: 978-3-031-20435-7. doi:10.1007/978-3-031-20436-4_14.
- [104] Walter Dick, Lou Carey, and James O Carey. 2005. The systematic design of instruction.
- [105] M David Merrill. 2002. First principles of instruction. *Educational technology research and development*, 50, (Sept. 2002), 43–59. doi:10.1007/bf02505024.
- [106] Robert Gagne. 1985. The conditions of learning and theory of instruction robert gagné. *New York, NY: Holt, Rinehart ja Winston*.
- [107] Charles M Reigeluth. 1999. What is instructional-design theory and how is it changing. *Instructional-design theories and models: A new paradigm of instructional theory*, 2, 5–29.
- [108] Richard E. Mayer. [n. d.] Principles of multimedia learning based on social cues: personalization, voice, and image principles, 201–212. ISBN: 978-0-521-54751-2. doi:10.1017/CBO9780511816819.014.
- [109] Richard E. Mayer. 2002. Multimedia learning. In *Psychology of Learning and Motivation*. Vol. 41. Academic Press, 85–139. doi:10.1016/S0079-7421(02)80005-6.

- [110] Philip J Guo, Juho Kim, and Rob Rubin. 2014. How video production affects student engagement: an empirical study of mooc videos. In *Proceedings of the first ACM conference on Learning@ scale conference*, 41–50.
- [111] Larry Lagerstrom, Petr Johanes, and Umnouy Ponsukcharoen. 2015. The myth of the six-minute rule: student engagement with online videos. In *2015 ASEE annual conference & exposition*, 26–1558.
- [112] Ahmad Manasrah, Mohammad Masoud, and Yousef Jaradat. 2021. Short videos, or long videos? a study on the ideal video length in online learning. In *2021 international conference on information technology (ICIT)*. IEEE, 366–370.
- [113] Juan C Castro-Alonso, Rachel M Wong, Olusola O Adesope, and Fred Paas. 2021. Effectiveness of multimedia pedagogical agents predicted by diverse theories: a meta-analysis. *Educational Psychology Review*, 33, 3, 989–1015.
- [114] Noah L Schroeder, Robert O Davis, and Eunbyul Yang. 2025. Designing and learning with pedagogical agents: an umbrella review. *Journal of Educational Computing Research*, 62, 8, 1907–1936.
- [115] James R Williams. 1998. Guidelines for the use of multimedia in instruction. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting number 20*. Vol. 42. SAGE Publications Sage CA: Los Angeles, CA, 1447–1451.
- [116] Ray Pastore. 2012. The effects of time-compressed instruction and redundancy on learning and learners’ perceptions of cognitive load. *Computers & Education*, 58, 1, 641–651.
- [117] Geneva Gay. 2002. Preparing for culturally responsive teaching. *Journal of teacher education*, 53, 2, (Mar. 2002), 106–116. doi:10.1177/0022487102053002003.
- [118] David Wiley and John Hilton III. 2009. Openness, dynamic specialization, and the disaggregated future of higher education. *International Review of Research in Open and Distributed Learning*, 10, 5, (Nov. 2009). doi:10.19173/irrodl.v10i5.768.
- [119] Thomas T Hewett, Ronald Baecker, Stuart Card, Tom Carey, Jean Gasen, Marilyn Mantei, Gary Perlman, Gary Strong, and William Verplank. 1992. *ACM SIGCHI curricula for human-computer interaction*. ACM.
- [120] Lauren Wilcox, Betsy DiSalvo, Dick Henneman, and Qiaosi Wang. 2019. Design in the hci classroom: setting a research agenda. In *Proceedings of the 2019 on Designing Interactive Systems Conference*, 871–883.
- [121] Siddharth Mehrotra, Carolina Centeio Jorge, Catholijn M Jonker, and Myrthe L Tielman. 2024. Integrity-based explanations for fostering appropriate trust in ai agents. *ACM Transactions on Interactive Intelligent Systems*, 14, 1, 1–36.
- [122] W Lewis Johnson and James C Lester. 2016. Face-to-face interaction with pedagogical agents, twenty years later. *International Journal of Artificial intelligence in education*, 26, 1, 25–36.
- [123] Chris Birmingham, Zijian Hu, Kartik Mahajan, Eli Reber, and Maja J Matarić. 2020. Can i trust you? a user study of robot mediation of a support group. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 8019–8026.
- [124] Kenny K. N. Chow. 2026. What AI pretends to be? Design rhetoric and trust in the appearances of AI applications. *Pragmatics and Society*. Forthcoming.
- [125] John D Lee and Katrina A See. 2004. Trust in automation: designing for appropriate reliance. *Human factors*, 46, 1, 50–80.
- [126] Kristina Höök and Jonas Löwgren. 2012. Strong concepts: intermediate-level knowledge in interaction design research. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 19, 3, 1–18.
- [127] Donald A. Schon. 1992. Designing as reflective conversation with the materials of a design situation. *Research in Engineering Design*, 3, 3, (Sept. 1, 1992), 131–147. doi:10.1007/BF01580516.
- [128] Tim Ingold. 2013. *Making: Anthropology, Archaeology, Art and Architecture*. Routledge, London, (Mar. 29, 2013). 176 pp. ISBN: 978-0-203-55905-5. doi:10.4324/9780203559055.

- [129] Mikael Wiberg. 2014. Methodology for materiality: interaction design research through a material lens. *Personal and Ubiquitous Computing*, 18, 3, (Mar. 1, 2014), 625–636. doi:10.1007/s00779-013-0686-7.
- [130] John Seely Brown, Allan Collins, and Paul Duguid. 1989. Situated cognition and the culture of learning. *1989*, 18, 1, 32–42.
- [131] Stefano Magistretti, Lorenzo Ardito, and Antonio Messeni Petruzzelli. 2021. Framing the microfoundations of design thinking as a dynamic capability for innovation: reconciling theory and practice. *Journal of Product Innovation Management*, 38, 6, (Nov. 2021), 645–667. doi:10.1111/jpim.12586.
- [132] Lisa Carlgren, Ingo Rauth, and Maria Elmquist. 2016. Framing design thinking: the concept in idea and enactment. *Creativity and innovation management*, 25, 1, (Mar. 2016), 38–57.
- [133] Robert G Cooper and Anita F Sommer. 2016. The agile–stage-gate hybrid model: a promising new approach and a new research opportunity. *Journal of Product Innovation Management*, 33, 5, (Sept. 2016), 513–526.
- [134] Pietro Micheli, Sarah JS Wilner, Sabeen Hussain Bhatti, Matteo Mura, and Michael B Beverland. 2019. Doing design thinking: conceptual review, synthesis, and research agenda. *Journal of Product Innovation Management*, 36, 2, (Mar. 2019), 124–148. doi:10.1111/jpim.12466.
- [135] Jaimie Lee Freeman and Amanda Nicole Curtis. 2023. Putting the self in self-tracking: the value of a co-designed ‘how might you’ self-tracking guide for teenagers. In *CHI ’23. Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, (Apr. 2023). doi:10.1145/3544548.3580938.
- [136] Mursyid Ibrahim, Penny Sweetser, and Anne Ozdowska. 2023. Tutorial level design guidelines for 2d fighting games. In *FDG ’23. Proceedings of the 18th International Conference on the Foundations of Digital Games*. Association for Computing Machinery, New York, NY, USA, (Apr. 2023). doi:10.1145/3582437.3582470.
- [137] Kirill Shatilov, Ahmad Alhilal, Tristan Braud, Lik-Hang Lee, Pengyuan Zhou, and Pan Hui. 2023. Players are not ready 101: a tutorial on organising mixed-mode events in the metaverse. In *MetaSys ’23. Proceedings of the First Workshop on Metaverse Systems and Applications*. Association for Computing Machinery, New York, NY, USA, (June 2023), 14–20. doi:10.1145/3597063.3597360.
- [138] Sark Pangrui Xing, Bart Van Dijk, Pengcheng An, Miguel Bruns, Yaliang Chuang, and Stephen Jia Wang. 2023. Puffy: a step-by-step guide to craft bio-inspired artifacts with interactive materiality. In *TEI ’23. Proceedings of the Seventeenth International Conference on Tangible, Embedded, and Embodied Interaction*. Association for Computing Machinery, New York, NY, USA, (Feb. 2023). doi:10.1145/3569009.3572800.
- [139] Alan Cienki and Cornelia Müller. 2008. Metaphor, gesture, and thought. *The Cambridge handbook of metaphor and thought*, 483, 501.
- [140] ELAN. [n. d.] Elan is an annotation tool for audio and video recordings. (). Retrieved Aug. 15, 2023 from <https://archive.mpi.nl/tla/elan>.
- [141] Johnny Saldaña. 2021. The coding manual for qualitative researchers.
- [142] OtterAI. [n. d.] Voice meeting notes & real-time transcription. (). Retrieved Aug. 15, 2023 from <https://otter.ai>.
- [143] Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O’Reilly Media, Inc".
- [144] Plask. [n. d.] Motion: ai-powered mocap animation tool. (). Retrieved Aug. 15, 2023 from <https://plask.ai>.
- [145] MoveAI. [n. d.] Motion capture in any environment. (). Retrieved Aug. 15, 2023 from <https://move.ai>.
- [146] Amy L Baylor and Jeeheon Ryu. 2003. The effects of image and animation in enhancing pedagogical agent persona. *Journal of Educational Computing Research*, 28, 4, 373–394.
- [147] Rudolph Flesch. 1948. A new readability yardstick. *Journal of applied psychology*, 32, 3, 221.

- [148] J Peter Kincaid, Robert P Fishburne Jr, Richard L Rogers, and Brad S Chissom. 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Tech. rep.
- [149] Victoria Clarke and Virginia Braun. 2017. Thematic analysis. *The journal of positive psychology*, 12, 3, 297–298.
- [150] Allison Druin. 2002. The role of children in the design of new technology. *Behaviour and information technology*, 21, 1, (Jan. 2002), 1–25. doi:10.1080/01449290110108659.
- [151] Wing Chee So, Colin Sim Chen-Hui, and Julie Low Wei-Shan. 2012. Mnemonic effect of iconic gesture and beat gesture in adults and children: is meaning in gesture important for memory recall? *Language and Cognitive Processes*, 27, 5, (June 2012), 665–681. doi:10.1080/01690965.2011.573220.
- [152] DeepBrain. [n. d.] Conversational ai humans humanize digital engagement. (). Retrieved Aug. 15, 2023 from <https://www.deepbrain.io/aihuman>.
- [153] Synthesia. [n. d.] Ai spokesperson video creator. (). Retrieved Aug. 15, 2023 from <https://www.synthesia.io/tools/video-spokesperson>.
- [154] Bilal Ahmed, Shagufta Zada, Liang Zhang, Shehla Najib Sidiki, Nicolás Contreras-Barraza, Alejandro Vega-Muñoz, and Guido Salazar-Sepúlveda. 2022. The impact of customer experience and customer engagement on behavioral intentions: does competitive choices matters? *Frontiers in Psychology*, 13, (May 2022), 864841. doi:10.3389/fpsyg.2022.864841.
- [155] Rui He, Huaxin Wei, and Ying Cao. 2024. An interactive system for supporting creative exploration of cinematic composition designs. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, 1–15.