Review

Character Design and Animation in RPG-Based Games: A Comprehensive Review of Traditional and AI-Driven Pipelines

Sri Kumaran R. P. *

Posted Date: 7 April 2026

doi: 10.20944/preprints202604.0296.v1

Keywords: RPG games; character design; character animation; generative AI; motion synthesis; facial animation; avatar customization; game development pipeline; identity preservation; human-in-the-loop systems

 doi:10.20944/preprints202604.0296.v1

Review

Character Design and Animation in RPG-Based Games: A Comprehensive Review of Traditional and AI-Driven Pipelines

Sri Kumaran R. P.

VIT-AP University, Amaravati, India; parthiban.23mis7191@vitapstudent.ac.in

Abstract

Role-playing games (RPGs) rely on character systems as the primary interface between narrative, player agency, and gameplay mechanics. Unlike static visual assets, RPG characters must preserve identity across customization, animation, equipment variation, and long-term progression. This paper presents a comprehensive review of character design and animation pipelines in RPG-based games, integrating perspectives from visual identity construction, motion synthesis, and AI-assisted generation workflows. The review synthesizes foundational studies on shape language, avatar representation, and expressive animation with recent advancements in generative artificial intelligence, including diffusion models, generative adversarial networks (GANs), dialogue-based editing systems, and speech-driven facial animation. The analysis highlights a transition from manual, artist-driven workflows toward hybrid pipelines where AI supports ideation, parameter editing, sprite generation, and animation synthesis. However, significant challenges remain, including weak semantic controllability, lack of identity preservation across transformations, inconsistencies between stylized and realistic character pipelines, and the absence of unified frameworks integrating design, rigging, and animation. This paper identifies these gaps and proposes future directions toward unified, human-in-the-loop character pipelines that ensure scalability, controllability, and consistency in next-generation RPG systems.

Keywords: RPG games; character design; character animation; generative AI; motion synthesis; facial animation; avatar customization; game development pipeline; identity preservation; human-in-the-loop systems

1. Introduction

In RPG-based games, characters are not simply visual assets. They are the main carriers of class identity, narrative role, emotional tone, and mechanical function. A healer, rogue, warrior, or mage is often recognized immediately through silhouette and form language [4]. Because RPGs often include long-term progression, outfit changes, branching story paths, and cosmetic customization, character design cannot be treated as a one-time illustration task. It must be designed as a flexible system that can survive transformation while remaining visually coherent and emotionally readable.

Recent work shows that the field is moving quickly. Earlier studies emphasized character design workflow as a foundational problem [2]. Character identity and shape language were also established as key visual cues in RPGs [4]. Motion creation had already begun to be addressed in earlier online-game character work [9]. More recent preprints and survey papers expand the field toward AI-assisted pipelines in which generative models and interactive editing systems support concept creation, parameter refinement, sprite-sheet generation, and animation synthesis [12]. This transition matters for RPGs because the genre has high asset volume and high identity demands at the same time. A single game may need many character archetypes, many expressions, many equipment combinations, and many animation states, all of which must remain internally consistent. Recent IEEE Access work on GenePixKolor (GPK) Fusion also shows how generative and optimization-

 doi:10.20944/preprints202604.0296.v1

driven systems can support game-related digital asset generation and rarity ranking in adjacent tokenomics ecosystems [1].

This review focuses on the intersection of character design and animation in RPG-based games. Its purpose is to consolidate what the literature establishes, explain how different studies fit together, and identify the unresolved problems that still limit adoption in real production pipelines. Rather than treating design and animation as independent stages, the paper argues that they should be viewed as one connected pipeline: concept art shapes the rig, the rig shapes the motion, motion shapes player perception, and player feedback shapes redesign. The review therefore emphasizes papers that are accessible as preprints or recent scholarly publications because those sources best reflect the active direction of the field [12].

2. Review Methodology

This paper follows a structured narrative review approach. Literature was gathered from Google Scholar, arXiv, IEEE Xplore, ACM Digital Library, and related open-access research repositories using keyword combinations such as “RPG character design,” “game character animation,” “avatar creation,” “facial animation,” “motion synthesis,” and “generative AI in games.” The review was intentionally broad at first so that both classical and recent work could be captured, then narrowed to studies that directly contribute to visual design, motion generation, or AI-supported character creation. The final selection emphasized papers that were accessible, methodologically clear, and relevant to the RPG character pipeline.

Inclusion criteria required that each study address game characters, character animation, or an adjacent pipeline used in interactive game environments. Papers were prioritized if they proposed a method, framework, or critical analysis that could be applied to RPG characters in particular. Exclusion criteria removed purely cinematic animation work without interactivity, papers without clear methodological contribution, and works unrelated to character generation, customization, or motion behavior in games. The final set was organized into three thematic layers: visual identity and design, motion and expressive animation, and AI-assisted generation and editing.

Because the goal is a review paper rather than an experimental benchmark, the emphasis here is on synthesis, comparison, and gap analysis. The literature is therefore organized by function and contribution rather than by publication venue alone. This allows the paper to move from foundational design theory to modern generative pipelines while keeping the RPG context central.

3. Literature Review

3.1. Character Design Workflow and Visual Identity

In RPGs, character design begins with form, but the form must communicate meaning [4]. Voimala’s workflow review is useful because it frames character design as a staged process that begins with ideation, continues through background development, and ends in visual realization [2]. In that view, the designer is not merely drawing a person; the designer is building a readable identity system. Shape language is one of the clearest visual tools for such identity construction [4]. RPGs use this principle constantly: rounded forms often suggest friendliness or support, angular structures can imply aggression or authority, and asymmetry may suggest improvisation, corruption, or uniqueness. The practical implication is that art evaluation should focus not only on aesthetics but also on communicative clarity.

A second issue is role continuity. RPG characters are expected to evolve through leveling, weapon changes, gear upgrades, and story-driven transformation while still remaining recognizable. The literature on character play also shows that avatars in multiplayer RPGs act both as personal expression and as social objects, so design must support self-identification and external readability at the same time [3].

3.2. Character Creation, Customization, and Identity Transfer

 doi:10.20944/preprints202604.0296.v1

A major line of research in game character creation studies how user input can be mapped into avatar appearance. Face-to-Parameter Translation for Game Character Auto-Creation proposed a pipeline that uses facial similarity and differentiable imitation of the game engine to create in-game characters from a face photo [5]. This is important for RPGs because it converts character creation from a manual modeling task into a personalization problem. In effect, the player is no longer constrained by fixed templates.

PokerFace-GAN extends this line by focusing on neutral-face generation and by disentangling identity from expression [6]. This distinction is especially relevant in RPGs, where a player may want the avatar to resemble a real face but still support a wide emotional range during gameplay. The paper's differentiable rendering and adversarial design show how modern character creation increasingly depends on machine-learning representations of the rendering process itself.

ICE advances the customization experience further by introducing multi-round dialogue-based editing [7]. This changes customization from a one-shot parameter search into an iterative conversation. For RPGs, that matters because players often refine avatars gradually: first the face, then the hair, then the body, and then clothing or armor. Dialogue-based editing matches that natural workflow and lowers the barrier for non-expert users.

3.3. Motion Creation and Expressive Animation

Animation gives the RPG character timing, weight, and behavior. Early game-oriented work on motion creation for online characters already recognized that the player perceives a character through motion as much as through appearance [9]. Even though that research predates current generative AI, the central insight remains current: motion creation must be integrated into the character pipeline rather than treated as an isolated finishing stage.

Emotion Capture shows why this matters for game experience [10]. Expressive characters do not simply move; they communicate mood through posture, gestures, and facial movement. In RPGs, emotional readability is essential in dialogue scenes, quest reactions, cutscenes, companion interactions, and combat response. A character that looks rich in concept art but moves stiffly in the engine will feel unconvincing even if the visual design is strong.

Recent survey work on human interaction motion generation broadens this topic by organizing motion synthesis around human-human, human-object, and human-scene interactions [14]. That taxonomy is highly relevant to RPGs because the genre constantly requires characters to interact with weapons, doors, objects, environmental obstacles, and other characters. This makes motion synthesis a core part of character design rather than a separate downstream feature.

3.4. Facial Animation, Lip Sync, and Emotional Realism

RPGs increasingly rely on voiced dialogue and cinematic cutscenes, which makes facial animation a central quality factor. Speech-driven facial animation research aims to synchronize lip movement while preserving emotional expression and identity [15]. The field has moved beyond simple lip-syncing toward emotionally aware motion, which is particularly relevant for RPG companions, quest givers, and player-controlled avatars who must react in believable ways.

The newer audio-driven systems also show a practical game-development orientation [15]. This is critical because RPG facial animation usually runs inside the game loop and must work at interactive frame rates.

Even with these advances, the literature still reveals a gap between research prototypes and production needs [12]. Game teams need facial systems that work across stylized and realistic characters, multiple rig configurations, different lighting setups, and many levels of expressiveness. Realism alone is not enough; the animation must also match the game's art direction. A hyperreal facial solver may fit one RPG, while a simplified stylized or anime-like pipeline may be better for another.

3.5. AI-Assisted Creation and Generative Pipelines

 doi:10.20944/preprints202604.0296.v1

Recent literature increasingly treats AI as a partner in the character pipeline rather than a replacement for artists [12]. The major takeaway is that AI is most useful when it supports ideation, variation, and customization while leaving control with designers and players.

Related preprints expand this idea into concrete tools. Spiritus combines natural language understanding, diffusion models, segmentation, and skeleton binding to produce 2D characters and animations from descriptions [13]. Sprite Sheet Diffusion addresses the manual burden of producing 2D sprite animations by using diffusion-based generation to automate frame creation and preserve style coherence [8]. Together, these works suggest that generative models can reduce repetitive production work, particularly in RPGs with many action states and costume variants.

A recent IEEE Access paper, GenePixKolor (GPK) Fusion, is also useful as an adjacent example because it combines genetic algorithms, image processing, and machine learning to optimize game-related digital asset generation and rarity ranking [1]. Although it focuses on NFT cards rather than RPG avatars, the paper shows how generative and optimization-based methods can be used to produce visually appealing assets while keeping ranking logic tied to both visual and trait-based criteria.

A Text-to-Game Engine for UGC-Based Role-Playing Games takes the idea even further by linking characters, narrative, visuals, and mechanics inside a text-driven pipeline [14]. This is important because RPG character design does not happen in isolation; the appearance of the character must match story logic, quest structure, and the interaction model of the game world. In this sense, character generation is becoming a multimodal system problem rather than a simple art-production problem.

3.6. Broader Game-AI Context

The survey literature on large language models in games helps explain why character design is moving toward dialogue, explanation, and procedural generation [12]. Although that work is not limited to visual design, it is relevant because many modern character tools now use language-based prompts to refine style, posture, background, and behavior. In practice, the character creator is becoming an interactive authoring environment.

Taken together, the reviewed studies show that RPG character design is converging with broader generative game research [14]. The boundary between concept art, customization interface, animation system, and narrative agent is thinning. The literature therefore points toward unified creative pipelines in which a player or designer can describe a character in words, refine the design conversationally, and receive both the static avatar and its motion assets as outputs.

3.7. Technical Pipeline from Concept Art to Runtime

A production-ready RPG character passes through several technical stages before the player sees it in-game [14]. Concept art establishes the silhouette, costume language, and emotional tone, but that concept must later be translated into topology, rigging, skinning, animation clips, and engine-readable assets. The reviewed literature implies that this chain is becoming more integrated.

In a conventional pipeline, the character model is built in a modeling or sculpting package, then rigged with bones and control handles, then weighted so the mesh deforms correctly, and finally connected to animation state machines and blend trees inside the engine. In RPGs, that pipeline must support multiple equipment states, combat poses, idle cycles, dialogue gestures, and facial expressions. A design that looks excellent in still art can fail if the rig collapses during crouch, if armor clips through the body, or if expressions become unreadable in close-up scenes.

The practical implication is that character design should be evaluated with downstream production constraints in mind [12]. Thick armor may simplify silhouette but reduce animation flexibility. Highly complex accessories may enrich personality but increase the risk of clipping and retargeting errors. For stylized RPGs, the pipeline must also preserve art direction under deformation, which makes style consistency not only an aesthetic issue but also an engineering one.

 doi:10.20944/preprints202604.0296.v1

Motion reuse is also strategically important. Walk cycles, idle cycles, combat recoveries, and interaction gestures can be shared across related character types if the rig is planned carefully. That makes modularity a key design goal. The more reusable the skeleton and animation library, the easier it becomes to support large casts of NPCs and playable classes without multiplying production cost.

3.8. Evaluation Metrics and Production Constraints

A recurring weakness in the literature is the absence of a single evaluation framework that captures the full quality of an RPG character [12]. Some studies focus on likeness or identity preservation, others on animation realism or emotional expressiveness, and others on usability or iterative editing. In practice, however, production teams need a multi-criteria evaluation that combines all of these concerns.

A more complete evaluation framework for RPG characters would include at least five dimensions: visual recognizability, identity consistency, motion plausibility, authoring efficiency, and player satisfaction. Production constraints also matter because RPG pipelines are usually built under tight asset budgets. A method that looks impressive in a research demo may still be difficult to deploy if it requires too much compute, too much manual cleanup, or too much special-case tuning. This is especially true for live-service or long-running RPGs, where teams must maintain consistency across patches, downloadable content, localization changes, and evolving art direction. Therefore, the most useful methods are not necessarily the most visually advanced ones; they are the ones that balance quality with predictability, scalability, and integration cost.

Table 1. Summary of the reviewed literature.

Study	Year	Domain	Technique/Method	Main Contribution	Limitation	RPG Relevance
Guruprakash et al.	2025	NFT / Game assets	Genetic algorithms + image processing + ML	Optimizes NFT asset generation and hybrid rarity ranking	Focused on NFT ecosystem, not full RPG pipelines	Relevant as adjacent AI-driven asset generation approach
Voimala	2023	Design workflow	Narrative literature review	Structures character creation into phases	Less emphasis on runtime animation	Useful for pipeline planning
Tychsen et al.	2008	Avatar play	Conceptual analysis	Explains avatar as identity medium	Predates modern AI tools	Foundational for RPG identity
Nasution & Rahmi	2024	Shape language	Case study	Links shape to personality reading	Focused on 2D design	Strong for silhouette design
Shi et al.	2019	Auto-creation	Optimization + imitation network	Maps face images to game parameters	Single-shot and constrained	Relevant to player avatars

 doi:10.20944/preprints202604.0296.v1

Shi et al.	2020	Auto-creation	Differentiable renderer + GAN	Improves identity/expression disentanglement	Parameter-space limitations remain	Useful for controlled customization
Wu et al.	2024	Interactive editing	Dialogue + parameter solver	Supports iterative refinement	Depends on prompt interpretation quality	Important for user-driven design
Hsieh et al.	2024	Sprite animation	Diffusion model	Reduces manual 2D animation labor	Style consistency remains difficult	Useful for 2D RPGs
Nakano & Hoshino	2006	Motion creation	Behavior synthesis	Early online motion framework	Older toolchain and weaker realism	Foundational motion work
Gallotta et al.	2024	Game-AI context	Survey and roadmap	Shows LLM roles in games	Broad rather than RPG-specific	Supports language-based tools
Wu et al.	2025	Generative AI review	Systematic review	Synthesizes applications and challenges	Survey-level, not a single method	Useful for state-of-the-art context
Sun et al.	2025	2D character generation	Text + diffusion + skeleton binding	Creates 2D characters and animations from text	2D-centric and tool-dependent	Strong for stylized RPG assets
Zhang et al.	2024	UGC RPG pipeline	Text-to-game engine	Connects story, visuals, and mechanics	Still emerging as a production pipeline	Highly relevant to RPG creation
Chung et al.	2025	Facial animation	Audio-driven realistic facial animation	Preserves lip sync, emotion, identity	Needs adaptation for stylized characters	Useful for cutscene realism

4. Discussion

 doi:10.20944/preprints202604.0296.v1

The literature suggests that RPG character design has moved from an art-only problem to a systems problem [12]. In earlier work, the central concern was how to represent the player, express role identity, and produce motion that matches the social and narrative context of the game. More recent studies broaden this by using generative models, dialogue systems, and differentiable rendering to compress the design process and increase customization. The design challenge is no longer only whether a character looks attractive; it is whether the system can preserve identity while supporting iterative editing, expressive motion, and engine constraints.

A recurring pattern is the tension between control and automation [12]. Manual workflows give artists strong creative control, but they are slow and difficult to scale. AI-assisted methods improve speed and accessibility, but they can reduce predictability unless the system offers fine-grained control. For RPG development, this trade-off is especially important because character assets must remain coherent across many states: facial close-ups, battle scenes, inventory portraits, dialogue windows, and gameplay animation. A useful pipeline must therefore support both high-level intent and low-level correction.

The same broader shift is visible in adjacent game-asset work. GenePixKolor (GPK) Fusion [1] shows that evolutionary algorithms, image processing, and machine learning can be combined to generate and rank digital assets by considering both trait-based and pixel-based signals. Although the domain is NFT cards rather than RPG avatars, the underlying idea is relevant: AI systems become more useful when they optimize assets while still preserving controllable visual criteria and ranking logic.

Another strong theme is semantic consistency [5]. The reviewed studies show that players want the avatar to feel like the same person across appearance, motion, and dialogue. That means face synthesis, costume design, emotion modeling, and animation timing should not be treated as separate modules that only meet at the end of production. Instead, the papers point toward a unified character model in which identity is encoded once and then retargeted across visual and motion outputs.

The literature also reveals an opportunity for human-in-the-loop workflows [7]. Dialogue-based editing, prompt-based generation, and assistant-guided refinement are appealing because they reduce the technical burden on users while preserving creative agency. For RPGs, this is a particularly good fit because the genre already invites players to co-author the character. The player is not merely consuming a preset asset; the player is choosing how the avatar should look, sound, and behave.

5. Research Gaps and Future Directions

5.1. Persisting Technical Gaps

First, there is no widely accepted benchmark that evaluates RPG character systems across both visual and behavioral dimensions [12]. Current studies may measure likeness, realism, editability, or animation quality, but they rarely combine these criteria in one framework. As a result, it remains difficult to compare character pipelines in a way that reflects actual RPG production requirements.

Second, semantic control remains weak in many generative pipelines [7]. While dialogue-based and text-based tools are promising, they still need better ways to translate abstract descriptions such as “more heroic,” “older but agile,” or “battle-hardened healer” into consistent visual and motion changes. This gap is especially important for RPGs because character meaning is often expressed through subtle combinations of costume, posture, and facial behavior.

Third, animation systems are still insufficiently integrated with character customization [14]. A body type or facial structure chosen during creation can affect the feasible range of movement, expression, and cloth behavior later in the game. Yet the literature often treats generation and animation as separate topics. Better research is needed on end-to-end systems where rigging, mesh deformation, and motion retargeting are co-designed.

Fourth, stylized RPGs remain underrepresented in many generative methods [12]. Some methods are oriented toward photorealistic faces or realistic avatar motion, but RPGs often use cel-

 doi:10.20944/preprints202604.0296.v1

shaded, low-poly, anime-inspired, or hybrid visual languages. Systems that work well for realistic digital humans may not transfer cleanly to these styles.

Fifth, there is limited work on long-term character consistency [12]. RPG characters frequently evolve through upgrades, narrative branching, and repeated interaction, yet many AI generation papers focus on one-off outputs. Future research should study how to maintain identity over time, across scenes, and across multiple creative sessions.

5.2. Future Research Directions

The most promising direction is a unified character pipeline that links concept generation, customization, rigging, facial animation, and motion synthesis in one workflow. Such a system would allow an artist or player to define identity once and then propagate it consistently into appearance, expression, and action. This would reduce manual rework and make RPG character production more scalable.

A second direction is real-time adaptive animation. Instead of static motion clips only, future systems could adjust posture, expression, and gesture based on gameplay state, emotional context, or dialogue choices. That would make non-player characters and player avatars feel more responsive and alive, especially in narrative-heavy RPGs.

A third direction is multimodal co-creation [6]. Text prompts, sketches, reference images, and user feedback should be combined inside one interactive environment so that non-expert users can refine characters without losing control. This would be particularly useful in UGC-based RPG tools and mod-friendly pipelines.

A fourth direction is cross-style generation. RPGs are highly diverse in visual language, so future systems should support style transfer between realistic, stylized, anime, pixel-art, and cel-shaded character types while maintaining identity and motion coherence. This remains one of the hardest open problems because style, topology, and animation constraints interact strongly.

A fifth direction is benchmark development. The field needs evaluation datasets and metrics that measure identity preservation, animation smoothness, semantic controllability, and user satisfaction together. Without such benchmarks, it will remain difficult to know whether a new method genuinely improves production quality or only improves a narrow technical metric.

6. Conclusions

Character design and animation in RPG-based games are now best understood as a connected pipeline rather than isolated tasks. The reviewed literature shows a clear progression from foundational work on avatar identity, shape language, and motion creation to recent AI-assisted systems for face-based character generation, dialogue-based editing, sprite animation, expressive facial motion, and text-to-game pipelines. This progression reflects a broader shift in game development toward hybrid workflows that combine artistic intent with algorithmic support.

The main conclusion is that future RPG character systems should aim for unified control. A strong system should allow a designer or player to define identity, refine appearance conversationally, and preserve that identity across animation, emotion, and gameplay states. At the same time, the system must remain compatible with engine constraints, stylistic variation, and production scale. The most promising direction is therefore not full automation, but assisted co-creation.

The field gaps identified in this review point to three priorities for future work: robust evaluation metrics, stronger semantic control, and tighter integration between design and animation. Addressing these gaps would make it possible to build RPG characters that are not only visually appealing, but also emotionally readable, technically stable, and genuinely editable over time.

References

 doi:10.20944/preprints202604.0296.v1

1. J. Guruprakash, P. Pradeep, and L. B. Krithika, "GenePixKolor (GPK) Fusion: A Novel Evolutionary Algorithm-Based Optimized NFT Card Generation and Rarity Ranking Method for Gaming Tokenomics," IEEE Access, vol. 13, pp. 65566-65581, 2025, doi: 10.1109/ACCESS.2025.3558833.
2. J. Voimala, "Creating a Video Game Character Design Workflow: A Literature Review on the Theories and Methods Used by Professionals in Video Game Character Design," Bachelor's thesis, Jamk University of Applied Sciences, 2023.
3. A. Tychsen, M. Hitchens, and T. M. Brolund, "Character Play: The Use of Game Characters in Multi-player Role-Playing Games Across Platforms," Computers in Entertainment, vol. 6, no. 2, article 22, 2008.
4. U. K. Nasution and A. Rahmi, "The Analysis of Shape Language in Character Design: A Case Study of 2D Animation 'Terjebak Hoaks'," Proceedings of the 11th International Applied Business and Engineering Conference, 2024.
5. T. Shi, Y. Yuan, C. Fan, Z. Zou, Z. Shi, and Y. Liu, "Face-to-Parameter Translation for Game Character Auto-Creation," in Proc. IEEE/CVF ICCV, 2019.
6. T. Shi, Z. Zou, X. Song, Z. Song, C. Gu, C. Fan, and Y. Yuan, "Neutral Face Game Character Auto-Creation via PokerFace-GAN," in Proc. ACM Multimedia, 2020.
7. H. Wu, M. Zhao, Z. Hu, L. Li, W. Chen, R. Zhao, C. Fan, and X. Yu, "ICE: Interactive 3D Game Character Editing via Dialogue," arXiv preprint arXiv:2403.12667, 2024.
8. C.-A. Hsieh, J. Zhang, and A. Yan, "Sprite Sheet Diffusion: Generate Game Character for Animation," arXiv preprint arXiv:2412.03685, 2024.
9. A. Nakano and J. Hoshino, "Motion Creation for Online Game Characters," in Advances in Computer Entertainment Technology (ACE), 2006.
10. C. Ennis, L. Hoyet, A. Egges, and R. McDonnell, "Emotion Capture: Emotionally Expressive Characters for Games," Proceedings of Motion on Games, 2013.
11. R. Gallotta, G. Todd, M. Zammit, S. Earle, A. Liapis, J. Togelius, and G. N. Yannakakis, "Large Language Models and Games: A Survey and Roadmap," arXiv preprint arXiv:2402.18659, 2024.
12. Z. Wu, Z. Chen, D. Zhu, C. Mousas, and D. Kao, "A Systematic Review of Generative AI on Game Character Creation: Applications, Challenges, and Future Trends," IEEE Transactions on Games, 2025.
13. Q. Sun, Y. Ni, T. Yuan, J. Zhang, F. Yang, Z. Yao, and H. Mi, "Spiritus: An AI-Assisted Tool for Creating 2D Characters and Animations," arXiv preprint arXiv:2503.09127, 2025.
14. L. Zhang, X. Peng, S. Yang, and F. Wang, "A Text-to-Game Engine for UGC-Based Role-Playing Games," arXiv preprint arXiv:2407.08195, 2024.
15. C. Chung et al., "Audio2Face-3D: Audio-driven Realistic Facial Animation For Digital Avatars," arXiv preprint arXiv:2508.16401, 2025.