BraintankAI ASIA

DISPATCH · 3D HAIR GENERATION · 3 MIN

BEIJING · Tsinghua University

Diffusion model builds animatable stylised game hair from a single picture

HairDeck treats each ribbon-like hair card as one token and generates them all at once, then adds the bones that make the hair sway, a preprint reports.

10 November 2026 · arXiv preprint, not yet peer reviewed

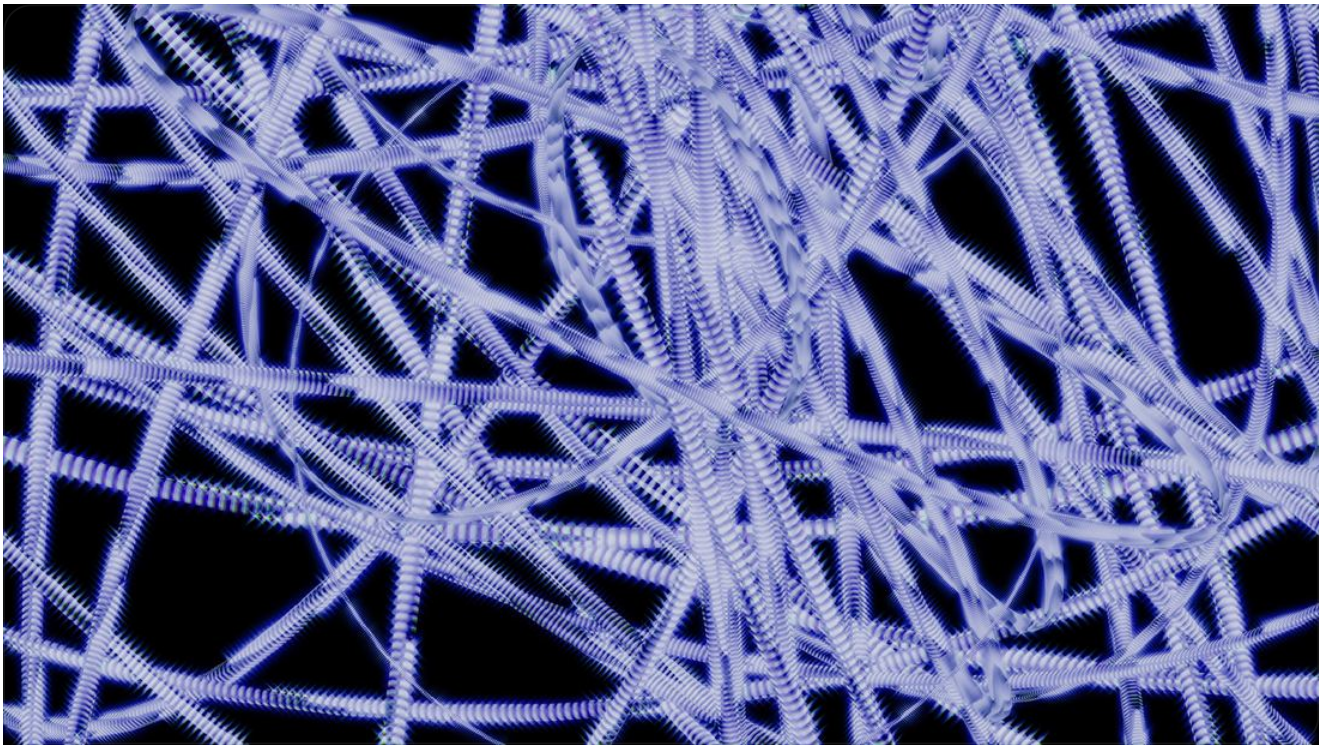


Photo: Logan Voss

In games and virtual avatars with an anime or cartoon look, hair is rarely made of individual strands. It is built from hair cards: a few dozen to more than a hundred ribbon-like 3D strips with rhombic or triangular cross-sections, each placed by an artist so that its shape is the hair the viewer sees. To make it move, artists also rig spring bones, short chains of joints hanging from the head that a game engine simulates every frame so the hair sways and settles. Both jobs are done by hand. A 2026 preprint from researchers at Tsinghua University, Carnegie Mellon University and the Chinese University of Hong Kong, which has not yet been peer reviewed, describes HairDeck, a system that produces the cards and their bone chains from a single image or a 3D point cloud.

The only earlier generator of stylised cards, CHARM, works autoregressively, producing output one piece at a time as a language model writes text, and spells out a hairstyle as hundreds to thousands of tokens describing points along each card. HairDeck's authors start from the observation that

cards have no natural order. They encode each card as a single token and generate the whole set at once.

Each card is first compressed by a variational autoencoder, a network that squeezes an object into a short vector and can decode it back, into 32 numbers. The token also carries the card's root position, its cross-section type and a root normal, the direction along which the ribbon's width runs where it leaves the scalp. A hairstyle occupies 192 slots, each with an existence flag whose sign says whether it holds a real card, so the model settles the card count itself. A diffusion transformer of 181.3 million parameters, trained to turn random noise step by step into data, refines all slots together; with no positional information, it treats them as an unordered set. A second, smaller transformer generates up to 48 spring-bone chains, and a fixed geometric rule attaches cards to them.

Training used 36,440 hairstyles from CHARM's dataset, which is built from characters made in VRoid, an anime character-creation tool. On CHARM's 100 test hairstyles, the authors report three improvements. Under CHARM's own protocol, Chamfer distance, the average gap between points on the generated and true shapes, fell from 0.0117 to 0.0066. Card-orientation error, measured against the artists' original meshes, fell from 40.0° to 16.8°. Generation time on a single GPU fell from 33.7 seconds to 1.01 seconds. They also report that simply retrieving the nearest training hairstyle scored 0.0119, almost level with CHARM. An autoregressive model using HairDeck's own one-per-card tokens reached 0.0078, short of HairDeck's 0.0066, and was marginally better on orientation. It tended to stop early, producing 60.8 cards on average against 68.9 for HairDeck and 67.4 in the artist originals. A 54.6-million-parameter version of HairDeck, smaller than CHARM, still reached 0.0072, far below CHARM's 0.0117.

From single images, HairDeck was compared with CHARM, the general image-to-3D generators TRELLIS.2 and Hunyuan3D-2.1, and the strand generator DiffLocks with its output converted to cards. On 40 anime illustrations from outside the dataset, it failed on 10.0% of inputs against 17.5% to 77.5% for the others, and took 1.33 seconds against 15.6 to 83.0 seconds.

The authors list four limitations. HairDeck learnt only from stylised VRoid hairstyles, so very short hair and the textured cards used for realistic hair lie outside its training. It produces geometry but no textures. Some characters model bows and hats as cards, so HairDeck generates them as cards when an input shows them. It also generates too few chains for unusual bone layouts. Chains built on artist cards matched artists' own chains better than retrieval or a procedural rig, but quality drops on cards generated from images, with a recall score of 0.625 falling to 0.390. The benchmark test sets are small and come from a single dataset, and no variance across runs is reported. The illustrations were judged only by failure rate, and the image-to-3D baselines were never designed to output cards.

The authors present the approach as a general rule for assets assembled from artist-defined parts, one token per part with no order among them, and suggest garments sewn from panels and drawings made of strokes as further candidates.

PAPER	*HairDeck: Animation-Ready Stylized Hair Cards from a Single Image*
AUTHORS	Yuze He, Yixiang Wei, Wang Zhao, Chuhao Chen, Yanning Zhou, Yong-Jin Liu et al.
VENUE	arXiv preprint — not yet peer reviewed · 29 October 2026
LINK	arxiv.org/abs/2610.40599
INDEX	stacks.braintank.ai
CODE	github.com/hyz317/HairDeck