Hertz

Question 1: Choice of Field and Future Expectations How did you choose your field and what are your primary expectations of your future career? If you are currently in your second or later year of graduate school, you should make your case here for why receiving a Fellowship would result in exceptional leverage in the kind and quality of your graduate work, including your ability to pursue promising new ideas. Please understand before continuing this application, that such a case will have to be very strong to be considered further, and that new Hertz Fellowships are very rarely granted to students currently in their second year of graduate study or beyond. Question 1 Essay - How did you choose your field and what are your primary expectations of your future career? (300 words) My interest in computer vision and graphics began with a simple curiosity—how can machines see and understand the world as humans do? This question led me to explore how visual data could be processed, interpreted, and generated. Over time, my curiosity evolved into a deeper research goal: to build AI systems that can understand complex visual distributions and dynamic processes through a data-centric approach.

I began my research journey at Tongji University during the pandemic, working on foundational vision problems such as feature matching and image super-resolution. Those early projects taught me the importance of data quality and structure in determining model performance. Later, at USC's iLab, I expanded into generative AI, creating synthetic datasets for detection and segmentation tasks. That experience convinced me that high-quality, well-designed data can be as powerful as model innovation itself.

My Master's research and internship at Microsoft Research Asia further connected AI with real-world impact, as I developed large-scale datasets for radiology report generation. At Stanford University, I extended these ideas to 3D and 4D motion understanding—transforming videos into dynamic reconstructions of animals in motion. This project reinforced my belief that data is the foundation of meaningful AI progress.

Now as a PhD student at UIUC, I am exploring generative models for scientific data such as eye-tracking, where privacy and scarcity are major challenges. Looking ahead, I aspire to work at a world-leading research institute or company, collaborating with scientists and engineers across disciplines to develop data-driven AI systems that advance science and benefit society. My ultimate goal is to make AI not only more capable, but also more accessible and equitable for people around the world.

Question 2: Proposed Field of Study Question 2 Essay - How does your proposed field of study and career constitute an application of the physical sciences or biological sciences or engineering? (300 words)

My research lies at the intersection of computer vision, graphics, and generative modeling—fields grounded in both the physical sciences and engineering. I aim to develop AI systems that can perceive, reconstruct, and simulate the physical world from data. This involves modeling the dynamics of real-world objects, understanding their geometry, motion, and interactions, and embedding physical reasoning into data-driven algorithms.

Many problems I study—such as inverse rendering, 3D and 4D reconstruction, and motion generation—are, at their core, questions about how to infer and represent physical reality. For instance, reconstructing an animal's motion from 2D video requires understanding projection geometry, light transport, deformation, and physical constraints of articulated bodies—all classical problems of physics and engineering reframed through the lens of machine learning. Similarly, generative models that synthesize realistic movements must capture the statistical and physical properties of motion distributions to ensure stability, continuity, and plausibility over time.

My approach combines data-centric engineering—building large-scale, structured datasets that reflect physical processes—with physics-integrated AI modeling, enabling models to learn both from examples and from the underlying mechanics of the world. The same principles that govern optical imaging, motion dynamics, and material behavior underpin the algorithms I design.

Ultimately, I hope to advance a new generation of AI systems that not only interpret data but also reason about the physics of the real world. Such systems could transform fields ranging from robotics and biomedical imaging to environmental monitoring and scientific simulation—bridging artificial intelligence with the core principles of the physical sciences.

Question 3: Choice of Graduate School Question 3 Essay - What are the considerations involved in your choice of graduate school? (300 words)

I chose to pursue my Ph.D. at the University of Illinois Urbana-Champaign (UIUC) because it offers one of the strongest and most supportive environments for research in computer science, especially in computer vision, graphics, and machine learning. My decision was guided by a combination of academic excellence, collaborative culture, and the research-focused environment that UIUC uniquely provides.

UIUC's Department of Computer Science consistently ranks among the top programs worldwide and is home to many pioneering researchers in artificial intelligence. I was drawn to the opportunity to work with Professor James Rehg, whose work in human-centered AI and visual understanding aligns closely with my interest in data-centric modeling and physical reasoning in vision. The breadth of expertise across the department—ranging from machine learning theory to large-scale systems—creates a rich environment for cross-disciplinary collaboration and innovation.

Beyond academic prestige, I was equally influenced by UIUC's community and culture. Although it is a large institution, students describe feeling deeply supported due to the accessibility of faculty, abundant resources, and a strong sense of academic community. Frequent seminars, guest lectures, and research symposia create constant opportunities to learn from and engage with leaders across computing and engineering.

The setting of Urbana-Champaign also played a key role in my decision. As a close-knit college town, it offers an ideal environment for focused research and meaningful connection among students and faculty. The collaborative and grounded atmosphere encourages both intellectual and personal growth.

Ultimately, I chose UIUC because it combines world-class research strength with a collegial and inclusive environment—one that will help me grow as an independent researcher and contribute to the broader scientific community.

Question 4: Chronological Synopsis Provide a concise summary, similar to a resume, in chronological order, with dates, recapitulating significant periods of technical and other creative activity since high school graduation. Omit activities only distantly related to your professional development. Include workshops, summer schools, a general description of all courses of study pursued (e.g. "3 quarters of Differential Equations") and degrees expected or awarded (dates, institutions, fields). Separate your undergraduate activities from your graduate activities (if/as applicable) with a single dashed line. Question 4 Essay - Chronological Synopsis (300 words)

Undergraduate Studies: University of Southern California (USC), BS in Computer Engineering and Computer Science, BA in Applied and Computational Mathematics, 08/2018-12/2022

Courses: 2 semesters of Calculus and Linear Algebra, 2 semesters of Probability and Statistics, 2 semesters of Discrete Math, 2 semesters of Numeric Methods and Modeling, 2 semesters of Algorithms, 4 semesters of Programming and Software Development, 4 semesters of AI and Machine Learning, 9 semesters of Computer/Embedded Systems and Architecture, 3 semesters of Physics, 3 semesters of Astronomy

Research: Tongji University, Visiting Student Researcher, 06/2020-12/2022, USC iLab, Student Researcher, 04/2022-05/2024, USC Information Sciences Institute, Research Intern, 06/2022-03/2023, Microsoft Research Asia, Research Intern, 05/2023-02/2024

Graduate Studies: University of Southern California (USC), MS in Computer Science 01/2023-05/2024, University of Illinois Urbana-Champaign (UIUC) PhD in Computer Science, 08/2025-Present,

Courses: 1 semester of Algorithm, 5 semesters of AI and Machine Learning,

Research: Stanford Vision and Learning Lab, Research Assistant, 05/2024-04/2025

Honors: UIUC Summer Predoctoral Emerging Scholar

Personal Essay Use this page to submit a personal essay that reveals your personal creativity beyond your research. You may choose to focus on activities since your graduation from high school, but these do not have to be related to the classroom or lab, to STEM or your future career—they can also be from your more informal pursuits and passions, e.g., your extracurricular and leisure activities. This is an opportunity for you to share your narrative with experiences that demonstrate your creativity, perseverance, curiosity and/or grit. You

may use more than the allotted space for your response, and any additional text will be included in your application.

Growing up, my ears were filled with two different rhythms: the intricate wordplay of American hip-hop and the sharp, tonal percussion of Chinese rap. They were the soundtracks to my two worlds, and from a young age, I was driven by a core question: how could music built on such different linguistic architecture still evoke the same fundamental emotions?

My connection to music began on stage. For four years in high school, I was immersed in our musical theater team and a cappella club, playing roles from King Oberon in A Midsummer Night's Dream to the lead in an original Chinese musical. In a cappella, I started as the bass but quickly found my voice as the lead rap and hip-hop vocal. What I loved most wasn't the spotlight, but the synthesis—the intricate challenge of weaving together different languages, sounds, and stories into a single, coherent piece of art.

In college, though my performance time dwindled, my curiosity only deepened. I pursued a minor in music recording not as a diversion, but to gain a new set of tools. Learning to mix, master, and produce sound gave me a technician's control. Suddenly, I could move from the stage to the architecture of the song itself, experimenting with vocal layering and sound design. For the first time, I could deconstruct the music I loved to its core components.

Behind the mixing board, I began to see the raw mechanics of rhyme. I realized English hip-hop builds its flow on fluid, multi-syllable combinations, while Mandarin rap functions more like a drum machine—each monosyllabic character, with its inherent tone, lands with a distinct percussive force. I saw how Chinese artists masterfully blend modern beats with traditional pentatonic scales, creating a sound that feels simultaneously ancient and futuristic. This fusion resonated with me; it sounded like my own experience.

This insight unlocked a new creative obsession: "translating" songs not literally, but culturally. To capture the essence of an American song about resilience, I wouldn't search for direct translations; I would rebuild it using classic Chinese metaphors—a single plum blossom blooming in the snow, or bamboo bending but never breaking in a storm. These symbols, deeply rooted in centuries of poetry, could carry the same emotional weight. The process was one of deconstruction and reconstruction: breaking a song down to its essential meaning, then reassembling it in a new framework that spoke to both halves of my identity.

This hobby has become a model for how I solve problems. It taught me that the most innovative ideas often come from connecting systems that don't seem to fit together. It taught me patience and the rigor of experimentation—of testing a rhyme or a beat dozens of times until it feels true in two different contexts. Most of all, it taught me to listen deeply—to sound, to meaning, and to the human experience that powers both.

Whether I am remixing a song or approaching a research problem, my method is the same: to act as a bridge between different worlds, find the universal patterns that connect them, and build something new in the space between. Somewhere between the rhythms of English and the tones of Mandarin, I learned how to blend structure with freedom and analysis with art. That is the creative space where I feel most at home.

In the space provided below, list, in chronological order, academic honors and distinctions (e.g., Phi Beta Kappa, Tau Beta Pi, or Sigma Xi; class rank recognition, etc.) which you have received and the time or time-interval of receipt. Separate your undergraduate from your graduate awards (if/as applicable) with a single dashed line (Include title, reason for award, and where/when received). Especially significant awards received in high school can also be included. Use no more than one line per award whenever possible (what, where/when received). You may use more than the allotted space for your response, and any additional text will be included in your application.

Undergraduate:

USC Academic Achievement Award, partial tuition award recognizing academic excellence and pursuit of multiple degree objectives, 07/2020 at USC, USC Associates Senior Scholar, partial tuition award honoring top academic performance among senior undergraduates, 09/2021 at USC, USC Steven and Kathryn Sample Renaissance Scholar Distinction, monetary award recognizing academic excellence across two or more distinct fields of study, 05/2022 at USC Dean's List, USC Viterbi School of Engineering, recognizing academic excellence, 2019-2023 Dean's List, USC Dornsife School of Letters, Arts and Sciences, recognizing academic excellence, 2019-2023

Graduate: UIUC Summer Predoctoral Emerging Scholar, monetary stipend recognizing academic excellence, leadership potential, and commitment to diversity, inclusion, and enriching the university community, 07/2025

Compelling example(s) of personal creativity are a very important factor in our selection process. We are more interested in what makes your personal contributions stand out and the creativity you employed than what your research group did as a whole, unless you contributed centrally to its leadership. *In addition to responding to the items below, please also submit copies of your most significant scientific publications/reports on the Supplemental & Other Information page of this application. Project 1: First, provide a brief overview of your research project for a non-specialist, technical audience. Second, describe YOUR direct contribution to this research that exemplifies your creativity and independent thinking.

DreamDistribution: Learning Prompt Distribution for Diverse In-distribution Generation

In machine learning and computer vision, images are often represented as latent vectors—compact encodings of visual information. To describe specific visual concepts that are difficult to capture in natural language (like a particular person or dog), a technique called textual inversion learns to embed that concept as a vector in the text space of a model. For instance, given photos of your pet dog, textual inversion can learn a token representing it, allowing prompts like "<dog> on a spaceship" to generate new images of your pet.

However, textual inversion struggles when the goal is to learn a broader concept—such as an artistic style, an object category, or a visual theme represented by a set of images. These concepts have natural variation, which cannot be captured by a single point in the text space.

In my project DreamDistribution, I proposed a simple yet effective solution: instead of learning a single embedding, we learn a distribution in the text space that models the variation within the image set. This approach allows generative models to produce diverse yet coherent outputs, enabling creative control and variety while remaining faithful to the intended concept. This methods further enabled new ability such as varying learned variance to control generation diversity, and interpolation between two distributions to generate mix-of-concept images.

As the project lead and first author, I initiated the idea, developed the theoretical framework, implemented all code, and conducted most of the experiments. The creativity lies in rethinking where and how diversity is represented in a generative model, not in the image latent space as others have done, but in the text embedding space, where such modeling had not been explored before. This conceptual shift opens a new direction for controllable, expressive image generation.

Project 2: (optional) First, provide a brief overview of your research project for a non-specialist, technical audience. Second, describe YOUR direct contribution to this research that exemplifies your creativity and independent thinking.

Compared to humans, animal motion study remains far behind. Human motion can be easily captured with motion capture systems—through sensors or visual markers—enabling accurate digital representation. For animals, however, this approach is often impossible or unethical. While domesticated species like cats and dogs can sometimes be recorded in controlled environments, most wildlife, such as elephants or bison, cannot. Even for tamed animals, artificial environments can disrupt natural behaviors, making traditional motion capture unsuitable.

I approached this long-standing problem from a new direction. Recent advances in image-to-3D reconstruction now allow accurate estimation of 3D shapes and poses from ordinary images and videos. This inspired me to explore whether we could reconstruct animal motion directly from in-the-wild videos—an ethical, scalable, and noninvasive alternative to motion capture.

To make this feasible, I designed and implemented a fully automatic data processing pipeline that collects and cleans large-scale animal motion data from online videos. The system filters out clips with occlusion, shot transitions, low resolution, or static motion, resulting in clean, animal-centric motion clips. Within a few weeks, this pipeline produced a dataset of nearly 30,000 high-quality videos, enabling large-scale 3D motion reconstruction.

This work was published at the Conference on Neural Information Processing Systems (NeurIPS 2025). I

am currently extending it to reconstruct 3D motion and train generative models for realistic animal motion synthesis.

My creative contribution lies in rethinking how motion data can be collected: transforming raw, unstructured internet videos into a powerful, ethical, and data-driven foundation for animal motion research—bridging the gap between human and animal motion understanding.