

# Kyle Berkson

California Institute of Technology | Pasadena, CA | [kberkson@caltech.edu](mailto:kberkson@caltech.edu) | [linkedin.com/in/kyle-berkson](https://www.linkedin.com/in/kyle-berkson)

## EDUCATION

<b>California Institute of Technology</b>	Pasadena, CA
<i>Bachelor of Science in Applied Computational Math, Minor in CS</i>	<b>Graduating June 2028</b>
<b>Clubs/Activities:</b> Division III NCAA Track and Field and Cross Country	<b>GPA: 4.2</b>
<b>Relevant Coursework:</b> Data Structures and Algorithms, Analytical Linear Algebra, Computational Methods, Software Design.	

## WORK EXPERIENCE

<b>Caltech Summer Undergraduate Research Fellowship</b>	Pasadena, CA
<i>Undergraduate Researcher</i>	<b>Jun 2025 - Present</b>
<ul style="list-style-type: none"><li>Developed a novel Bayesian uncertainty algorithm for NeRF-based Event Horizon Telescope tomography in JAX/Flax, giving per-voxel confidence maps for black-hole reconstructions in both 2D and 3D scenarios.</li><li>Delivered ensemble-level uncertainty calibration with 100x less GPU time, enabling quick analysis of low-confidence geometry in reconstructions. understand geometric uncertainty in their reconstructions.</li></ul>	
<b>Nexa</b>	Pasadena, CA
<i>Co-Founder</i>	<b>November 2024 - Present</b>
<ul style="list-style-type: none"><li>Secured \$600k+ in pre-seed funding, got accepted into YCombinator (S25) to build a state of the art microchip and software platform for livestock monitoring, acquiring \$400k in LOIs from target cattle operations.</li><li>Led a 4-engineer software team: owned architecture and roadmap, ran weekly sprints, set coding/testing standards, and mentored others, shipping MVP in 3 weeks.</li><li>Built real-time disease &amp; estrus detection and classification pipelines blending signal processing with time-series ML, outperforming research baselines by 15% recall + 10% AUROC on internal validation and reduced false alerts by 26%.</li><li>Designed an end-to-end evaluation and testing suite: synthetic data generators, backtests, ablations, shadow-mode trials, and CI with unit/integration/regression tests; added data-drift checks and load tests that simulate bursty device traffic.</li><li>Architected an end-to-end platform—Django + Celery on PostgreSQL/Redis with versioned REST APIs powering a React/Tailwind UI—ingesting 100k+ readings/day and delivering herd/cow analytics with latency &lt; 2 min.</li></ul>	
<b>Shibolet</b>	Pasadena, CA
<i>Machine Learning Engineer</i>	<b>February 2025 - Jun 2025</b>
<ul style="list-style-type: none"><li>Cut classification spend by 82% and latency by 200% by leveraging NLP techniques to build an asynchronous FinTech-to-Bank taxonomy mapper that auto-labels thousands of complaints per second, raising flags for mislabels or ambiguous cases to increase accuracy by 10%.</li><li>Developed a RAG library, improving precision@5 by 12% through cross-encoder reranking and configurable cleaning/filtering.</li><li>Wrote a dual-output preprocessing layer that strips PII, generating LLM-ready JSON and normalized token streams for embeddings, enabling transparent audits and downstream analytics.</li></ul>	

## SKILLS

<b>Programming Languages:</b> Python, C, Java, JavaScript, HTML/CSS	<b>Hardware:</b> Microcontrollers, I2C communication, Control Theory
<b>Frameworks:</b> Django, Celery, React, Tailwind	<b>Infra/Tools:</b> PostgreSQL, Git, Linux, Docker, AWS, REST APIs, Redis
<b>ML/Data:</b> Machine Learning Algorithms, JAX/Flax, Pytorch, LLMs and RAG, Tensorflow, Uncertainty Quantification, Signal Processing.	<b>Soft Skills:</b> Technical writing, Research collaboratio