Raymond A. Yeh

raymond-yeh.com

Education

University of Illinois at Urbana-Champaign   **Urbana, IL**
Doctor of Philosophy in Electrical Engineering   *2016–2021*
Advisors: Alexander G. Schwing and Minh N. Do

University of Illinois at Urbana-Champaign   **Urbana, IL**
Master of Science in Electrical Engineering   *2014–2016*
Advisor: Minh N. Do

University of Illinois at Urbana-Champaign   **Urbana, IL**
Bachelor of Science in Electrical Engineering   *2010–2014*
Advisor: Mark Hasegawa-Johnson

Professional Employment

Purdue University   **West Lafayette, IN**
Assistant Professor in the Department of Computer Science   *2022-present*

Toyota Technological Institute at Chicago   **Chicago, IL**
Research Assistant Professor   *2021-2022*

Google   **Mountain View, CA**
Research Intern   *Summer 2018,2019*

Google   **Seattle, WA & Mountain View, CA**
Software Engineering Intern   *Summer 2015,2016,2017*

Johns Hopkins University, Human Language Technology Center of Excellence   **Baltimore, MD**
Research Intern   *Summer 2013, 2014*

Awards & Recognition

Best Paper (Runner up) at CVPR Workshop on AI for Content Creation   2024
○ Awarded to Towards Safer AI Content Creation by Immunizing Text-to-image Models.

The Seed for Success Acorn Award   2024
○ Given to investigators in recognition of their accomplishment in obtaining an external sponsored award of $1 million or more.

UIUC Graduate College Dissertation Completion Fellowship (Declined)   2020
○ Awarded to help outstanding students complete the doctoral degree.

Robert T. Chien Memorial Award   2020
○ Awarded to a doctoral graduate student who has demonstrated excellence in research.

NeurIPS High-Scoring Reviewers   2020
○ Awarded to top 10% reviewers based on the quality of reviews as rated by the Area Chairs.

ECCV Outstanding Reviewers   *2020*
- Awarded to top 215 reviewers based on the number of papers they reviewed and based on the quality as rated by the Area Chairs.

Mavis Future Faculty Fellowship   *2019*
- Awarded to facilitate the training of the next generation of great engineering professors.

Google Ph.D. Fellowship (Machine Perception)   *2018*
- Awarded to recognize and support outstanding graduate students doing exceptional research in Computer Science and related disciplines.

CVPR Outstanding Reviewers   *2018*
- Awarded to reviewers who contributed at least two reviews noted as excellent by area chairs.

James M. Henderson Fellowship   *2015*
- Awarded to recognize an outstanding first-year graduate student in the Department of Electrical and Computer Engineering.

Teachers Ranked as Excellent   *2014*
- Awarded to instructors and teaching assistants with outstanding teaching ratings based on student reviews.

Henry Ford II Scholar Award   *2014*
- Awarded to an undergraduate student with a 4.0 grade point average and intends to pursue graduate studies.

Graduation Highest Honors   *2014*
- Awarded to any student with at least a 3.8 grade-point average at graduation (A= 4.0) and completed a special project of superior quality.

Bronze Tablet Award   *2014*
- Awarded to students ranked in the top 3% of the students in their graduating class.

Henry O. Koehler Merit Scholarship   *2014*
- Awarded to juniors or seniors with outstanding scholastic records.

Publications

Google Scholar. Total Citation: 4551; h-index: 18; i10-index: 28

G : Graduate student, U : Undergraduate student.

Refereed Conference Publications

[C1] W. Chen^G, **R. A. Yeh***, S. Mou, and Y. Gu*. Leveraging perturbation robustness to enhance out-of-distribution detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. (8 pages, acceptance rate 22.1%).

[C2] A.Y Zheng^{G}and **R. A. Yeh**. Multi-concept model immunization through differentiable model merging. In *The AAAI Conference on Artificial Intelligence (AAAI)*, 2025. (7 pages, acceptance rate 23.4%).

[C3] H. Zhang^G, C. Y. Yang^G, and **R. A. Yeh**. Multi-object 3D grounding with dynamic proposals and language informed spatial attention. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. (10 pages, acceptance rate 25.8%).

[C4] M. A. Rahman[G], R. J. George, M. Elleithy, D. Leibovici, Z. Li, B. Bonev, C. White, J. Berner, **R. A. Yeh**, J. Kossaifi, K. Azizzadenesheli, and A. Anandkumar. Pretraining codomain attention neural operators for solving multiphysics PDEs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. (10 pages, acceptance rate 25.8%).

[C5] A.Y Zheng[G]and **R. A. Yeh**. IMMA: immunizing text-to-image models against malicious adaptation. In *European Conference on Computer Vision (ECCV)*, 2024. (8 pages, acceptance rate 27.9%) **Best paper runner-up and oral presentation at CVPR workshop for AI for Content Creation. Selection rate: 5.26%.**

[C6] A.Y Zheng*[G], C. Y. Yang*[G], and **R. A. Yeh**. Learning to obstruct few-shot image classification over restricted classes. In *European Conference on Computer Vision (ECCV)*, 2024. (8 pages, acceptance rate 27.9%).

[C7] C. Y. Yang[G], Z. Liu, and **R. A. Yeh**. Deep nets with subsampling layers unwittingly discard useful activations at test-time. In *European Conference on Computer Vision (ECCV)*, 2024. (8 pages, acceptance rate 27.9%).

[C8] J. Lee[G], B. Li[G], S. Beery, J. Huang, S. Fe, **R. A. Yeh**, and B. Benes. Tree-d fusion: Simulation-ready tree dataset from single images with diffusion priors. In *European Conference on Computer Vision (ECCV)*, 2024. (8 pages, acceptance rate 27.9%).

[C9] R. A. Rojas-Gomez[G], T. Y. Lim, M. N. Do, and **R. A. Yeh**. Making vision transformers truly shift-equivariant. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. (8 pages, acceptance rate 23.6%).

[C10] J. Ahn*[U], H. Wang*[G], **R. A. Yeh**, and G. Shakhnarovich. Alpha invariance: On inverse scaling between distance and volume density in a neural radiance field. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. (8 pages, acceptance rate 23.6%).

[C11] E. Kim[G], T. Bui[G], J. Yuan[G], S. C. Mouliand, B. Ribeiro, **R. A. Yeh**, M. P. Fassnacht, and M. B.G. Jun. Online real-time machining chatter sound detection using convolutional neural network by adopting expert knowledge. In *SME North American Manufacturing Research Conference (NAMRCs)*, 2024.

[C12] M. A. Rahman[G] and **R. A. Yeh**. Truly scale-equivariant deep nets with fourier layers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. (10 pages, acceptance rate 26.1%).

[C13] Y.-T. Hu, A. G. Schwing, and **R. A. Yeh**. Surface snapping optimization layer for single image object shape reconstruction. In *International Conference on Machine Learning (ICML)*, 2023. (8 pages, acceptance rate 27.9%).

[C14] H. Wang*[G], X. Du*[G], J. Li*[G], **R. A. Yeh**, and G. Shakhnarovich. Score jacobian chaining: Lifting pretrained 2D diffusion models for 3D generation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. (8 pages, acceptance rate 25.8%).

[C15] A. Firoze[G], C. Wingren[G], **R. A. Yeh**, B. Benes, and D. Aliaga. Tree instance segmentation using temporal structured images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. (8 pages, acceptance rate 25.8%).

[C16] C.-A. Yang[G], M.-L. Wu, **R. A. Yeh**, and Y.-C. F. Wang. Consistent and multi-scale scene graph transformer for semantic-guided image outpainting. In *The IEEE International Conference on*

Image Processing (*ICIP*), 2023.

[C17] R. A. Rojas-Gomez[G], T. Y. Lim, A. G. Schwing, M. N. Do, and **R. A. Yeh**. Learnable polyphase sampling for shift invariant and equivariant convolutional networks. In *Advances in Neural Information Processing Systems* (*NeurIPS*), 2022. (10 pages, acceptance rate 25.6%).

[C18] W. M. Gao[G], A. Wang[U], G. Metzer, **R. A. Yeh**, and R. Hanocka. TetGAN: A convolutional neural network for tetrahedral mesh generation. In *British Machine Vision Conference* (*BMVC*), 2022.
Oral Presentation. Selection rate 3.1% (24 papers selected out of 770 submissions, 8 pages).

[C19] R. A. Rojas-Gomez[G], **R. A. Yeh**, M. N. Do, and A. Nguyen. Inverting adversarially robust networks for image synthesis. In *Asian Conference on Computer Vision* (*ACCV*), 2022. (8 pages, acceptance rate 33.3%).

[C20] **R. A. Yeh**, Y.-T. Hu, Z. Ren, and A. G. Schwing. Total variation optimization layers for computer vision. In *The IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2022. (8 pages, acceptance rate 25.3%).

[C21] **R. A. Yeh**, Y.-T. Hu, M. Hasegawa-Johnson, and A. G. Schwing. Equivariance Discovery by Learned Parameter-Sharing. In *International Conference on Artificial Intelligence and Statistics* (*AISTATS*), 2022. (8 pages, acceptance rate 31.7%).

[C22] I.-J. Liu*, Z. Ren*, **R. A. Yeh***, and A. G. Schwing. Semantic Tracklets: An Object-Centric Representation for Visual Multi-Agent Reinforcement Learning. In *International Conference on Intelligent Robots and Systems* (*IROS*), 2021. (8 pages, acceptance rate 45%).

[C23] I.-J. Liu, U. Jain, **R. A. Yeh**, and A. G. Schwing. Cooperative Exploration for Multi-Agent Deep Reinforcement Learning. In *International Conference on Machine Learning* (*ICML*), 2021.
Long Talk. Selection rate 3.0% (166 papers selected out of 5513 submissions, 8 pages).

[C24] Y.-T. Hu, J. Wang, **R. A. Yeh**, and A. G. Schwing. SAIL-VOS 3D: A Synthetic Dataset and Baselines for Object Detection and 3D Mesh Reconstruction from Video Data. In *The IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2021.
Oral Presentation. Selection rate 4.2% (295 papers selected out of 7039 submissions, 8 pages).

[C25] J. Zhu, **R. A. Yeh***, and M. Hasegawa-Johnson. Multi-decoder dprnn: Source separation for variable number of speakers. In *The IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*), 2021. (4 pages, acceptance rate 48.0%).

[C26] Z. Ren*, **R. A. Yeh***, and A. G. Schwing. Not all unlabeled data are equal: Learning to weight data in semi-supervised learning. In *Advances in Neural Information Processing Systems* (*NeurIPS*), 2020. (8 pages, acceptance rate 20.1%).

[C27] I.-J. Liu, **R. A. Yeh**, and A. G. Schwing. High-throughput synchronous deep reinforcement learning. In *Advances in Neural Information Processing Systems* (*NeurIPS*), 2020. (8 pages, acceptance rate 20.1%).

[C28] **R. A. Yeh***, Y.-T. Hu*, and A. G. Schwing. Chirality nets for human pose regression. In *Advances in Neural Information Processing Systems* (*NeurIPS*), 2019. (8 pages, acceptance rate 21.1%)
Contributed Talk at Sets & Partitions Workshop.

[C29] I.-J. Liu*, **R. A. Yeh***, and A. G. Schwing. PIC: Permutation invariant critic for multi-agent deep reinforcement learning. In *The Conference on Robot Learning* (*CoRL*), 2019. (8 pages, acceptance rate 34.7%).

[C30] K. C. Mac, D. Joshi, **R. A. Yeh**, J. Xiong, R. S. Feris, and M. N. Do. Learning motion in feature space: Locally-consistent deformable convolution networks for fine-grained action detection. In *The IEEE International Conference on Computer Vision* (*ICCV*), 2019.
Oral Presentation. Selection rate 4.6% (200 papers selected out of 4303 submissions, 8 pages).

[C31] **R. A. Yeh**, J. Huang, A. G. Schwing, and K. Murphy. Diverse generation for multi-agent sports game. In *The IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2019.
Oral Presentation. Selection rate 5.6% (288 papers selected out of 5160 submissions, 8 pages).

[C32] **R. A. Yeh**, M. N. Do, and A. G. Schwing. Unsupervised textual grounding: Linking words to image concepts. In *The IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2018.
Spotlight Presentation. Selection rate 6.8% (224 papers selected out of 3303 submissions, 8 pages) .

[C33] T. Y. Lim*, **R. A. Yeh***, Y. Xu, M. N. Do, and M. Hasegawa-Johnson. Time-frequency networks for audio super-resolution. In *The IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*), 2018. (4 pages, acceptance rate 49.7%).

[C34] **R. A. Yeh***, T. Y. Lim*, C. Chen, A. G. Schwing, M. Hasegawa-Johnson, and M. N. Do. Image restoration with deep generative models. In *The IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*), 2018. (4 pages, acceptance rate 49.7%).

[C35] **R. A. Yeh**, J. Xiong, W.-M. Hwu, M. N. Do, and A. G. Schwing. Interpretable and globally optimal prediction for textual grounding using image concepts. In *Advances in Neural Information Processing Systems* (*NeurIPS*), 2017.
Oral Presentation. Selection rate 1.2% (40 papers selected out of 3240 submissions, 8 pages) .

[C36] Z. Liu, **R. A. Yeh**, X. Tang, Y. Liu, and A. Agarwala. Video frame synthesis using deep voxel flow. In *The IEEE International Conference on Computer Vision* (*ICCV*), 2017.
Oral Presentation. Selection rate 2.1% (45 papers selected out of 2143 submissions, 8 pages) .

[C37] **R. A. Yeh***, C. Chen*, T. Y. Lim, A. G. Schwing, M. Hasegawa-Johnson, and M. N. Do. Semantic image inpainting with deep generative models. In *The IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2017. (8 pages, acceptance rate 29.9%).

[C38] **R. A. Yeh**, M. Hasegawa-Johnson, and M. N. Do. Stable and symmetric filter convolutional neural network. In *The IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*), 2016. (4 pages, acceptance rate 47.0%).

Refereed Journal Publications

[J1] O. Pearl[U], I. Lang[G], Y. Hu[U], **R. A. Yeh**, and R. Hanock. Geocode: Interpretable shape programs. In *Computer Graphics Forum*. Wiley Online Library, 2025.

Refereed Workshop Publications

[W1] A.Y Zheng[G] and **R. A. Yeh**. Towards safer AI content creation by immunizing text-to-image models. In *CVPR Workshop on AI for Content Creation*, 2024.

[W2] B. Zhao[U], R. Hanocka, and **R. A. Yeh**. AmbiGen: Generating ambigrams from pre-trained diffusion model. In *CVPR Workshop on Graphic Design Understanding and Generation*, 2024.

[W3] Q. Jiang[G], X. Zhang, Deming D. Chen, Minh N M. N. Do, and **R. A. Yeh**. EH-DNAS: End-to-end hardware-aware differentiable neural architecture search. In *Differentiable Almost Everything Workshop at International Conference on Machine Learning* (*ICML*), 2023.

[W4] **R. A. Yeh***, Y.-T. Hu*, and A. G. Schwing. Chirality nets: Exploiting structure in human pose regression. In *Sets and Partition Workshop at Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Talks

Immunization and Test-Time Augmentation for Pre-Trained Computer Vision Models
○ Invited Talk at TTIC Seminar Series — *2024*

IMMA: Immunizing text-to-image Models against Malicious Adaptation
○ Invited Talk at Midwest Computer Vision Workshop at Indiana University Bloomington — *2024*

Tree-D Fusion: Simulation-Ready Tree Dataset from Single Images with Diffusion Priors
○ Invited Talk at Purdue Institute for Digital Forestry Annual Retreat — *2024*

Towards Truly Equivariant Vision Models
○ Invited Talk at National Taiwan University (NTU) — *2024*

Distilling Pretrained Diffusion Model for Content Creation
○ Invited Talk at Purdue Institute for Control, Optimization, and Networks (ICON) — *2024*

Adapting CLIP For Phrase Localization Without Further Training
○ Invited Talk at UIUC's Computer Vision External Speaker Series — *2022*

Invariance and Equivariance in Computer Vision
○ Research at TTIC Talk Series — *2021*

Incorporating Equivariance for Human Pose Regression
○ Invited Talk at IEEE M2VIP — *2021*

Chirality Nets for Human Pose Regression
○ Oral Presentation at CSL Student Conference — *2020*
○ Contributed Talk at NeurIPS Sets and Partitions Workshop — *2019*

Diverse Generation for Multi-agent Sports Games
○ Oral Presentation at CVPR — *2019*

Unsupervised Textual Grounding: Linking Words to Image Concepts
○ Spotlight Presentation at CVPR — *2018*

Interpretable and Globally Optimal Prediction for Textual Grounding using Image Concepts
○ Oral Presentation at CSL Student Conference — *2018*
○ Oral Presentation at NeurIPS — *2017*

Teaching Experience

CS47100 — Introduction to Artificial Intelligence (280 students; maxed capacity) — *Spring 2025*
○ **Evaluation:** In progress.

CS59300CVD — Computer Vision with Deep Learning (39 students; maxed capacity) — *Fall 2024*
○ **Evaluation:** Course: 4.6, Prepared: 4.7, Fair: 4.3, Caring: 4.4, Questions: 4.7

CS58700 — Foundations of Deep Learning (53 students) — *Spring 2024*
○ **Evaluation:** Course: 4.7, Prepared: 4.6, Fair: 4.4, Caring: 4.5, Questions: 4.6

CS47100 — Introduction to Artificial Intelligence (179 students) — *Fall 2023*
○ **Evaluation:** Course: 4.4, Prepared: 4.6, Fair: 4.4, Caring: 4.3, Questions: 4.3

CS47100 — Introduction to Artificial Intelligence (186 students) — *Spring 2023*

- **Evaluation:** Course: 4.4, Prepared: 4.5, Fair: 4.3, Caring: 4.4, Questions: 4.4

CS69000-TMP — Topics in Machine Perception (17 students) *Fall 2022*
  - **Evaluation:** Course: 5.0, Prepared: 5.0, Fair: 5.0, Caring: 5.0, Questions: 5.0

Services

Professional Service

Senior roles in program committees:

Area Chair: Computer Vision and Pattern Recognition (CVPR), 2025, 2024, 2023

Area Chair: Neural Information Processing Systems (NeurIPS), 2025, 2024, 2023

Area Chair: International Conference on Learning Representations (ICLR), 2025, 2024

Area Chair: International Conference on Machine Learning (ICML), 2025

Area Chair: Asian Conference on Machine Learning (ACML), 2024

Area Chair: International Joint Conference on Artificial Intelligence (IJCAI), 2023

Associate Editor: IET Computer Vision, 2023-present

Grant Panels:

NSF Panelist: Robust Intelligence (RI), 2024

Program committees:

Journal Reviewer: IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)

Journal Reviewer: IEEE Transactions on Image Processing (TIP)

Journal Reviewer: IEEE Transactions on Visualization and Computer Graphics (TVCG)

Journal Reviewer: International Journal of Computer Vision (IJCV)

Journal Reviewer: Special Interest Group on Computer Graphics and Interactive Techniques (ACM SIGGRAPH)

Journal Reviewer: SIGGRAPH Asia

Journal Reviewer: Pattern Recognition (Pattern Recognit.)

Journal Reviewer: Machine Learning (Mach. Learn.)

Conference Reviewer: Neural Information Processing Systems (NeurIPS)

Conference Reviewer: International Conference on Machine Learning (ICML)

Conference Reviewer: International Conference on Representation Learning (ICLR)

Conference Reviewer: Computer Vision and Pattern Recognition (CVPR)

Conference Reviewer: International Conference on Computer Vision (ICCV)

Conference Reviewer: European Conference on Computer Vision (ECCV)

Outreach

Judging for Start-up x AI Hackathon: Catapult Hacks
  - Invited judge for Catapult Hacks, Purdue's FIRST AI x Start-up hackathon. *2024*

Purdue University's Superheroes of Science
  - Invited interview for the Superheroes of Science podcast on the topic of AI. *2023*

Speed Mentoring at Conference on Computer Vision and Pattern Recognition (CVPR)
  - Invited mentoring and discussion on topics of interest to students — Vancouver *2023*

○ Invited mentoring and discussion on topics of interest to students — New Orleans   *2022*

Machine Learning— Making Computers See and Talk

○ Invited Talk with Girls Who Code Program at McClellan Elementary School — Chicago   *2022*