Nancy (Xin Ru) Wang

www.nancyxwang.com
nancywang1991@gmail.com | +49 15751369819 | Berlin, Germany

EDUCATION

UNIVERSITY OF WASHINGTON

PhD in Computer Science and Engineering
Dec 2018 | Seattle, WA

UNIVERSITY OF BRITISH COLUMBIA

BSc (Honours) in Computer Science
May 2014 | Vancouver, BC
Dean's Honour List, all semesters

SKILLS

LANGUAGES

Primary: Python
Working: JavaScript • TypeScript • SQL
Familiar: Java • C++ • Matlab

MACHINE LEARNING

PyTorch • computer vision • vision-language models • LLM agents and tool use • judge-model evaluation • speech and real-time audio

SYSTEMS

FastAPI • React • PostgreSQL • WebSockets • LiveKit • Docker • Kubernetes • Google Cloud • Cloudflare • Stripe

HUMAN LANGUAGES

English and Mandarin (native) • German (B1)

LINKS

LinkedIn: nancy-wang-phd
Scholar: Nancy Xin Ru Wang

EXPERIENCE

CO-FOUNDER AND CTO | FluenTea

May 2025 – Present | Berlin, DE

- Co-founded FluenTea, a video classroom with an AI speaking coach for language schools, and built the product alone as the technical co-founder, from the first commit to Google Play. Bought managed real-time video, a hosted speech model and a managed database so that one person could take it from concept to launch in three months and keep a practice call at about 1.4 cents a minute.
- Designed the voice agent's tool use with safety first and tested it the way a research system is tested. Every tool is listed in a manifest that tests enforce, each role sees its own list, nothing that reaches another person happens without a confirmed second turn, and text the learner did not type is treated as data. A red-team suite, a second model grading real answers against a baseline per error type that each release has to match or beat, and weekly reads of real transcripts.
- Ran production alone on Kubernetes for eighteen months, with deploys behind a test gate and a live-schema check, CI on seeded fixtures, contract tests for route authentication and security headers, and property-based timezone tests.
- Worked directly with teachers, coordinators and a non-technical team. Rebuilt the dashboard from user interviews (monthly active users up 247% between April and August 2026), then had ChaCha run the interviews herself (eight scripted questions, voice or text; eleven fixes from the first round). Built the age gate, parental consent, GDPR deletion, recordings, attendance and per-student speaking time a university required before enrolling minors.
- Launched Sep 2025, three months after the concept; more than 16,000 registered users from 136 countries; three university clients, the first a university in Mexico with 500 students; an Android app on Google Play. Consumers signed up in large numbers and schools were the ones who paid, so the 2026 roadmap followed the schools.

APPLIED SCIENTIST, COMPUTER VISION | Amazon

Apr 2021 – Nov 2026 | Berlin, DE

- Technical lead on three production launches on amazon.com and Alexa with multi-million dollar impact from increased sales and reduced returns; five papers at Amazon's internal computer vision conference (one oral, one best-paper shortlist) and D-Extract at WACV 2023.
- Dimension extraction from product images: conceived a text-VQA approach (a TAP-style multimodal transformer that reads measurements, arrows and tables in the image), wrote the annotation protocol and ran a five-person annotation team across 130 product types, productionised the model as an inference pipeline that backfilled the catalogue, and led the public dataset release on amazon-science. Sizes filled in for more than ten million products; Door Desk award.
- Interest and theme prediction for social shopping images and video: zero-shot and low-shot CLIP (open_clip ViT-L) against a Swin baseline with a custom loss for missing labels, Tip-adapter and class-description augmentation, four times faster inference than Swin. Designed the crowd annotation (qualification tasks, model-proposed candidates) over 270 interests and 30,000 videos and the audit-based per-interest thresholds at 90% precision. Launched into Amazon's interest mapping and used by two recommenders.
- Visual similarity and refinement for Alexa Shop the Look: embedding retrieval with sibling and near-duplicate filters for diverse results, index coverage raised from 15% to 95% by tracing the data pipeline to one attribute, and a CLIP-style composed-retrieval model (reference image plus modification text). Launched at 100% in June 2023; shown at re:MARS.
- 2025: generative video and 3D feasibility studies: frame selection for structure-from-motion and Gaussian splatting, video transitions and product swap with diffusion models, image editing with Flux Kontext.
- Supervised four research interns; organised the Berlin ML Workshop 2023 (180 attendees); won the Berlin GenAI hackathon 2024 with auto-generated buying-guide videos; Amazon Berlin Champion of Diversity and Inclusion award, 2021.

TEACHING

Teaching assistant at UW and UBC:

Artificial Intelligence
Algorithms and Data Structures
Models of Computation
Computer Vision
Machine Learning
(graduate)

SERVICE

Lead organiser, SemEval 2021 Task 9

Tutorials on table extraction, ICDM 2019 and VLDB 2020

EXTRA-CURRICULAR

2016–2017: Founder and President, UW High Performance Computing Club, Cloud Chapter

2010–2011: Elected Science Councillor, UBC Alma Mater Society

INTERESTS

Brains and computers, any cognitive system. Languages and living inside other cultures. Role-playing and strategy games (Dragon Age, XCOM).

EXPERIENCE, CONTINUED

RESEARCH STAFF MEMBER | IBM Research Almaden

Mar 2019 – Mar 2021 | San Jose, CA

- Deep learning and computer vision for enterprise document understanding in the Scalable Knowledge Intelligence department: algorithms to find and read tables, and active-learning tools that let users build models for their own documents.
- Technical lead and main developer of GTE, an end-to-end table extraction system (WACV 2021), and of TableLab, the labelling interface built on it, from concept to launch.
- Lead organiser of SemEval 2021 Task 9 (SEM-TAB-FACTS) on fact verification against tables. Co-author of CORD-19 and of D2S (NAACL 2021). Tutorials on table extraction at ICDM 2019 and VLDB 2020.

GRADUATE RESEARCH ASSISTANT | University of Washington

Sep 2014 – Dec 2018 | Seattle, WA

- Sensor-fusing deep learning models that decode and predict human behaviour from hundreds of hours of simultaneous video, audio and electrocorticography recordings (AAAI 2018, oral). Advised by Bingni Brunton and Rajesh Rao, with Jeff Ojemann and Ali Farhadi.
- Founded the project to automate the annotation of naturalistic neuroscience data; it became a multi-year effort funded by NSF and DARPA, and parts of the dataset were released.

PHD INTERNSHIPS | Zillow · Rakuten Institute of Technology

Jun – Sep 2018, Seattle, WA · Mar – Jun 2016, Tokyo, Japan

Computer vision for Zillow housing products; unsupervised user modelling for search ranking and page design at Rakuten.

SELECTED PUBLICATIONS

Pushpendu Ghosh, Nancy Wang, Promod Yenigalla. D-Extract: Extracting dimensional attributes from product images. *WACV 2023*.

Edward Sun, Yufang Hou, Dakuo Wang, Yunfeng Zhang, Nancy X. R. Wang. D2S: Document-to-slide generation via query-based text summarization. *NAACL 2021*.

Xinyi Zheng, Douglas Burdick, Lucian Popa, Xu Zhong, Nancy Xin Ru Wang. Global Table Extractor (GTE): A framework for joint table identification and cell structure recognition using visual context. *WACV 2021*.

Nancy X. R. Wang, Diwakar Mahajan, Marina Danilevsky, Sara Rosenthal. SemEval-2021 Task 9: Fact verification and evidence finding for tabular data in scientific documents (SEM-TAB-FACTS). *SemEval 2021*.

Nancy Xin Ru Wang, Ali Farhadi, Rajesh Rao, Bingni Brunton. AJILE movement prediction: Multimodal deep learning for natural human neural recordings and video. *AAAI 2018*, oral, 25% acceptance.

Nancy Xin Ru Wang, Jared D. Olson, Jeffrey G. Ojemann, Rajesh P. N. Rao, Bingni W. Brunton. Unsupervised decoding of long-term, naturalistic human neural recordings with automated video and audio annotations. *Frontiers in Human Neuroscience*, 2016.

AWARDS

2016–2019	Doctoral stipend	NSERC Postgraduate Scholarship (the Alexander Graham Bell CGS-D was declined in its favour)
2016–2017	$20,000 Azure credits	Microsoft Azure Research Award
2015–2017	Full tuition and stipend	Washington Research Foundation Innovation Graduate Fellow in Neuroengineering and Data Science
2014	Full tuition and stipend	Wilma Bradley Endowed Fellowship
2014	Research grant	NSERC Undergraduate Student Research Award
2011	Research scholarship	DAAD RISE, Germany
2009–2013	Full undergraduate scholarship	Hugh M. Brock Award