


Human and Machine Cognition Lab



Thesis topics

Prof. Dr. Charley M. Wu
Computational Cognitive Science
hmc-lab.com
Landwehrstr. 50A
charley.wu@tu-darmstadt.de







TECHNISCHE
UNIVERSITÄT
DARMSTADT



Centre for
Cognitive
Science



hessian.AI




Human and Machine Cognition Lab

Guiding questions of the lab


1. How do we **generalize** in novel environments?
2. How does intelligence **develop** over the lifespan?
3. How do we learn from each other in **social** and **cultural** settings?

Human **generalization** as Bayesian function learning





[Wu, Meder, Schulz \(AnnRevPsy 2025\)](#)

Developmental change as stochastic optimization



[Giron*, Ciranka*,... & Wu \(NatHumBehav 2023\)](#)

Cultural transmission through **social** interactions




[Velez, Wu, Gershman & Schulz \(PsyArxiv 2024\)](#)

Thesis topics are organized around the three guiding questions

**An orange box like this indicates a potential
thesis topic**

Generalization as function approximation


Law of Generalization



Shepard (Sci 1987)


+

Function Approximation



Silver et al., (Nature 2016)

Generalization as function approximation



Spatially Correlated Bandit

7	5	10	22	32	32	28	24	22	26	33
6	11	19	29	38	41	42	40	37	36	40
22	27	30	35	43	50	53	53	51	49	46
45	44	38	36	40	46	47	49	54	55	48
61	55	46	40	37	32	27	31	44	52	44
62	59	57	54	44	27	14	17	33	46	45
53	59	68	71	59	36	17	15	28	45	51
46	57	71	77	67	47	26	18	27	45	56
45	56	65	67	60	46	29	20	27	42	55
51	57	58	53	47	40	30	23	28	40	49
60	62	58	47	39	38	35	31	35	41	46



Click tiles on the grid to maximize rewards



Each tile is an arm of the bandit, giving noisy rewards




Reward expectations are spatially correlated



Limited search horizon (20-40 clicks) requires balancing explore-exploit


Framework and models replicated across a variety of tasks and contexts

Spatial and conceptual search




[Wu, Schulz, Garvert, Meder & Schuck \(PLOS Comp Bio 2020\)](#)

Graph-structured rewards




[Wu, Schulz & Gershman \(CBB 2021\)](#)

Forgetful generalization



[Ten, Sakaki, Breit, Chandrasekaran, Murayama & Wu \(PsyArXiv, 2025\)](#)

Hippocampal gradients




Wong, ... Hauser* & Wu* (in prep)

Previous Masters Thesis Example


Can human hippocampal gradients also be found in non-spatial tasks? Reanalysis of fMRI data from Garvert et al., (2017)



Anna-Lea Beyer



Garvert et al., (2017)



Brunec & Momennejad (2022)

Thesis topic: Structured Generalization in Dance

- **Goals:**

- Adapt an existing online experiment for learning & completing melody sequences to movement sequences in Labanotation
- Define movement primitives in combinatorial logic
- Pilot the experiment

- **Past work:**





- [Zhou et al., \(2026\)](#) developed a model of melody learning as “program induction” under resource constraints
 - Program induction: learning as inferring a distribution over generative programs
- Note: the computational model is quite complex, but the focus here will be on experiment design

- **Skills you will learn:**

- Online experiment development in JavaScript
- Version control with Git/GitLab
- Experimental design for behavioural studies in humans



Anna-Lea Beyer



Note: these topics require a very high level of independence, because no current PhD or Postdocs in the lab are working on this topic

Other potential thesis topics

- Reanalyze published data to answer open questions
 - Some examples:
 - Do people gradually adapt how much they generalize to the task at hand?
 - How can we design more effective curriculum for learning (i.e., by ordering environments)?
 - Would graph kernels ([Wu, Schulz & Gershman, 2021](#)) provide a similarly good account of hippocampal gradients?
- Implement your own experiment and conduct new analyses
 - More challenging and may not be feasible in a single semester
 - Needs a very compelling research question
- **How to get started**
 - Read [Wu et al., \(AnnRevsPsych 2025\)](#) for an overview
 - Read other relevant papers (see <https://hmc-lab.com/publication/> for open-access pdfs of all papers)
 - Identify open questions and identify hypotheses
 - Develop an initial roadmap

Development as Stochastic Optimization



- Children initially perform high-temperature search, which gradually “cool offs” as they grow older
- The cooling schedule induces a gradual shift from exploration —> exploitation, preventing the algorithm from getting stuck in a local optimum





Anna Giron
Uni Tübingen

Previous Masters thesis

First direct test of the “cooling-off” hypothesis: human vs. optimization algorithms

Human vs. SHC-fast



Current directions

- How do (adversive) childhood experiences influence learning strategies in adulthood?
- Would the predicability of your childhood environment influence how we generalize and explore as an adult?

Potential thesis topics

- Simulation-based study of how strategies optimized for one set of environments generalize to new environments
- Implement an experiment to relate childhood background to normative predictions

How to get started

- Read [Giron et al., \(Nature Human Behavior 2023\)](#) for an overview and recent papers (e.g., [Schechtl, NatComms 2025](#); [Vermeent et al., JEP:G 2025](#))
- Develop theory-driven predictions and propose a plan for how to map them to the model




Dr. Simon Ciranka
MPI Berlin






Dr. Anna Thoma
MPI Berlin / Uni Frankfurt

High autocorrelation
(predictable, benign environment)



Low autocorrelation
(unpredictable; adverse environment)




Recent Masters Thesis Example



Intrinsically Motivated Learning in Open-Ended Environments:
Understanding the Interplay of Empowerment and Information Gain



Fryderyk Mantiuk




Lava Wall Ice Agent





Human intelligence is defined by our capacity for **social learning** and **cumulative culture**



Physical





Social




Hierarchy of Social Learning

Theory of Mind (ToM)



[Wu, Dale & Hawkins \(2024\)](#)




[Wu, Veléz & Cushman \(2022\)](#)

Cost-based arbitration


Social transitive inference task



Value hierarchy




[Witt et al., \(2023; in prep\)](#)



Flexible social learning despite individual differences


[Witt et al., \(2024\)](#)

PNAS




Visual-spatial dynamics of adaptive social learning

[Wu, Deffner et al., \(2025\)](#)






Cultural evolution in virtual societies



~ 60 minutes real time



Born to a random player
Dependent on online player community to survive until adulthood
Incremental improvements lead to largescale technological progress over successive generations



Recent Masters' Thesis Example


Comparing Global and Local Advice in Human-AI Interaction



Orsolya Szócs

Global vs. local AI assistance in home design


Global AI: assigns functional roles to rooms within a home



Example for global advice




Local AI: optimises furniture arrangements within a room




Example for local advice



Experimental procedure overview



Average performance across phases






Current Masters' Thesis Example

Can simple social learning mechanisms give rise to the social transmission of theory-of-mind like representations?



Silja Keßler



Potential extensions as a thesis topic



- How would these results generalize to more dangerous environments?
- How would different teaching strategies influence performance?

Thesis topic: Decision-making in coordination games as a function of action costs

- **Research Question:** Do people consider the potential cognitive/movement costs for a partner and/or the dyad when selecting solutions for a game?
- **Past work:**
 - In a simple object matching task with visual stimuli ([Török et al., 2021/ PhD thesis, unpublished Study 4](#)) we found that people tend to consider the aggregate costs of a dyad that would have to be expended to collect objects by both co-actors in a coordination game
 - Participants tended to make choices to coordinate with an assumed remote partner such that the joint object collection costs would be minimized
 - However, we cannot fully rule out that choices were made based on the visual layout of the scenes and not only based on the movement costs, as visual arrangements and movement costs strongly correlated.
 - This project would provide follow-ups to the original two experiments via conceptual replications.



Georgina Török





Thesis topic: Social inference in collective search tasks

- **Research Question:** Does social inference in social learning occur to the same extent when people believe they interact with a machine as when they believe they interact with a human? In other words: in contexts of uncertain collective search, do humans treat machines as if they were other human agents and copy them similarly selectively?
- **Past work:**
 - rying to read other minds (i.e. inferring rewards based on observed behavior) where there are in fact none may not be worth the computational costs
 - This project would test if heuristic models might better explain human behavior in collective search tasks than a social inference model when people know they interact with a bot
 - Alternatively, people might treat known artificial agents as if they were humans in such contexts, and infer rewards in the same way as they would with human others
 - This project would be directly based on [Hawkins et al. \(2023, see images\)](#) but would also build on and connect to research on ascribing minds to artificial agents and treating them as intentional agents



Georgina Török



Other potential thesis projects

Social Learning

- **Agent-based simulations** to study which social learning strategies perform best in different environments ([Keßler et al., Cogsci 2026](#))
- **Implement experiments** with AI-Human collaboration ([Szócs et al., Cogsci 2025](#))
- [Infrastructure in progress]: **online multiplayer experiments**

Cultural Transmission

- **Agent-based simulations** to study the dynamics of social transmission over multiple generations
 - Eg., [Ammar, Fogarty & Kandler \(2023\)](#)
- [Infrastructure in progress]: **model cultural transmission using LLMs**
- [Infrastructure in progress]: **online transmission chain experiments**

How to get started

- Read [Wu, Dale, & Hawkins \(Open Mind, 2024\)](#) for an overview of some open questions
- Read related literature to identify current theoretical accounts and hypotheses
- Propose a plan for how to test them empirically

General information

- Working language is English
- Strong computational skills are required for all lines of research
 - Please make sure to *include a link to public code repositories* when contacting me about a thesis project
- My lab members are quite new when it comes to supervising theses in Darmstadt
- You will be expected to take part in weekly lab events (lab meeting, journal club, and talks) and contribute to discussions
- A high-degree of independence and self-motivation is required, including willingness to learn new skills and read relevant papers on your own initiative

Thank you for your attention

- Lab website: hmc-lab.com
 - 📌 PUBLICATIONS for open access pdfs of all papers & links to open code/data
 - 📌 RESOURCES for these slides as a pdfs
- Offices: @ Hessian.AI, Landwehrstraße 50A
 - Flex office for Bachelor/Masters students
- Contact:
 - Email: charley.wu@tu-darmstadt.de
 - Bluesky: [@thecharleywu.bsky.social](https://bsky.app/profile/thecharleywu.bsky.social)

