

Savvas Dimitrios Petridis

e-mail: savvas@cs.columbia.edu

Personal Website: <https://savvaspetridis.github.io/>

LinkedIn : www.linkedin.com/in/savvaspetridis/

Github : <https://github.com/savvaspetridis>

Education

Columbia University, New York, NY USA

PhD, Computer Science

October 2022

Research Area: Human-AI Interaction, HCI, AI, NLP

Thesis: Designing Exploratory Search Systems that Stimulate Memory and Reduce Cognitive Load

Advisor: [Lydia B. Chilton](#)

Committee: [Michael Terry](#), [Nicholas Diakopoulos](#), [Eugene Wu](#), [Brian A. Smith](#)

MS, Computer Science (completed during PhD)

May 2018

BA, Computer Science

May 2017

GPA: 3.78/4.0, Dean's List all semesters

Research Experience

Google People + AI Research (PAIR) / Responsible AI

Visiting Research Scientist (PostDoc)

November 2022 - now

with Carrie J. Cai and Michael Terry

- Developed and led research on ([ConstitutionalExperts](#), in submission at *ACL 2024*), an automatic method for improving prompts via fine-grained edits of principles and mixture of experts.
- Developed ([ConstitutionMaker](#), IUI 2024), a method for interactively improving LLM prompts via natural language feedback.
- Developed ([PromptInfuser](#), in submission at *DIS 2024*), a Figma plugin which enables designers to quickly create *functional* AI prototypes, by connecting LLM prompts to UI elements.
- Helped design and build ([LinguisticLens](#), VIS 2023) an interactive tool for visualizing the diversity of LLM-generated datasets.

Columbia University

PhD Candidate

September 2017 - October 2022

with Lydia B. Chilton

- Developed an interactive tool that uses a large language model (LLM) to generate angles for journalists given a press release ([AngleKindling](#), CHI 2023).
- Designed and developed an interactive web tool, which enables novice designers to explore a word-association network to find multiple, diverse visual symbols for abstract concepts ([SymbolFinder](#), UIST 2021).
- Conducted a study on Mechanical Turk to understand how people interpret complex visual messaging, like visual metaphors, and distilled four types of errors in their interpretations (Creativity & Cognition 2019).
- Developed and studied an interactive web tool, which helps novice designers create visual blends, images that combine symbols to convey a greater meaning ([VisiBlends](#), CHI 2019).

Spotify

Research Intern

Summer 2021

with Nediya Daskalova, Sarah Mennicken, Sam Way, Paul Lamere, & Jenn Thom.

- Designed and developed a web tool that helps users understand and explore a music genre by generating a personalized artist-graph with guide. ([TastePaths](#), IUI 2022).

Adobe Research

Creative Intelligence Lab Intern

Summer 2019

with Valentina Shin, Joy Kim, Mira Dontcheva, & Karrie Karahalios.

- Studied how audio can be used to animate graphics. Created a web-prototype interactive storybook, where the graphics could be animated by the reader’s voice in real-time. Conducted a pilot study, with a few parents and their kids, to understand how users would interact with such a book.

IBM

Research Intern

Summer 2017

with Seraphin Calo and Maroun Touma.

- Created an interactive web-tool to experiment with different features to train a suite models designed to classify devices communicating in a wireless network. Users could specify different models, parameters, and features to include, as well as the training and test set.

Research Intern

Summer 2016

with Seraphin Calo and Maroun Touma.

- Scraped data from multiple online newspapers and blogs to train models designed to classify the political bias of news articles.
- Designed and developed an interactive web-based strategy game end-to-end in order to study how individuals make complicated decisions.

Wireless & Mobile Networking Lab at Columbia University

Undergraduate Research Assistant

June 2015 - February 2016

with Prof. Gil Zussman.

- Developed an interactive web-based application, which demonstrates the efficacy of feedback node selecting algorithms in a WiFi multicast network. Featured: **NYC Media Lab Summit (won second best demo)**, IEEE Local Computer Networks conference, IEEE INFOCOM16, and the GENI NICE conference.

Publications

Savvas Petridis*, Benjamin Wedin*, Ann Yuan*, James Wexler, Nithum Thain. “ConstitutionalExperts: Training a Mixture of Principle-based Prompts.” *In Submission at ACL 2024*. * denotes equal contribution.

Savvas Petridis, Benjamin Wedin, James Wexler, Aaron Donsbach, Mahima Pushkarna, Nitesh Goyal, Carrie J. Cai, Michael Terry. “ConstitutionMaker: Interactively Critiquing Large Language Models by Converting Feedback into Principles.” *IUI 2024*.

Savvas Petridis, Michael Terry, Carrie J. Cai. “PromptInfuser: How Tightly Coupling AI and UI Design Impacts Designers’ Workflows.” *In Submission at DIS 2024*.

Emily Reif, Minsuk Kahng, **Savvas Petridis**. “Visualizing Linguistic Diversity of Text Datasets Synthesized by Large Language Models.” *VIS 2023*.

Savvas Petridis, Michael Terry, Carrie J. Cai. “PromptInfuser: Bringing User Interface Mock-ups to Life with Large Language Models.” *CHI 2023 LBW*.

Savvas Petridis, Nicholas Diakopoulos, Kevin Crowston, Mark Hansen, Keren Henderson, Stan Jastrzebski, Jeffrey V. Nickerson, Lydia B. Chilton. “AngleKindling: Supporting Journalistic Angle Ideation with Large Language Models.” *CHI 2023*.

Sitong Wang, **Savvas Petridis**, Taeahn Kwon, Xiaojuan Ma, Lydia B. Chilton. “Pop-Blends: Strategies for Conceptual Blending with Large Language Models.” *CHI 2023*.

Savvas Petridis, Nediyan Daskalova, Sarah Mennicken, Samuel F. Way, Paul Lamere, Jenn Thom. “TastePaths: Enabling deeper exploration and understanding of personal preferences in recommender systems”. *Intelligent User Interfaces (IUI) 2022*.

Savvas Petridis, Hijung Valentina Shin, Lydia B. Chilton. “SymbolFinder: Brain-

storming Diverse Symbols using Local Semantic Networks”. UIST 2021.

Savvas Petridis, Lydia B. Chilton. “Human Errors in Interpreting Visual Metaphor”. Proceedings of the 2019 on Creativity and Cognition, 187—197, 2019.

Lydia B. Chilton, **Savvas Petridis**, Maneesh Agrawala. “VisiBlends: A Flexible Workflow for Visual Blends”. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019.

Tariq Alhindi, **Savvas Petridis**, Smaranda Muresan. “Where is your Evidence: Improving Fact-checking by Justification Modeling”. Proceedings of the First Workshop on Fact Extraction and VERification (FEVER) at EMNLP, 2018.

Lydia B. Chilton, **Savvas Petridis**, Maneesh Agrawala. “An Interactive Pipeline for Creating Visual Blends”. The 31st Annual ACM Symposium on User Interface Software and Technology Adjunct Proceedings, 188—190, 2018. (Demo)

Varun Gupta, Raphael Norwitz, **Savvas Petridis**, Craig Gutterman, Gil Zussman, Yigal Bejerano. “AMuSe: Large-scale WiFi video distribution-experimentation on the ORBIT testbed”. IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPs), 2016. (Demo)

Varun Gupta, Raphael Norwitz, **Savvas Petridis**, Craig Gutterman, Gil Zussman, Yigal Bejerano. “WiFi Multicast to Very Large Groups-Experimentation on the ORBIT Testbed”. IEEE LCN, Oct. 2015. (Demo)

Mentoring Experience

Advised

[Sitong Wang](#) (MS at Columbia), now PhD student at Columbia.

[Terry Kwon](#) (MS at Columbia), now software engineer at Foliage.

Thesis Committee

[Terry Kwon](#) (MS Thesis).

Professional Service

Paper Reviewing

ACM CHI 2021-2024, IUI 2022-2023, UIST 2022-2023, DIS 2021

Program Committee

ACM IUI 2023

Technical Skills

Languages: Python, JavaScript, HTML, CSS

Databases: MySQL, PostgreSQL, Neo4j

Tools/Framework: Flask, Django, Node.js, Keras, NumPy, spaCy, D3.js

Familiar: Java, C++, Matlab