The Structure of Human-AI Collaboration

Large tasks can seem impossible until they are broken down into smaller parts. My research demonstrates how to decompose hard problems into small, actionable steps. I operationalize theories of human cognition to generate insights into how people work and how machines can support them. In 2010, I showed how crowd algorithms could organize people to solve complex tasks through simple actions [1, 2]. My current work breaks down large problems into small tasks that people and generative AI can tackle together [6,7,8,9,10,11].

To enable collaboration on hard problems, we must first understand them. My VisiBlends [3] system showed how cognitive non-routine (or creative) work can be decomposed and executed. Drawing from foundational psychological concepts of divergent, convergent, and fluid thinking, I produced mechanized models of the creative processes – including ideation, prototyping, and iteration (see figure). In this model, rather than ideating multiple complete solutions, the system generates multiple *options for each part of the solution* (in a visual blend it generates multiple symbols for each concept to be blended). Rather than trying all possible combinations in the prototype phase, *the system uses abstract design patterns to guide the search* for elements that can fit the pattern (such as blending symbols that contain the same main shape). The system *computationally annotates images, finds matches, then synthesizes prototypes*. People evaluate the prototypes and iterate on the best ones by *repeating the process*.

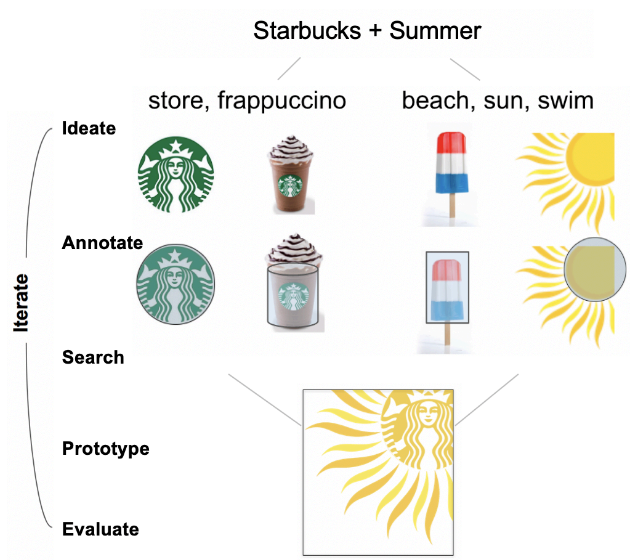

Given this model, I have shown that AI can facilitate many parts of the process: finding elements in ideation [5], applying design patterns to create prototypes [3], and iterating to improve prototypes [4]. User studies show people are 10 times more effective at making successful blends with the system than without it because it allows people to focus on high-level tasks like critical evaluation and directing the search as AI does low-level work like searching, optimizing, and prototyping.

One of the biggest impediments to solving a problem is not actually solving it - it is understanding the problem. This is called problem framing. My recent work shows that given a set of *abstract problem framings* mined from previous solutions, AI can quickly apply each framing to the problem for people to quickly evaluate. My ReelFramer work shows that AI can help journalists pick from multiple narrative framings to translate different types of print news stories into short social media videos with the appropriate balance of news and fun [8].

My current research shows my approach to human-AI collaboration works for many types of tasks [8,9,10,11], including making software applications to solve users' problems [10]. *AI can help ideate for the key dimensions of an app* – understanding users' needs, proposing theories to guide the solution approach, and suggesting user interactions to implement it. Once the user selects options that are appropriate for the problem, *AI can synthesize code* to make a working prototype. Compared to people using state-of-the-art AI, the system helps people achieve double the performance with half the mental effort [10]. To help AI gather the context necessary for problem-solving, I introduced an *AI-powered context curation* method. The method detects when problems are underspecified, and selects relevant context from local files; if that is not sufficient, the AI also asks users short questions to elicit necessary context that is not yet written down [11]. Our user studies show AI context curation is 90% accurate and increases users' efficiency by 100%. To be truly helpful, people need AI integrated into their tasks, their workflow, and their context. My future research will provide that seamless collaboration.

My mechanized models of complex tasks help us understand and augment human cognition. Both humans and AI are imperfect; my research shows how to structure human-AI work that allows people to solve problems faster and better than they can alone.

1. Greg Little, Lydia B. Chilton, Max Goldman, Robert C. Miller. TurKit: Human Computation Algorithms on Mechanical Turk. *UIST 2010.*
2. Lydia B. Chilton, Greg Little, Darren Edge, Daniel S. Weld, James A. Landay. Cascade: Crowdsourcing Taxonomy Creation. *CHI 2013.*
3. Lydia B. Chilton, Savvas D. Petridis, Maneesh Agrawala. VisiBlends: A FlexibleWorkflow for Visual Blends. *CHI 2019.*
4. Lydia B. Chilton, Ecenaz Jen Ozmen, Sam Ross, and Vivian Liu. VisiFit: Structuring Iterative Improvement for Novice Designers. *CHI 2021.*
5. Savvas Petridis, Hijung Valentina Shin, Lydia B. Chilton. SymbolFinder: Brainstorming diverse symbols using local semantic networks. *UIST 2021.*
6. Vivian Liu, Han Qiao, and Lydia B. Chilton. Opal: Multimodal Image Generation for News Illustration. *UIST 2022.*
7. 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*.
8. Sitong Wang, Samia Menon, Tao Long, Keren Henderson, Dingzeyu Li, Kevin Crowston, Mark Hansen, Jeffrey V Nickerson, Lydia B Chilton. ReelFramer: Human-AI Co-Creation for News-to-Video Translation. *CHI 2024.*
9. Samia Menon, Sitong Wang, Lydia Chilton. MoodSmith: Enabling Mood-Consistent Multimedia for AI-Generated Advocacy Campaigns. *ICCC 2024.*
10. Jenny Ma, Karthik Sreedhar, Vivian Liu, Pedro A. Perez, Sitong Wang, Riya Sahni, Lydia B Chilton. DynEx: Dynamic Code Synthesis with Structured Design Exploration for Accelerated Exploratory Programming. *Under submission to CHI 2025.*
11. Sitong Wang, Xuanming Zhang, Jenny Ma, Alyssa Hwang, Lydia B Chilton. JumpStarter: Getting Started on Personal Goals with AI-Powered Context Curation. *Under submission to CHI 2025.*