

This document outlines my vision on research, valorization, education, leadership, and organization. It complements my CV, which includes my background and experience in each of these five aspects. I start by describing my vision of sociotechnical systems, which is at the core of my research, teaching, and valorization efforts.

1 Research

Engineering sociotechnical systems (STSs) is the overarching theme of my research. I envision an STS as a system that supports rich interactions between *principals* (humans or organization) and computational *agents*, enabling a variety of individual and societal applications. Figure 1 depicts my vision, which includes an STS at the center.

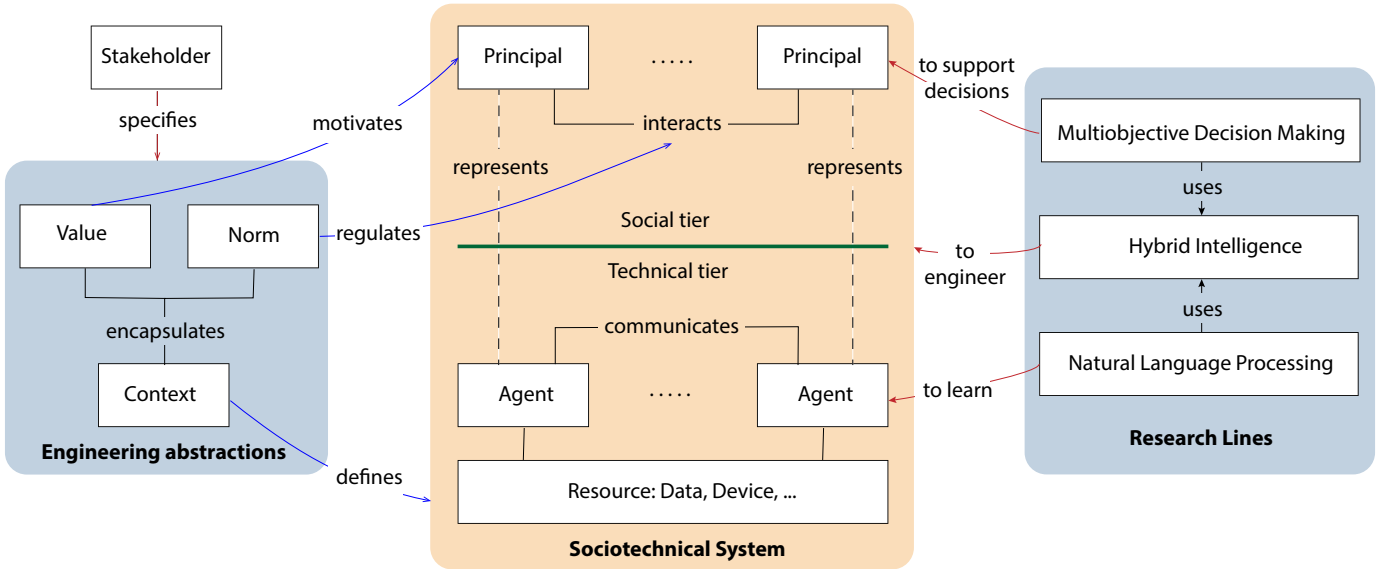


Figure 1: A sociotechnical system or STS (at the center), key abstractions necessary to engineer an STS (on the left), and the methods I research and develop (on the right) to realize an STS.

Figure 1, on the left, shows three key abstractions that guide my research.

A value is a principal’s deep-rooted motivation to act a certain way.

A norm represents a principal’s (directed) expectation of another principal.

A context defines the micro-society an STS represents, and it makes values and norms concrete.

I have carved three technical research lines to investigate context-specific values and norms.

Natural Language Processing (NLP). To act inline with a principal’s values and norms, an agent must first learn those values and norms. Traditional approaches (e.g., value surveys or hard-coded norms) yield broad-brush, context-agnostic, values and norms, which are often not useful in concrete applications. I investigate NLP techniques for learning context-specific values and norms based on stakeholders’ natural language interactions.

Multiobjective Decision Making (MODM). Most practical applications involve multiple principals, and consequently, require balancing several competing values and norms during decision making. The well-established field of Multiobjective Optimization (MOO) considers such problems, but has two key limitations. Most MOO methods (1) start with well-formulated objectives, and (2) end by producing a solution set (e.g., a Pareto set). I seek to address these two limitations. To address the first, I explore methods formulate objectives from values and norms (extracted via NLP). To address the second, I investigate methods to explain MOO output to stakeholders.

Hybrid Intelligence (HI). In an STS, a principal is paramount. Principals act autonomously (based on values) and are accountable to each other (as specified by norms). Agents, in contrast, support decision-making by principals. Thus, principals and agents must collaborate in an STS, making it a HI system. I investigate systematic methods for developing and evaluating HI systems. This paradigm is inherent in my works on NLP and MODM.

Figure 1 shows the positioning of my PhD and PD advisees across the three research lines.

- The NLP and HI research lines have shaped well in the last five years. I have found a niche—NLP techniques for inferring contextual values and norms—which is novel, challenging, and has potential for impact.
- The MODM research line is developing. My interests in MODM—pre-optimization (how objectives come about) and post-optimization (how to convey MODM output to stakeholders)—are largely unexplored topics. I have good collaborators (especially, in the Hippo Delft AI lab¹) on this endeavor.

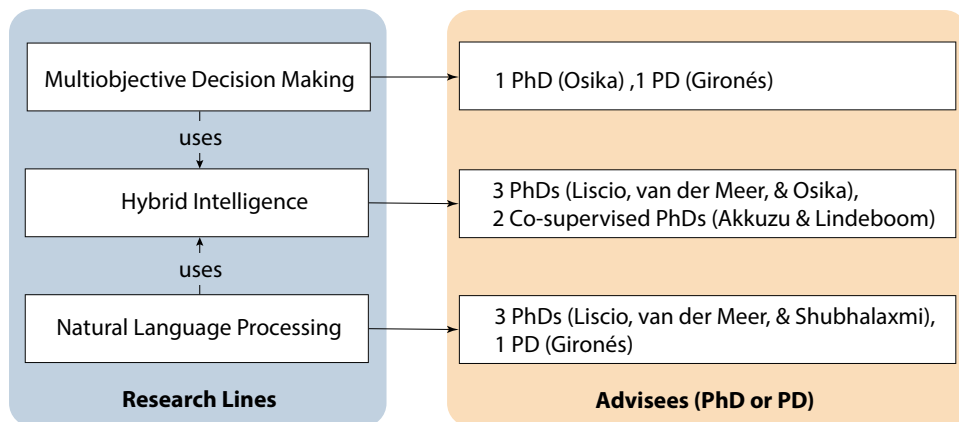


Figure 2: A mapping of my PhD and PD advisees into my research lines.

Commitment. My research vision is a result of years-long thinking. I am committed to realizing this vision because it (1) involves challenging topics that I enjoy working on, (2) contains a variety of topics, which keep me curious, and (3) has potential for significant scientific and societal impact, which I find fulfilling.

2 Valorization

Figure 2 show how my current projects contribute to my research lines and how I plan to develop those further.

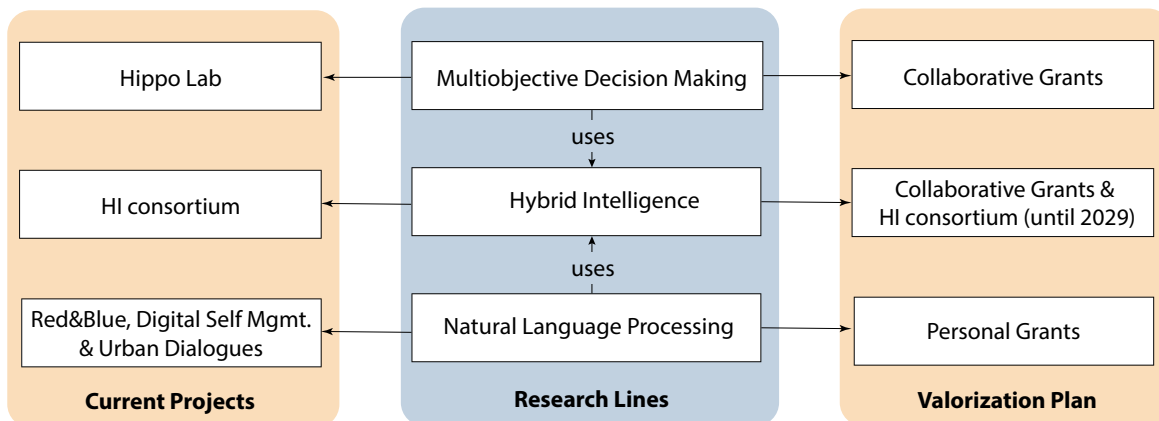


Figure 3: A mapping of my research lines to my current projects and valorization plan. Most projects involve more than one research line (since the research lines are intersecting). I only show one representative arrow in the mapping.

Current Projects. I am fortunate to have found adequate funding and talented students for developing my research.

- The INSY department supported me with one PhD position to start with.

¹<https://www.tudelft.nl/ai/hippo-lab>

- The Hybrid Intelligence center², funded by a 10 year Zwaartekracht grant, started just after I joined TU Delft. I became an affiliated member of the center since my research agenda and the center’s missions have a strong overlap. I am the daily supervisor of one PhD and co-supervisor of two PhDs funded by the HI center.
- I co-direct the Hippo Lab, one of the 24 TU Delft AI Labs, funded by the TU Delft AI initiative. In the Hippo lab, I supervise two PhD students at EEMCS, and collaborate with two other PhD students at TPM.
- I acquired an NWO project, RED&BLUE³, as a co-applicant. I co-developed a work package in this project, and I am supervising a (three-year) Postdoctoral Associate on that work package.
- I recently acquired another NWO project on Digital Self-Management⁴ as a co-applicant. I will co-supervise a Postdoctoral Associate and a Research Engineer as part of this project.

Valorization Plan. The following is an outline the valorization plan I hope to execute in the next three years.

Personal grant on NLP and Values. Of my research lines, I have the strongest track record at the intersection of NLP and HI. I am working on these lines since I started at TU Delft. I would like to first start with an ERC Consolidator grant application on this topic, followed by an NWO Vici application after a few years.

Collaborative on MODM. I would like to submit collaborative grant applications on the topic of MODM. This is an interdisciplinary topic and I have a good network to pursue such an application. I am already working on two avenues in this regard: (1) a joint-application with my Hippo lab co-director from TPM, and (2) a joint-applications with colleagues in Leiden University, whose expertise complements mine.

Collaborative grant on HI. The HI Zwaartekracht grant lasts until 2029. The consortium will recruit the third cohort of PhDs and PDs in the next year or two. I am playing several active roles in the consortium and will try my best to get a PhD/PD in the next cohort. In addition, I and a few other colleagues in the HI center are working on a EU COFUND project for HI. If we succeed, I hope to get an additional PhD/PD from it.

I would like to mention that I decided not to apply for a personal grant during my tenure as an Assistant Professor for two main reasons. (1) I acquired adequate resources via other projects to develop each of my research lines. Applying for a personal grant would not only have cost me significant time, but also I would not have had sufficient time to supervise additional PhD students or PDs if the personal grant was granted. (2) I started developing the NLP and HI research lines after I started at TU Delft, and I did not have a strong record on that topic at that time to submit an NWO Vidi application (I was not eligible for a Veni application when I started at TU Delft).

3 Education

Brief Teaching Statement. Learning Computer Science (CS) can be both exciting and challenging. Perhaps, a major problem CS students face is information overload. With so much to learn and so much of information (and misinformation) easily available, learning CS concepts, staying in the *flow channel*—neither getting bored nor panicking—can be extremely difficult. As a teacher, I strive to show my students that learning CS is about grasping a relatively few fundamental ideas and principles that underlie a plethora of application-level technologies.

University Teaching Qualification (UTQ). I successfully completed my UTQ at TU Delft in September 2022. I found the UTQ trajectory to be quite helpful. When I look back, I do see differences in the way I teach and mentor before and after UTQ. Although most ideas presented in UTQ modules seemed intuitive and familiar, the trajectory forced me to reflect on and challenge those ideas. This process made a difference in the way I approach teaching and mentoring. For a detailed version of my reflection on teaching, please see my UTQ dossier⁵.

Current Teaching. I am currently at my full capacity for teaching. I teach in three courses (two BSc and one MSc), and coordinate two of those. I am developing a new course for the upcoming MSc in DSAIT program, and I intend to coordinate and teach in it. In addition, I supervise two–three MSc theses and two–three Bachelor projects in a year.

Teaching and Research Alignment. As Figure 3 depicts, my teaching and research align very well. This is not coincidental, but a result of my strategic effort—I developed three out of the four courses shown here.

²<https://www.hybrid-intelligence-centre.nl/>

³<https://redblueclimate.nl/>

⁴<https://www.tilburguniversity.edu/current/news/more-news/grant-self-management-tool-severe-mentally-ill>

⁵<https://ii.tudelft.nl/~pradeep/doc/Murukannaiah-UTQ-Final-Dossier.pdf>

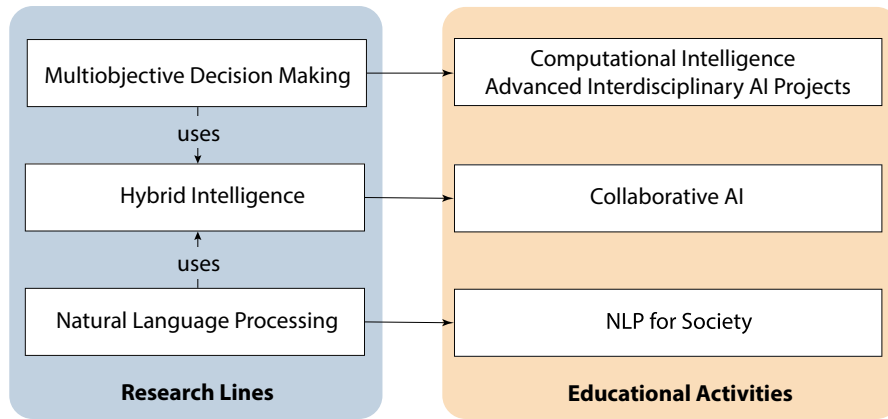


Figure 4: A mapping between my research lines and teaching activities.

Short-Term Plans. The following is my concrete short-term planning of educational activities.

Collaborative AI: Maintain. I developed this course five years ago. It is one of the largest third year BSc electives (250–300 students). The course content has become largely stable, but this is an emerging topic. I will continue to improve this course based on student feedback and new developments in the field.

Advanced Interdisciplinary AI Project (AI2P): Grow. I developed this course three years ago. It is a 15 ECTS project-based MSc course available to students of all TU Delft faculties. In the first two editions, the student enrollment has been low (10–15). I believe this is mainly because many students (especially those outside the MSc CS program) do not know about it. I am working with the education team to grow this course.

NLP for Society: Develop. I am currently developing this course (to teach in Q4, 2024-2025). With the popularity of Large Language Models (LLMs), NLP is a topic of growing interest among students. In this course, I will teach how to develop and use NLP responsibly and for societal good.

4 Leadership

I strive to develop as an academic leader along three avenues.

Scientific Visionary. Two processes I enjoy immensely are *abstracting* a problem or a solution, i.e., generalizing or simplifying it to the extremes, and *learning* new things. These processes have led me to explore several computing disciplines, including AI (a few of its sub-fields), HCI, and Software Engineering. This exploration has led me to a realization that despite an increasing variety computing sub-fields, we seek to solve very similar fundamental problems. As a concrete example, I have studied goal models, requirements models, and value models. Despite different fields of origin, all these models seek to answer one fundamental question: *what does a user want?*

This abstraction bent has led me to successfully publish visionary ideas as Blue Sky papers, research agendas, and spotlight articles. I will continue this type of work since it can have significant scientific impact.

Role Model for Students. I am where I am today because of the great mentors I have had. Thus, I strive to be a good mentor, myself. The crux of my mentoring philosophy is: *always put students first*. To do so:

- First, it is necessary to understand what motivates a student to act (or not act) in a certain way. I believe that a student would do their best if my expectations align with their motivations. In the STS lingo, this is akin to aligning norms (expectations) with values (motivations). Thus, I always try to elicit a student’s motivations (this is often nontrivial) and carve a direction for the student inline with their motivation.
- Second, I believe in leading by example. This becomes challenging, sometimes. For example, I work very closely with the students on their initial projects, but some students seem to expect the same for all their projects. I am learning to better handle such situations by setting appropriate expectations.

- Third, I do not micro-manage. I have learned, from personal experience, that working independently and making mistakes may slow you down in the short-term but serves well in the long-term. I embrace this philosophy and encourage my students to develop independent research lines, even if that deviates from my own research lines.

Trustworthy. Trustworthiness is an important trait of a leader. I serve on important leadership roles both in the HI consortium and in the broader research communities, e.g., being a Program Chair for a conference. I believe that I am entrusted for such roles because of my diligence—if I accept a role, I will do my best to exceed the expectations.

At this career stage, I have *so much time and so little to do. Wait a minute. Strike that. Reverse it!*⁶. So, I am learning to choose my roles wisely. My ideal roles are those that challenge me and lead to significant impact.

5 Organization

I currently have three key organizational roles. Below, I outline my ambitions for each.

Coordination of MSc in DSAIT. I intend to serve as one of the two Coordinators for the MSc in DSAIT program. I am humbled to have been entrusted with this major responsibility, considering that it is a new program both for the staff members and the students. This is a step-up for me since I have never had such a large educational responsibility. I will do my best support running this program as smoothly as possible.

Delft AI Labs & Talent Management Team. I serve as the education coordinator in this team. One of our main goals is to enhance the AI education capacity of TU Delft. To this end, my team strives to connect AI L&T (supply) with various education teams (demand). However, this is a challenging task since education, across faculties, is administered in various silos. My ambition is to make this an effective and efficient process.

Coordination of AI2P. I coordinate Advanced Interdisciplinary AI Project (AI2P), a 15 ECTS course, available to all TU Delft students. I list this as an organizational role because it requires coordinating with multiple faculties (where students come from) and several AI Labs & Talent (where projects come from). In addition, my ambitions is to bring industry projects (with academic supervision) AI2P. I am working with the Innovation and Impact team on this.

I started this document with a research vision and let me end it with another. Figure 5 shows how a well-designed STS and its members interact. A key aspect of this design is that the norms of the system and the values of the individuals reinforce each other. I believe that this applies to all organizations—effective governance requires a continuous interplay between the organizational norms and the values of the members of the organization.

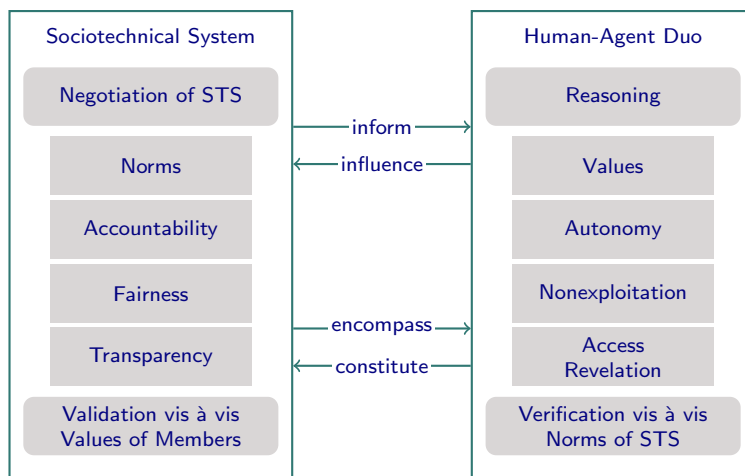


Figure 5: The interplay between an STS and its members

⁶A quote from *Willy Wonka & the Chocolate Factory*