Hybrid Intelligence

Mid-term Self-evaluation

1 Participating researchers →

2 Carrying out the research programme →

3 Organisation structure and policies of the consortium →

4 Knowledge use and transfer activities; interaction with knowledge users →

5 SWOT analysis →

6 Budget situation →

Annexes →

VU Vrije Universiteit Amsterdam · TU Delft · University of Twente. · Utrecht University · Universiteit van Amsterdam · rijksuniversiteit groningen · Universiteit Leiden

Project title:	Hybrid Intelligence (HI): augmenting human intellect
Principal investigator:	Prof. Frank van Harmelen
Dossier #:	024.004.022

Abstract

By creating intelligent machines that interact with humans, the Hybrid Intelligence (HI) Centre aims to give people new, intelligent artificial collaborators for joint reasoning to optimise decision-making and problem-solving. This interaction has the potential to amplify both human and machine intelligence by combining their complementary strengths. Hybrid Intelligence (HI) focuses on the assistive and collaborative role of Artificial Intelligence, emphasising its potential to enhance human intelligence instead of replacing it.

The Hybrid Intelligence (HI) Centre is a collaboration of top AI researchers from across The Netherlands in areas such as machine learning, knowledge representation, natural language understanding & generation, multi-agent systems, collaborative AI, cognitive psychology, multimodal interaction, social robotics, AI & law, and ethics of technology.

The HI Centre operates as a virtual research centre between 6 Dutch Universities: TU Delft, Univ.'s of Groningen, Utrecht, Amsterdam and Leiden, under leadership of the VU Amsterdam. The HI Centre started in January 2020 funded from the OCW/NWO "Gravitation" programme with a budget of 20m € for 10 years. Per 1 April 2023 the HI Centre comprises 44 PhD/postdoc positions, of which 30 are funded from the Gravitation grant, and some 50 faculty members. We have the ambition to be a national and international centre of gravity for research on all aspects of Hybrid Intelligent systems.

Results achieved

Answering the question: "What has been the added value of the Gravitation grant beyond the individual projects?"

- A leading position in the rapidly growing area of Hybrid Intelligence. This is due both to the size of the Gravitation grant, as well as its timing, which was well ahead of the curve.
- Creation of a national community on Hybrid Intelligence. This is due to our consortium's strong integrative and collaborative policies.
- Supporting a European community on Hybrid Intelligence. This is due to our initiatives such as the HHAI conference series, and our alignment with other European initiatives such as Humane AI Net.
- Tangible progress towards HI as a scientific discipline, by publishing a high-impact research agenda, and by developing a shared conceptual basis between subfields that have different approaches and metrics: behavioural vs. technical, symbolic vs. statistical, empirical vs. philosophical.
- A magnet for partners further down the innovation chain (such as TNO) who are interested in the application potential of HI. Again, this is due to sufficient mass under a single organisational umbrella.
- An environment in which young and mid-career scientists are given the opportunity to take leadership roles.

1 Participating researchers

1.1 Attracting personnel and promoting diversity

Attracting personnel

The HI consortium hired the first cohort of 22 PhD candidates and 5 postdocs. Despite COVID adversity, we succeeded in an early start, with almost all of the PhD/PDs starting between the summer of 2020 and early 2021. Additionally, 8 PhD candidates, 1 postdoc (PD) and 5 staff members joined the HI Centre, as an investment from the participating universities. These include two staff members and their 2 PhD students of the University of Twente, who joined in 2022 as a collaborating partner (see our page with 36 Funded Projects). A further 8 PhD projects have associated themselves with our programme with their own funding. This total volume of PhD candidates, postdocs and staff exceeds our initial ambitions. In the first half of 2023 we are hiring the second cohort of 20 PhD/PDs of which 2 already started. In 2025 we plan to hire our third cohort of some 30 PhD/PDs from the Gravitation grant.

Diversity

At the start of the project Prof. Catholijn Jonker was appointed as Diversity Officer, and we formulated a Diversity Policy, included in our HI recruitment procedure. We adopted an ambitious 50/50 gender-balance target for the PhDs and PDs, and strived for ethnographic diversity. We advertised vacancies internationally and we ensured visibility at Women in AI meetings at international conferences and at our partner institutes, and at activities such as the Digivita summer camps. Positions were immediately offered if the top candidate was female, whereas male top candidates only received an offer when gender balance was restored across the cohort. This did not unduly delay the hiring process, and everyone was able to offer the position to the highest-ranking candidate. This resulted in a gender balance of 22 female and 22 male PhDs and PDs across all funded and associated projects. This is a substantially better gender balance than what is typically seen in Computer Science and AI. Furthermore, our teams foster 14 nationalities, attracted from universities around the globe.

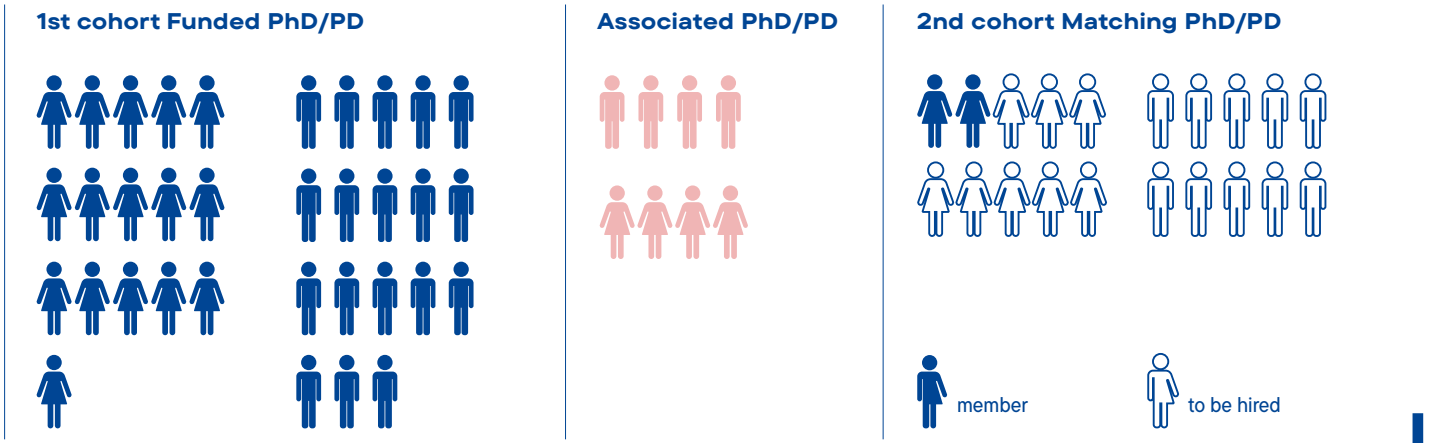

1.2 Fostering new research generations

Our total staff (95) comprises approx. 50% senior scientists (with a PhD from 1998 or before), and 50% early and mid-career scientists. The original (co-) PI(s) and other senior members use their experience to provide opportunities for the mid-career members to advance their career and to keep The Netherlands a leading nation in HI research. Mid-career scientists have been given leading roles in the programme; each of our Research Lines is co-led by a mid-career scientist, all the second cohort projects are written by early- and mid-career scientists, and prof. Yolum, a mid-career researcher, has joined the Executive Board (EB). Such new positions in EB are both a way to offer opportunities and to mitigate original (co)-PIs leaving the programme. To mitigate the risk of overloading our early- and mid-career researchers, we are setting up a mentoring programme, and the management team meets monthly with mid-career researchers in leading roles to support their progress. These policies will put our mid-career scientists in a good position to take over leadership of the field during and after the runtime of the Gravitation grant.

Seven of our early and mid-career staff have seen promotions: Mike Ligthart, Erman Acar and Karine Miras from PD to assistant professor; associate professors Pinar Yolum, Bart Verheij, Antske Fokkens and Annette ten Teije to full professor. We are pleased that 4 of these 7 promotions concern female colleagues. As further signs of the career development of our mid-career staff members, dr. Birna van Riemsdijk was one of the 5 nominees for Women in Dutch AI, dr. Hayley Hung was awarded a prestigious ERC Consolidator grant on using AI systems to infer social intentions from non-verbal behaviour, dr. Eric Nalisnick acquired a prestigious NWO Veni grant on continual learning under human guidance, prof. Annette ten Teije is co-PI on an NWO grant on personalised oncology care that is closely aligned with the HI programme, and Dr. Davide Grossi was awarded a fellowship at the Netherlands Institute for Advanced Study to study "Algorithms for Large-Scale Deliberative Democracy". We are pleased to see a strong female representation among these recipients of grants and awards.

1.3 Second term perspective

In the original proposal, we intended to hire scientific personnel in two cohorts. However, to improve knowledge transfer between the cohorts, and with approval of our governing board we introduced an extra cohort that bridges the two originally planned cohorts. Knowledge transfer will be further facilitated by a "buddy system" in which first-cohort researchers collaborate with second cohort researchers. If successful, we will continue this for the third cohort. All of this will ensure the continuity of our research programme across its 10 year funding period.

Besides the "buddy system" for the second cohort, we are setting up an exchange programme. Furthermore, we establish connections with industry for collaboration and internships with the help of TNO (the Dutch national organisation for applied research) and through the national Innovation Center for AI (ICAI). PhD exchange possibilities are in place with the Humane AI Net funded by the European Commission. This will ensure that our PhD students are well positioned for future career steps either in industry or in international academia.

For our mid-career researchers, we will initiate a mentoring program, and we will continue to offer positions of trust and leadership. The new matrix lines and special interest groups (see 2.3) are led by early- and mid-career researchers. Futhermore, we will include mid-careers in the EB when current members retire or step down (expected in 2024 or 2025), and extend the Management Team (MT) with a mid-career researcher. This will provide our early- and mid-career scientists with valuable research management experience. Because diversity is important for innovation and to create hybrid intelligence for everyone, we will continue to deploy our diversity policy for recruiting. (see Sections 1.1 and 3.9)

2 Carrying out the research programme

2.1 How does the research programme progress?

The overall research programme

At the start of our research programme, we published an agenda-setting paper in IEEE Computer. This paper has been very influential: it received the paper of the year award for the IEEE Computer journal, and has been cited 140 times since it appeared in 2020). The paper operationalised the overall HI research challenge ("*How to build adaptive intelligent systems that augment rather than replace human intelligence, that leverage our strengths and compensate for our weaknesses?*") into four objectives: (i) systems that **collaborate**, (ii) systems that **adapt** to changes in their environment, (iii) systems that can **explain** themselves, and (iv) systems that behave **responsibly** according to awareness of norms and values. Below, we report on the progress on each of these objectives in terms of the "core capability tables" we defined in our original proposal for each of the four Research Lines (**C**ollaborative, **A**daptive, **R**esponsible, **E**xplainable; CARE)

IEEE COMPUTER SOCIETY

Collaborative HI

(prof. Hindriks (VUA) & dr. Hung (TUD))

Aim: To develop HI systems that collaborate effectively with humans by: understanding social contexts, norms and practices; recognizing and reasoning about the abilities, beliefs and goals of other agents through a computational theory of mind; communicating appropriately, including via physically embodied cues and signals.

Progress: work in this Research Line was spread among almost the entire capability table, with award-winning work on establishing shared situational awareness (Putra Manggala et al., IGLU challenge winning entry; Selena Baez best paper at the COLING Workshop on Customised Chat Grounding Persona and Knowledge) and personalised multimodal interaction (Mike Lighthart: HRI best paper; Davide Grosse best paper at the AAMAS conference on computational social choice; Selene Baez, Lea Krause and Taewoon Kim: EMISSOR platform for capturing multimodal interactions, see our code and data page). For studying *group interactions*, a substantial multimodal dataset was collected by Maria Tsfasman and dr. Catharine Oertel (ICMI paper, see the MEMO entry on our code and data page).

Core capabilities
Initiating relationships
Establishing shared situational awareness
Personalised multi-modal user interaction
Collaborative group support

Adaptive HI

(dr. Oliehoek (TUD), dr. van Hoof (UvA) & prof. Plaat (UL))

Aim: To enable systems to adjust to previously unknown situations, changing circumstances, new user requirements and goals, and variable team configurations. To reconcile such behavioural adaptation with requirements on safety, transparency and predictability.

Progress: Progress was made on all four core capabilities. Progress on learning through interaction is illustrated through award-winning work by a team of our PhD students using argument mining (COLING workshop best paper award, see the HyEnA entry on our code and data page) and by Bram Renting using negotiation (AAMAS paper). Results on learning how to interact were obtained by Robert Loftin (studying how two agents that have not previously interacted can settle on a joint solution, ICML paper), and by Erman Acar (showing how two agents that have not been instructed to interact can learn to interact under reward pressure from the environment, see the Space Canons entry on our code and data page). For incremental adaptivity, Andreas Sauter studied how targeted interventions can reduce the uncertainty in causal discovery and Anna Kuzina studied the behaviour of deep generative models in the presence of distribution shift (ICLR paper and two NeurIPS papers). Finally, for integrating symbolic constraints, Niklas Höpner showed how an agent can learn rich meanings of symbols by connecting relations in a knowledge graph to the agent's experience from interaction with the world (IJCAI paper).

Core capabilities
Learning through interaction
Learning how to interact
Incremental adaptivity
Integrate symbolic constraints

Responsible HI

(prof. Verheij (UG) & dr. van Riemsdijk (UT))

Aim: Ensure that behaviour of artificial agents in HI systems is aligned with ethical and legal values and social norms.

Progress: a natural distinction emerged in this Research Line between the questions "how do we want the AI to behave" and "how do we make the AI behave like that". Progress was made mostly on the second question. This includes work on Validating whether legally or morally acceptable behaviour is learned (by Cor Steging on Discovering the rationale of decisions, and by Annete Onnes on Monitoring AI systems), and on Reasoning about the legal or ethical acceptability of intended behaviour, with work by Wijnand van Woerkom on justifications in case-based reasoning, and two empirical studies by Enrico Liscio and Michiel van der Meer: a hybrid method for extracting arguments from opinions (see HyEnA on our code and data page), and a collaborative platform for identifying context-specific values (AAMAS paper, see AXIES on our code and data page).

Core capabilities
Critically examining decisions of big-data applications
Validating whether legally or morally acceptable behaviour is learned
Reasoning about the legal or ethical acceptability of behaviour

Explainable HI

(prof. Vossen (VUA) & dr. Fokkens (VUA) & dr.Murukannaiah (TUD))

Aim: Methods to generate appropriate explanations in different circumstances and for different purposes, even for systems whose internal representations are vastly different from human cognitive concepts. Methodology to evaluate explanations in a scalable and reproducible manner.

Progress: Progress was made on *Transferability of shared representations* with Sharvaree Vadgama's work on Interpretable ECG classifications and on interpretable shapes in a generative setting (spotlight paper at NeurIPS workshop). An example of progress on *Quality of explanations* and on *Interactive explanations* is Lea Krause's work on explanation triggers in human-agent interaction, Erman Acar's results on using hyperbolic manifolds for producing explainable hierarchical image labelling (CVPR paper) and Enrico Liscio's and Michiel van de Meer's work on making values explicit in deliberation processes involving agents (What values should an agent align with?)

Core capabilities
Transferability of shared representations
Quality of the explanations
Interactive explanations

A full overview of our 90+ publications to date is available at our publication page. These include papers at top conferences (AAAI, IJCAI, NeurIPS, AAMAS, ICML, CVPR, SIGIR), including 6 best paper awards, 1 paper of the year award, and 2 challenge-winning entries. Our code and data page lists 7 open source software packages and datasets.

We list below two projects as examples of the research spectrum in the HI programme. The first example (project nr. 8) studies AI algorithms, and aims to extend deep reinforcement learning with human-interpretable background knowledge in order to learn more explainable strategies for agents that execute sequential tasks. The second example (project nr. 35) studies what is important in dialogues between human subjects: what are important moments in such dialogues, and which multi-modal features can be used to recognise such moments. These are just two examples from our 44 projects that are producing new scientific insights into Hybrid Intelligence.

Example 1

Project 08: Hybrid approaches to learn explainable sequential behaviours

Researchers: Niklas Höpner (PhD candidate), dr. Herke van Hoof, dr. Ilaria Tiddi (supervisors)

Project Description: We investigate the combination of (deep) reinforcement learning with prior knowledge from knowledge graphs to learn explainable strategies for complex sequential tasks. Many problems of practical interests are concerned with optimising sequential decision-making, such as finding optimal trajectories for robots or deciding which medical tests to run. Methods for classical symbolic planning are nicely explainable to human collaborators, but rely on an exact problem description; while data-driven (e.g., reinforcement learning) approaches do not rely on a provided problem specification, but are data hungry and inscrutable. We investigate the combination of the complementary strengths of these approaches.

Publication (IJCAI 2022): Enabling reinforcement learning (RL) agents to leverage a knowledge base while learning from experience promises to advance RL in knowledge intensive domains. However, it has proven difficult to leverage knowledge that is not manually tailored to the environment. We propose to use the subclass relationships present in open-source knowledge graphs to abstract away from specific objects. We develop a residual policy gradient method that is able to integrate knowledge across different abstraction levels in the class hierarchy. Our method results in improved sample efficiency and generalisation to unseen objects in commonsense games, but we also investigate failure modes, such as excessive noise in the extracted class knowledge or environments with little class structure.

Example 2

Project 35: Memory for Conversational Agents

Researchers: Masha Tsfasman (PhD candidate), dr. Oertel, prof. Jonker (supervisors)

Project Description: On the path to creating a socially-aware memory system, we are creating a data-driven model to predict what is memorable from conversations for different participants in a meeting. For that, we are exploring different modalities of human-human interactions. This topic is relevant for HI since we are looking into better understanding of humans to advance the quality of long-term human-agent interactions and collaborations.

Publication (ICMI 2022): Previous research on group interactions has addressed topics such as dominance detection, group engagement, and group creativity. Conversational memory, however, remains an unexplored area in the study of group interaction. The ability to track what each participant or a group as a whole finds memorable from each meeting would allow a system or agent to continuously optimise its strategy to help a team meet its goals. In this paper, we therefore investigate what participants take away from each meeting and how it is reflected in group dynamics. We recorded a multimodal longitudinal meeting corpus (MEMO), which comprises a first-party annotation of what participants remember from a discussion and why they remember it. We investigated the use of non-verbal multimodal features to predict what groups are likely to remember. We devised a coding scheme to cluster participants' memorisation reasons into higher-level constructs. We find that low-level multimodal cues such as gaze and speaker activity can predict conversational memorability, including start and end of memorable moments.

2.2 Interaction and coherency between the research lines

Interconnections and research

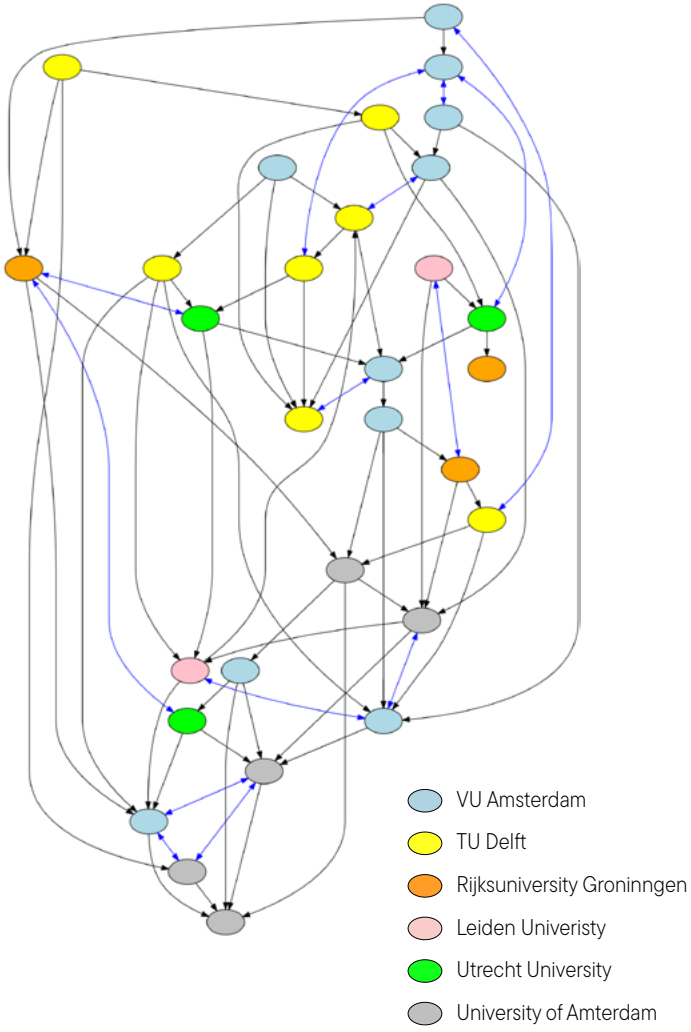

Coherence

We designed for tight interconnections between the Research Lines. Each Research Line combines teams from 4 or 5 different institutions and is jointly coordinated by a senior and a mid-career scientist from different institutions. All PhD students meet bi-weekly (on-line, interspersed with in-person meetings) in a joint seminar/reading group. They discuss literature of shared interest and present their own work to each other to increase awareness. The three courses we organised for our PhD student population (jointly with SIKS, one of The Netherlands Research Schools, see 3.5) cover our Research Lines and help develop a shared knowledge base. All Research Lines organised plenary activities at our consortium meetings and hosted consortium-wide workshops, including workshops at our HHAI 2022 conference. The success of our strategies is illustrated by the collaboration diagram, which lists for all PhD/PDs their three closest collaborators, colours indicate universities. The absence of single-colour clusters in this diagram shows the healthy interconnection between PhDs/PDs at different universities.

To further strengthen cross-line collaborations, we initiated a consortium-wide **Taskforce** for HI Methodologies, Integration and Implementation. This Taskforce ran several activities at consortium meetings on definition and measurement of Hybrid Intelligence, initiated the development of a shared HI vocabulary, and initiated a Hackaton organised by TNO at the HHAI2022 conference.

More information on our activities can be found in our annual reports.

Shared supervision and facilities

A crucial integrative Supervision Policy of the HI Centre is that all PhD students and postdocs have co-supervisors from two institutions. This enhances collaboration both at the level of the PhD candidates and PDs, and at the staff level. The diagram shows the shared supervision network between the institutions (the weight of edges indicates number of co-supervised candidates). We consider this shared supervision policy a crucial ingredient to guarantee real and substantial collaboration between teams located at different universities across the country.

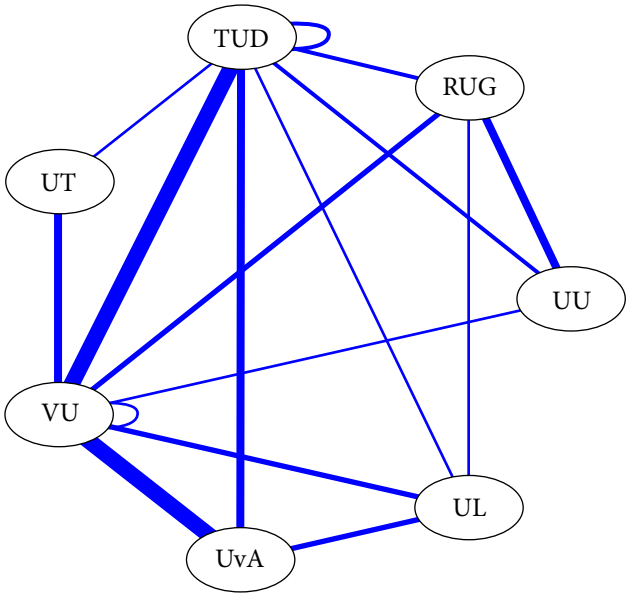

The VU hosts a joint facility used by researchers from the partner universities to jointly work on their research. The Heat map HI meetings on the next page shows the frequency of our scientific meetings (excluding management meetings) substantially increased over time. Across the consortium, and with different participants, there are now multiple scientific meetings each week in which the scientists exchange ideas, knowledge, and insights and brainstorm on their research. We consider these frequent interactions among our researchers, across locations, across research lines and across generations a major indicator of the success of the HI Centre.

Joint activities

See hybrid-intelligence-centre.nl/meetings for more information on the activities below:

Internal meetings

- Kick-off meeting with all researchers involved in the programme in December 2019
- 5 online consortium meetings (COVID enforced)
- 4 face-to-face consortium meetings (2x one day, 2x two days including overnight stay)
- Bi-weekly Friday reading group and colloquium for the PhD students online and regularly face-to-face
- Research Line meetings every 6 weeks for each Research Line
- Monthly meetings between Management Team and all Research Line leaders
- Taskforce meetings across the Research Lines every 2 weeks
- Debate on defining HI vocabulary/Ontology Engineering (F2F)

Meetings with other consortia

- Workshop with ESDiT “The good, the bad and the ugly” (online)
- Workshop with ESDiT on Agency (online)

National courses

- Graduate School course on Neuro-Symbolic AI (Combining Learning & Reasoning)
- Graduate School course on Explainable AI
- Graduate School course on Responsible AI

Community events

- Debate on embodiment of AI (F2F)
- First international conference on Hybrid Human-Artificial Intelligence (HHAI 2022), Amsterdam
- Second international conference on Hybrid Human-Artificial Intelligence (HHAI 2023), Munich

Heat map HI meetings

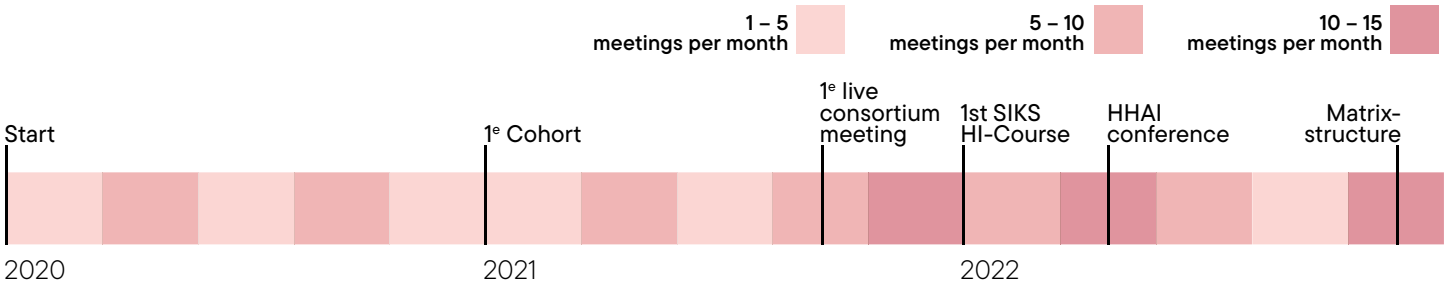

2.3 Second term perspective

At the end of year 3 (Fall 2022) we revisited the original organisational structure to optimise it with respect to the four CARE research objectives of the project: to understand and develop systems that **collaborate**, **adapt** and **explain** themselves and do so in a **responsible** way. We assessed the progress against the capability tables of the original proposal and evaluated the continued suitability of these tables as the right milestones to pursue. This process led to a new structure, designed during our Fall 2022 consortium meeting, and was driven by early- and mid-career researchers. By putting them in the driver seat we provided opportunities for these researchers to co-lead the HI Centre and the HI research field in the coming years. The objectives and capability tables were kept, but progressive insight showed that a matrix structure in which researchers can participate in more than one focus point would be beneficial.

This new Matrix structure, effective per February 2023, is an evolution of the original four CARE Research Lines into three perspectives: (i) four core scientific HI *Challenges*; (ii) which are crosscut by *Special Interest Groups (SIG)* that focus on specific AI techniques, and (iii) four *Case Studies* to anchor our research in concrete societal situations. The Taskforce (see 2.2) continues its integrative work and coordinates the Case Studies (Dr. Pradeep Murukannaiah (TUD)). Per February 2023, the Taskforce is supported by a dedicated PD (dr. Davide Dell'Anna (UU)) and led by a team of our mid-career researchers.

CARE: The new Matrix structure gives us a scalable organisation to accommodate the growth of the HI Centre with the 2nd cohort, while still addressing the CARE objectives we identified at the start of the programme.

- The research objective of *collaborative* systems will be addressed through the Challenges on *Collaboration & Synergy* and *Dialogue* as well as the new SIG on *Deliberation and Argumentation*.
- The SIG on Reinforcement Learning is dedicated to the most relevant machine learning technique for *adaptive* HI systems.
- The new Challenges on *Systems Design in Context & User Literacy* and *Assistants & Trust* will focus on *responsible* systems. The connection with ethicists will be significantly further strengthened, as both Gravition-funded projects ESDiT and HI have committed to intensifying the HI-ESDiT collaboration in the *Ethics* SIG.
- *Explainability* will be addressed through the new Challenge on *Dialogue* as well as through the SIG on *Deliberation & Argumentation*.

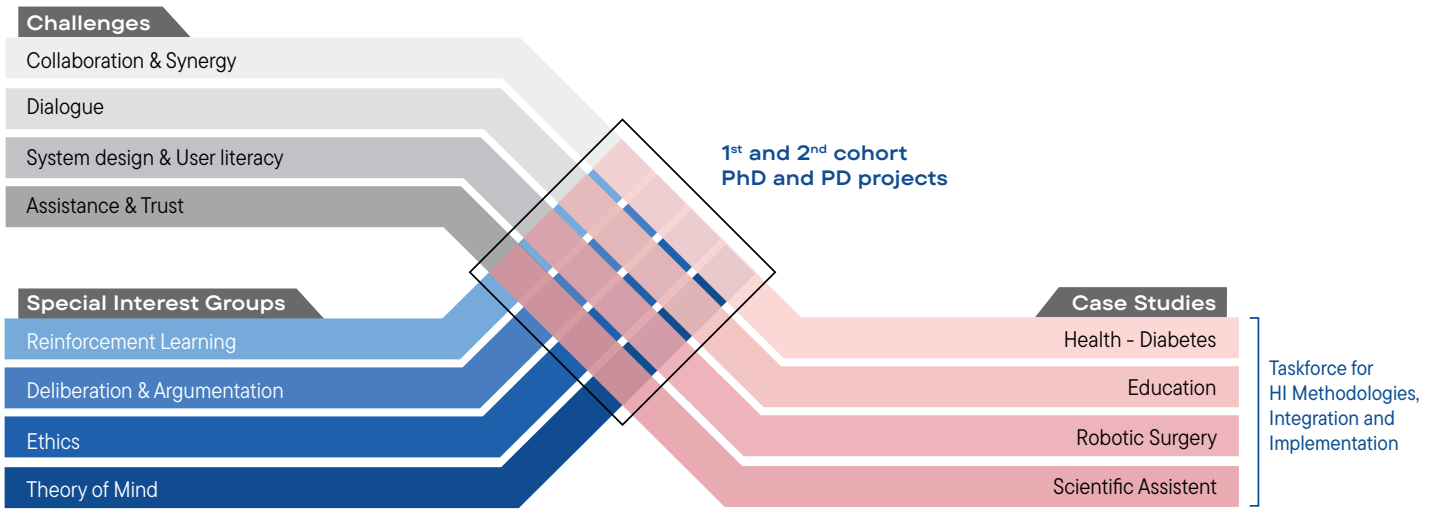

All of the second cohort projects connect to one of the new Challenges, and to at least one of the Case Studies, and the majority of our PhDs and PDs participate in at least one of the SIGs, while others are preparing to propose new SIGs. This shows that the Matrix is already functioning as a fruitful organisational structure for the next phase of the Hybrid Intelligence Centre.

The recent turbulent developments in AI around Large Language Models (LLMs) has also had its impact on the thinking in all Challenges and Case Studies, as well as in the SIGs on *Deliberation & Argumentation* and *Reinforcement Learning*. Our consortium day in February 2023 was partly dedicated to the impact of LLMs on HI; we explicitly asked the leaders of the second cohort projects to reflect on these recent developments, and we are currently writing a white paper on the impact and challenges of LLMs for Hybrid Intelligence.

Case study Health - Diabetes: Diabetes is a chronic lifestyle related disease. A change of lifestyle requires intensive, personalised support involving both the patient and their social environment. A hybrid intelligent coach should help diabetes patients to adopt a healthier lifestyle while at the same time lowering the workload of healthcare professionals. A key challenge is the creation of multi-party dialogues between HI system, patient and healthcare professional that create long-term engagement. Explainability is crucial in this domain.

Case study Education: The aim is to develop a Hybrid AI robot tutor and to evaluate this tutor in field experiments at schools with teachers and students. The robot tutor should engage and motivate, using techniques such as interactive storytelling and mental models. A key challenge is to create long term personalised interactions, using a memory and retrieval mechanisms that should adapt to students, provide explanations and give feedback in a responsible way.

Case study Robotic Surgery: The aim is to create a Hybrid Intelligent system in which a robotic microscope and a surgeon complement each other's skills on the level of micrometre surgery. The robotic microscope has to learn how best to align its angle and zoom of vision with the activities of the human surgeon. This requires a mutual understanding of each upcoming surgical procedure, which is to be acquired during practice sessions. The scientific challenges are human-robot dialogues, theory of mind, shared mental models, shared planning, and collaborative ontology determination.

Case study Scientific Assistant: The aim of the scientific assistant is to support one or more steps in the scientific method cycle: formulating research questions, analysing the literature, formulating a hypothesis, designing an experiment, analysing data, and drawing conclusions. This will require a combination of symbolic and subsymbolic AI techniques, ranging from domain ontologies to deep learning, as well as theory of mind and shared planning to support the collaboration.

2.4 Future research embedding

Beyond 2028, we aim for a sustainable Hybrid Intelligence Centre, with (i) integration in undergraduate and graduate teaching, (ii) strong connections with national and international research institutes and (iii) collaborations with industry to integrate the HI principles and technology into practice. As funding will have ended, the HI-Centre should be self-supporting and financed by the HI-partners and second and third money streams. Progress is made on all these ambitions:

(i) Integration with undergraduate and graduate teaching programs is progressing, with BSc and MSc programs at VU, and through the national SIKS graduate school.

(ii) Internationally, collaboration with Humane AI Net has been established, and discussions are initiated on continuation beyond the grant period for both programmes. Another potential future connection is the Swedish Wallenberg Foundation's AI programme led by Prof. Virginia Dignum that is also considering its future after 2030.

(iii) Finally, a firm collaboration has been established with TNO which is expected to lead to contacts with further industrial partners. Collaboration has also been initiated with Erasmus Medical Centre. Further ambitions are formulated in Section 3.6.

The above progress on teaching, research collaborations and industrial collaborations shows that even at this stage of the consortium we are actively working towards sustainability of the HI Centre beyond the Gravitation funding period.

3 Organisation structure and policies of the consortium

3.1 Organisation structure and management

The HI organisation is structured in two levels: strategic management and operational management (see figure in 3.2):

At the strategic level the entire programme is overseen by the **Governing Board** (GB), consisting of the Rectores of each institution. It is chaired by the Rector of the VU, and meets annually. The **Executive Board** (EB) is the highest decision making body of the programme, with a representative from each founding institution. GB and EB are supported by the international **Scientific Advisory Board** (SAB, see section 3.4). The EB meets monthly and is chaired by the programme coordinator, Prof. Frank van Harmelen, with Prof. Catholijn Jonker as vice-coordinator. The EB monitors scientific progress, and decides on the scientific, organisational and impact strategy of the programme. Any changes affecting the Consortium Agreement (redistribution of budget, addition of new partners) require approval of the GB.

At the operational level, the **Management Team** (MT) consists of the programme coordinator, the vice-coordinator, Dr. Stefan Schlobach, and Dr. Wendy Aartsen from Hands4Grants b.v. The MT supports all activities of the consortium, and its members communicate on a daily basis, meet every week, as well as quarterly in additional strategy meetings.

The scientific operational management is in the hands of the Research Line coordinators and the Taskforce for HI Methodologies, Implementation and Integration. The Taskforce covers the Case Studies and the collaboration with external partners, replacing the Professional User Panel that we envisaged in our grant proposal. Monthly meetings are organised to coordinate between all Research Line coordinators. As of February 2023, the new Matrix Line structure supplants the Research Lines, see 2.3, the meetings with coordinators continue. Our management strategy shows power of flexibility and the capacity to foster engagement at all levels of scientific seniority.

3.2 Organogram of the organisation structure

International Scientific Advisory Board

Governing Board

VU Vrije Universiteit Amsterdam · TU Delft · Utrecht University · Universiteit van Amsterdam · rijksuniversiteit groningen · Universiteit Leiden

Executive Board

Frank van Harmelen, Catholijn Jonker, Piek Vossen, Pinar Yolem, Rineke Verbrugge, Joost Batenburg, Maarten de Rijke

Management Team

Frank van Harmelen, Wendy Aartsen, Stefan Schlobach, Catholijn Jonker

Research Lines

Collaborative
Koen Hindriks, Hayley Hung

Responsible
Bart Verhey, Birna van Riemsdijk

Adaptive
Herke van Hoof, Aske Plaat, Frans Oliehoek

Explainable
Pradeep Murukannaiah, Piek Vossen, Antske Fokkens

3.3 Monitoring progress, risks and conflicts; quality assurance

Core capabilities	Level 1	Level 2	Level 3
Learning through interaction	Ask humans for suitable and sufficient feedback	Rich interaction with one human partner	Rich interaction with multiple AI and human partners
Learning how to interact	Predict what human partner knows, wants and needs (ToM)	Anticipate how human reacts in human-agent collaboration	Anticipate how humans react in larger, mixed teams
Incremental adaptivity	Detect uncertainty of performance due to changing situations, goals and preferences; take suitable action	Predict likely upcoming situations, switching between pre-learned models	Online learning for predicted and surprising changes in situations, goals, and preferences
Integrate symbolic constraints during learning	Identify and ask for meaning of new symbols	Learn rich meaning of symbols	Shared meaning and use of symbols

Progress monitoring

We monitor scientific progress through the "HI capability tables" provided in 2.1, and already present in sections 2.2.1-2.2.4 of the original proposal. These four tables, one for each Research Line, each identify a number of "core capabilities", with three levels of increasing sophistication for each capability.

No widely accepted metrics for HI exist yet, and an entire consortium meeting (July 2021) was devoted to "how to measure Hybrid Intelligence" through metrics, benchmarks and test-environments. Part of a consortium meeting a year later (September 2022) was devoted to scoring each of the projects against the relevant capability table. The Taskforce with participants across all Research Lines has the development of metrics as part of its remit, and we have appointed dr. Davide Dell'Anna as a PD with the specific task to develop a shared HI methodology and metrics.

Annual progress reports are written by each Research Line. Daily monitoring of progress in each of the individual projects is the responsibility of the supervisor and co-supervisor, based on our supervision policy, as well as the supervision standards and policies of the participating universities. All projects submit an annual progress report (2021). Beyond the metrics per Research Line and per individual project, in the coming years we will increasingly use our four case studies to execute lab- and field-experiments with integrated HI systems.

Risk management

We maintain a risk registry that is reviewed and updated by the MT. Operational risks and issues are addressed by the MT, potential high impact risks and their mitigation strategies are discussed with the EB. One particular risk is caused by several senior members nearing the end of their career, and withdrawing from EB and MT. We are mitigating this risk by already having mid-career scientists into leading positions in EB and as leaders of Challenges, SIGs, and Case Studies. A particular financial risk is caused by the increasing gap between current salary levels and the original grant, depending on indexation by the funding agency.

Highest risks from Risk Registry

Title	Likelihood	Impact	Risk level	Strategy
Upcoming changes in key individuals	5	3	15	Senior PI members that are nearing the end of their careers will withdraw over time from the EB and MT to make room for mid-career people. The consortium has a wealth of talented mid-career researchers to provide ample opportunity to take over from these senior PIs.
Lack of compute power	4	3	12	This risk is brought forward due to the increasing usage of data-driven machine learning, and the environmental concerns associated with high-performance computing. All partner institutions (still) have enough compute power available, in fact, it is a shared concern with our partner institutions and dealt with at that level and with applications for grants by the computer science departments of the universities in The Netherlands. Therefore, for now, our strategy is to monitor the situation.
Lack of high quality data	3	4	12	Careful selection of case studies, insisting on early commitment of collaborating institutions and companies to ensure that data will be available. For the individual projects this is left to the responsible project supervisors. Set up formal terms of reference and agreements.

Quality assurance

The HI consortium has an extensive Impact Plan with the aim to assure that a) researchers are recruited in an ethical way in line with our diversity statement, b) the quality of the programme's output will be in line with the project strategy, c) the output is effectively disseminated and communicated to engage with the programme's stakeholders. To assure quality, the programme's strategy is being updated following the SAB and EB meetings. An example of such an update is the introduction of mentoring programmes for both junior and mid-career scientists. The related policies e.g. supervision policy, have been generated and approved by the EB and based on these policies we created a project manual, several procedures, e.g. the recruitment procedure, and a communication plan to engage with stakeholders mentioned in our stakeholder analysis.

Data management

A Data Management Plan is available at consortium level and submitted to NWO. Furthermore, conform our data management Policy, all researchers are requested to generate a DMP per project following the institutional guidelines as each individual project has distinct needs for data securing and handling.

Research ethics

Our ethics policy states that research projects gathering empirical data should receive prior approval by the ethics committees of the respective university. It is the responsibility of the project supervisors to ensure that prior approval has been obtained. Note that this is ethics at the operational level. Ethics as a research topic is led by the SIG Ethics.

Composition of the international scientific advisory board (SAB)

Name	Affiliation	Expertise
Prof. Roni Reiter-Palmon	Univ. of Nebraska at Omaha	Organisational psychology, creativity in teams
Prof. Wendy MacKay	INRIA and the Univ. Paris-Saclay	Mixed human-machine systems
Prof. Ana Paiva	University of Lisbon	Social agents
Prof. Paul Lukowicz	DFKI Germany	Leader of the European **Humane AI Net**
Prof. Yolanda Gil	ISI and Univ. of Southern California	Intelligent interfaces for knowledge capture and discovery
Prof. Petra Ahrweiler	Johannes Gutenberg University Mainz	Social consequences of technologies and innovations
Prof. Juanzi Li	Tsinghua University	Knowledge engineering

3.4 Scientific and organisational oversight

The Governing Board and Executive Board are supported by a multidisciplinary Scientific Advisory Board (SAB), consisting of **7 members** from the EU, US and China. The SAB meets once a year, provides strategic advice, and ensures that Research Lines and results are critically tested. Due to COVID-restrictions, the first SAB meeting in 2021 was fully online. The SAB meeting in 2022 was co-located with the HHAI conference, with 3 members physically present, and the other members online. Topics that were discussed included the balance between science and case studies, bottom-up vs. top-down organisation, the Research Lines vs. the risk of silos, the use of demonstrators as experimental testbeds, the importance of experiments involving human subjects, the effectiveness of hackathons, diversifying the training programme by career stages, and the unique selling points of the HI Centre. The minutes are available online (**2021**, **2022**). The SAB met online in March 2023 to provide feedback on this mid-term review report.

We will install a Societal Advisory Board when our case studies reach a minimal level of maturity. We expect this to be the case for the first case studies in 2024. Members of this board will reflect the domains of our case studies.

3.5 Contribution to training of talented researchers

To train and educate our PhD candidates and postdocs, policies have been designed and are overseen by our Education Director, Prof. Rineke Verbrugge.

Intake: Verbrugge interviews every PhD student together with their supervisor(s) in their first year. The goal of these meetings is two-fold: first, to support the PhD candidate and supervisors in designing a training plan that complies with local policies and HI policies; and second, to ensure that PhD candidates are well supervised and to establish herself as a person of trust.

PhD meetings: Verbrugge organises bi-weekly PhD meetings, first online (due to COVID), later intermittently online and on-site. The goal of those meetings is to create a common conceptual basis as well as a social community among PhD students from diverse fields ranging from CS to philosophy, and to train them in leading group discussions about HI-literature. As work progressed, PhD students also gave presentations about their research. Students receive training credits (represented on annual personal certificates) for participation, leading a reading group session, presenting their research, and co-organizing events. The consortium-wide HI meetings and Research Line meetings regularly contain training elements, such as possibilities for presentation and collective problem-solving.

The appointment of a dedicated and active Educational Director, the intake interviews and the bi-weekly PhD meetings have resulted in a PhD cohort that is both socially coherent and that has a shared understanding of Hybrid Intelligence.

Courses: Our PhD students follow courses in the The Netherlands Research School SIKS (School for Information and Knowledge Systems) and the Kurt Lewin Institute (Social Psychology). There is a close collaboration with SIKS: most PhD students are members and follow several SIKS-courses a year. In addition, three bi-yearly courses have been developed by HI seniors on Neuro-symbolic AI, Explainable AI, and Responsible AI. Currently, a fourth HI + SIKS course on Adaptive AI is being developed, and a new course on Collaboration & Synergy is foreseen. At HHAI 2023, together with ESDiT, we will provide a tutorial Exploring Responsible AI by Design.

With three courses already developed and three further courses foreseen, we are well under way towards providing a substantial body of graduate-level educational material, which will be an important asset when training new generations of HI students and researchers.

3.6 National and international research embedding and collaboration

National embedding and collaboration

- We extended our national collaboration when the University of Twente joined the consortium in 2022 with 2 self-funded PhD students and staff members. A similar expansion is currently being discussed with the University of Maastricht, which would extend our membership to 8 of the 12 Dutch universities.
- We closely collaborate with the Dutch national graduate School for Information- and Knowledge-Systems (SIKS). We developed three new courses on neuro-symbolic AI, explainable AI and responsible AI, to be taught bi-annually, the standard course frequency in SIKS so that new generations of PhD students can benefit.
- We initiated a firm collaboration with the Gravitation-funded project "Ethics of Socially Disruptive Technologies" (ESDiT): Prof. Ibo van der Poel from ESDiT co-supervises a HI PhD student. ESDiT post-doc dr. Michael Dale is co-supervised by Prof. Jonker (member of both HI and ESDiT) and Prof. IJsselstein (ESDiT). Dr. Dale is a member of our Taskforce and participates in our consortium meetings. We organised two shared workshops with the ESDiT consortium, in December 2020 and October 2021.
- We are collaborating with TNO on the use of HI to support Diabetes II patients and their social environment. Dr. Jasper van der Waa from TNO has been added as a member of our Taskforce.

As a result of these activities, we have succeeded in placing Hybrid Intelligence firmly on the Dutch research landscape, including a prominent place in the Dutch national AI research agenda AIREA-NL

National ambitions

- We are extending our alliance with TNO into a signed partnership and expect to extend the current collaboration on Diabetes II management to other case studies with third parties who collaborate with TNO.
- We initiated a collaboration with Dr. Dalibor Vasilic, a plastic surgeon specialising in microsurgery at the Erasmus medical centre in Rotterdam. He uses a robot assistant for training and operational purposes, and we will work with him and his team to increase the collaborative capabilities of that robot assistant in a pre-operation training setting.
- We will initiate a collaboration with the Kurt Lewin Institute, a graduate school for Social Psychology, to expand the training opportunities for our PhDs and to reach out to the young generation of social psychology researchers.
- We will establish collaborations with industrial and governmental organisations through ICAI labs, LTP projects, and case studies through our Taskforce. This will strengthen our potential for securing second and third money stream projects and provide an ecosystem for hybrid intelligence that is open for BSc, MSc and PhD students for internships.
- Strategic collaborations with other consortia will be established to strengthen our scientific basis in disciplines such as the social sciences and life sciences. A particularly relevant candidate is the Gravitation-funded consortium ALGOSOC, of which three of our scientists are members.

International embedding and collaboration

- The first international conference on Hybrid Human-Artificial Intelligence (HHAI) which we hosted in Amsterdam in summer 2022 attracted 260 participants, including 90 from other European countries.
- We set up an internship agreement with the EU-funded Humane AI Net arranging that our PhD students can apply for Humane AI Net's micro-project funding for an internship of 6 months at one of their 53 academic and industrial partners across Europe.
- We signed letters of support and letters of intent to collaborate with a number of proposals for related initiatives across Europe (Aachen, Stuttgart, Copenhagen), to be followed up when grants are awarded.

International ambitions

- We aim to establish the Hybrid Human-Artificial Intelligence (HHAI) conferences as an annual international conference series. The 2023 edition is co-chaired by dr. Tiddi from the HI consortium and is hosted by DFKI in Munich. The HHAI 2024 edition has already been secured. All this shows that the HHAI conference series is a fruitful new addition to the international HI research community.
- Profs. Virginia and Frank Dignum, co-founders of the Hybrid Intelligence Centre, have moved to Umeå University, where they are leading a large Swedish initiative funded by the Wallenberg foundation, which is closely related to Hybrid Intelligence. We will set up an exchange program with this initiative for PhD, PDs and staff members. Furthermore, they have agreed to organise HHAI2024.
- We aim to engage with new and closely related initiatives in Europe (Oxford, Aarhus, St. Gallen, Bielefeld, and Oulu) and the US (Microsoft, Stanford).
- A longer term option that we are considering is to establish a pan-European HI network of academic and industrial partners. Possible partners for this are the Humane AI Net programme, its requested successor proposal in which we play a role, as well as the aforementioned Swedish programme.

3.7 Open Access publishing

Our project manual stipulates the requirement to implement the NWO Open Access policy for all publications. The manual guides our researchers to three indices on Open Access policies of publishers. Our researchers have greatly benefited from the Open Access agreements between Dutch universities and publishers. As a result, all our 90+ papers have been published as Open Access, in either green or grey form, with a small number under a temporary Open Access embargo. The success of the Open Access arrangements by the Dutch Universities meant that in the first three years we didn't spend any of the funds for Open Access fees, and we have re-allocated this funding towards other HI activities. All publications are listed on the HI Centre's website, including links to the full text of the publications. This shows our full compliance with NWO's Open Access publishing policy.

3.8 Attention to data management, ethical aspects of the research

All researchers comply with local university regulations, and our supervision policy states our compliance with The Netherlands Code of Conduct for Research Integrity. Experiments with human participants require approval from the ethics board of the local university. A data management plan is in place for every project, in line with our data management policy and our ethics policy. All PhD candidates receive training in research ethics at their local university.

Data management & Ethical aspects

Every project has submitted a Data Management Plan following the NWO template. We maintain a "code and data" collection, but publication of datasets is always through the institutional repositories of the participating universities, and must comply with local policies. Similarly, ethical approval is arranged through the relevant ethics board of the participating universities. All experiments carried out in the first period of the project have asked for (if applicable) and obtained approval by their ethics board and have worked with Data Stewards of the partner organisations.

Our lab experiments produced the HyEnA Covid corpus and the Memo corpus. The code underlying the HI annotation methods of AXIES and HYeNA and the pointers to our data sets are published on our Code and Data page.

3.9 Attention to (gender) diversity

In section 1.1 we described our successful procedure for hiring a diverse researcher population, resulting in an equal male/female balance, which is substantially better than the field average.

We will deploy a similar procedure when recruiting our second cohort, with added attention to a wider notion of diversity.

All efforts regarding diversity are coordinated by our Diversity Officer Prof. Catholijn Jonker. We pay special attention to invited speakers at HI workshops and consortium meetings to ensure diversity. When appointing positions of leadership the Diversity Officer checks gender balance and overall diversity. The diagram shows that our gender balance is close to 50/50 across the different leadership roles, which is again substantially better than the field average. In 2023 we installed a Diversity Committee (Dr. Pradeep Murukannaiah, PhD candidate Annet Onnes) to support the work of the Diversity Officer.

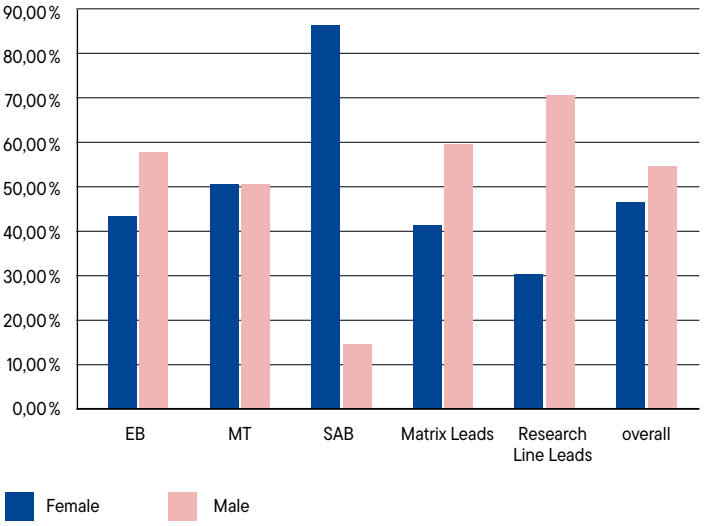

4 Knowledge use and transfer activities; interaction with knowledge users

To promote the spread of knowledge to societal and economical stakeholders, we adopted a communication plan which addresses publications, meetings, events, and other activities to interact with target groups of scientists, societal institutions, the general public, industry as well as other public and private organisations. Besides, the consortium is actively pursuing collaborations with parties that can support the spread of HI insights and results (see 3.6). In 2022 our website has daily attracted 400+ visits.

Scientists in related fields

The strategy for communication is as usual through

- Publications in scientific journals and proceedings of conferences and workshops. See our publications page, listing 90+ publications over the first 3 years.
- Presentations at conferences and workshops. See our page on keynotes and invited presentations, listing 15 keynotes and invited presentations in 2022 alone.
- Releasing open source code and data, see our Code and Data page, listing 7 open source software packages and datasets that we released in the first 3 years. All available from the website that has between 400+ visitors a day in the last 12 months.
- Scientific workshops, tutorials (e.g., on Exploring Responsible AI by Design) and courses offered together with The Netherlands Research School SIKS. See our education page, listing courses we developed and taught for SIKS.
- Through our personal networks, e.g., this is how the connections with HI-ESDiT and Humane AI Net were created, and how we are considering collaborating with other consortia such as the Gravitation-funded projects ALGOSOC, and Stress in Action, and the Wallenberg programme in Sweden.
- We reach out to and collaborate with scientists in related fields, for more information see 2.1 and the Annex.

General public

We actively reach out to the media to contribute or initiate a debate on HI related topics, and journalists successfully reach us to contribute to a topic in the media. Our Media page lists 28 contributions, including podcasts, interviews in national newspapers, articles in trade magazines and appearances on national television. Some highlights are the AI.nstein video on YouTube, the Boerhave podcast on Spotify, a contribution to the Dutch National AI course, and an interview on Tweakers.net. An upcoming popularising science book on "AI as turbobooster" devotes an entire chapter to our programme.

Governmental institutions

Our researchers participate in sessions organised by governmental institutions. We presented twice for committees of the Upper House in Parliament. We participate in ELSA labs (Ethical, Legal and Social Aspects of AI) that collaborate with municipalities. This creates the opportunity to inform policy makers on the HI perspective on Artificial Intelligence. Hackathons are a further way of interacting with other institutions and industrial partners. Our first hackathon, H3AI at HHAI 2022, was hosted by TNO, and sponsored by Netwerk Mediawijsheid, by Nieuws Checkers, and by Nederlands Forensisch Instituut.

Industry

Industrial collaboration is primarily fostered through our case studies (see 3.1). Our collaboration with TNO is intended to connect us to other industrial partners. Similarly, the Robotic Surgery case study involves a company providing the software for interacting with the robotic arm. The intended close collaboration with TNO will allow HI researchers to anchor their research in industrial case studies, to increase the applicability of (parts of) our research. It will simplify co-organisation of further hackathons, create internship opportunities for HI researchers, and help to connect more industry partners to the HI Programme. We expect a strategic collaboration to be signed in Spring 2023. Moreover, we will investigate collaborations with the ICAI Labs in the recently funded ROBUST programme through our co-PI Prof. de Rijke, who is leading the ROBUST programme, and through members of our consortium, e.g, Dr. Frans Oliehoek as co-founders of the Mercury Machine Learning ICAI Lab with Booking.com.

Private & Public organisations in application domains

Through our individual researchers and through our collaboration with TNO we have strong ties with health care (hospitals, psychiatric care institutions, RIVM, helplines), incident management - search and rescue (police, fire departments, ambulance), education (elementary schools), and governmental organisations (municipalities, ministeries, Juridisch Loket). Each of these domains of application provides a wealth of opportunities for (future) knowledge use and transfer, comparable to our initial case studies Health - Diabetes, Education and Robotic Surgery.

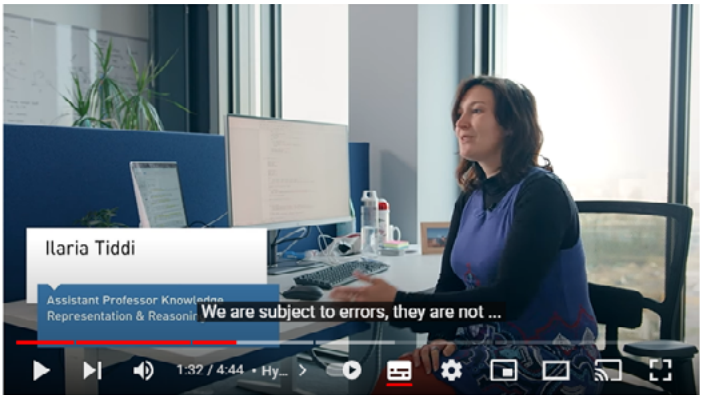

5 SWOT analysis

Strengths

1 Team

- Diversity of our consortium in gender, nationality and scientific generations
- Engagement across all generations
- Growth beyond the original grant
- Three cohorts instead of two

2 Governance

- Strong strategic management with high engagement
- Bottom-up joint research vision development
- New generation leadership
- Dual supervision structure
- Institutional commitment

3 Research & Impact

- Healthy volume and high quality published work (prizes, awards)
- Grant acquisition and promotions by early- and mid-career staff
- Coherence between projects is increasing
- HHAI conference series initiated
- Strong outreach beyond scientists

4 Education

- Dedicated Education Director
- External courses for The Netherlands Research School
- Bi-weekly reading Club & colloquium with great attendance and junior leadership

Weaknesses

1. Current Management Team is reaching its limits because of consortium growth
2. International outreach programme has not been defined yet
3. Difficulty of aligning projects towards Case Studies
4. Low on expertise in HCI, Cognitive Psychology
5. Low awareness of the impact of HI research and deployment on the environment

Opportunities

1. Strong interest in HI: greater visibility, easier recruiting, opportunities for growth
2. Increasing demand for human-centred AI (industry, academia, education)
3. International interest in an HI education programme
4. Many new international academic initiatives to form new partnerships
5. New developments in AI (e.g., Large Language Models as an HI-enabling technology)

Threats

1. Loss of front-runner position and unique selling point due to global increase in related initiatives
2. Fast moving external world of AI
3. Increasing pressure and stress on all academic levels
4. Funding will be eroded by inflation-induced wage increases

Actions needed & intended

- W1: Bring management capacity and budget back in line with the size of the consortium
- W2: Invest time in international networking and outreach
- W3: Continuation of our efforts to bring science and case studies together
- W4: Attract expertise in HCI and Cognitive Psychology through joint funding applications and strategic collaborations with other consortia
- O1-O3: Apply for funding for international training & teaching programme
- T4: Renegotiate the size of third cohort if the grant will not be inflation-indexed.

6 Budget situation

Personnel (in FTE and k€) (amounts rounded per cell)

The total grant in 2019 was k€19874, which became k€20381 after indexation in 2022. The first half year (01-01-2020 start date) has been dedicated to writing the 1st cohort proposals and recruiting. Due to the pandemic, spending on some parts of the non-personnel HI budget has been somewhat delayed while the staff budget is on schedule. The first request for programmers has been approved in 2022 and will be spent in 2023. **Unclarity around the indexation and the current inflation has forced us to forecast the budget on the most negative scenario for the coming years.**

Science staff	FTE cum	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	Total
Postdocs	5.6	57	140	233	444	533	533	533	533	533	533	40705
PhDs	48	228	1011	1193	1193	1193	1193	1072	1072	1072	1072	10300
Education coordinator	0.6	31	34	35	37	39	41	43	45	47	49	401
Total scientific staff	54.2	315	1185	1462	1674	1765	1767	1648	1650	1652	1654	14771

Non-science staff	FTE cum	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	Total
Manager	0.9	42	12	39	40	40	40	40	40	40	40	373
Programmers		-	-	-	251	351	351	295	234	234	234	1950
Knowl. transfer & dissem. coord.	0.9	21	12	20	30	30	40	40	40	40	40	313
Total non-science staff	1.8	62	24	59	321	421	431	375	314	314	314	2634
Total staff	56	378	1209	1521	1995	2186	2198	2023	1964	1966	1968	17406

Investment costs (in k€) (amounts rounded per cel)

So far most of our work has been possible using existing infrastructures. We however foresee the risk of additional costs to extensive computing expenses. This risk is logged in the risk registry.

Description	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	Total
Servers	-	-	-	21	21	21	21	21	21	21	150
Robots	-	-	-	21	21	21	21	21	21	21	150
Total investments	-	-	-	43	43	43	43	43	43	43	300

Other costs (in k€) (amounts rounded per cel)

Due to the corona pandemic most of the budget for travel and physical meetings had to remain in the account.

Description	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	Total
Travel and subsistence	0.2	0.5	31	66	82	95	95	95	100	100	633
Communication/dissemination	-	-	-	60	85	85	85	85	85	85	570
Training/meetings/symposia	5.2	8	58.2	34	40	40	50	50	60	60	279
Exchange progr.	-	-	-	12	18	18	18	18	18	18	117
Consumables	9.8	7.5	41	25	25	25	25	25	30	30	185
Experiment fees	-	2.4	-	16	22	22	25	25	25	25	160
Advisory boards	-	-	-	10	12	12	12	12	15	15	88
Competitions	-	-	-	33	40	40	50	50	60	60	333
Total other costs	15	18	130	256	324	337	360	360	393	393	242

Annexes

Annex 1
Joint key publications

(See our **Publications** page for the full collection of joint publications)

1 Z. Akata et al., "A Research Agenda for Hybrid Intelligence: Augmenting Human Intellect With Collaborative, Adaptive, Responsible, and Explainable Artificial Intelligence", IEEE Computer, vol. 53, iss. 08, pp. 18-28, 2020. (**best paper award**) **https://ieeexplore.ieee.org/document/9153877**

2 M. Ligthart, M. Neerincx, and K. Hindriks, "Memory-Based Personalization for Fostering a Long-Term Child-Robot Relationship", in ACM/IEEE International Conference on Human-Robot Interaction, HRI 2022, Sapporo, Hokkaido, Japan, March 7 – 10, 2022, 2022, pp. 80–89. (**best paper award**) **https://repository.tudelft.nl/islandora/object/uuid%3A99db22f6-19cb-4d21-85c2-bc1250a9fa01**

3 T. Kim, M. Cochez, V. François-Lavet, M. Neerincx, and P. Vossen, "A Machine with Short-Term, Episodic, and Semantic Memory Systems", in Proc. of the **AAAI** Conf, 2023, to appear. **https://arxiv.org/abs/2212.02098**

4 D. Grossi, W. van der Hoek, and L. B. Kuijer, "Reasoning about General Preference Relations," **Artificial Intelligence**, vol. 313, iss. C, 2022. **doi:10.1016/j.artint.2022.103793**

5 R. Loftin and F. A. Oliehoek, "On the Impossibility of Learning to Cooperate with Adaptive Partner Strategies in Repeated Games," in International Conference on Machine Learning (**ICML**), 2022, pp. 14197–14209. **https://arxiv.org/abs/2206.10614**

6 B. M. Renting, H. H. Hoos, and C. M. Jonker, "Automated Configuration of Negotiation Strategies," in Proc. of the 19th International Conference on Autonomous Agents and MultiAgent Systems (**AAMAS**), 2020, pp. 1116–1124 **https://ifaamas.org/Proceedings/aamas2020/pdfs/p1116.pdf**

7 E. Liscio, M. van der Meer, L. C. Siebert, C. M. Jonker, and P. K. Murukannaiah, "What values should an agent align with?", in Proc. of **AAMAS**, 2022, vol. 36, iss. 23, p. 32, **https://link.springer.com/article/10.1007/s10458-022-09550-0**

8 N. Kökciyan and P. Yolum, "Taking Situation-Based Privacy Decisions: Privacy Assistants Working with Humans," in Proc. of 31st Int. Joint Conf. on AI, (**IJCAI**), 2022, pp. 703–709. **https://www.ijcai.org/proceedings/2022/0099.pdf**

9 N. Hopner, I. Tiddi, and H. van Hoof, "Leveraging Class Abstraction for Commonsense Reinforcement Learning via Residual Policy Gradient Methods," in Proc. of 31st Int. Joint Conf. on AI, (**IJCAI**), 2022, pp. 3050–3056. **https://www.ijcai.org/proceedings/2022/423**

10 Ilir Kola, Pradeep K. Murukannaiah, Catholijn M. Jonker, M. Birna van Riemsdijk, "Toward Social Situation Awareness in Support Agents". **IEEE Intell. Systems** 37(5): 50-58 (2022) **https://ii.tudelft.nl/~pradeep/doc/Kola-2022-IS-SocialSituation.pdf**

Annex 2
List of key pages on our website

- **Team**: Members and affiliated members
- **Funded** and **Associated** PhD projects
- **Publications**
- **Newsletter**
- **In the Media**
- **Keynotes and invited presentations**
- **Code and Data**
- **Educational activities**

Annex 3
Workshops and conferences we (co-)organised

External workshops

- The First International Hybrid Human-Artificial Intelligence conference (**HHAI 2022**)
- The Second International Hybrid Human-Artificial Intelligence conference (**HHAI 2023**)

Workshops we (co-)organized at HHAI2022:

- **Heterodox Methods for Interpretable and Efficient Artificial Intelligence**
- **Knowledge Representation for Hybrid Intelligence**
- **Human-centered Design of Symbiotic Hybrid Intelligence**
- **The (eco)systemic challenges in AI**
- **H3AI Hackaton**
- **Representation, sharing and evaluation of multimodal agent interaction**

Annex 4
Scientific Advisory Board

- **Report**

Annex 5
Impact Plan

- **Project Strategy**
- **Communication Strategy**
- **Quality Strategy**