TU Delft

Self-Assessment Report Computer Science TU Delft 2015–2020

22 September 2021 — *Public final version with selected appendices*

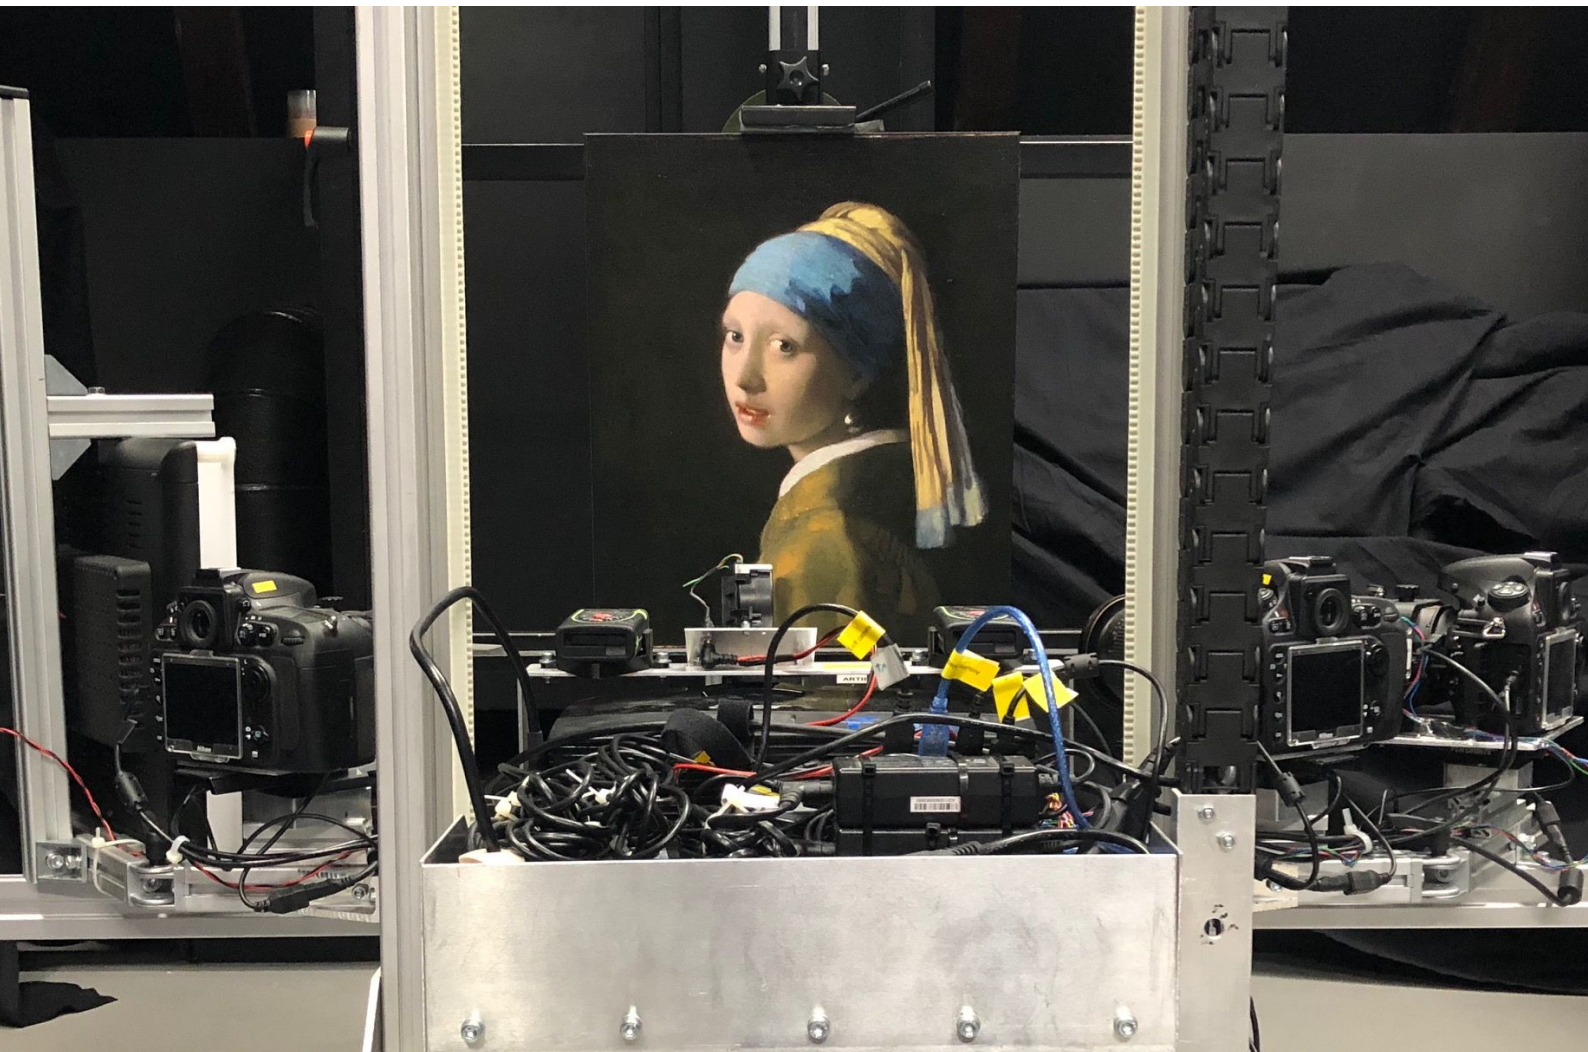

Note: ***Public version in which appendices with personal information about staff have been removed.***

Faculty of Electrical Engineering, Mathematics and Computer Science
Van Mourik Broekmanweg 6
2628 XE Delft
The Netherlands
https://www.tudelft.nl/en/eemcs

Cover image: Tom Callewaert, Jerry Guo, Guusje Harteveld, Abbie Vandivere, Elmar Eisemann, Joris Dik, and Jeroen Kalkman. Multi-scale optical coherence tomography imaging and visualization of Vermeer's *Girl with a Pearl Earring*. In *Optics Express* Vol. 28, Issue 18, pp. 26239-26256 (2020). https://doi.org/10.1364/OE.390703 (Open Access). Press release:
https://www.tudelft.nl/en/stories/articles/looking-over-vermeers-shoulders

Contents

Preface

Computer science is a game changer and plays a crucial role in the ubiquitous use of advanced sensing, communication and computing technologies. These technologies have fuelled the transformation into a digital society, of which we are witnessing great progress as well as concerns. Society calls upon computer science for intelligent algorithms and reliable software systems that facilitate responsible use and trust in a secure cyberworld. Academia plays an important role, because universities have the task to combine advancement of scientific progress by means of ground-breaking research with educating the next generation of responsible computer scientists and future leaders in this field.

The faculty of Electrical Engineering, Mathematics and Computer Science (EEMCS) has fully embraced this challenge. We offer high-quality bachelor's and master's degree programs in computer science to a large international student population. We have grown substantially in the number of faculty positions and have the ambition to grow even further. To compete for talent, we are in the process of further expansion of our resources and office space. At the same time, we continue to strive to become an even better employer by promoting an inclusive and diverse community of staff and students.

Computer Science at TU Delft is the vibrant heart of the faculty of EEMCS. It is organised in two large departments – Intelligent Systems and Software Technology – that collaborate seamlessly within the university as well as with external partners. We foster crossovers with mathematics in creating the foundations for data science and machine learning, and crossovers with electrical energy in computer and embedded systems engineering, network science and digital energy systems. It is EEMCS' mission to provide our society with technology to create a better world. We focus on key technologies such as artificial intelligence, next generation of communication, sensing and computing, as well as the emerging field of quantum technologies. Together with our partners we work in multidisciplinary teams to provide solution for major societal challenges such as climate action, the energy transition, health and well-being as well as the digital society.

TU Delft has embraced artificial intelligence as a university-wide theme and implemented two large collaborative programs: the Delft AI Labs and Convergence with social sciences as well as life sciences and health. The DAI-Labs connect fundamental advances in methodology and algorithms in computer science with the advanced application thereof in all engineering sciences and design on our campus. Convergence provides the interdisciplinary platform to co-create societal impact with predominantly our strategic partners in Rotterdam and Leiden. We are leading and participating in public-private partnerships that accelerate innovations and societal impact.

It is my pleasure to present the self-assessment of Computer Science at the TU Delft. It provides an honest insight in our strategy – mission, ambition and approach – as well as our research accomplishments, organisation and policies. The world is changing and we take our responsibility to provide technology to turn it into a better place. We continuously strive to improve ourselves in virtually all aspects. All input from the expert panel that will help us, is greatly appreciated.

Prof. dr. ir. Lucas van Vliet
Dean of the Faculty EEMCS
Delft University of Technology

Summary

The phenomena of *datafication* and *AI-zation* reflect the increasing tendency to quantify everything through data and to automate the decision-making processes that are also largely based on data. Since these phenomena have entered all segments of our lives and since research in computer science (CS) is at the heart of the technological developments underlying these phenomena, CS as a research field has gained strategic importance. TU Delft Computer Science operates at the forefront of these developments with the aim to help society at large, by enabling it to maximally benefit from these phenomena, while protecting it from potential risks. To that end, inspired and driven by the TU Delft core values of *Diversity*, *Inclusion*, *Respect*, *Engagement*, *Courage* and *Trust* (DIRECT), our mission includes (1) conducting world class research in selected computer science core areas; (2) maximizing opportunities for societal impact of our research; (3) providing rigorous, research-inspired engineering education in computer science; and (4) contributing to an international academic culture that is open, diverse and inclusive, and that offers openly available knowledge.

We are organized in two departments, *Intelligent Systems* and *Software Technology*, consisting of 5 and 6 *sections* respectively. Sections are small-scale organizational units, typically headed by a full or associate professor and marking a particular CS disciplinary scope. While the departments are separate units, they work closely together in research and education, and collaborate for societal impact. The convergence between the departments in terms of alignment and joint pursuit of strategic and operational goals has even become so strong over recent years that we can speak of an increasingly recognizable CS entity in Delft organizing its research into five main *themes* transcending the departmental and section boundaries: (1) decision support; (2) data management and analytics; (3) software systems engineering; (4) networked and distributed systems; and (5) security and privacy. The themes offer critical mass in order to achieve substantial impact, and each theme involves many researchers with various CS backgrounds and expertise.

Award-winning research in these themes achieved during 2015-2020 include a novel cross-modal (e.g., combining text and images) retrieval method based on adverserial learning; genetic algorithms for the automatic reproduction of software crashes to facilitate automated debugging; and Trustchain, a permission-less tamper-proof data structure for storing transaction records of agents with applications in digital identity. International recognition of our expertise is reflected by numerous leadership roles, e.g., as general or program chairs in numerous flagship conferences, such as AAAI, EuroGraphics, ACM/IEEE ICSE, ACM OOPSLA, ACM RecSys and ACM Multimedia. In the same time period, several staff members also received the highest (inter)national recognition in their fields, such as IEEE Fellow, membership of the Young Academy of the Royal Dutch Academy of Sciences, or the Netherlands Prize for ICT Research. Our scientific reputation also brought us into the consortia of two prestigious NWO Gravitation Projects ("NWO Zwaartekracht") of the Dutch Research Council, Hybrid Intelligence and BRAINSCAPES - the consortia that *"belong to the world top in their field of research or have the potential to do so"*.

To maximize societal impact, we embrace eight key *sectors*: transport and logistics, energy, health and well-being, safety and security, finance, online education, creative industry, and smart cities. To enable and support us in making substantial interdisciplinary impact in these sectors, we have built up expertise, a network of collaborators and societal partners, and established the necessary organizational structures. Prominent examples of our impact in these sectors include the *NGS sequencing analysis pipeline* we designed and implemented as part of the NIPT test, which is used routinely by hospitals in several countries; *Cytosplore*, a software system for interactive visual single-cell profiling of the immune system; and *SocialGlass*, a tool suite for integration, enrichment, and sense-making of urban data. Our close ties with society are also reflected in our strategic collaborations with socio-economical partners, such as ING, DSM, Booking.com, Adyen, Ripple, Erasmus Medical Center and Leiden University Medical Center, leading amongst other things to strategic investments in the form of three large industry-funded labs (with ING, DSM and Booking.com) setup in the assessed time period for a duration of five years. Furthermore, we have invested extensive effort in public outreach, explaining and discussing science with a broad audience, and in particular in the context of complex societal debates in the domain of AI and blockchain. Finally, we play a leading role in regional, national and European initiatives, most notably in the Dutch AI Coalition (NLAIC).

In addition to scientific excellence and strong impact in the selected societal sectors, we are committed to (a) meeting the increasing societal need for highly skilled CS experts, (b) development of human capital

in our organization, leading to a new generation of international academic leaders, and (c) advancing the organization and academic culture, with the key pillars of open science, diversity and inclusion.

Regarding (a), we embraced an over 100% increase of our student population, but also aim at securing the highest possible level of their knowledge, skills and academic forming despite scaling up. Therefore, we value a close connection between research and education, and let both MSc and BSc students participate actively in our research. We also formulated an ambitious strategy, the realization of which would enable us to manage this education scale-up efficiently and effectively, leaving sufficient room to our staff for further developing scientific excellence and deploying it for societal impact. Part of this strategy is the growth of our academic staff towards 100 FTE by 2024 to meet the stabilization of the student numbers (due to numerus fixus). Between 2015 and 2020, we already achieved a net growth from 54 to 72 faculty members (+33%), with more to come in the upcoming years.

Next to BSc and MSc students, we are committed to delivering highly skilled CS experts at the PhD level. The number of PhD students grew from 105 to 165 (+57%) in the assessed time period, reflecting our ability to successfully acquire research funding in the present landscape. For our PhD students, the Graduate School defines a framework in which they can develop their skills next to conducting their thesis research. We strive towards completion of PhD theses within four years and organize our supervision, official moments of assessment, requirements on the volume and quality of the conducted research, as well as evidence of scientific impact through publications, accordingly.

Regarding (b), development of human capital: as computer science expertise is in high demand across the globe, finding strong new people as well as retaining our current staff proved highly challenging, especially given the high teaching load due to our record student intake. Therefore, acquiring, developing and retaining academic talent has been one of our most important goals. Dedicated actions, such as devising of a *Development Track Plan*, serve to empower each staff member to provide contribution to the organization in his/her own way, based on individual interests, talents and ambitions, and in view of our joint ambition as organization.

In view of (c), our organization, we embrace open science, with a substantial percentage (80% in 2020) of our articles available as open access, and by making numerous software tools and data sets openly available. We are a highly international organization with employees and students from all over the world. We strive to be an inclusive organization, where staff and students feel at home and valued, regardless of their background, age, gender, sexual orientation or functional disability. In terms of female faculty, we realized a net growth from 11 to 14 faculty members. As the number of men employed also increased, the percentage of female faculty stayed stable at around 20%. We consider this too low. We are committed to addressing this, for which we will take a long-term approach with, amongst other means, dedicated budget reserved for continued openings for female faculty in the upcoming years.

We are proud of our scientific successes and societal impact in the core computer science disciplines as well as in interdisciplinary research in our target societal sectors. This is especially so as those were achieved in a period that was transformational for CS@Delft, characterized by substantial growth and development across our organization and activities. We anticipate an even stronger societal demand for our research and expertise in the future. We will therefore continue to initiate, participate and take on a key role in effective and interdisciplinary partnerships at the university (TU Delft AI), regional (LDE), national (ICAI, IPN), and European (ELLIS, CLAIRE) levels. Furthermore, we will continue the growth path for our staff, in order to build up capacity enabling us to further develop our scientific excellence and offer our strongly increased student population the world-class research-intensive education they deserve. To achieve this, we center the next steps in our ongoing transformation around *people*, *organization*, and *profiling* and identify seven key actions for the upcoming years that aim at (1) improving our attractiveness as an employer; (2) improving diversity and inclusion; (3) improving the execution of the PhD program; (4) expanding our staff capacity; (5) aligning our office space with the optimal way of working; (6) articulating the scientific profile; and (7) boosting our scientific and societal impact.

Prof. dr. Alan Hanjalic
Head, Intelligent Systems Department
Delft University of Technology

Prof. dr. Arie van Deursen
Head, Software Technology Department
Delft University of Technology

1 Introduction

1.1 Context

This research assessment of Computer Science at the Delft University of Technology (TU Delft), further in this document referred to as CS@Delft, has arrived at a good time for its departments to support them in their ongoing efforts to position themselves in what could be the most dynamic and critical development period to date. One key aspect of this development is *datafication*[1] - a trend seen in an increasing number of societal, industrial, and scientific domains and referring to quantifying more and more aspects of our lives through data. Because of its intensity and extent, this trend is already radically influencing the way people, companies, societies, and governments exist and operate, by creating new opportunities as well as new risks. With the increased use of algorithmic and software tools that can deploy data for automating decision-making processes, we can also speak about another trend that has accompanied datafication in recent years: *AI-zation*. While datafication has raised the discussion about the volume, variety, velocity and veracity of data (the *four V's of Big Data*), AI-zation has pointed to the explainability and trustworthiness[2] of the automated decision-making systems.

Focusing at the heart of the datafication and AI-zation phenomena, computer science (CS) has the responsibility to play a leading role in the data- and AI-induced transformation of our society, and in helping the society through the *fourth industrial revolution* that embodies this transformation and that is characterized by a '*fusion of technologies that is blurring the lines between the physical, digital, and biological spheres*'[3]. This requires CS@Delft to (a) position itself successfully in terms of its scientific, educational and societal ambition, and (b) find ways to equip itself to pursue its ambition effectively with maximized long-term impact. The current growth and transformation of CS@Delft is serving precisely this purpose. Being already the largest in the Netherlands, it is on its way to growing even further in order to better meet the increasing societal demand for socially responsible CS engineers and to contribute to CS education of all engineers. At the same time, it is aiming at further expanding its disciplinary scientific frontiers, crystallizing its strong and recognizable scientific focus and profile, and strengthening its role in multidisciplinary collaboration with all engineering and other sciences. The material provided in this self-assessment document serves to show to what extent we believe CS@Delft is equipped to pursue its growth and transformation in the best possible way. We expect that the interaction with the international assessment committee will give us valuable pointers towards further improving our position and capabilities in this respect, helping us to find and fix our blind spots and empowering us to approach the future successfully.

1.2 Computer Science at TU Delft

1.2.1 Organization and Leadership

Research and education in CS@Delft is carried out in two departments: **Software Technology (ST)** and **Intelligent Systems (INSY)**. These departments are part of the Faculty of Electrical Engineering, Mathematics and Computer Science (EEMCS), being one of the eight faculties of TU Delft (Figure 1.1).

Each department consists of a number of smaller organizational units called *sections*. While each section has a disciplinary origin, some sections have evolved towards larger units, allowing new disciplinary seeds to develop, which could eventually grow into new, separate sections. Ensuring full academic freedom and autonomy to each faculty member, the main benefit of the section-based organization is to provide a supporting and stimulating small-scale local environment for (new) faculty members to create and maintain their research momentum in their familiar disciplinary context. The sections are headed by an associate or

1 K.N. Cukier and V. Mayer-Schoenberger: The Rise of Big Data: How It's Changing the Way We Think About the World, Foreign Affairs, May/June 2013

2 Ethics guidelines for trustworthy AI

3 K. Schwab, World Economic Forum

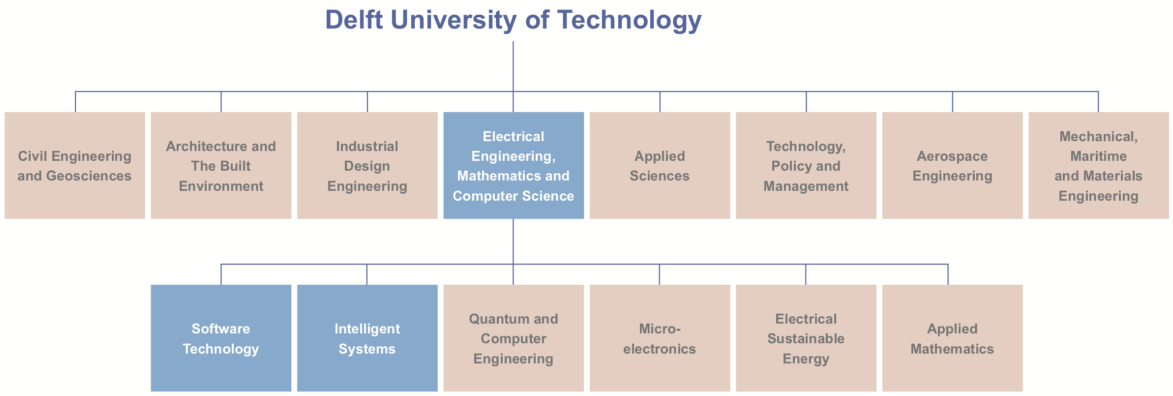

Figure 1.1: Overview of the university's organizational structure in which CS@Delft is positioned

full professor, and typically comprise 5-8 academic staff members (tenured or in tenure track), PhD candidates, postdocs, software engineers and support staff. The INSY and ST departments are roughly similar in size and have developed similarly over recent years. Appendix I provides an overview of the 11 CS@Delft sections. The figures in Appendix B illustrate the development of the staff composition per category and per department over the assessed time period.

The departments are led by the department heads - full professors appointed by and reporting to the dean of the EEMCS faculty. They are responsible for the personnel and budget in their departments, for building up and maintaining the education-related effort of the department's academic staff, and for the overall quality, impact, and viability of the department's research. The appointment term is four years, with the possibility for one extension. The department head chairs the department's management team consisting of the section heads. The management team meets on a weekly basis, to discuss policies and decisions concerning, e.g., hiring, promotions, project acquisition, finances, research strategy, education, and the relevant (inter)national developments. Each department is supported by a manager, who takes over part of the administrative and budget-related responsibilities.

The departmental budget covers all staff and their activities (e.g., equipment, travel). Following the classification common to the Dutch academic system, the budget is fed from two sources, internal and external, divided into three *money tiers*. The primary source (direct funding, first money tier) is an internal one and used to cover all structural costs of a department, in particular the salaries of the faculty members and the permanent support staff. It comes from the faculty, which in turn is funded by the subsidy TU Delft receives from the government. This subsidy is divided between faculties based on an allocation model defined by the TU Delft's Executive Board. This model includes a fixed (structural) component, and a variable one depending on the teaching effort of the department's staff (expressed by the ECTS credit points earned by students) and on the number of completed PhD theses. The second funding source consists of research grants and funding from industrial and other external parties, also referred to as contract research. The research grants provided by the Dutch Research Council (NWO) or other governmental institutions fall under the second money tier, while other grants (including EU projects) and the direct industrial funding form the third money tier. Both the second and the third money tiers are deployed to build up and maintain temporarily employed research and development capacity, consisting of PhD students, postdocs or software engineers, but also to finance the acquisition of extra equipment and to provide more financial room for the operations of the faculty members beyond what is possible from the primary funding source. The revenues and expenditures per department over the assessed time period are presented in Appendix B.

1.2.2 From TWO departments to ONE Computer Science

The division of CS@Delft into two departments is primarily due to practical reasons, as it helps preserve the span of control and secure balance in academic staff volume across the EEMCS departments. Evolved as such over many years, the two departments have marked two major research lines addressing in a complementary fashion the trends of datafication and AI-zation. Regarding the first research line, mainly anchored at the ST department, the development and adoption of the technology underlying these trends brings with it many challenges that point to growing complexity of designing distributed, embedded and autonomous systems

and software for data/information gathering, communication and processing. This points to open issues at the system and software layers of the technology, such as complexity, reliability, scalability, evolution and adaptability. However, these challenges also point to the issues at the application layer related to the practical impact of datafication and AI-zation, such as how to understand and manage, as well as reason upon the complex relationships that are contained in the sheer volume of data being produced by a wide variety of physical and social sensing systems, how to handle the data and information overload, how to optimally deploy the new technology for enrichment of individuals and society, and how to protect ourselves from the negative consequences of this technology becoming an integral part of our lives. Regarding the latter category of issues, which form the second research line and that are mainly anchored at the INSY department, typical challenges are related to utility, quality of experience, explainability, responsibility, security and privacy. The split of CS@Delft into two departments never presented barrier for collaboration and joint ventures, as can be seen by many joint initiatives to set up or influence broader interdisciplinary research programs and institutes (see Appendix G). Hence, while the two research lines may be anchored at different places in the two departments, neither of them is necessarily restricted to a particular department, and increasingly they have been pursued across the department boundaries.

The ongoing boundary fading between the departments is also a consequence of the need (a) to respond to the increasing complexity and dynamics of our research targets and the strategic developments related to CS on the (inter)national stage, and (b) to scale up our organization to address the demand for growth (Section 3.2), while keeping the excellence level we have built up over the years. Positioning us well in this respect requires efficiency in using our resources and effectiveness in pursuing our goals. This can only be achieved by further tightening the connection between INSY and ST and operating as one well-aligned and organized CS@Delft team. Regarding the efficiency of scaling up our organization, the two departments built on already strong collaboration on education and in 2018 also formally established the joint CS Teaching Team. This team works together and coordinates with the academic staff from both departments in shaping up and executing our joint CS curriculum. Regarding the effectiveness of pursuing our goals, we have given more weight to our joint CS@Delft mission, especially regarding the aspects related to impact in key societal sectors, teaching-research nexus, academic organization and culture, development of human capital and PhD program (Section 2.2). This joint mission unifies and enhances the individual departmental missions[4] and is critical for us to operate successfully on the increasingly competitive CS funding and personnel market.

1.3 Approach to assessment

The strategy sketched above towards building a strong and recognizable CS entity in Delft serves the purpose of improving internal consolidation and deployment of resources across the two departments and give them more leverage in influencing the faculty, university, regional, national and international developments related to CS. The success of the strategy directly depends on our ability to address challenges that are department agnostic and can best be pursued through a joint effort. This ability, however, depends on the effectiveness of the strategic choices the management of INSY and ST departments have made and the instruments they have put in place. It is therefore the implementation of this strategy (Section 3.3) that we wish to place centrally in this research assessment.

In view of the above, we present the mission and strategic aims we defined in the assessed time period (Chapter 2), as well as the ways of addressing these aims (Chapter 3), from the combined CS@Delft perspective. This holds conceptually also for the overview of past performance (Chapter 4), with the difference that we make the results, achievements and contributions traceable per department. The latter also holds for Chapter 5, where we highlight the key scientific results, their scientific and societal relevance and future prospects. Then, in Chapter 6, we reflect as CS@Delft on our strengths, weaknesses, as well as on the major external threats and opportunities, and identify the main pointers for adjusting our current strategy for the time period of 2021–2026. This approach, centered around joint positioning, reflection, ambition and strategy of the INSY and ST departments as one CS organization has been suggested by the previous research assessment committees (Appendix *app:responses omitted*). The approach, however, also allows to assess the performance of each individual department and their individual contribution to the joint CS@Delft strategy, as requested by the TU Delft Executive Board.

This document is optimized for online reading. In addition to footnotes providing links to URLs (where applicable), the text pointers indicated in blue provide direct links to chapters, sections, figures and tables.

4 Information about this can be found in the documentation submitted for the research assessments in 2015 and 2018.


/b æ d/ / b æt/
/b ɛ d/ /b ɛ t/

2 Mission and Strategic Aims

2.1 Mission

Over the assessed time period, the INSY and ST departments have focused on their joint CS mission to help them operate more effectively on the international CS stage (Section 1.2.2). This mission is:

1. *to conduct world-class research in selected computer science core areas;*
2. *to maximize opportunities for socio-economic impact of our research;*
3. *to provide rigorous, research-inspired engineering education in computer science;*
4. *to contribute to an international academic culture that is open, diverse and inclusive, and that offers openly available knowledge.*

While the mission formulation itself largely follows those from the previous research assessments, the critical difference with the past lies (a) in setting up a joint strategy to pursue this joint mission that aligns the efforts from INSY and ST around what they see as their main strategic aims (Section 2.2), and (b) in the coordinated joint CS@Delft-wide approach to addressing these aims (Chapter 3).

Regarding the first item of the mission above, inspired and stimulated by the increasing collaboration, diverse expertise and synergistic potential across the departments, but also by the need for mutual positioning of the Dutch CS institutes in the context of the *Sectorplan* (see Section 3.2), we have focused on the scientific profiling of CS@Delft as a whole, beyond individual departments, and especially in the cross-departmental domains, like AI, data analysis and management, distributed systems and security and privacy. This led to five research themes (Table 2.2) that build on the two research lines anchored in the two departments (Section 1.2.2) and their individual research missions (Table 2.1). Each theme is supported by a critical mass of research capacity and has emerged from a significant track record of excellence over the past years by multiple INSY and ST sections (Table 2.3, Appendices C, *app:outcomes omitted*, and D). Through these themes, CS@Delft aims at crystallizing its profile in the (inter)national CS landscape and indicates where it wishes to maximize its scientific excellence and societal impact. For this reason, we also organized the presentation of our scientific highlights in Chapter 5 around these themes.

Intelligent Systems. The research mission of the INSY department is *to enable human and machine, in close cooperation with their environment, to deal with the increasing volume and complexity of data and with the increasing role that the AI-driven decision-making systems are taking up*. The department pursues its mission by integrating fundamental research, engineering and design in the interlocking fields of processing, interpretation, visualization of and interaction with data using model- and knowledge-based methods and algorithms. The INSY research is inspired largely by challenges from the domains of social media and entertainment, medical and health sciences, security and privacy and cultural heritage.
Software Technology. The research mission of the ST department is *to advance the design, engineering and analysis of complex, distributed and data-intensive software and computer systems*. Society is on its way to be completely dependent on such systems, which include online payment systems, smart grids, smart cities, and self-driving cars. In its research, the ST department wants to contribute both to the foundations of software and computer systems, and to the understanding of the operation of such systems. Therefore, its research approach is both fundamental — designing algorithms and abstractions, developing new concepts and theories, and identifying principles, as well as experimental — engineering and analyzing the behaviour of actual systems.

Table 2.1: Research missions of the INSY and ST departments.

2.2 Strategic Aims

We have pursued our mission by following the overarching strategic aims given below.

World-class research portfolio Scientific excellence and reputation come with research quality and scientific leadership at the international level. To achieve this, we set as our first strategic aim *to maximize our international visibility and our scientific impact within the joint research themes described above*, along the current focus dimensions (see topics per theme in Chapter 5) or some new ones in the future. This requires us *to articulate our scientific profile* built upon these themes in the best possible way and *to empower ourselves to influence the international scientific agenda* on these themes.

Strong population of PhD students In order to increase our research volume, maintain the dynamics and momentum in our research, and enable our staff to achieve more impact, a good number of PhD students, typically 2.5-3 per faculty, is essential. Given the funding constraints of the Dutch academic system, maintaining the PhD-student population is only possible by the external research funding. It is therefore our strategic aim *to empower our academic staff to meet the challenges related to generating high-quality project proposals and building up networks with societal partners and to become and remain successful in acquiring research funding*. We also consider it essential to deliver to academia and industry the future leaders in research and innovation in computer science. Therefore, another strategic aim is *to maximally exploit the possibilities given to us to educate and train PhD students, not only in carrying out world-class research, but also in acquiring professional skills and competences*.

Development of human capital Visibility and impact are achieved by thought leaders in their scientific fields and their engagement and initiative in broader societal discussions. Academic leadership builds on idealism, ambition and drive, but needs skills and competences to grow and become effective. It is also essential to stimulate academic leadership development as early as possible, starting with our tenure-track faculty members. It is therefore our strategic aim *to foster academic leadership among our academic staff and to offer them all means possible to develop and exploit their leadership potential.*

Tight connection between research and education We place great value on the relationship between our research and our education. From one side, having an international scientific expert working with MSc students and coaching them in understanding current scientific developments and in pursuing scientific discoveries boosts the quality of the engineers we wish to deliver to society. From the other side, discussions with students in the context of specific courses or assignments may also inspire or inform the research of our staff members. It is therefore our strategic aim *to strengthen the relationship between research and education and maximize their mutual benefit*.

Strong impact in key societal sectors Our mix of foundational and use-inspired research allows us to be active in a range of application domains. In view of our research profile, however, we focus our collaboration with industry and societal partners on a number of areas where we can maximize inter-disciplinary impact. These areas fall into the sectors of *transport and logistics*, *energy*, *health and well-being*, *safety and security*, *finance*, *online education*, *creative industry* and *smart cities*. In this respect, our strategic aim is *to empower ourselves to maximize visibility and impact in these sectors*.

Advancing the organization and academic culture An important precondition to research excellence and leadership is organizational excellence and an academic culture that can serve as role model to others. It is therefore our strategic aim *to arrange human resource policies, life-work balance, career promotions protocols and coaching such to secure excellent work and development conditions for our scientific staff*. In this regard, it is our strategic aim *to keep nourishing the CS@Delft academic culture where research integrity, open science, diversity, and inclusion are the main pillars*.

Decision Support. This theme encompasses devising, implementing and evaluating fundamental (classical and data-intensive) AI solutions to improve the machine-level processes involving humans and supporting them in making decisions. We have focused on challenges related to *interactive and hybrid intelligence*, *intelligent agents and multiagent systems*, *computer vision and machine learning*, *user modeling* and *search and recommendation*.

Data Management and Analytics. This theme focuses on enabling humans not only to handle large quantities of data efficiently, but also to gain insight into and analyze complex (big, unstructured, distributed, heterogeneous, web) data and the AI decisions, actions and their consequences forthcoming from data collection and processing. This research also generates scientific foundations of data-driven, human-centric AI technology. We have focused on the challenges related to *data management and engineering*, *bioinformatics*, *content generation and data visualization*, and *data analysis and process modeling in complex networks*.

Software Systems Engineering. This theme brings together the expertise in the areas of software engineering, programming languages, and software reverse engineering and is aimed at developing new ways to safeguard key quality attributes of complex software systems, such as robustness, reliability, and evolvability. To that end, this theme has combined foundational (e.g. program semantics), empirical (software analytics), and AI (model learning, search-based software engineering) research. In particular, we focus on the challenges related to *software analytics*, *AI for software engineering*, *language engineering*, and *software verification*.

Networked and Distributed Systems. This theme aims at devising, implementing and analyzing fundamental algorithmic and systems concepts for distributed and networked systems. Focus has been on *distributed trust*, in particular on the problems related to scalability and consistency in existing blockchain-based approaches, *Internet of Things*, with special attention to low-power devices, *big data processing*, which addresses the complexity and scalability challenges in the data streaming and data processing components of these systems, and *visible light communication*, aiming at unlocking the bandwidth of the visible light spectrum for wireless communication.

Security and Privacy. This theme targets foundational computer science challenges in security and privacy, namely rigorously showing that a computer system is inherently (in)secure, and the maturing of design-for-security and privacy methodologies. We have focused on the following aspects of this theme: *AI for intrusion detection and prevention*, *AI-based testing for secure systems*, *security of AI*, and *privacy protection*.

Table 2.2: An overview and description of the joint CS@Delft research themes.

		Algorithmics	Interactive Intelligence	Multimedia Computing	Web Information Systems	Pattern Recognition and Bio-Informatics	Computer Graphics and Visualization	Programming Languages	Embedded and Networked Systems	Software Engineering	Cybersecurity	Distributed Systems
Themes	Decision Support	*	*	*	*	*						
	Data Management and Analytics			*	*	*	*					
	Software Systems Engineering							*	*	*		
	Networked & Distributed Systems								*		*	*
	Security and Privacy									*	*	*

Table 2.3: Coverage of the themes by the main contributing sections from INSY (green) and ST (blue).

3 | Strategy 2015-2020

3.1 Strategic Process

The strategic aims stated in the previous chapter transcend departmental boundaries. Therefore, our strategy to pursue these aims is defined at the level of CS@Delft. This also holds for the four specific evaluation aspects of the SEP Protocol, namely *Open Science*, *PhD Policy and Training*, *Academic Culture* and *Human Resource Policy*. Shaping up this joint CS@Delft strategy is a continuous process, triggered by initiatives from either the INSY or ST department, inspired by the comments received by international peers, influenced by our local university context, and steered by the management teams of the two departments.

Feedback on our strategy from our international peers is collected regularly through the six-year cycle of the *TU Delft Research Assessments Protocol*, which includes an assessment at the end and midway the reporting period. In Appendix *app:responses omitted*, we present the main points brought up by previous assessments from 2015 and 2018 and how we responded to them.

Within TU Delft, the Executive Board, as part of its *Planning and Evaluation Cycle*, discusses strategic matters with the management team of the EEMCS faculty during the annual *Strategic Spring* and *Strategic Fall* meetings. Then, within the faculty, the CS@Delft strategy is discussed in the EEMCS Management Team under the leadership of the dean, and aligned with the faculty's multi-year plans and long-term budget, again in yearly spring and fall meetings.

Devising the CS@Delft strategy, ensuring its consistent execution, and adjusting it as new insights and opportunities arise, is the joint responsibility of the heads of the INSY and ST departments. To that end, they meet on a weekly basis and are informed through frequent discussions at multiple levels, including the two departmental management teams, faculty meetings, faculty lunches, section meetings, as well as the yearly faculty retreat of the joint departments. Through the local dynamics of the operational contexts in two departments, different aspects of the joint strategy may receive different priorities in a given time span, resulting sometimes in different pace of progress of implementing these aspects in INSY and ST. The close collaboration between the departments ensures, however, that the priorities remain 'fresh' across CS@Delft, inspiring each other by their individual ideas and initiatives.

3.2 Strategic Context

During the assessed time period, we have witnessed a huge increase in the societal demand for computer science knowledge and expertise and academically qualified computer science engineers. We embrace these developments and have been committed to meet this societal demand.

From the societal perspective, the interest in and significance of such topics as cybersecurity, blockchain, and artificial intelligence has dramatically increased. We invest considerable resources in playing a proactive, agenda-setting role in such topics, in close collaboration with key societal partners (see Section 3.3.10). A prominent success is the Dutch AI Coalition (NLAIC[1]), in which CS@Delft has played a leading role since its inception, resulting in an investment of €276M through the National Growth Fund (contributions from many, with key role for Inald Lagendijk, INSY/CYS).

From the educational perspective, we are proud that between 2012 and 2020 the total number of bachelor and master students in our international educational programs has tripled (from 750 to over 2200), and from 2015 to 2020 (the assessed time period) this number has doubled (from 1100 to over 2200). We presently ensure the intake in the BSc and MSc programs of 500 and 300, respectively, for a stable population of 2500 students in the upcoming years. With our research tightly connected to our educational programs (see Section 3.3.8), we made it our top priority to organize sufficient funds and attract new faculty members so that research and educational responsibilities are in balance. In particular, to ensure sufficient opportunities for students to interact with active researchers, we formulated the ambition to achieve a student-staff ratio of 25:1 by the year 2024, which implies growth towards 100 faculty members in total.

To accommodate this growth we rely on three income sources that form the first money tier (see Section 1.2.1). The first source is the variable part of the faculty income, which directly depends on credit points

1 Dutch AI Coallition

earned by students. Within the faculty, we have been active in ensuring that the growth of this income that can be attributed to increased student numbers and teaching effort in Computer Science (around €3.5M) is indeed used to support CS@Delft. The second income source is structural support from the TU Delft Executive Board of €2M yearly for our growth plans from 2020 onwards (from the TU Delft *Van Rijn means*[2], a national reallocation of university funding to strenghten disciplines with high economic impact). Lastly, we successfully participated in the *Sectorplan*[3] (a national plan to strengthen fields including computer science), from which CS@Delft obtained structural funding for 10 new faculty positions (€1.2M per year).

The development of our research capacity and staff in light of this demand for growth is shown in Appendix B. From 2015 to 2020, we realized a net growth from 54 to 72 faculty members (Table B.4), a 33% increase. The student-staff ratio was 21:1 in 2015, and has deteriorated to 36:1 in 2020. Our target of 25:1 by 2024 would help us get back to normal and corresponds to the TU Delft average. Reaching this target is challenging due to a highly competitive academic job market in CS, a strong pull on CS experts working in academia by industry, and reduced attractiveness of an academic position in CS@Delft due to the currently high teaching load. This challenging context has largely marked the assessed time period, especially 2018–2020, has influenced the way and the results of how we have addressed our strategic aims, and will remain relevant in the (near) future, influencing our future strategy.

3.3 Addressing the Strategic Aims

As stated in Section 2.2, our strategic aims are: (1) *World-class research portfolio*, (2) *Strong population of PhD students*, (3) *Development of human capital*, (4) *Tight connection between research and education*, (5) *Strong impact in key societal sectors*, and (6) *Advancing the organization and academic culture*. Our strategy to address these aims encompasses different aspects, described in different subsections below. Each aspect serves to address one or more aims. Also, a strategic aim can be addressed in different ways by multiple aspects of our strategy.

3.3.1 Our Values

The definition and implementation of our strategy is largely inspired and driven by the six TU Delft values: **Diversity**, **Integrity**, **Respect**, **Engagement**, **Courage**, and **Trust**, collectively known as **DIRECT**[4]. For us as CS@Delft, *diversity* means that we welcome different perspectives, backgrounds, and identities. We foster this by stimulating cultural and gender diversity (Section 3.3.4), as well as diversity in academic backgrounds of the staff we hire, including non-CS backgrounds. *Integrity* means that we adhere to the TU Delft *Code of Conduct*[5], as well as those of relevant organizations like ACM and IEEE. For *research* integrity in particular, we strive to be as transparent as possible through *open science* (see Section 3.3.7). Concerning *respect*, we encourage all employees to bring their own unique skills and contributions to the table through a high level of independence and academic freedom we create for our academic staff (Section 3.3.2). This is directly in line with the San Francisco Declaration of Research Assessment (DORA[6]). We acknowledge each other's freedom and ambition to excel, which is also reflected in our HR processes (Section 3.3.3). In terms of *engagement*, we seek to maximize involvement of staff in organizational processes, as well as our *societal* engagement and impact (see Section 3.3.10). With respect to *courage* we welcome risk appetite and initiative. We encourage research off the beaten track, as reflected in our publication and project acquisition strategies (Sections 3.3.6 and 3.3.11). Lastly, *trust* refers to the way we work with each other and with our external partners. We are not afraid to make ourselves vulnerable, and put people over process.

3.3.2 Managing the Research

Related to the strategic aims *World-class research portfolio* and *Development of human capital*, we made a set of choices on how to develop, direct, organize, and maintain the research momentum in the dynamic transition period described in Section 3.2.

2 Van Rijn Funding Program Higher Education

3 Sectorplan to strengthen scientific research in technical and natural sciences

4 TU Delft Strategic Framework

5 TU Delft Code of Conduct

6 San Francisco Declaration on Research Assessment (DORA)

Striving for international excellence in the scope defined by the five joint research themes (see Section 2.1) already brings some level of focus (for example, in our themes we choose not to conduct research on purely theoretical computer science). Within these themes, however, we choose not to restrict the scope of the research, nor to top-down decide to stop with a particular line of research. Instead, we leave it to our faculty members to direct, focus, shape up, or eventually change their research in the way they find it to be most effective in their specific domain of expertise. We count here on the individual ambition and strong embedding of our faculty members in their international communities, based on which they can best judge which research direction or topic is most valuable for them and for the CS@Delft as a whole. At the same time, we strive to provide in this way a work environment in which every faculty member can enjoy the maximum of academic freedom and be maximally motivated, productive, creative, and inspired to search for new possibilities for own or collaborative research.

While allowing the research themes to develop organically as explained above, the management teams of the two departments make strategic choices to facilitate this development. For example, our ambition to excel in all five themes determines how we allocate new positions across CS@Delft given a growth context described in Section 3.2. Since all sections contribute to these themes, we balance the staff coverage per contributing discipline (section) with the needs to develop a new discipline (like, for instance, speech and natural language processing) of interest for one or more CS@Delft research themes or to further stimulate a research cluster that has already built a strong strategic momentum (like, for instance, Computer Vision) benefiting a particular research theme.

3.3.3 Hiring, promoting and developing talent

While new positions are typically advertised internationally, we also actively scout for talent on the international academic market. New appointments are usually at the level of a tenure-track Assistant Professor, which allows us to maintain a balanced composition of academic staff in terms of age, and to continuously acquire new fresh perspectives on the field of computer science. In order to select top talent, when hiring new scientists we follow a rigorous procedure, involving a broad scientific committee and including multiple interview rounds, scientific presentations and teaching trial lectures.

An overview of the newly appointed staff during the reporting period is provided in Table *tab:appointed omitted*. We appointed a total of 48 new faculty members – 16 during the first and 32 during the second half of the reporting period, illustrating our increased focus on attracting new people. The vast majority (39) was appointed at the assistant level (in tenure track): five were appointed at the associate level, and four as full professor. The academic backgrounds were highly diverse and from a wide range of strong institutions, thus further enriching our international scientific network. Also noteworthy is that five faculty members moved from industry back to academia, illustrating and strengthening our connection to the (tech) sector (in these cases IBM, SAP, and Amazon).

Related to the aims *Development of human capital* and *World-class research portfolio*, we have devised a strategy towards developing and promoting our faculty members towards becoming not only scientific, but also academic leaders in a broader sense. Promotion from tenure-track up to the full professor level is based on personal performance, not limited by the capacity plan and evaluated by a faculty/university career development committee. Recently, we have adjusted the implementation of the tenure-track process to allow for more diversity in the development pace of a staff member to better take into account the specific aspects and the context of each individual academic career. This led to a reduced role of the quantitative performance indicators, and stronger role of the general development potential when deciding to award tenure to a scientific staff member. Beside yearly evaluation, the tenure track comes with a *mid-term* (after 2.5 years) and an *end-term* (after five years) assessment by the departmental management team and HR. Since 2018, based on several years of experience with the tenure track system, permanent positions are granted upon positive evaluation at the mid-term (upon approval by the dean).

We strive for T-shaped academic staff, who are recognized specialists in their scientific discipline, but can also connect to other disciplines and application domains and can recognize and exploit strategic aspects of their research. To stimulate this, reflection on and coaching in academic leadership has become central in career development discussions and monitoring of our staff. In order to assess the development potential of a staff member for a nomination for tenure and identify the best means of support and coaching towards exploiting this potential and climbing the academic career ladder, we recently devised a *Development Track Plan* document (Appendix F), in which reflection is asked about various considerations and choices a staff member encounters in their daily professional life, and which influence their career development. Upon

approval by the dean and the HR department, this document is currently deployed as a pilot within the INSY department. Furthermore, with support of the HR department, we provide systematic training of academic staff towards teaching qualifications, life-long learning, leadership and management skills.

Since 2019, new faculty members in the computer science departments are provided with a startup package consisting of a single funded PhD student, thus helping the faculty members to bootstrap their research. The section provides the context in which the tenure tracker "lands": The section leader often acts as a coach, and is responsible for the yearly evaluation. This typically also entails assisting in getting to understand the Dutch funding landscape. As an example of the success of this, in the Programming Languages section (ST) three consecutive new hires, Krebbers, Poulsen and Cockx, each managed to acquire an NWO Veni personal grant in their first year (in 2018, 2019, and 2020).

Promotions to the ranks of associate and full professor are fully aligned with the evaluation criteria during the tenure track, with focus on leadership and impact. Proposals for such promotions are additionally evaluated by the faculty's *Career Committee* ('loopbaancommissie'), as well as, for promotions to full professor, a university-wide committee installed by the Executive Board. During the assessed time period, 15 assistant professors were promoted to the associate level, and 2 associate professors were promoted to full professor (see Table *tab:promotions omitted*). As the number of promotions is directly related to the staff composition, in the upcoming years we expect an increase in such promotions, given the many newly started assistant professors.

During the assessed time period, in addition to retirement (7 faculty members), 25 decided to pursue their career elsewhere (see Table *tab:departed omitted*). Of these 25, 6 switched to industry. Of the 19 that stayed in academia, 4 moved abroad, and 15 stayed in The Netherlands. Of these 15, 11 combined their move with a promotion, and 4 stayed in the same rank. Reasons to switch to another (Dutch) university vary, and include new collaboration opportunities, better topical fit, or personal reasons (the "two-body problem"). We are proud of the strong positions our staff members take at other places and consider such career switches natural, healthy and refreshing for an academic environment. However, we are even more proud if our staff members feel comfortable and empowered and do not have the incentive to leave. The activities and initiatives described above, aiming at allowing more job security, creating *'room for everyone's talent'*[7] and helping everyone to maximally develop and exploit own talent, are geared towards achieving precisely that.

3.3.4 Improving Diversity, and Inclusion

An important aspect of our strategic aim of *Advancing the organization and academic culture* is to improve diversity and inclusion. Diversity is the first of the six TU Delft DIRECT values (Section 3.3.1) and features prominently in the TU Delft Strategic Framework 2018-2024 (see footnote on p.10), describing it *'as a cornerstone for excellence and innovation'*, a position we fully support.

The composition of our academic staff is highly international, with a large diversity in cultures and nationalities. An impression of the background of the newly appointed staff is provided in Table *tab:appointed omitted*. Around 20% have a former affiliation in The Netherlands (8 at INSY, 2 at ST), 60% come from various European countries (including Germany (6), UK (4), Portugal (2), and Switzerland (2)), and 20% come from the Americas and Australia (USA (4), Brazil (3), Canada (1), Australia (1)). The many different backgrounds enrich our departments and help to make us aware of cultural differences in, e.g., communication, providing and receiving feedback, and leadership.

The gender composition of our tenured staff is displayed in Table B.4. The overall percentage of female staff members fluctuated between 19%-23% during the assessed time period. In absolute numbers the number of women grew from 11 to 14. During 2015–2020, we were able to hire 11 new female faculty (6 at INSY and 5 at ST – see Table *tab:appointed omitted*). During those same 6 years, in total 8 decided to pursue their career elsewhere (5 from INSY and 3 from ST – see Table *tab:departed omitted*), giving a net growth of 3. The reasons for departure are largely similar to those for all staff discussed in Section 3.3.3. An additional reason, though not explicitly expressed often, will be that 20% is simply a low number, and that many of our female faculty were the first or only woman in a section. This may make switching to a group at another university with more women attractive. To mitigate this problem, despite the highly competitive labor market, we are committed to appointing new female faculty at an increased pace each year in the upcoming years.

7 TU Delft Recognition and Rewards Perspective 2021-2024

A key instrument to attract strong female faculty members is the Delft Technology Fellowship[8] (DTF) program. This annual program seeks to recruit exceptional female talent at all academic ranks and across all TU Delft disciplines. In every round a total of 10 women is appointed. The program is highly competitive, across all disciplines, and not restricted to specific topics, allowing female candidates to put their unique expertise on the table. As Computer Science, we have succeeded in appointing at least one fellow in each round. In 2017 we were able to appoint four, of which two at the level of associate professor. We are committed to play an active role in the DTF recruitment process and incorporate the faculty growth via DTF in our overall growth strategy. This is needed because the financial means for appointing a DTF candidate comes only for 50% from the TU Delft executive board and only for the first five years of each appointment. Since hiring is also ultimately limited by the percentage of women active in the field, we also focus on attracting female PhD students and postdocs, and preparing them for an academic career. Table B.7 shows our intake of new PhD students over the last decade. The average percentage of female PhD students is 22%, with a growth towards 23%, 29%, and 34%, respectively, over the last three years.

In 2017 TU Delft established a Diversity Office. Soon after that, Wioletta Ruszel (department of mathematics), Felienne Hermans, and Arie van Deursen (both SE/ST) took the initiative to launch the *EEMCS Diversity and Inclusion Team* (EDIT).[9] The EDIT proposal was approved and financially supported by the dean in November 2018. While the responsibility for strengthening diversity and inclusion remains within the departmental leadership, EDIT has recurrent meetings with the dean on all issues related to diversity and inclusion, has organized awareness events on, e.g., International Women's Day, and has ensured that diversity and inclusion is sufficiently visible in the faculty-wide employee monitor ('medewerkersmonitor'). CS@Delft continues to play a key role in EDIT, through the involvement of several faculty members including Odette Scharenborg (MMC/INSY, chair), Jan Rellermeyer (DS/ST), and Luciano Cavalcante Siebert (II/INSY). We wholeheartedly share EDIT's mission that we *'want that our students and staff feel at home and valued, regardless of their sex, age, background, sexual orientation or functional disability.'*

3.3.5 Getting the most from PhD training

Related to our aims of *Development of human capital* and *Strong PhD population*, we take the best possible care of our PhD students to become independent researchers and future scientific leaders. We do this through our PhD program, which is organized within the TU Delft's Graduate School. In Appendix E, we describe the way PhD students are embedded in the Graduate School, and explain the PhD development cycle in more detail, including the means of supervision and support and the *Doctoral Education Program*.

We strive toward completion of PhD theses within four years and organize our supervision, official moments of assessment, and requirements on the volume, quality and impact of the conducted research accordingly. This helps to stimulate the effectiveness of the research and supervision processes, enforce focus and maintain a sense of urgency with both the PhD student and the supervisors regarding the timely development of the knowledge, skills and competences. We seek the most suitable publication venues for a PhD student to acquire feedback from the peers (Section 3.3.6), but also help them to learn to publish in high-reputation venues. As an indicator of the quality of their work, the vast majority of our best publications as listed in Appendix C are written with one or more PhD students, usually as first author.

We foster a culture of helping and giving advice to PhD students. In regular group meetings, PhD candidates are encouraged to speak up and engage in discussions, thus learning to present their research and receive constructive comments, in order to equip them to get the most from their participation in (inter)national conferences. We help to the best of our ability in arranging and/or facilitating external internships as a part of the PhD training. We also encourage community building among PhD students and staff members. Once a year, the Faculty Graduate School organizes a faculty-wide PhD event with poster presentations by PhD candidates for their colleagues and staff members and invited talks on aspects of doing research. Furthermore, students can be enrolled in the national Research Schools ASCI[10], SIKS[11], or IPA[12] which organize multiple events per year targeted specifically at PhD students (see also Appendix E).

8 Delft Technology Fellowship

9 EEMCS Diversity and Inclusion Team, EDIT

10 Advanced School for Computing and Imaging, ASCI

11 School for Information and Knowledge Systems, SIKS

12 Institute for Programming Research and Algorithmics, IPA

Table B.3 gives an overview of the PhD students that started four or more years before the end (December 2020) of the assessed time period. There are 22 PhD students (18%) who discontinued, on average after 26 months. In four cases this was after a negative go/no-go advice (see Appendix E for details), while in most other cases the student resigned. Most (16) of the 25 students listed as "not yet finished" started their PhD in 2016, and are on track to defend their PhD thesis in 2021. The TU Delft-wide target for the duration of PhD projects is that 70% of all standard PhD students should defend their thesis within five years. For CS@Delft, this number was on average 48% for PhD students who started in the period 2011–2015, who could finish within five years during the period 2015–2020.[13] As common reasons for delay we see personal circumstances (health, pregnancy), struggling in finishing or publishing a final chapter and/or paper, switch of supervisor (also related to staff turnover), time devoted to project obligations, and doing a PhD in a part-time appointment. Note that the teaching load for PhD students is generally light (at most 10-15% of the appointment), so we do not see this as a common cause of delay. In Chapter 6 we identify a number of measures we are currently implementing in order to move CS@Delft towards the target of 70% completion within five years.

3.3.6 Publication strategy

Critical for pursuing our strategic aim *World-class research portfolio* is our publication strategy. This strategy has two main aspects. The first aspect is that we primarily target high-reputation publication venues that maximize impact and visibility of our research results. The list of our key publications in Appendix C is an indicator for the type and level of the scientific journals and flagship conferences we typically publish in.

The second aspect of our publication strategy is our value-based philosophy. Looking beyond the widely used quantitative indicators, we see great publication opportunities in specialized (often smaller) venues, which allow for dissemination of early results and detailed discussions with peers working on similar topics and pointing to new scientific directions. This is also related to the fact that most of our publications are written together with our PhD students, leading to venue choices that are most helpful to PhD students at the given moment in their PhD projects. For example, a starting PhD student may benefit more from a workshop paper with rapid feedback than from an attempt to directly target the flagship conference or a top journal with a possibly long publication pipeline. Furthermore, for our interdisciplinary research, we seek to publish our results in computer science venues as well as in domain-specific scientific venues, for instance in the areas of health, transportation, or energy. Thanks to our collaborations with internationally leading domain experts, we can target top venues in these domains as well.

We have also been raising awareness among our faculty members on the impact of different author roles in articulating their own academic profile. In this way we aim at helping our faculty members to realize whether they prefer to profile themselves as the main contributing principal investigators (first or single author), as contributors to larger scientific teams (middle author) or scientific coaches, supervisors and managers (last author). Thereby, we urge our junior faculty members to use the room they have during their early career stages to establish themselves in their peer community through first- or single-author papers. This will help them in building up their scientific reputation, especially if they intend to apply for personal grants, like NWO VENI/VIDI/VICI or ERC grants. Naturally, there is no place for "free riding" in our academic culture, and all authors must meet common criteria for authorship, such as those formulated by the ACM.[14]

3.3.7 Open Science

Related to the aim *Advancing the organization and academic culture*, an important aspect of the academic culture at CS@Delft that we nourish is open science. We consider it a key ingredient to scientific success, also contributing to our aim of a *World-class research portfolio*. We therefore embrace the ambition of the TU Delft Open Science Program[15] as our own, namely to be an international open science frontrunner.

Regarding **Open Access**, we use the TU Delft Research Repository[16] to offer (links to) open access versions of our papers. TU Delft has expressed the ambition that 60% of peer-reviewed publications should

13 Note that the standard SEP Table B.3 includes the starting year 2016, which for a goal of 70% within *five* years leads to incomplete data since the fifth year (2021) is not yet finished at the cutoff date for the table (ultimo 2020). This is why we offer this data for the period 2011–2015, and why the number (48%) is slightly higher than that in the table (43%).

14 ACM Criteria for Authorship

15 TU Delft Open Science Program

16 TU Delft Research Repository

be available through open access in 2018, which should grow to 100% by 2024. In case a publisher is not open access, authors offer a link to a self-archived version, often via arXiv, but also by directly uploading it onto the institutional repository. With this, we also fulfill the mandates of funding agencies (NWO, EU, ...). While we aim at fulfilling the ambition of TU Delft regarding open access (Section 4.1), we do not adopt a policy of publishing in open access venues exclusively. We believe researchers should select the venue that matches their discipline best, also considering the aspects of our publication strategy (Section 3.3.6). In the upcoming years, however, we hope not to be required to make this choice. The ambitious open access plans of the *cOAlition S*[17] founders will help to move more and more publishers to an open access offering. To speed up this process, we exercise influence on publishers that are not yet open access to change their course, like for example at ACM via ACM SIGPLAN (Eelco Visser, PL/ST) and ACM SIGSOFT (Arie van Deursen, SE/ST).

Regarding **Open Source**, we share many of our research prototypes and demonstrators as open source. Since our intent often is to involve external developers, we typically do so on shared platforms like GitHub.[18] When relevant, we submit our systems to (ACM) *Artefact Evaluation Tracks* to badge them as *functional*, *reusable*, or *available*. We abstain from open source availability, for example, if we see opportunities for a start-up requiring protected software, or in case the software is developed in close collaboration with industry.

Regarding **Open Data** we encourage our researchers to make curated datasets available either in the 4TU Data Center,[19] or in widely used repositories like Zenodo,[20] so that the data remains Findable, Accessible, Interoperable, and Reusable (FAIR). To support researchers with data management, in 2019 the EEMCS faculty has appointed a *data steward*, in line with the general TU Delft policy to have one data steward available for each faculty.

Open data and open source artefacts offer a strong impulse to **replication** and **reproducible** AI.[21] We seek to provide our papers with artifacts to make replication as easy as possible. Furthermore, we conduct replication ourselves, either as standalone studies or as baselines for our experiments. Replication also forms a natural ingredient in our work with bachelor and master students, either in courses or in their thesis projects. A prominent example is `ReproducedPapers.org`, a collection of papers reproduced by students from a variety of TU Delft courses.

3.3.8 The Teaching-Research Nexus

Related to our aims *Tight connection between research and education*, *World-class research portfolio*, and *Strong impact in key societal sectors*, we organize and conduct our education at both the bachelor and master level such to be inspired by and related to the research we do. In this way, we can guarantee that the students we educate are optimally equipped with the newest knowledge and insights in computer science. At the same time, we can help our staff members to let their teaching activities benefit their research.

The first aspect of our strategy to achieve this is to let our students *participate* in research, starting at the bachelor level. In their individual thesis project (10 weeks, 15 EC)[22] at the end of their third year, students work in groups of 4-5 under the supervision of a faculty member and often a PhD student. Assignments include replications of related papers, alternative solutions to a similar problem, and addressing different stages of a shared problem. For strong students, this scheme permits publishing a paper while at the undergraduate level, for which we can also appoint the student as a *research assistant*, for example in the summer period. At the master's level, the thesis project takes nine months (45 EC) and has a strong research focus. The assignment is more substantial, and fits in a bigger research program carried out by a faculty member and one or more of their PhD students. For around 20% of the theses, the results are published in a refereed venue, ranging from simple workshops to the most prestigious conferences and journals.

Furthermore, we aim at assigning courses to faculty members such that they best match their expertise and interest. This helps to reuse the deepened insights obtained through teaching in their research. For educational responsibilities with less emphasis on research, our faculty members can work with the newly

17 cOAlition S

18 GitHub

19 4TU Data Center

20 Zenodo

21 AAAI Reproducibility Checklist

22 This was introduced in the curriculum renewal 2018/2019, and run first as a pilot in academic year 2019/2020. Results from 2021 can be viewed at https://cse3000-research-project.github.io/2021/Q4.

established CS Teaching Team (see Section 1.2.2). Lastly, given Delft's strong position in computer science education as well as in online learning through its participation in EdX,[23] we have adopted education itself as a topic of research. During the assessed time period, we appointed Marcus Specht (WIS/ST) as a professor of digital education, and as director of the Leiden-Delft-Erasmus Center for Education and Learning (LDE-CEL). Through this appointment, CS@Delft has expanded its research scope to include learning analytics and EdTech.

3.3.9 Internal collaborations

Many key challenges computer science is addressing nowadays are interdisciplinary and can only be addressed successfully by combining computer science with other expertise. Related to the aims *World-class research portfolio* and *Strong impact in key societal sectors*, our strategy to foster interdisciplinary research is twofold.

First, we have invested significant effort in building up and participating in a number of joint research programs, which we describe in more detail in Appendix G. These research programs, sometimes embodied in a formal institute, are either at the level of the EEMCS faculty, where CS@Delft collaborates with the departments of applied mathematics or electrical engineering, or at the level of the university, in which we combine our forces with the disciplines from other faculties.

Second, we took the initiative to strengthen the visibility of our data- and AI-related research, and to align this research across the university. While started bottom up, this eventually helped to shape a coordinating initiative for AI-related research and education (AIDU) in Delft, out of which the TU Delft AI initiative evolved. Within the realm of this organization, 24 interdisciplinary TU Delft Artificial Intelligence Labs (so-called DAI-Labs) are emerging, combining the research 'in-AI' with research 'with-AI' to address various engineering and societal challenges. Currently, 16 of these labs have been realized, with CS@Delft participating in 6 labs (Appendix G.3). Substantial funding dedicated to establishing the DAI Labs is an indicator of the strong support of the TU Delft Executive Board for AI across the university. Recent appointment of Geert-Jan Houben (WIS/ST) as the Pro Vice Rector for AI, Data and Digitalization, reflects a pivotal role CS@Delft has played in shaping up AIDU and implementing the TU Delft AI initiative.

3.3.10 External collaborations

Building up and maintaining an effective portfolio of external collaborations with societal partners is critical to reach our aim of *Strong impact in key societal sectors*. Collaborating with academic institutions helps us build up larger momentum and critical mass for influencing the national and international research agenda as well as the national and international research funding mechanisms. Through the collaborations with non-academic parties, the "use" for our research results can be identified and our results can be verified. Furthermore, this is also the way to realize effective technology transfer. Our vision on technology transfer is that the most promising route to success is by closely working together, right from the start, with societal partners. Transfer is not a step at the end of a research project, but a continuous process. It involves problem selection, data availability, problem solving in context, as well as result evaluation in that context. It requires a relation of trust, and a long-term perspective spanning multiple years.

To build up and maintain external relations, we invest in building up sector-specific expertise in **key sectors** mentioned in Section 2.2: *transport and logistics*, *energy*, *health and well-being*, *safety and security*, *finance*, *online education*, *creative industry* and *smart cities*. We do this either through the institutes (Appendix G) that also link to external partners (e.g., TU Delft Bioengineering and AgTech institutes), or through dedicated projects with sector-partners, typically through direct research funding in public-private collaborations (Appendix *app:outcomes omitted*).

The expertise per sector is then deployed in building up collaborations at various levels. At the **regional** level, we highlight here different collaboration lines involving TU Delft, Leiden University, Erasmus University Rotterdam and the two university medical centers, Leiden University Medical Center (LUMC) and Erasmus Medical Center (EMC). Within the *health* sector, TU Delft, Erasmus University and EMC established the *HealthTech convergence*[24] (key role of Marcel Reinders, PRB/INSY) that offers unique possibilities for our work on bioinformatics and AI for medical applications. Also in the health sector, and involving in addition Leiden University and LUMC, *Medical Delta*[25] program was established (Boudewijn Lelieveldt and

23 EdX online education platform

24 HealthTech Convergence

Marcel Reinders, both PRB/INSY). Looking broader than the health sector, the *AI, Data and Digitalization convergence* (key role of Geert-Jan Houben, WIS/ST and Arie van Deursen, SE/ST), was initiated, where the combination of 'in-AI' research and 'with-AI' research in various domains is taken up involving these institutions. These collaborations build on the regional *Leiden-Delft-Erasmus (LDE)*[26] alliance that has a long history of collaboration, based on complementarity of research and education expertise and with a strong focus on societal impact. This alliance has also brought up other joint initiatives in the areas of smart cities (Geert-Jan Houben, WIS/ST), learning analytics (Marcus Specht, WIS/ST), and cybersecurity (Zekeriya Erkin and Inald Lagendijk, both CYS/INSY). We take an active role in making the different collaboration lines and programs a success.

At the **national** level, our strategy is as follows:

- We maintain our strong institutional collaboration and leadership in the 4TU federation with TU Eindhoven, University of Twente and Wageningen University in the field of ICT (NIRICT) and Human-Technology Interaction (HTI). More information about the objectives and strategy of the 4TU federation and the 4TU institutes we contribute to can be found in Appendix G.
- We engage and take leadership roles in the process of public-private research coordination and programming and institutional collaboration, such as through NWO program committees and board memberships, ICT Research Platform Netherlands (IPN) and national programs board memberships (COMMIT/, NBIC, DTL).
- The Amsterdam Metropolitan Solution (AMS) initiative of TU Delft together with Wageningen University and MIT (USA) provides us future means to link well to ICT-entrepreneurial activities in Amsterdam.
- Since 2019, we have been active partner in the ICAI (Innovation Center for AI[27]), by establishing an ICAI Lab in the *finance* sector funded by the ING Bank (Arie van Deursen, SE/ST, as PI). In the course of 2020, the process of setting up another two ICAI Labs has been finalized (though starting only in 2021) in the domain of biotechnology, funded by the DSM company (Marcel Reinders, PRB/INSY, as PI), and with Booking.com and the University of Amsterdam in the domain of recommendation systems (Frans Oliehoek, II/INSY, Matthijs Spaan, ALG/ST).

Successful collaboration with industry requires a clear vision on **intellectual property** (IP). In line with our open science policies (Section 3.3.7), where possible we make software tools and data openly available. For many of our collaborators (e.g., Google, Facebook, and Microsoft), open source software is a natural form of sharing knowledge and increasing impact. In case consortia or individual companies require explicit protection of IP, we work with the faculty's project office and the TU Delft Innovation & Impact Center[28] to arrive at terms that are beneficial to all parties. In line with our research mission, while research results can be embargoed for a short period of time to organize, e.g., patent protection, in all contracts it is explicit that researchers (such as PhD students) have the right to eventually publish their results.

We collaborate with the Yes!Delft incubator to create a fertile environment for entrepreneurship initiatives. An important instrument for early stage ideas are the NWO Take Off Grants. Furthermore, several faculty members have previous experience as co-founder, and are always willing to support new ventures. Successful earlier startups (like SIG, SocialGlass, Infotron) continue to be research partners in various projects.

Finally, in order to strengthen the collaboration with external institutes, strategic part-time (associate) professorships were established, as well as visiting professorships (*Table Staff exchange omitted*). Under such constructions, employees from external organizations become affiliated with CS@Delft for a fraction of their FTE. Such appointments, especially at the level of full professor, are realized following standard strict selection and appointment procedures, similar to the full-time employed faculty members.

3.3.11 Research Funding Portfolio

Sufficient income via the second and third money tier are critical for building up research capacity in the form of PhD students, postdocs and software engineers. In this sense, successful research funding acquisition

25 Medical Delta program

26 Leiden-Delft-Erasmus alliance

27 ICAI - Innovation Center for AI

28 Formerly called the TU Delft Valorisation Center

is the necessary condition for reaching the aim *World-class research portfolio*. Since our projects are often established with external, societal partners, our efforts in this direction also serve the aim *Strong impact in key societal sectors*.

Since we are part of a technical university, we have a focus on the engineering aspects of computer science. Hence, our research approach, while fundamental, is also largely use-inspired. We therefore encourage and formally support (with a strong top-down financial, lobby, marketing and communication assistance) initiatives to form or participate in national and international project consortia, like for instance those funded by EU (e.g., H2020, Horizon Europe) that require fundamental computer science research inspired from generic needs or from modern industrial and societal challenges, often in a systemic and technological setting. In line with this, we also help and encourage our staff members towards acquiring direct funding from companies, either on ad-hoc basis or as a part of a larger (strategic) framework (such as ICAI Labs discussed in the previous section). Furthermore, we engage in curiosity-driven research. We stimulate projects that can accommodate out-of-the-box research ideas introducing new or redefining existing research directions in the field. Such projects are typically those funded by the ERC or NWO, either targeted by individual staff members (personal grants) or by larger consortia (e.g. NWO Gravitation Program).

An essential aspect of our strategic process on research funding acquisition is balancing the second and third money tier projects in order to (a) allow faculty members to deepen their research and grow their reputation through personal grants, (b) influence the national research agenda by introducing new fundamental research directions via top national scientific consortia, and (c) build up and maintain collaborations with societal partners. Clearly, (a) and (b) push towards the second money tier and (c) towards the third money tier projects. Strategically, focusing on (a) and (b) serves building up and maintaining our scientific reputation and impact, while (a) in addition gives individual research space for our staff members. Focusing on (c) gives valuable access to interesting contexts, which can serve as both problem provider and experimental play ground ("industry as a laboratory"), and which often also offer relevant data.

Ideally, all three categories (a)-(c) of projects should be represented by good shares, both in terms of the number of projects and funding volume. This is because of our strategy to develop and contribute along all three dimensions. However, achieving this faces challenges, due to which developing our research funding acquisition strategy is an ongoing process. These challenges are as follows:

- Working on personal grants requires substantial amounts of time and energy, which is difficult to find in view of overfull task portfolio of an average staff member. While we have, despite this, succeeded in increasing the numbers of NWO personal grants at the junior and mid-career level (NWO VENI and VIDI, respectively), more effort is needed to create conditions and improve success rates of the personal grant at the highest level (ERC, NWO VICI). In order to help this, we have introduced various support mechanisms, including the help of our TU Delft Innovation & Impact Center and an external advisor with writing proposals and helping each other by establishing internal ad-hoc reading and assessment committees mimicking the selection process at NWO.

- Building on the strategic roles of our staff in various national committees and on their visibility as prominent scientists, we invested a long-term effort in building top national research consortia, which brought us participation in two prestigious Gravitation Projects (Hybrid Intelligence and BRAINSCAPES) (see footnote on page vii). We learned, however, that in order to repeat this success and expand it towards leading such consortia requires enormous level of dedication of staff members, which again is difficult to realize in terms of the available time and capacity in view of the complexity of their task portfolios.

Table B.2 summarizes the funding of the INSY and ST departments during 2015–2020. Despite the challenges, the financial figures and analysis provided in this table and other tables in Appendix B, supported by the projects listed in Appendix *app:outcomes omitted*, reflect our success in attracting research funding across the scope. While the balance between the funding through the second and third money tier (research grants and contract research, respectively) has varied over the years, it is important to note that since 2016, INSY and ST departments have attracted each more research funding per year than the direct funding (first money tier) they receive. Achieving this while growing in our staff, together with the highest-ever project intake in 2020 (Table B.6), and the strategic investments prepared in the assessed time period to become effective in 2021 (e.g., ICAI Labs with DSM and Booking.com), give us confidence in building up this success even further in the upcoming years. We reflect on the financial development of CS@Delft and the INSY and ST departments individually in more detail in Chapter 6 and point there also to our future ambition and strategy regarding project intake.

4 Key Performance Indicators

In this chapter we give an overview of the key indicators through which we assess our successes, as well as the main scientific accomplishments in terms of these indicators. The indicators are listed in Appendix A. Following the *Strategy Evaluation Protocol*[1], they fall into the two categories **research quality** and **societal relevance**. These in turn are grouped into demonstrable products and their use, as well as marks of recognition, for either peers (research quality) or societal target groups (societal relevance). Complementary to this, the appendices B, C, and *app:outcomes omitted* provide lists of key outcomes in terms of the indicators. In Chapter 5, we illustrate our scientific accomplishments by highlighting key results across our research themes, with a focus on their scientific and societal impact and outlook. We made all quantitative and qualitative evidence on the past performance traceable to either the INSY or ST department.

4.1 Research Quality

Research products for peers Our research products for peers include (open access) scientific publications, dissertations, as well as software platforms, datasets and tools. According to our publication strategy, we target venues that match our values (Section 3.3.1), but that are preferably also of (objectively) high-reputation. Our understanding of high reputation can be quantified reasonably well by the Google Scholar Metrics[2] parameters. We therefore used these parameters to mark our key publications in Appendix C. For each venue of a key publication, we report its rank in the relevant research area as well as its h5-index[3] (if available). For each paper, we report its citation count as one indicator of impact, unless a publication is too recent to attract a significant number of citations.

The list of our key publications is representative for our overall publication output and shows that we have been consistently successful in publishing in major venues across the CS domain, in major venues in other domains (through our interdisciplinary research, e.g., in transportation, education and learning), but also in major scientific venues in general (e.g., Nature, Scientific Reports, Cell). We also monitor and foster our adherence to our open science strategy and the ambition of TU Delft in this respect (see Section 3.3.7). Since this ambition was introduced in 2016, the percentage of our open access papers thus available has grown from 46% (2016) to 80% (see Table B.5). Open access applies to all our key publications (Appendix C). Furthermore, since our dissertations are typically based on publications, their impact can be inferred from that of their underlying publications, and they are all available via open access through the TU Delft Repository.

The two CS departments have built up a strong track record of making their software platforms, datasets and tools available to the international scientific community (*Table Software and datasets for peers omitted*). Several of these are widely deployed in the community, reflecting our ambition and success in making our scientific findings tangible and facilitating and boosting scientific research in general.

Use of research products by peers Use of publications is visible through citations, which we report for our key publications (Appendix C). While the use of software platforms, datasets and tools delivered in the assessed time period is indicated in Appendix *app:outcomes omitted*, their impact can best be assessed over a longer time period. Furthermore, many of such products demand a long-term investment – work on some of our most successful tools started over a decade ago. These effects may not yet be sufficiently visible for all the reported results in this category reported in Appendix *app:outcomes omitted*, especially for those, that were generated recently. Therefore, next to the results we delivered over the past six years, we find it illustrative of our performance and international scientific standing to also show here examples of the lasting impact on the scientific community we achieved through the tools we made available further in the past, but which continued to be downloaded, used and/or cited over the assessed time period:

1 Strategy Evaluation Protocol

2 Google Scholar Metrics

3 The venue's h-index over the past five years, which is the basis for the rank

- PRtools[4] is a widely used MATLAB toolbox for statistical pattern recognition (received over 100,000 downloads),
- The tSNE (more than 20,000 citations of the original paper[5], available from GitHub[6]is a data analysis and visualization tool that detects and projects high dimensional data into low dimensions. The tSNE is used in a wide range of research fields and has become a standard tool in many data exploration environments, such as R or scikit-learn.
- Our open source peer-to-peer research platform Tribler has received almost 2 million downloads. Its development started in 2005, and it continues to be a key experimental platform for blockchain research (see Section 5.4.1).

Marks of recognition from peers In terms of personal recognition, in the assessed time period, Alan Hanjalic (INSY/MMC) became IEEE Fellow and Pablo Cesar (INSY/MMC) became ACM Distinguished Member. Furthermore, both Cynthia Liem (MMC/INSY) and Alexandru Iosup (DS/ST) became members of the Young Academy of the Royal Dutch Academy of Sciences (*Table Professional standing omitted*). An important achievement in this respect also relates to the Netherlands Prize for ICT Research, which TU Delft researchers have won three times since 2015 (Alexandru Iosup, DS/ST, Elmar Eisemann, CGV/INSY and Pablo Cesar, MMC/INSY). These and other research awards and honors can be found in *Table Awards and honors omitted*.

At the international level, a key mark of recognition comes from the general and program chair responsibilities in flagship conferences that are entrusted on our staff by the research community. As shown in *Table General/program chair of top conferences omitted*, many conferences took place during the reporting period with CS@Delft staff in leadership roles, including ACM Multimedia, EUROCRYPT, Eurographics, CIKM, ESEC/FSE, AAMAS, and SPLASH/OOPSLA. The list of chair roles is a sign that this recognition is not concentrated with a few senior members, but spread across all sections in CS@Delft.

Appendix *app:outcomes omitted* highlights additional achievements, including keynotes, editorial roles in major scientific journals (editorial boards, steering committees), best paper awards, challenge awards and test of time awards for papers 10 years after publication. Particularly relevant are personal grants, which not just recognize expertise and leadership, but also the ability to formulate and execute a ground breaking research programme. Many of our personal grants come from the NWO Veni/Vidi/Vici scheme, as listed in *Table Personal grants omitted*.

When compiling the lists in Appendix *app:outcomes omitted*, we were selective and included only the evidence of recognition from the major venues and organizations, while aiming at a list that is representative of the breadth of our activities. It is important to see that, next to the fact that a large number of our staff members have received (inter)national recognition, this also holds for many of our recently hired tenure track assistant professors.

4.2 Societal Relevance

Research products for societal groups Possibly the largest impact on society is made by the people we train, be it at MSc, PhD or at postdoc level. A large majority of our PhD students have their next job in industry, both at large companies (e.g., Amazon, Facebook, Airbus, Microsoft, Philips, Google, ING) as well as small and medium enterprises, including our spin-offs. Even though our intake at the PhD level is highly international, many stay in The Netherlands strengthening the potential for innovation here.

Furthermore, our research products have become a part of the societal and economic infrastructure. Regarding the former, we mention from the examples listed in Appendix *app:outcomes omitted* the flooding visualization and simulation software we co-developed, which was used to illustrate and commemorate the 1953 flooding in Zeeland. It is currently an installation in the Science Center in Delft and the 'Watersnood-museum' in Zeeland. Regarding the latter, the CS departments encourage staff members to pursue science and technology transfer by means of spin-offs. An example is Apta Technologies, co-founded by Sicco Verwer (CYS/INSY), which helps companies understand software through their log data and identify anomalies, bugs, malicious events all the way down to root causes. Another example is Generatrix, co-founded by Lydia Chen (DS/ST), which offers tools and services in the area of privacy preserving synthetic data generation.

4 PR Tools
5 L. van der Maaten, G. Hinton: Visualizing data using t-SNE, *J. of Machine Learning Research* 9 (2579-2605), 85, 2008
6 t-SNE @ GitHub

Finally, there has been extensive effort towards public outreach. The coverage included major newspapers, TV and radio channels, but also unconventional settings, such as the event series 'Professors in the Theater' (Geert-Jan Houben, ST/WIS, Alan Hanjalic and Cynthia Liem, MMC/INSY)[7], where the links between societal phenomena and Computer Science is explained to and discussed with a broad audience. While we see our senior staff being contacted regularly as experts for various commentary in media, we also see large enthusiasm at our junior staff to share their views with a broad audience. They do this to discuss research results, as well as to contribute to societal debates, for example in the areas of fairness in AI (Cynthia Liem, MMC/INSY, Nava Tintarev, WIS/ST) or blockchain (Johan Pouwelse, DS/ST). Furthermore, we are active on social media such as Twitter (several faculty members have thousands of followers), LinkedIn, personal blogs, and Q&A sites like StackOverflow (where we have reached millions of software developers).

Use of research products by societal groups A key means to enable the use of scientific results in society is close collaboration with industry and societal organizations. Taking our research in Bioinformatics as an example, we highlight here our ability to bridge a technical gap in clinical sciences, which has contributed to major scientific breakthroughs, amongst others, the discovery of a new Alzheimer-associated gene[8], or the development of tools, like WISECONDOR[9], which is used as standard analysis tool for NIPT analysis in Dutch, French and Korean hospitals.

Another means to stimulate and facilitate the use of research in society are public-private partnerships and contract research. We have had many public-private collaborations, through EU projects (H2020, Marie Curie, ERASMUS+) or Dutch national collaborations (NWO/NWA). For example, in the NWO Perspective program CAS on sports, we have collaborated with many companies and research groups to prevent sports injuries and to improve performance. Our focus was on artificial intelligence, building mental models of sporters and providing personalized feedback. We are particularly proud of the projects (e.g. H2020 projects TROMPA, FASTEN), in which we came up with the initial idea and played a leading role via forming and coordinating the project consortium to a successful completion.

Finally, we consider professional education as an important tool in transferring our knowledge to a non-academic audience. A good example of this is our leading role in the Cybersecurity Academy in The Hague (involvement of CYS/INSY led by Zekeriya Erkin), in which we organize (together with Leiden University) a professional cybersecurity master. We have also leveraged the TU Delft Extension school[10], where we have offered popular online courses on Automated Software Testing (Arie van Deursen, Mauricio Aniche, SE/ST), data analysis with Excel (Felienne Hermans, SE/ST), AI in Practice (with ICAI, Arie van Deursen, SE/ST, Asterios Katsifodimos, WIS/ST, Cynthia Liem, MMC/INSY), Unix data analysis (Diomidis Spinellis, SE/ST), big data processing (Jan Rellermeyer (DS/ST), Asterios Katsifodimos (WIS/ST)), and Globally Distributed Software Engineering (Rini van Solingen, SE/ST). Collectively, these courses have reached tens of thousands of practioners.

Marks of recognition by societal groups Staff exchange serves as a clear mark of recognition and as a fruitful way to embed our own staff in industry or societal organizations, as well as vice versa. This has resulted in various part-time professor appointments (from such organizations as LUMC, NKI, TNO, and Keygene) as well as industry appointments of our own staff (e.g., TNO, Dutch Railways) (*Table Staff exchange omitted*).

Another important mark of recognition is formed by the advisory and leadership roles we fulfil in industry, government and various other societal bodies. As the most prominent example, Inald Lagendijk (CYS/INSY) has fulfilled a leading role in the Dutch AI Coalition[11], a massive national effort aimed at bringing together knowledge institutes, industry, and societal partners to advance the development and application of artificial intelligence in The Netherlands. This builds upon Lagendijk's long experience bringing academic and societal partners together nationally, in particular through his role as Captain of Science of the Dutch Top-Team ICT (Dutch Digital Delta)[12]. While an analysis of the past reveals that the strategic roles have been

7 Professors in the Theater

8 Discovery of an Alzheimer gene

9 WISECONDOR

10 TU Delft Extension School

11 Dutch AI Coalition

12 Dutch Digital Delta

pursued mainly by senior faculty members, *Table Memberships of management and advisory bodies omitted* show that in the meantime, many more of our staff members have assumed roles in various management and advisory committees and societal bodies, as a result of, but also stimulating their growth in terms of academic seniority. This ongoing transfer of strategic responsibilities across generations reveals a strong and broad potential of CS@Delft to keep influencing the societal context of CS for many years to come.


Best Reviewer Award
presented to
Jaehun Kim
for outstanding reviewing
steinberg

5 Scientific Highlights

In order to highlight the quality and societal relevance of our research over the period 2015–2020, we offer narratives for our five joint *research themes* (see Section 2.1), that integrate the accomplishments from different sections across the two departments. Each narrative is organized around specific research topics of our current focus mentioned in Table 2.2, for which we pick a number of representative results and discuss their origin, scientific and societal impact, and outlook. The corresponding key publications are listed in Appendix C, hyperlinked to their open access version in the TU Delft repository. Furthermore, the narratives are supported by the data collected in Appendix *app:outcomes omitted* and the four case studies discussed in Appendix D.

5.1 Decision Support

5.1.1 Interactive and hybrid intelligence (II/INSY, PRB/INSY)

The research on interactive and hybrid intelligence focuses on fostering long-term interaction and collaboration between humans and artificial intelligent agents in so-called human-agent-robot teams. The main challenge lies in the question how agents can learn about and effectively use the value systems of humans to support the humans in achieving their goals.

Representative results: The most representative result is the strategic NWO Gravitation Project *Hybrid Intelligence* (19 M€ subsidy, 41.9 M€ in total) funding 10 years of research (2020-2030). The work on acquiring this project involved significant effort of our scientific staff, not only regarding the shaping of the research program, but also regarding national coordination by Catholijn Jonker (II/INSY) as Co-PI.

Scientific impact: We mainly target publication venues that set the future research agenda of the community. A good example of such a venue is the Blue Sky track of AAMAS, to which we have contributed regularly and successfully over the past years. For example, the contribution of Murukannaiah et al. (key publication [1]), in which we sketch ethical AI as a multiagent (as opposed to single agent) challenge, highlighting the sociotechnical challenges of embedding AI in societal applications, was awarded the AAMAS'20 Blue Sky paper award. Furthermore, the position paper (key publication [2]) we contributed to, has put forward a new perspective on 'Hybrid Intelligence' which details components of AI that are necessary to work with and augment what humans can achieve on their own. Also, our work on interactive and affective robots and agents was published in major journals and conferences, including IEEE Transaction on Affective Computing, Machine Learning, and IJCAI. Finally, our software and data sets are impactful. For many years we set the research agenda on automated negotiation through our open source negotiation framework GeniusWeb[a]. Another example is the new dataset (MatchNMingle[b]), which is the largest collection of annotated human social behavior in a romantic setting.

Societal impact: We highlight here our projects and collaborations with external parties. The first example is *Virtual eCoaching and Storytelling technology* for post-traumatic stress disorder treatment (VESP), in which we investigated a virtual coach that can assist in the execution of treatment of PTSD patients. In the Sleep care project we designed and evaluated mobile phone-delivered cognitive behavioral therapy for insomnia Furthermore, we collaborated with Latent Logic (now Waymo, UK) on a project that focussed on learning from demonstration, in which we learned realistic driver behavior based on camera data of intersections.

Outlook: Interaction, and specifically the combination with AI, such as in hybrid intelligence settings, is expected to be of increasing importance. In order to position ourselves optimally in this respect, we have recently expanded our scientific staff with Frank Broz and Luciano Cavalcante Siebert to address the challenges of, respectively, interactive AI for social robotics and the development of practical methods to keep AI under meaningful human control.

a GeniusWeb
b MatchNMingle dataset

5.1.2 Intelligent agents and multiagent systems (II/INSY, ALG/ST)

The aim of this topic is to design and understand fundamental properties of planning and coordination protocols and algorithms, and to develop agents and methodologies for social simulations.

Representative results: Partially Observable Markov Decision processes (POMDPs) form a canonical framework for planning problems when an agent is uncertain about the environment's state. In 2017 we presented the fastest optimal solver for POMDPs currently available (key publication [3]). The corresponding PhD thesis was selected as a runner-up for the ICAPS-21 Best Dissertation Award. Another way to speed up solving such planning problems is via abstraction, such as for Monte Carlo tree-search methods (key publication [13]). Intelligent coordination of the plans of agents can lead to higher efficiency, as shown by coordinating the en-route charging of electric vehicles (key publication [5]). Often, our work is simulation-based (rather than based on datasets) and we released an open source toolbox[a] to support such developments.

Scientific impact: Catholijn Jonker (II/INSY) is president of the International Foundation of Autonomous Agents and Multi-Agent System (IFAAMAS) and Mathijs Spaan (ALG/ST) and Frans Oliehoek (II/INSY) are IFAAMAS board members and Associate Editors of the Journal of AI Research. We organized the main international conference on automated planning and scheduling (ICAPS) in Delft in 2018. Spaan will be program co-chair for AAAI in 2022.

Societal impact: Our expertise in decision making is particularly sought after. For instance, electricity grid operators need AI decision making to enable the energy transition, which we elaborate in Case Study D.4. The Dutch railways use algorithms to schedule the routing and servicing on shunting yards. In the Horizon 2020 project MEDIATOR we collaborate with developers of advanced automotive automation systems and human-factors experts to build an AI system that can safely switch from automated to manual driving and vice versa. Our team provides its algorithmic decision-making core. Our interdisciplinary TU Delft team won in the Second International Autonomous Greenhouses International Challenge in 2019-2020, where optimisation, learning and control algorithms in the form of a user-assistive agent automated the control of a real-world greenhouse climate over six months.

Outlook: In the recently founded ICAI Labs with DSM and Booking.com (formally starting in 2021) we aim to develop novel algorithms for planning problems for agents in dynamic situations. We expect these strategic projects to give a new boost to our research on this topic, which we will leverage through the upcoming initiatives, enabled by the Dutch AI coalition and the Leiden–Delft–Erasmus AI convergence (Section 3.3.10), bringing together researchers from different fields to address algorithmic as well as socio-technical challenges of embedding AI in society, and of making automated decision making more reliable.

a MADPToolbox

5.1.3 Computer vision and machine learning (PRB/INSY)

With our machine learning research we aim to gain deeper understanding of conditions in which learning can work. Deep learning is dominating our computer vision research where we focus on adding prior knowledge to deep learning. Through new hires, we gained critical mass on (deep) reinforcement learning.

Representative results: Six years ago, we developed the first semi-supervised learner ever to come with guaranteed performance improvements (key publication [6]) and recently, we did the same for the field of domain adaptation. An important finding at the foundations of AI is our work on learning curves and empirical risk minimisation that showed that more data does *not* always lead to better performance (key publication [7]). For computer vision applications we have shown that visual inductive priors allow a reduction in valuable training data by replacing data by knowledge (CVPR20, ECCV20). An example application area is non-contact visual human physiology estimation, where we proposed Eulerian motion representations in video that can magnify tiny motions in the presence of camera motion or depth discontinuities (key publication [8]). On reinforcement learning we developed formal models and approaches for multi-agent active perception, interactions between intelligent vehicles and traffic light control.

Scientific impact: The Computer Vision Lab consistently publishes in top venues, fulfills area chair roles (CVPR/ICCV), receives best reviewer awards and best paper awards, and has organized five international

workshops. We combined research with education by promoting reproducibility of research results to students[a], which has already led to around 100 reproductions of deep learning papers. We published several open source packages, e.g. on motion magnification[b],[c] and knowledge-based deep learning[d] and provide continuing contributions to PRtools (see Section 4.1), which is also at the basis of our machine learning educational activities.

Societal impact: Our Computer Vision research generates media attention, such as national newspaper articles and national popular science ('Universiteit van Nederland'). We are proud to have supported a science refugee through the NWO Hestia Impulse for Refugees in Science.

Outlook: As causal inference is gaining momentum, especially in the health domain, we have expanded our machine learning research in this direction with the hiring of Jesse Krijthe. Computational efficiency and sustainable learning are gaining traction within the deep learning community, which we strengthed with the hiring of Silvia Pintea. To bridge the Socially Perceptive Computing Lab and the Computer Vision lab, we invested in recognizing gestures and gaze with the hire of Xucong Zhang. We are also exploring spiking neural networks as well as time series data analysis with application to detect, for example, cardiovascular events.

---

a reproducedpapers.org
b Video Acceleration Magnification
c Phase based video motion magnification
d Active Decision Boundary Annotation with Deep Generative Models

5.1.4 Search and Recommendation (MMC/INSY)

The guiding research question here is how to provide effective, intuitive and (socially) responsible access to and interaction with information stored in large (multimedia) data collections or available on social media platforms in a variety of use scenarios involving both search and recommendation.

Representative results: We further strengthened our reputation in the domain of video search by pioneering the effort of making video search engines aware of the user's search intent - the reason *why* a user consults a search engine in the first place, which would lead to a higher overall utility of the search results for the user. Furthermore, we have intensified research on cross-modal retrieval, generation and captioning (key publication [9]). Boosted through the collaboration with the University of Electronic Science and Technology, Chengdu, China (Alan Hanjalic), and with the University of Illinois at Urbana-Champaign, USA (Odette Scharenborg), this research aims at automatically inferring semantic links between pieces of information conveyed by different modalities, like between text and image/video or between speech and images.

Scientific impact: Our research on video search intent brought our PhD student, Christoph Kofler the ACM SIGMM Best PhD Thesis Award in 2016: *'This Ph.D. Thesis On Deciphering User Intent Could Be the Next Breakthrough In Multimedia Search Engines'* [a]. Furthermore, our research on cross-modal retrieval has already led to numerous and well cited publications in major conference and journal venues, attracting close to 900 citations and bringing two best paper awards (ACM Multimedia 2017 and IEEE International Conference on Multimedia Big Data 2019). Finally, our strong track record, visibility and reputation in the domain of search and recommendation has contributed to bringing three major conferences to the Netherlands, ACM Multimedia 2016, ISMIR 2019, and ACM Recommender Systems 2021 with Alan Hanjalic, Cynthia Liem and Martha Larson as General Co-Chairs, respectively.

Societal impact: We have initiated several EU projects in the *creative industry* sector, in which we brought international consortia together to enable the transfer of our scientific results in multimedia search and recommendation towards industrial and societal partners involved, among which the Royal Concertgebouw Orchestra. Specifically we mention the FP7 projects PHENICX and CrowdRec (both reaching into the assessed time period) and the H2020 project TROMPA, which were initiated and (scientifically or completely) coordinated by our staff members (PHENICX and TROMPA by Cynthia Liem and CrowdRec by Martha Larson). In addition, based on our expertise in and experience with search and recommendation, we have taken the initiative and have been increasingly engaged in collaborative efforts with external scientific CS and non-CS partners towards responsible data analysis and AI, as illustrated by the project with Erasmus University

(Case Study D.3) with the involvement of Cynthia Liem, and the MediaEval Benchmarking Initiative[b] which was initiated and has been coordinated for many years by Martha Larson.

Outlook: Over the past six years, our research on developing recommender systems has increasingly shifted from optimizing the pure performance (effectiveness, relevance) to address aspects such as user's values, trustworthiness, explainability and fairness. A recently published work of potentially high impact is about neutralizing bias in recommendations to better serve less active and/or non-mainstream users (key publication [10]). To further strengthen this shift in strategic focus, two new staff members have been appointed: Elvin Isufi, to apply his expertise on graph signal processing to address biases in the networks of users and items, and Julian Urbano, to apply his expertise on evaluation of information retrieval systems to help improve trustworthiness and explainability of recommender systems.

a Tech at Bloomberg: Article about the award of Christoph Kofler
b MediaEval evaluation benchmark

5.1.5 User modeling (WIS/ST)

A large part of our scientific work around *user modeling* is concerned with techniques for extracting semantics and making sense of user and interaction data to support the engineering of large-scale Web-based systems in a variety of contexts.

Representative results: One line of work is on search and sensemaking, in particular as ubiquitous parts of human learning. Learners turn to the Web seeking additional information, and over time create mental representations of the material. This work is aimed at investigating search and sensemaking activities of learners and subsequently enhancing them in order to improve learning outcomes. One key result is *SearchX* (key publication [11]), an open-source, modular, scalable search system for interactive information retrieval experiments that has recently been extended by our framework-agnostic logging framework *LogUI*[a]. SearchX is used in a series of published experiments at ACM CIKM and ACM SIGIR. In our work on digital learning and educational technologies, we use interaction and tracking data to build user and learner models for assessment, reflection, and intelligent feedback (key publication [12]). Main results are configurable *Learning Analytics Dashboards* and our *edX logging analysis* framework.

Scientific impact: Our work on instructional scaffolding for complex searches has received the Best Student Paper Award at CHIIR. The LogUI framework received the Best Demo Award at ECIR. Our work on digital learning and educational technologies won best full paper awards at LAK [12] and CTSE. Our members (e.g., Claudia Hauff, Geert-Jan Houben, Marcus Specht, Nava Tintarev) fullfil leadership roles as program, track, or general chair at key conferences (ACM Hypertext, ACM WebSci, ECIR, ICWE, CIKM, ACM Learning@Scale) and editorial boards (UMUAI, ACM TWEB, Information Retrieval Journal, JWEB).

Societal impact: For our search and sensemaking work, we collaborate with TU Delft's Extension School which is responsible for the 100+ MOOCs run by TU Delft. Thus as part of our research uptake, learning-oriented interventions that have been shown to be beneficial for human learning in experimental interventions can be employed to large cohorts of MOOC learners. Our research is sponsored largely by NWO (Vidi and Top grants), NWA, and ING. For our work on interactive explanatory interfaces, we have been collaborating with IBM, Twitter, Blendle and are part of an ITN NL4XAI. This work won the TU Delft Mekelprize on Responsible Innovation 2020. For our work in digital learning and educational technologies, we collaborate with the SURF Acceleration Program and run joint research projects with several EdTech startups.

Outlook: We will continue investing in this research topic and its lines, through NL4XAI and the collaboration with IBM, ING and other partners, and with new hires in explanations (complementing the part-time (0.2 FTE) contributions of Nava Tintarev after her promotion to full professor at University of Maastricht) and in sensemaking (strengthening our NLP expertise related to search and sensemaking). We will strategically extend and broaden the topic of digital learning and educational technologies, for example through the research group's collaboration with LDE-CEL and 4TU-CEE, focusing on AI-Ed technology.

a LogUI framework

5.2 Data Management and Analytics

5.2.1 Data Management and Engineering (WIS/ST)

With data being a cornerstone of modern AI and data-driven systems, our work on data management and engineering addresses the scientific challenges associated to scalable, semantically-rich, and human-enhanced management and engineering of large volumes of data.

Representative results: Our scalable data management research aims at designing systems and methods to cope with the ever-increasing demand for storage and processing power. Our results address this challenge by scaling data operations in large clusters and the cloud, for example in *Apache Flink*, the state-of-the-art stream processing system, including the *Rho* software for stateful functions as a service, and in *OpertusMundi*, a European effort building a marketplace for spatial data. Our human-enhanced data management research aims at efficiently and effectively integrating human and artificial intelligence at scale for data (and metadata) creation, integration, enrichment, analysis, and sensemaking. The key outcome is our approach to combining crowdsourcing and expertise modeling to extend data- and knowledge-intensive systems with large-scale human interpretation (key publication [13]). This resulted in systems that found successful application in several domains: *SocialGlass*, for integration, enrichment, and sensemaking of urban data; *eCrowd*, for enterprise crowdsourcing for training data generation for AI; and *SmartPub*, for extraction and exploration of semantically-rich metadata from scientific publications.

Scientific impact: Our research on scalable stream processing on Apache Flink has received more than 1150 citations and received the best paper award in EDBT 2019 (key publication [14]). In 2019 Asterios Katsifodimos co-organised the prime database conference (ACM SIGMOD) in Amsterdam. Our human-enhanced data management research received best paper and honorable mention awards (ACM HT 2017, AAAI HCOMP 2020, CSCW 2020).

Societal impact: We constantly work with industry partners in H2020 projects, but also directly: data integration with ING Bank, stream processing with Hazelcast Inc, scalable distributed databases with Adyen, video analytics with Cognizant, and medical data integration with Erasmus Medical Center. Asterios Katsifodimos created an online training program "Big Data to Insights" for working professionals. Alessandro Bozzon and Christoph Lofi founded the Future Libraries Lab together with the Dutch Royal Library. Our human-enhanced data management work attracted substantial funding from various research programs (H2020, EIT Digital, JPI, COMMIT/, NWO, NWA VWData, NWA JOIN) and industrial partners (IBM, KPN, Telecom Italia, Cognizant). Alessandro Bozzon was a faculty fellow at IBM and received an IBM Faculty Award in 2017, and Geert-Jan Houben and Alessandro Bozzon served as PI of Urban Data Science at AMS Institute.

Outlook: We invest in our scalable data management work, through our intention to hire a new colleague in the area of data-intensive systems, and by building an open-source data integration system (Daisy, https://github.com/delftdata/daisy) as part of our OpertusMundi efforts. The scientific relevance and excellence of the human-enhanced data management research resulted in the promotion of Bozzon to full professor (Human-Centred AI) at the IDE faculty (he now works with us for 0.2). At the same time, two new assistant professors, Gadiraju and Yang, joined this program in 2020. They already established one of the TU Delft AI Labs, "Design@Scale", to investigate Human-Centered AI for societal-scale design. The team is central in the development of a cross-faculty program on Human-Centred AI Systems.

5.2.2 Bioinformatics (PRB/INSY)

Molecular biology has become a data-driven research field. The strength of our bioinformatics research is that we use our expertise in machine learning and algorithmics to develop new methodologies to analyse and interpret genomic data.

Representative results: The genomics field changed drastically with the advent of single cell sequencing. We have been at the forefront of these developments (a.o. in collaboration with the Allen Brain Institute for Brain Sciences (US)). In addition to the efforts on single cell visualisation (see Data Visualisation achievements in Section 5.2.3), we analyzed and interpreted spatial-temporal gene expression data ((key publication [14]). Furthermore, the PhD thesis of Ahmed Mahfouz, Winner of the BIOSB Young Investigators award), proposed several clustering approaches, as well as a much cited single cell classifier comparison paper in genome biology

(key publication [16]). Another revelation in genomics has been the possibility to do long-read sequencing. We introduced novel genome representations and graph alignments. This led to the characterization of the first aneuploid yeast species crucial in the beer brewing industry, as well as variations in Mycobacterium tuberculosis (key publication [17]). We contributed to a better understanding of the impact of guide RNA (gRNA) and genomic locus on CRISPR-Cas9 activity, crucial to design effective gene editing assays. We also introduced new algorithms to distinguish closely related organisms in complex microbial communities with applications in infectious disease, industrial biotech and environmental biotech.

Scientific impact: Our impactful work on single cell analysis led to prizes (Princess Beatrix Spierfonds, KHMW Pfizer Life Sciences Price) and resulted in a partnership with the NWO Gravitation Project BRAIN-SCAPES (10MEuro over 10 years, 1MEuro in Delft). Our biotechnology endeavours over the last years have led to a strategic public-private collaboration in the form of the ICAI Lab, funded by the DSM company (2.5MEuro over 5 years, formally starting in 2021). We are one of three non-US partners in HuBMAP (3.7 MEuro, 0.7MEuro in Delft), a consortium funded by the NIH (US) to decipher molecular distributions in human tissue, as a result of our work on temporal algorithmics and CRISPR-Cas9.

Societal impact: Our NGS sequencing analysis pipeline is now routinely used in Non-Invasive Prenatal Screening throughout hospitals in various parts of the world. Our genome analyses of Mycobacterium tuberculosis, the number two killer infectious disease, has led to the first characterization of an elusive host interaction gene family. Our sequence analysis work in the context of Alzheimer revealed a number of novel markers (such as risk variants in SORL1, protective role of PLCG2, and Polygenic Risc Scores that predict predisposition for AD).

Outlook: With Jasmijn Baaijens, a new hire (in 2021), we will expand our efforts in complex genome analysis predominantly to capture variations in microorganisms. Within the context of the DSM ICAI, we will hire a new tenure-track assistant professsor with a focus on metagenomics and nutrition (aligning with our research on life style intervention and aging). We are part of a consortium together with the VUMC and Pacbio (world leader in long read sequencing) in which we will generate a globally unique dataset of 300+ humans. Our work on CRISPR-Cas9 continues along the lines of studying DNA repair mechanisms in the context of cancer as well as proton radiation therapies.

5.2.3 Content generation and data visualization (CGV/INSY, PRB/INSY)

We aim at efficiently and effectively generating visual representations to produce and simulate virtual worlds, as well as to analyze, understand, inspect, and interact with real-world data. We develop fundamental concepts, algorithmic solutions, and transfer the resulting technologies to various application areas. A main focus is the (bio-)medical sector, but also other areas like architecture, culture or entertainment.

Representative results: In the assessed time period, we contributed to various areas of computer graphics, visualization and beyond. We developed novel technology for physical simulations in virtual worlds (including detailed object representations, efficient display and complex editing), which were presented at top-tier venues in the field, such as ACM SIGGRAPH, ACM SIGGRAPH Asia, Eurographics, and IEEE Visualization. We also developed a new approach to geometric deep learning on surfaces (key publication [18]), which addresses a fundamental problem in the design of CNNs on surfaces and provides substantial improvements over the state-of-the-art on several benchmark tests. In visualization, our work on interactive visual analysis has led to several key contributions that have had a profound impact both within and outside our field. For example, our proposed Hierarchical Stochastic Neighbor Embedding (HSNE) formed the basis of an article in Nature Communications (key publication [19]), which presented our solution for mass cytometry data analysis.

Scientific impact: Two PhD students have received prestigious international doctoral-thesis awards; N. Pezzotti (CGV/INSY) received the IEEE VGTC 2019 Best Dissertation Award and C. Brandt (CGV/INSY) the Eurographics 2020 PhD Award. More than 20 papers were published in the journals of the leading graphics and visualization venues (ACM TOG and IEEE TVCG), and more than ten paper awards have been received. Cytosplore, a software system for interactive visual single-cell profiling of the immune system, in collaboration with LUMC, received the Dirk Bartz Prize for Visual Computing in Medicine. The Cyctosplore system has had a high impact on the single-cell analysis community and has already been downloaded by more than 4000 individual users and has led to numerous collaborative publications (e.g., Nature, Nature

Immunology (2x), and Science Translational Medicine). The original dimensionality reduction algorithm tSNE is now a built-in functionality of Matlab and SPSS, and our recent variant has been integrated into Tensorflow. In 2018, the Eurographics conference, a top-tier conference for computer graphics, took place in Delft with Elmar Eisemann (CGV/INSY) as General Chair.

Societal impact: Our visualization and display technologies can be found in professional products (e.g., NVIDIA GIWorks or MeVisLab), AAA games (e.g., Tomb Raider), or world-leading engines (e.g., Unreal). We also developed online platforms ourselves, such as the Online Anatomical Human (OAH), an education system based on real human anatomy. It is used in a MOOC by the University in Leiden, which is followed by almost 20K students per year.

Outlook: One of the fundamental aspects of our research in data visualization and visual analytics is the focus on large-scale data sets. Consequently, the connection towards data-driven solutions and the integration of data as an algorithmic component are natural evolutions. The department invested into these developments by strengthening the areas of data analysis and visualization through several new faculty positions in computer vision with links to computer graphics, as well as AI-driven image synthesis.

5.2.4 Data analysis and process modeling in complex networks (MMC/INSY)

Data collected from generic complex systems can be represented in the form of networks, ranging from online social networks, physical contact networks to critical infrastructures. We aim to develop methodologies to predict, model and control processes (such as information, disease, failure propagation and social or financial contagion) on networks.

Representative results: We further strengthened our leading role in network science via the following progress. We have developed methodologies, both theoretical and empirical, to understand how the underlying network topology influences a dynamic process. We addressed the challenge that the underlying network is evolving over time (key publication [20]) and that the spreading process is heterogeneous. The results enable the optimization of network topology to be robust against virus and failure propagation, or efficient in information diffusion. Furthermore, strategies to mitigate the spread of epidemics via e.g. reducing time-evolving network interactions and information spread have been developed (key publication [21]). During this time period, we have also intensified our research in the domain of *network data science*, aiming to integrate network science and data science approaches. For example, our expertise in spreading process has been utilized to improve the design of temporal network embedding algorithms for general link prediction and classification tasks.

Scientific impact: A series of papers by Huijuan Wang and Alan Hanjalic has been published in top-tier journals such as Scientific Reports, Plos One and New Journal of Physics. Together with other scholars who have been publishing in venues of Nature, we have contributed to the book "Temporal Network Theory", a milestone in the field of temporal networks. Huijuan Wang has co-founded and been chairing the Dutch Network Science Society and she is also currently chairing the Netherlands Platform Complex Systems.

Societal impact: We transform and disseminate scientific results towards industry and society via projects ranging from EU FP7-FET CONGAS, NWO-TOP to KPN-TUDelft project on AI networking that is financed by KPN and via the activities that we lead through research communities. Moreover, our methods that predict late payment and default of companies by analyzing the monetary transaction network among companies have been launched in the products of Exact. Finally, a group of our bachelor students developed a web application that visualizes the prediction and mitigation of Covid19 to facilitate the knowledge transfer towards researchers from diverse fields and the public in general.

Outlook: We will keep investing in the further development of our research on network data science, particularly aiming at understanding and optimizing complex socio-physical systems by combining network science and data science. The necessary expertise has been strengthened by the recent appointment of a new staff member, Elvin Isufi, who has a strong background in graph signal processing.

5.3 Software Systems Engineering

5.3.1 Software Analytics (SE/ST)

In *software analytics* we investigate data originating from the software development process to obtain actionable insights to improve the effectiveness of development processes.

Representative results: One key result of the assessed time period is `ghtorrent.org`. This is a service and curated data set that we began offering in 2012, which collects the continuous stream of data about GitHub open source systems in an accessible data base. Since then we (and many other researchers) have used GHTorrent in a series of experiments to understand pull-based software development process as adopted in open (and closed) source engineering teams around the world. As an example, in a series of 3 papers that appeared at ICSE 2014–2016, we assessed the way in which *integrators* decide to merge in changes offered by *contributors* (key publication [22]). Mixing quantitative and qualitative research methods, we established that reducing response time, maintaining awareness, improving communication both in content and in form, and quality assessment are key to contribution acceptance in the pull-based model. We particularly focus on the testing process in our work surrounding WatchDog (key publication [23]) and TravisTorrent. Here we involve test execution data from both the developer's IDE and the Continuous Integration server, to investigate software testing processes as they actually take place (as opposed to what testers are supposed to do). Findings include that developers tend to considerably over-estimate their testing effort, and that most local test executions fail.

Scientific impact: Various TU Delft papers on software analytics on GitHub have collectively attracted over 2000 citations, and have sparked world-wide interest in this area. For this, Georgios Gousios received the Foundational Contribution Award at MSR 2018. The selected key publication [22] received an ACM SIGSOFT Distinguished Paper Award at ICSE 2016. Our PyDriller tool (published at MSR 2018) to support git repository analysis has been downloaded over 250,000 times. Additional marks of recognition include the many conference chair roles in software engineering held by the group, including ICSE 2021 (van Deursen, program chair) and ESEC/FSE 2021 (Spinellis, Gousios, general chair).

Societal impact: The work on GHTorrent has been financially supported by Microsoft, who are using the service to monitor project health. As prominent public examples Deloitte has used GHTorrent to assess the evolution of blockchain technology,[a] and the European Commission has used GHTorrent in its study on the "Impact of Open Source Software and Hardware on technological independence, competitiveness and innovation in the EU economy". We are presently conducting software analytics research within the Dutch ING bank to help them understand and mitigate delays in software deliveries (case study D.2).

Outlook: In 2020, we strengthened our presence in this field by appointing Diomidis Spinellis as part time (0.2 fte) full professor. We aim to continue to play a world-leading role in this area, for example in the context of the ING ICAI lab (2020–2024, Case Study D.2) headed by Arie van Deursen, and the NWO VICI personal grant (2019–2023) obtained by Andy Zaidman.

a Deloitte Insights: Evolution of blockchain technology

5.3.2 AI for Software Engineering (SE/ST)

Given the advances in AI and wide availability of software analytics data, our work mainly revolves around the use of AI techniques like machine learning and evolutionary computing to support software engineering.

Representative results: An important line of work in this area relates to search-based *crash reproduction*, which aims at automating the first step of debugging, reproducing a failure. We define a fitness function and mutation operators that allow us to generate and refine test cases that try to reproduce a given crash (key publication [24]). To seed the process, we use models learned from source code and test case execution The resulting approach, dubbed `Botsing`, is able to beat the state-of-the-art in reproducing a large collection of crashes of open source systems. A second line of research initiated by Aniche involves the use of supervised learning to guide refactorings.Using the data of over 2 million refactorings in over 10,000 projects, we establish that random forest models can accurately recommend refactorings, and that process and ownership metrics

are the strongest predictors.

Scientific impact: The work by Annibale Panichella on search based software engineering is widely used, in particular his recent extension DynaMOSA to the evolutionary testing tool EvoSuite.[a] The work on crash reproduction has led to two members being invited as program co-chair of the Symposium on Search Based Software Engineering (SSBSE). Furthermore, it has led to a best paper award at SSBSE 2020. In addition, we have created Botsing, an open source crash replication tool, and the associated Java crash reproduction benchmark JCrashPack.

Societal impact: Botsing is a technology that came out of the H2020 STAMP project, where it was applied directly by the companies involved to make their software systems more reliable. Also in our collaboration with ING Bank, we investigate how to improve reliability, through investigating how search-based test generation techniques can be brought from the unit to the integration and system test levels. Within our sponsored collaboration with Ripple (2018–2022), we investigate how we can make blockchain applications more secure through the automatic generation of security tests. Finally, we are working on making automated driving systems more reliable through automated repair (key publication [25]).

Outlook: We will continue investing in this research topic, through the ING ICAI lab (2020–2024) (Case Study D.2) headed by Arie van Deursen, and the H2020 COSMOS (2021-2023) project acquired by Annibale Panichella and Andy Zaidman.

a EvoSuite

5.3.3 Language Engineering (PL/ST)

In *language engineering* we aim to support the high-level, declarative definition of programming languages and the derivation of safe and performant programming environments from these definitions.

Representative results: In the evaluation period we advanced the state of declarative definition of syntax, static semantics, and dynamic semantics of programming languages. A key result is a language independent *theory of name resolution* (key publication [26]) based on *scope graphs*, a language independent representation of the name binding facts in a program, with support for reasoning about name resolution, scope, shadowing, renaming, and capture. The scope graph representation provides a *lingua franca* for talking about name binding patterns in a wide range of languages and provides the foundation for language-agnostic support for a range of operations on programs including type checking, renaming, and code completion. We have developed the Statix language for declaratively describing the static semantics of programming languages with scope graphs at its core. Its operational semantics is safe with respect to its declarative semantics, ensuring stable answers to name resolution queries in an evolving scope graph. Query safety is based on a theory of *critical edges* (key publication [27]), which provides the basis for automatic derivation of type checkers and automatically parallelizing and incrementalizing type checkers.

Scientific impact: The theory of name resolution was published in ESOP 2015 and received the overall EAPLS best paper award at the umbrella conference ETAPS 2015. Further results were published in top tier venues including ACM OOPSLA, ACM POPL, and ECOOP. The work on Spoofax was recognized with the OOPSLA 2020 Most Influential Paper Award for the 2010 paper that introduced the workbench. We organized the ACM SPLASH conference for the first time outside North America (2016, Eelco Visser as general chair), chaired the review committee of ACM OOPSLA (2019, Eelco Visser), and organized a Dagstuhl Seminar on Domain-Specific Languages (2015, Sebastian Erdweg and Eelco Visser).

Societal impact: The research results are integrated in the Spoofax language workbench, which is made available under an open source software licence. Spoofax is applied in a long-term collaboration with Oracle Labs for the development of a suite of languages for graph analytics. Oracle Labs contributes funding to the open source development of Spoofax. Recently, Canon Production Printing systems started using Spoofax for the development of domain-specific languages. The WebDSL web programming language developed with Spoofax is used to develop several production web applications including ones for (programming) education and conference organization used by thousands of students and researchers.

Outlook: Sebastian Erdweg became a professor at Mainz university. We hired Casper Bach Poulsen as assistant professor, who was awarded an NWO VENI grant in 2019 to work on composition of intrinsically typed definitional interpreters. We are starting work on an NWO/AES project we acquired to develop programmable software restructurings for legacy software modernization in collaboration with TU Eindhoven, TNO, and Philips with scope graphs as key component.

5.3.4 Software Verification (PL/ST)

In our software verification research we seek to develop methods and techniques to support rigorous reasoning about the correctness of programs. Our interest emerged out of our work on language engineering, and was given a boost in the evaluation period with the appointment of Robbert Krebbers.

Representative results: A key result is Iris, a framework for machine-checked verification of programs and programming language theory in languages like Rust, Go, OCaml, C, and Scala (key publication [28]). Iris started as a collaborative effort involving MPI-SWS (Germany), Aarhus University (Denmark), and TU Delft, and has grown into an active open source project with collaborators and users at numerous institutes and companies world-wide. One of the flagship applications of Iris is RustBelt (key publication [29]), providing the first machine-checked proof of soundness of the Rust type system and a collection of its standard libraries.

Scientific impact: The RustBelt paper is the most cited paper of POPL'18. The version in CACM reached the top 3 of most popular CACM articles during the first month (more than 15K downloads). The Iris software is used in 45 papers and 7 PhD theses, including papers in systems and security conferences. Iris is used by Aarhus University, Boston College, CNRS/LRI, Delft University of Technology, Groningen University, INRIA, ITU Copenhagen, KU Leuven, Microsoft Research, MIT, MPI-SWS, NYU, Radboud University Nijmegen, Saarland University, and Vrije Universiteit Brussel. Krebbers was awarded NWO VENI grant in 2018.

Societal impact: The (open source) Iris software is used by Google as well as by BedRock Systems, specialized in trustworthy computing with a minimal attack surface. Rust is an increasingly popular systems programming language originally developed by Mozilla. Since the Rust community puts a strong emphasis on reliability, it has embraced the RustBelt work.

Outlook: We have strengthened the software verification program by hiring Jesper Cockx, an expert in dependently-typed programming and core contributor to the Agda programming language. He started at the end of 2019 and was awarded an NWO VENI grant in 2020. After Robbert Krebbers was appointed an Associate Professorship at his alma mater Radboud University in Nijmegen, we hired Soham Chakraborty, who works on the semantics of concurrency, and Benedikt Ahrens, who works on categorical semantics of programming languages to further strengthen the verification research program.

5.4 Networked and Distributed Systems

5.4.1 Distributed Trust (DS/ST)

Our long-term, high-level research challenge in *distributed trust* is how to digitize trust for society and the economy, for which we are building solid theoretical, experimental and technical foundations.

Representative results: A blockchain of our own design called *Trustchain* (key publication [30]) is a key result in the assessed time period. This alternative blockchain recognizes that global consensus on all transactions worldwide as in many crypto-blockchains is not required for many application areas. Rather, in Trustchain every participant has their own chain of transactions which are linked according to common transactions. Trustchain represents a novel principle to organize any economic activity, while addressing the key Bitcoin and Ethereum deficiencies. Trustchain's scalable graph of bilateral agreements can in principle serve as an incorruptible ledger to underpin the (European) digital economy. In another vein and at a more theoretical level, in our research on payment channel networks that offer higher transaction throughput to blockchains, we have investigated how nodes can maximize their financial gain by judiciously establishing channels and choosing fees for participating in payments (key publication [9]).

Scientific impact: The publication on Trustchain led by Johan Pouwelse (key publication [30]) has obtained over 120 citations. He organized the *Workshop on Distributed Infrastructure for Common Good* with the ACM/IFIP Middleware 2020 conference, which will be continued as a series of workshops in the coming years. Trustchain is part of our open-source, product-quality Tribler system[a], which serves as our research vehicle in leaderless self-organizing systems for the common good. Since its initial creation in 2005, Tribler has been installed 1,900,000+ times, and through its regular releases and its active user community of tens of thousands of users, it provides a prime example of open science.

Societal impact: Our work on distributed trust is conducted in partnership with several Dutch ministries and government regulators. In 2020 the Dutch Central Bank (DNB) gave us special permission to conduct field trials with our Trustchain-based digital Euro, which is in fact real money. The Dutch Ministry of the Interior and the European Commission are committed to providing a trustworthy digital identity to all Dutch and European citizens, and we are a key academic partner in these efforts with our Trustchain-based, passport-level self-sovereign identity framework.

Outlook: We have hired over the last few years two new staff members (Stefanie Roos and Jérémie Decouchant) in order to achieve a better balance towards a more theoretical approach to blockchain, consensus and anonymity. When we started this topic in 2017 with the Delft Blockchain Lab (see Appendix G.4), we planned to be in this research area for at least ten years. For the next five years we have ample funding from NWO and the Dutch Ministry of the Interior.

a Tribler

5.4.2 Internet of Things (ENS/ST)

To advance the Internet of Things (IoT) field we have contributed to the fundamentals (intermittent computing paradigm, and LoRa security), extended the application domain (by taking it to space), and established a fieldlab to foster experimental research in a multi-disciplinary setting.

Representative results: Devices running on ambient energy suffer from frequent (but unpredictable) power-outages; we have extended the standard checkpointing approach to include sensor data (which may become stale during recharging) by developing a new software framework called *Ink* (key publication [32]) that connects the field of intermittent computing and event-based processing common in IoT applications. LoRa (Long Range) communication is quickly becoming a corner stone of outdoor IoT deployments, but this relatively young product is not as safe as more mature standards. In particular we have identified a number of fundamental security flaws (key publication [33]) that need to be addressed before LoRa can be used in privacy-sensitive settings. In Sensys 2017 (the leading IoT/WSN conference) we put forward the radical idea of using IoT components for operating swarms of pico satellites, in which power consumption is the main limitation. Following suit –by means of collaboration with the Indian Space Research Organization– we have developed, tested, and launched several novel low-power satellite components. Most notable is the *Hummingbird* GPS receiver that features a novel duty-cycling algorithm reducing the energy consumption with a factor of 25.

Scientific impact: The collective set of IoT papers published in the time frame 2015-2020 has already received well over 2,000 citations, while the Hummingbird work by Narayana and Venkatesha Prasad [34] stands out individually as it received the prestigious MobiCom 20202 best paper award. We are also proud of regularly hosting premiere conferences, including Sensys (2017, Koen Langendoen as general chair) and SIGCOMM (2021, Fernando Kuipers as general chair).

Societal impact: Our academic standing has helped in connecting with local governments and organizations, and has culminated in the founding of the *Do IoT fieldlab* supported by a grant from the European Regional Development Fund (ERDF), and the Metropoolregio Rotterdam Den Haag (MRDH) amongst many others, with the ambition of accelerating IoT innovation by providing test sites equipped with the latest generation of mobile communication (currently 5G). A second highlight was winning the Airbus Fly Your Ideas challenge (285 entries) with a battery-free IoT device for aircraft monitoring developed by Venkatesha Prasad and a team of students, leading to a spinoff company.

Outlook: To expand our academic activities, we have attracted three new faculty members, who all have a link to AI (in particular machine learning). The main focus will be to explore the possibilities of AI, which is generally computationally complex, in an embedded context where resources are limited. Early results show that even the learning component may be executed on embedded devices if done collectively (federated learning).

5.4.3 Big Data Processing (DS/ST)

With the growing datacenters and the amounts of data being collected and analyzed, the problem of efficiently putting to work the large-scale infrastructures is constantly becoming more difficult. The flavor of the research we conduct in big data processing has changed over the years. Whereas we were originally focused on resource management and scheduling from the performance and reliability perspectives, we moved to supporting application areas such as graph processing and more recently, to distributed machine-learning systems.

Representative results: As a prime exponent of the performance-oriented research, we designed and experimentally evaluated autoscaling policies for workflows, which are widely used for data analysis in many scientific domains, in clouds that dynamically allocate resources to workflows using the elasticity of cloud resources (key publication [35]). As an exponent of the later research on the support for and use of machine learning, we have created and analyzed a two-layer learning framework for anomaly detection in various types of computer systems under the assumption that the labels of the features may be corrupted, either by errors in the system or by adversaries (key publication [36]). We show that with this framework, the accuracy of anomaly detection increases significantly and the dependability of datacenters is enhanced.

Scientific impact: In the first half of the reporting period, TU Delft was one of the top universities for research in scheduling and performance analysis for large-scale distributed infrastructures. The two key publications were published in the top journal *ACM TOMPECS* in 2018 and the *IEEE DSN* conference in 2019, and sofar received 21 and 19 citations, respectively. Alexandru Iosup was the General Chair of the *7th ACM/SPEC Int'l Conference on Performance Engineering* (ICPE) held in Delft in 2016, Lydia Chen was Program Co-chair of the *IEEE Int'l Conference on Autonomic Computing* (ICAC) in 2019, and Jan Rellermeyer was the General Chair of the *ACM/IFIP Middleware Conference* held in Delft in 2020.

Societal impact: Alexandru Iosup was the chair of the Cloud Working Group of the Standard Performance Evaluation Corporation (SPEC), which has many companies as its members. As a result of a collaboration with the Dutch insurance company Aegon, and using an NWO Take-Off grant, Lydia Chen has created a startup company called Generatrix. The purpose of the company is to build, market and sell a "beyond the state-of-the-art" decentralized synthetic data generator that allows for safe big-data knowledge discoveries without compromising privacy.

Outlook: Lydia Chen (as a DTF fellow) and Jan Rellermeyer joined this cluster in the reporting period. Supporting big-data processing and machine learning from a system's perspective remains an important topic for the future that we will continue to pursue.

5.4.4 Visible Light Communication (ENS/ST)

To address the ever increasing demand for wireless communication, Visible Light Communication (VLC) targets unlocking the massive bandwidth of the visible light spectrum. While most VLC systems are built on the active modulation of a LED, we have embraced a passive approach in which ambient light (e.g., from the sun or office fixtures) by means of LCD shutters is transformed into a modulated signal that can be decoded with a simple photodiode enabling low-power operation.

Representative results: We have established ourselves as pioneers in the area of passive VLC by publishing some radical new approaches. We highlight here the work by Wang and Zuniga (key publication [37]) that proposes to have mobile objects (e.g., cars) "wear" patterns consisting of distinctive reflective surfaces (much like a barcode) that can be decoded using reflected ambient light. A second novelty is the step beyond simplistic on/off keying proposed in key publication [38], which leverages frequency modulation to increase the signal-to-noise ratio (SNR); the functional prototype, named LuxLink, has been measured to cover distances up to 60 meters on a cloudy day in Delft, setting a new standard for low-cost passive VLC

systems.

Scientific impact: The above-mentioned key papers [37, 38] have been published in top venues (CONEXT and SenSys, respectively) in the field of low-power wireless communications. Furthermore two major research grants on passive VLC have been awarded to Marco Zuniga as PI: (i) "LuxSenz", 2019-2024, NWO (TOP Module 1), and (ii) "ENLIGHT'EM", 2019-2023, EU (MSCA-ITN), with a collective budget of 1.2 M€.

Societal impact: The idea of using sunlight for wireless communication has been well received by (local) industry with Marco Zuniga winning the Accenture Innovation Award, Perfect Cities track (out of more than 140 concepts) in October 2017. To valorize his knowledge of VLC, he has entered into collaboration with the lighting company Zumtobel (now Tridonics) resulting in multiple MSc projects with corresponding papers and a patent application.

Outlook: To take the research on passive VLC further we are working on theoretical contributions (on the optimal use of LCD shutters) as well as pursuing the ambitious, green goal of integrating VLC into standard window panels to realize the first "talking" office buildings that can send data back and forth using only sunlight (zero-energy communication).

5.5 Security and Privacy

5.5.1 AI for Intrusion Detection and Prevention (CYS/INSY)

We perform research in white-box machine learning for security goals such as intrusion detection. By focusing on interpretable models, we help security analysts to understand why a software file or network connection is labeled as malicious and enable root-cause analysis.

Representative results: The TABOR intrusion detection system (key publication [39]) is an example of interpretable AI. We demonstrate that data patterns in cybersecurity are regular and contain little noise. Our system learns models that can succinctly and accurately capture such patterns. We show that they perform better and significantly faster than opaque neural-network based solutions. We also developed tools for fingerprinting and analysing compromisation attempts over SSH (key publication [40]). Advanced attackers try to hide their intrusions by distributing their attacks to originate from multiple sources under their control, but they reuse software on all these sources. We show how to obtain identifying fingerprints for uncovering specific adversarial behaviors, tactics, and procedures.

Scientific impact: Our open source Flexfringe[a] learning tool (which includes RTI+, the algorithm used in TABOR (key publication [39])), can learn models for large amounts of network traffic and host-based logs in real-time. This tool is under continuous development and is actively used in both research and industry.

Societal impact: We collaborate with ASML and Adyen through Sicco Verwer's VIDI grant on learning state machines from software data. ASML and Adyen use the developed tools for understanding and debugging their software systems. In the H2020 project AssureMoss, we collaborate with Thales on detecting malicious behaviour in Kubernetes clusters. He founded the Apta.tech spin-off in this domain. Much of our societal impact also comes from our master students and alumni. Many of these find an internship and later a position as (cyber) data analyst or security analyst.

Outlook: With the new appointment of Apostolis Zarras, we strengthened our expertise in network security and in particular AI-based solutions for detection of cyber-attacks, for which he recently acquired two Horizon 2020 projects. Also, George Smaragdakis joined us in 2021 as new group chair, bringing significant expertise in networks and large-scale network data analysis. Together with Sicco Verwer and Mauro Conti, they will continue developing AI solutions for detecting and analysing threats, which continues to be important as attacks are becoming better and better at circumventing existing rule-based detection systems.

a Flexfringe

5.5.2 AI-based Testing for Secure Systems (CYS/INSY, SE/ST)

Vulnerability and bug discovery in software and hardware can be viewed as a search problem for the right inputs. Our work in this field is focused on improving this search process and generating insights from the searched paths, both by using and developing AI solutions.

Representative results: In hardware, the main security threat are side-channels, which attackers can use to extract confidential information such as private keys. We significantly improved the state-of-the-art by proposing a new neural network architecture for side-channel analysis which includes noise as a regularization factor (key publication [41]). With it, even previously developed architectures managed to achieve significantly better attack performance.

Scientific impact: We are recognized as one of the top centers for AI-based security analysis of hardware. For example, two years after publication, our approach (key publication [41]) still represents the architecture of choice for side-channel attacks when there are more than one dataset involved. We also won the Rigorous Examination of Reactive ystems (RERS) challenge (2016) and its Linear Temporal Logic (LTL) track (2020).

Societal impact: Our side-channel analysis toolbox is an example of our tools deployed by industry. Through projects with ING, RIPPLE, and ST Microelectronics, we apply our research in automated vulnerability discovery. We also collaborate with companies specialized in security, such as Riscure and Brightsight.

Outlook: In 2021, the principal investigator in side-channel analysis, Stjepan Picek, will transfer to Radboud University, part-time continuing his research at Delft University of Technology for a period of four years. The consequences of this transfer for preserving the corresponding research momentum will be addressed in consultation with the new full professor in cybersecurity George Smaragdakis. Automated Vulnerability Research (AVR), including discovery and patching, is a key focus point in the near future. Sicco Verwer and Annibale Panichella (SE/ST) are developing the AVR roadmap for the Dutch ministry of economic affairs and received funding for a first PhD position in this area.

5.5.3 Security of AI (CYS/INSY)

In security *of* AI, the goal is to design AI that is adversarially robust, i.e., to make AI itself secure. Relying on our expertise in learning algorithm design, we aim to make AI faster, more flexible, more effective, and provably robust against specified attackers. Since machine learning is commonly not designed to be secure, there also exist many other attack vectors to modern AI, such as poisoning, backdoors, model stealing, membership inference, and reprogramming. We develop technology to reverse-engineer information such as input data and hyperparameters from state-of-the-art AI implementations and build defences that make such attacks more difficult to perform.

Representative results: In the key publication [42], we present the first generic side-channel attack on neural networks. While previous work requires specific implementations or design choices, our work requires just the standard side-channel setting: the ability to measure the side-channel information. Furthermore, our team won the Robust Malware Challenge in 2019. We demonstrated how to design efficient algorithms for attacking neural networks that take the system calls made by an executable as input in order to detect malware programs. The discrete search space and constraints on the kinds of possible modifications made this design challenging. Our method outperformed the competition both in attacking unknown models and in defending against unknown attackers.

Scientific impact: Our work on side-channel attacks on neural networks (key publication [42]) was published at Usenix 2019, one of the top security conferences. It inspired a significant number of researchers to design even more powerful attacks or novel countermeasures, and is taught at several universities around the world as state-of-the-art attack on neural networks. It was selected as one of the *Top Picks in Hardware and Embedded Security* for 2020.

Societal impact: Trustworthy AI is essential to society given the wide-spread use of AI and the degrees in which we are becoming dependent on AI algorithms. We should not make the same mistakes as when we developed the internet, leaving open backdoors for attacks and threats unforeseen in those early days. Society will rely on academic and industrial researchers to not only develop useful AI tools, but also tools that safeguard our privacy and security. The societal relevance of the secure functioning of AI, protecting

humans and society from all sorts of harm, will only increase in the years to come.

Outlook: Although the field of secure AI is still young, our publications already made a significant impact in academia and industry. For instance, our reverse engineering attack was the first of its kind. We continue to build on our strengths in algorithm design to tackle this issue. Currently, most of our work is focused on protecting AI after deployment, but there is a strong push towards ensuring that AI is secure by design. To gain traction in this growing field, we aim for increased collaborative work with key strengths of other TU Delft research groups on static analysis, secure programming, and security testing.

5.5.4 Privacy Protection (CYS/INSY)

Digital services nowadays heavily rely on privacy-sensitive data. Unfortunately, services like genomic analyses and online personalized advertisement services are vulnerable to privacy attacks from malicious entities. Also use of data in machine learning is problematic due to privacy aspects during learning, and possibly private data exfiltration after releasing the model. Our goal is to design secure (machine learning) algorithms that provide privacy protection while they are accurate and efficient in terms of run-time, bandwidth and storage.

Representative results: In 2018 and 2019, we participated in iDash Challenge[a], where genomic data were processed under encryption for inferring diseases' predisposition using machine learning algorithms. Those challenges also included decentralized networks for sharing data (blockchain), and we developed new ideas as in key publication [43], to protect privacy using cryptographic tools.

Scientific impact: In iDash challenges, we competed against teams like Microsoft, IBM, MIT and were ranked 6th and 5th in 2018 and 2019, respectively. Zekeriya Erkin and Mauro Conti carry several responsibilities in the IEEE Information Forensics and Security community. Erkin, as general chair, brought IEEE International Workshop on Information Forensics and Security (WIFS) to Delft in 2019. ESORICS, one of the best and biggest computer security conferences, will be hosted in Delft in 2023, with Kaitai Liang as general chair.

Societal impact: We have been involved in NWO-funded projects, *Blockchain and Logistics Innovations* (2016-2018) and *Spark! Living Lab Project* (2020-2024), where we design decentralized privacy-preserving machine learning algorithms on supply chain operational services. A prototype that enables municipalities to share data about economically disadvantaged citizens with other governmental bodies (Financial Emergency Brake[b]) developed for the Ministry of Justice and Safety in 2019 will be tested by the Municipality of Amsterdam in fall 2021. We contributed to dissemination through visibility in hacking conferences like BlackHat USA and major international media, like Forbes and BBC.

Outlook: As GDPR and soon AI Regulation will demand privacy protection for EU citizens, we consider privacy protection, particularly in AI and machine learning, to be one of the core research domains. Our position in this domain has been strategically strengthened by appointing Kaitai Liang as a new staff member.

a iDASH secure genome analysis competition
b Financial Emergency Brake

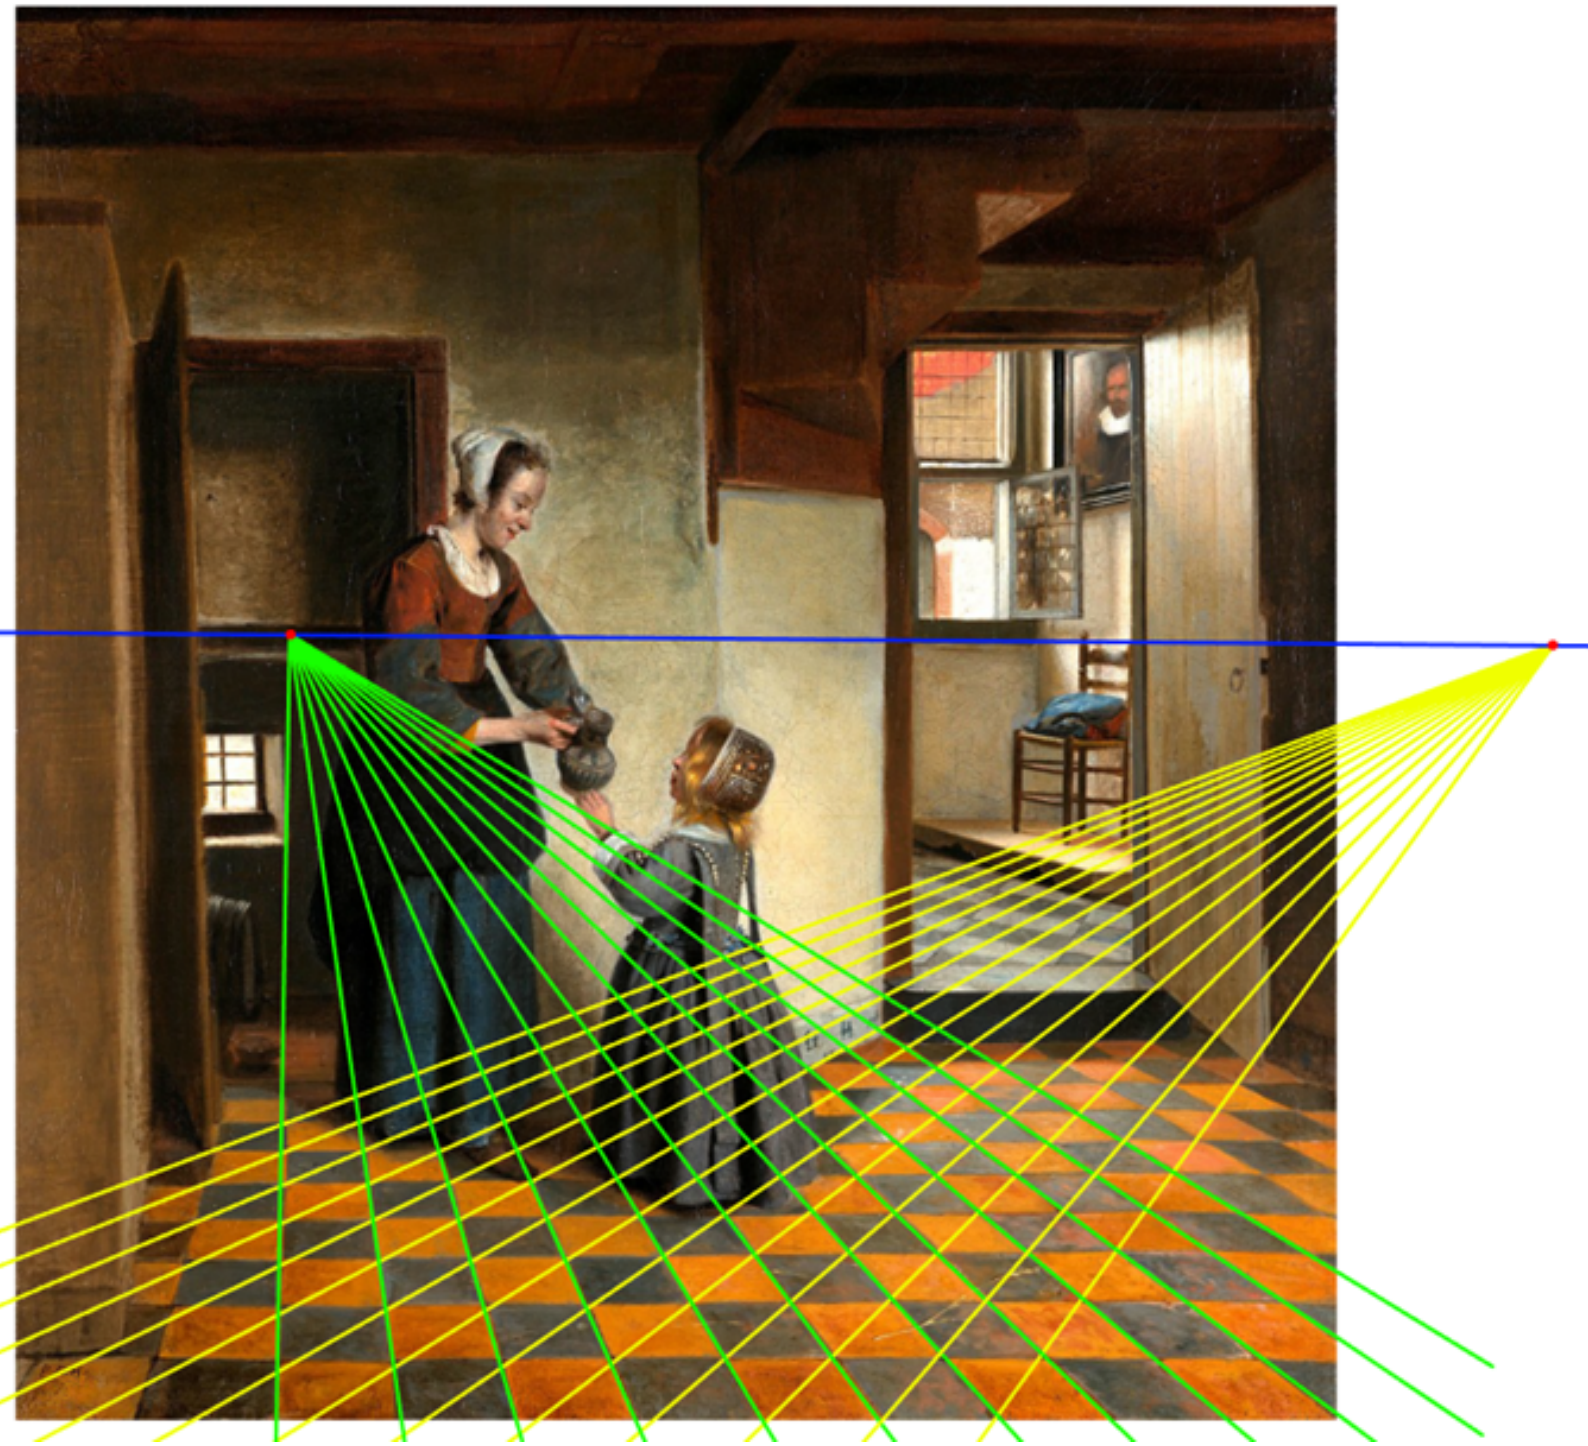

6 | Strategy 2021–2026

In the previous chapters, we presented CS@Delft through its integral mission and strategic aims, as well as the ways how we addressed these aims over the assessed time period. We also gave an overview of our past performance, from which contributions of the individual INSY and ST departments are clearly visible. Thereby, we highlighted the key scientific results, their scientific and societal relevance and future prospects. We are proud of our scientific successes and societal impact, which encompass core computer science disciplines as well as inter-disciplinary work in our target societal sectors. Given society's increasing demand for our expertise, the assessed time period was transformational for CS@Delft, characterized by substantial growth and development across our organization and activities. In this final chapter we reflect on our strategy to address this (ongoing) transformation and formulate seven key actions to successfully continue it in the years ahead of us.

6.1 Reflection

Talent-driven research strategy CS@Delft has shown the ability to cover a relatively broad research scope and still achieve excellence across this scope (Chapters 4 and 5). The breadth of activities defined by our five research themes (Chapter 2) enables us not only to address a large number of critical aspects of datafication and AI-zation phenomena, but also to capitalize on findings at the plethora of intersections of these aspects, including those with other disciplines. This has led to many valuable synergistic scientific initiatives, across section, department and even faculty boundaries. The INSY and ST department management has stimulated and facilitated creative research at these intersections by securing a high level of academic freedom and autonomy for our scientific staff, embracing new refreshing, largely interdisciplinary research directions. It is the role of the department to set a high level research scope and to provide the necessary means and context (career paths, work environment, resources, research/teaching balance). Within this scope, the development of the research directions relies entirely on the interest, drive and ambition of our people. The quality of the results and the intensity of the activities (Appendix *app:outcomes omitted*) demonstrate the effectiveness of this talent-driven research strategy, and its robustness to the dynamics and challenges of the context in which CS@Delft has been and will be operating in the upcoming years (Section 3.2).

Value-inspired perspective Our talent-driven research strategy directly relates to our value-inspired perspective on excellence and impact. We resist the temptation of the *"numbers game"* warned against by David Parnas.[1] Instead, our strategy is based on the DIRECT values, fostering the development of our staff according to DORA, and creating *room for everyone's talent* (Sections 3.3.1 and 3.3.3). This strategy is directly in line with the perspective of Mark Guzdial,[2] who, in the context of discussing the widely spread practice to assess research performance in computer science based on rankings, emphasizes the need for a CS research organization to know its own values and act upon them. He states that *"values would be hard to include in a ranking system based entirely on computable data"*. We fully subscribe to his perspective that *"we should promote CS faculty who* ***take risks****, who* ***explore new areas****, who* ***work with people in other disciplines*** *and* ***publish in those disciplines' venues****, and who* ***engage with other parts of the world***".

This value based perspective is reflected by our active engagement with other sciences and societal domains towards, e.g., adopting AI in a responsible manner (Case Study D.3). Further examples include our deep involvement in the health domain through our research in bioinformatics (Case Study D.1), our research on coordination algorithms for smart power systems (Case Study D.4), and our software analytics research conducted in the context of the financial sector (Case Study D.2). Risk taking can be seen from our yearslong investment in pioneering the platforms, such as Tribler, which eventually served as the platform for our work on distributed digital trust (Trustchain, Section 5.4.1), and our award-winning research on search intent proposing a disruptive innovation to the standard paradigm of the relevance-based video search (Section 5.1.4). Finally, we encourage new tenure track faculty to engage in inter-disciplinary research

1 D. Parnas: Stop the Numbers Game, CACM 50(11):19-21, 2007

2 M. Guzdial, Why I Don't Recommend CSRankings.org: Know the Values You are Ranking On, BLOG@CACM, 2020

through the Delft AI Labs (Section G.3). This is aligned with our modernized tenure-track policy focusing on potential rather than on quantitiative results, and our efforts to foster the development of academic leadership at our staff, reaching far beyond their scientific excellence (Section 3.3.3).

Viability An analysis of the tables B.4, *tab:appointed omitted* and *tab:departed omitted* in Appendix B points to a steady net growth of both departments since 2015 regarding their staff capacity, facilitated by building up significant income across the three money tiers, and despite a notable number of retired (7) or departed (25) staff members. The direct funding (first money tier) has grown by intensifying our educational activities as a consequence of our rapidly increasing student population, through which we respond to the increasing demand for CS specialists. As explained in Section 3.2, this line of income enabled us to work towards improving the balance between the teaching and research capacity as the prerequisite for maintaining a healthy task portfolio per faculty member, hence securing the conditions for both excellent teaching and excellent research. We have shown to be able to effectively exploit this income growth by hiring 48 new staff members over the assessed time period across both departments. While we were not able to influence all the reasons for the departure of 25 staff members, we realize we should better influence those that we can in order to reduce the incentive to leave. The value-inspired strategy we put in place (Section 3.3.1) serves precisely this purpose.

At the same time, we have demonstrated our capability of attracting and maintaining a healthy portfolio of external funding, leading to a rapid increase in the intake of PhD students in both departments (Table B.7) and in this way to a steady growth of our overall research capacity. This growth is even larger than observable in this table, because of the increase in the number of postdoctoral researchers over the assessed time period, and then especially due to a rapid build-up of the 'postdoc' population at the ST department (Table B.1). The significance of the project intake over the past years can be seen in Table B.2, which shows the development of the shares of the direct funding (only the part deployed to fund our research capacity) and the external funding divided over research grants and contract research. The key insight from this table is that both departments have managed to build up (ST from 40% in 2016 up to 62% in 2019) and maintain (INSY varying between 53-60%) the relative influx of external funds larger than its direct funding over the assessed time period, creating in this way a clear continuous multiplier effect in terms of 'return on investment'. The 'all-time high' in project intake in 2020 (Table B.6), combined with the intake results prepared in the assessed time period that will visible only in 2021 (e.g., ICAI Labs with DSM and Booking.com), reflect a strong position of CS@Delft as a whole, as well as of each individual department, to respond well to the increasing demand for CS scientific innovation (research grants) and its transfer to socio-economic sectors (contract research) in the years ahead.

The growth described above, especially for the second half of the assessed time period (2018-2020), has been achieved during a highly dynamic transition period, characterized, a.o., by a rapid increase in student numbers and the COVID pandemic. Our young faculty, from either the ST or INSY department, not only have shown themselves to be capable of acquiring research funding, even under unusually difficult conditions, but have also intensively been assuming responsibilities in various national and international scientific programs and consortia. This shows that the growth of our academic staff has also initiated the desired natural spread of responsibilities across generations, ensuring continuity of our organization. While the tables in Appendix *app:outcomes omitted* illustrate the role of younger faculty members in projects, committees and consortia, research highlights in Chapter 5 and the key publications in Appendix C show that the new generation is also picking up scientific leadership, securing scientific excellence and impact of CS@Delft in the coming years.

6.2 Future Strategy

The reflection above indicates that our strategy (Chapter 3) for further reinforcing the position of CS@Delft as a top international CS research organization is largely effective. We therefore choose to continue to follow this strategy conceptually also for the upcoming time period. The results of this strategy should further be facilitated by the expected stabilization of student numbers and the growth of our faculty formation towards 100 FTE (Section 3.2), securing the optimal student-staff ratio and empowering our faculty to create scientific excellence and impact even more than in the assessed time period.

However, we also see room for further improvement, achieving which could make our strategy even more effective. This is based on the findings from the SWOT analysis displayed in Table 6.1, and the benchmark

Table 6.1: SWOT Analysis CS@Delft

Strengths	**Weaknesses**
S1: Scientific reputation and impact	**W1:** Inadequate gender balance
S2: Socioeconomic embedding and impact	**W2:** Understaffed for educational responsibilities
S3: Financial potential through structural financial impulse, large numbers of students and increasing project income	**W3:** High staff turnover
S4: Organizational alignment of two CS departments with joint mission and ambition	**W4:** Ad hoc, scattered computer infrastructure; little structural funding for research software engineers
S5: Open access frontrunner	**W5:** Office shortage
	W6: Visibility of some of our research themes and directions
	W7: Efficiency and effectiveness of the PhD program
Opportunities	**Threats**
O1: Rapidly increasing societal demand for CS expertise and experts	**T1:** Foundational research under pressure
O2: Growing (inter)national strategic initiatives around CS (NLAIC, ICAI, ELLIS, CLAIRE, IPN)	**T2:** Pull of people by both industry and academia
O3: Global push for open science	**T3:** COVID pandemic
O4: Rapidly increasing opportunities for interdisciplinary research with CS as a key contributor	**T4:** Critical dependence on powerful computing, data, and software infrastructure and support

analysis, for which we engaged in a reflective discussion with our colleagues from the *Fachbereich Informatik*, TU Darmstadt, Germany (Appendix *app:benchmark omitted*). We outline the focus of the envisioned improvements below into three main aspects: people, organization and profiling.

- **People:** Our academic staff are our greatest asset, which makes attracting, coaching and retaining talent one of the most important pillars of our strategy. We will continue to strengthen this pillar with our scientific reputation and impact (S1) and deploy increasing opportunities for interdisciplinary research (O4) to attract talent. This will be facilitated by healthy financial position (S3) to invest in talent, our future-oriented CS@Delft mission (S4) and the emerging strategic initiatives (O2) involving CS in stimulating and accommodating talent development. The pillar will serve in particular to address the issues of gender balance (W1), under-staffing (W2), high staff turnover (W3) and efficiency and effectiveness of PhD research (W7), and to mitigate the effect of people being poached by external parties (T2).
- **Organization:** In order to stay strong, while growing and operating in an increasingly complex and competitive context, the CS@Delft organization needs to be reviewed and placed on stronger foundations. To achieve this, we will build on the already strong alignment of the INSY and ST departments and their joint mission and ambition (S4), in order to address the issues related to infrastructure (W4) and office space (W5). In the same way, we will mitigate the threats of the influence of the COVID pandemic on the work experience and work-life balance of staff members (T3), and the critical dependence of CS on software and hardware infrastructure and the related support (T4).
- **Profiling:** With the development of our staff and our organization, and also in view of the developments around us (O2), it is critical to improve our profiling along several dimensions: (a) articulating new promising research lines (W6), by building on our already strong reputation (S1) and further exploiting the opportunities regarding interdisciplinary research (O4), (b) creating more room and resources for foundational research (T1), e.g., by means of the NWO/ERC personal grants, and (c) open science, building on our extensive efforts from the past years (S5) and targeting an even greater role of it in our academic culture, in line with the global push in this direction (O3).

Table 6.2: Strategic actions and how they use, address, leverage, and mitigate the SWOT items

	Action	Use	Address	Leverage	Mitigate
A1	Improving attractiveness as employer	S3, S4	W1, W3	O2, O4	T2
A2	Improving diversity and inclusion	S1, S3	W1, W2, W3	O4	T2
A3	Improving the execution of the PhD program	S1, S4	W7	O1	
A4	Expanding the staff capacity	S3, S4	W2, W3, W4		T2, T4
A5	Aligning office space with the optimal way of working	S4	W5		T3
A6	Keep articulating the scientific profile	S1, S4	W6		T1
A7	Boosting the scientific and societal impact	S2, S5		O1, O2, O3, O4	

An overview of the actions towards our future strategy, indicating which strengths we will use, which weaknesses we will address, which opportunities we seek to leverage, and which threats we intend to mitigate is provided in Table 6.2. We discuss these actions in more detail below.

6.2.1 People

A1: Improving attractiveness as employer Within the assessed time period, we have shown ourselves to be capable of attracting (inter)national academic talent in substantial numbers (Appendix B), despite our disadvantageous position in terms of financial start-up and support packages, which are common at our international peer-institutes (see Benchmark: *Growth, Composition and Retention*, Appendix *app:benchmark omitted*). In the same time period, however, many faculty members have left, either to industry or to other universities. While some of the staff departures were due for personal reasons (e.g., two-body problem), we realize that we can still do better in improving the work conditions and promotion opportunities of our staff in order to minimize the incentive for them to leave due to career development reasons. We have therefore put in place several instruments that should help us reduce staff turnover and mitigate the attraction of external parties:

- A customized tenure-track process in which tenure can be awarded already after 2-2.5 years, and then primarily based on the development and impact potential. In view of the positive experience with this practice started in 2018 (Section 3.3.3), it is our intention to introduce this structurally.
- Continuous investment in the development of our academic staff at all levels, by means of courses, mentoring or personal coaching. While these instruments have been deployed before, we will deploy them in the future in a more effective fashion, informed by the *Development Track Plan* (Appendix F) for tenure-track Assistant Professors and Academic Leadership Profile for more senior faculty members.
- Fostering autonomy and leadership for assistant and associate professors. We will continue to stimulate and empower them to form own research clusters, to acquire the *"ius promovendi"*, to team-up with colleagues across section, department and faculty boundaries, and to contribute to the department, faculty or university-wide processes, collaborative research programs and institutes, on a both national and international level. In this way, they will receive the recognition and exposure they deserve, which will help mitigate limited recognition of assistant/associate professors in the Dutch university model and from external (non-academic) parties.
- Introducing more flexibility in the academic career development, not requiring each scientific staff member to excel along all dimensions. In this way, career progress from tenure-track towards full professorship should also be made possible through innovation and (international) impact in scientific education, and not only via top international research.

A2: Improving diversity and inclusion While cultural diversity has been achieved to a large extent (Section 3.3.4), the issue of gender imbalance persists. A substantial number of the faculty members who departed over recent years consisted of female scientists. Despite our efforts and success to maintain gender diversity through our staff acquisition, the share of female faculty members in CS@Delft is only around 20%.

Since we believe in the critical role of gender-diverse faculty population on productivity and excellence, we have introduced several plans to improve this percentage and reach the target of 25% by 2025. The first plan is to benefit more from the DTF rounds (Section 3.3.4). We will do this by working more closely with professional recruiters, by investing more in creating our own CS campaign as part of the overall DTF campaign, and by negotiating more financial room in the DTF budget for the prospective CS candidates (see Benchmark: *Growth, Composition and Retention*, Appendix *app:benchmark omitted*). Pursuing diversity, however, is not sufficient, unless we also work on inclusion. For this, we will intensify our efforts on mentoring and coaching of our staff to overcome biases (e.g., by deploying the *Harvard Bias Test*), as well empowering EDIT (EEMCS Diversity and Inclusion Team) to influence the development of our academic culture in this respect as much as possible.

A3: Improving the execution of the PhD program Over the past years, we have developed a strategy for optimizing our PhD program, which aims at (a) scientific excellence of PhD research, (b) development of skills and competences of our PhD graduates to make them qualify as independent scientists in the broadest possible sense, and (c) completing and defending the PhD thesis within five years. As the efficiency of the PhD research for CS@Delft remains as important as the PhD research quality (see Benchmark: *Performance measures*, Appendix *app:benchmark omitted*), we believe that we can improve our current strategy as explained in Section 3.3.5 through several measures that we will explore in the coming period. First of all, we will work closely with the EEMCS Graduate School to monitor and, where possible, help improve the current implementation of the PhD mentoring system (Appendix E). Secondly, we will set up discussion venues within and across the two CS departments to exchange best practices in PhD recruiting, supervision and assessment. Thirdly, we will work out a systematic procedure on what to do in case a serious delay in completing the PhD research occurs, especially when the reasons for the delay are not objectively justifiable.

6.2.2 Organization

A4: Expanding staff capacity The main prerequisite for excellence and impact is to create and maintain sufficient staff capacity enabling adequate achievement of an academic task portfolio across its scope (see Benchmark: *Growth, Composition and Retention*, Appendix *app:benchmark omitted*). As we have still not reached the faculty size stated in our ambition (Section 3.2), continuing this growth in the upcoming time period is of utmost importance, with the plan to complete it by 2024, synchronized with the stabilization of our student population. Since the funding situation enabling this growth has been developing positively, we expect to achieve this ambition as planned with the support of the EEMCS Faculty and the TU Delft Executive Board.

Furthermore, modern CS research depends increasingly on compute, storage, and software infrastructure, which needs continuous investments to facilitate research at scale. Merely investing in hardware infrastructure (see Appendix H) is not enough. What is also needed is substantial software development capacity, to support researchers in creating, maintaining, and operating (open source) software tools and services, for the research community and society alike. Therefore, in the upcoming time period, we will build up and structurally maintain a pool of software engineers as part of our staff formation, aiming at software development and consolidation for research, building up and maintaining educational software infrastructure and facilitating the use of the growing compute infrastructure for researchers.

A5: Aligning office space with the optimal way of working At the time of this research assessment, CS@Delft is in the process of reviewing its premises. This process was started in view of the significant growth of CS@Delft, and will rely on extra available space in the current building and new space in the adjacent ECHO building that is currently under construction. Part of this process is to take lessons from the current COVID crisis into account. These include the need to facilitate hybrid (on campus, remote) meetings, the importance of the office as a social meeting place as well as the place to do concentrated work, the role of the office in fostering team spirit, and the university campus as the place where students engage with faculty. Since our work is likely to still be largely influenced by the COVID pandemics, the management of the INSY and ST departments will initiate a discussion with the CS@Delft staff to discover what balance in working on and off campus to choose and for whom, and how to organize the offices to facilitate this. The discussion will be approached from both the perspective of an individual staff member and the perspective of the organization, with the objective to come up with a solution that serves us all best.

6.2.3 Profiling

A6: Keep articulating the scientific profile In the upcoming assessment time period, CS@Delft will continue to profile itself along the five research themes as presented in the previous chapter (see Benchmark: Thematic profile, Appendix *app:benchmark omitted*). The five themes were recently established and need time to mature. However, following our bottom-up, talent-driven research strategy (Section 6.1), topical evolution could emerge within the themes, for example through hiring of academic talent of a potential high impact, but with a new expertise, e.g., quantum software or Natural Language Processing (NLP), or with new research perspective on an existing topic. Also likely is the growth of the current research 'seeds', like around conversational AI, security & AI, and data management, which encompass complementary activities by young tenure-track assistant professors across section and even department boundaries.

An important other aspect of our scientific profile is the foundational research, typically approached by means of personal grants (NWO or ERC). As stated in Section 6.1, we see the need to strengthen this type of research in the CS domains, in which the role of academia is critical for resolving the bottlenecks related to the datafication and AI-zation phenomena. The developments in the EU and Dutch research funding landscape are, however, not in favour of foundational research (see Benchmark: *(Inter-)disciplinary research and collaboration*, Appendix *app:benchmark omitted*), with as a result that the project-submission channels are strongly overcrowded. In order to maximize the success probability, a strategy is needed, as well as allocation of resources, which can help our staff members to optimize their activities along the entire trajectory, from thinking about a suitable topic, via refining the topic through the discussion in various internal ad-hoc committees and writing the proposal with the help of experts, possibly from external companies, until practicing for an interview with the committee. Defining such strategy and reserving sufficient resources to facilitate personal-grant applications and improving their effectiveness is one of our highest priorities in the upcoming time period.

A7: Boosting the scientific and societal impact Over the past six years, we achieved a quantum leap in our scientific and societal impact, which can be seen in a number of (inter)disciplinary projects of strategic importance we acquired or (co-)initiated. Two Graviation Projects (Hybrid Intelligence and BRAINSCAPES) and three ICAI Labs (DSM, ING, Booking.com) are witness of our visibility and capability to exploit new strategic funding and networking opportunities. It is now our ambition to draw the maximum benefit from these projects to achieve scientific breakthroughs and maximize the societal and economic impact of our research. Furthermore, facilitated by expanding our research capacity (action A4) and learning from past experience, it is our ambition to motivate and empower our scientific staff to pursue new strategic initiatives, leading to a growing portfolio of impactful projects and increased scientific and societal reputation of CS@Delft. This will further be facilitated by our strong participation in European networks, such as ELLIS (AppendixG) and CLAIRE, which provides opportunities to interact and collaborate with the top AI researchers. This allows our faculty members to further expand their scientific network, to add to the visibility of their research, and enable them to influence the European agenda on AI and related topics. Finally, our efforts towards consolidating our education to large numbers of students, our leading role in setting up AI education for other disciplines at TU Delft (AIDU program, Appendix G) will secure continuous delivery of well-educated CS experts across the scope to satisfy the growing societal demand (see Benchmark: *Profile within the university*, Appendix *app:benchmark omitted*).

A | Performance Indicators

Research Quality

1. **Research products for peers**
   (a) (Open access) articles in refereed journals and conference proceedings
   (b) Open access dissertations
   (c) Technical products (algorithms, (open) data sets, (open source) software prototypes, algorithms, theories, methods, designs, on line services)

2. **Use of research products by peers**
   (a) Citations (e.g., Google Scholar)
   (b) Use of technical products (see 1c) in research

3. **Marks of recognition from peers**
   (a) Major awards and (personal) grants
   (b) Positions in research-related organizations (steering committees, program / general chair of major conferences, editorial boards, learned societies)
   (c) Research leadership (leadership or coordination of national and international research programs and projects)
   (d) Keynotes at leading conferences
   (e) Honors, professional standing, external appointments, best / most influential paper awards

Societal Relevance

4. **Research products for societal target groups**
   (a) Academically qualified engineers and scientists, at MSc, PhD, and postdoc level.
   (b) Technical products (algorithms, (open) data sets, (open source) software prototypes, algorithms, theories, methods, designs, on line services)
   (c) Start-up companies
   (d) Outreach activities (blogs, social media, presentations at industry conferences, press appearances)

5. **Use of research products by societal target groups**
   (a) Downloads (tools), monthly users (services) or GitHub stars (open source libraries)
   (b) Use of technical products (see 4b) in societal organizations.
   (c) Use of research in (online) professional education
   (d) Cooperation with industry or societal organizations (government) in consortia or directly funded by industry

6. **Marks of recognition by societal target groups**
   (a) Public awards (research-related appreciation, by industry or by the public)
   (b) (Invited) positions in management and advisory bodies.
   (c) Staff exchange (senior-level staff exchange, industry sponsored appointments)

B | Facts and Figures

Table B.1: Research staff for the INSY and ST departments, and for Computer Science overall, conform Table E2 of the SEP. The #-column refers to average headcount over a year, where one month counts as 1/12=0.083. Column entries are rounded to whole number

Conform SEP, the FTE column represents the *research* responsibility of the staff, using 0.4 for faculty members, and 0.8 for postdocs and PhD students.

Visiting fellows have an appointment of at least 0.5 FTE. Support staff includes developers.

Note that counting staff by the percentage of months employed makes it somewhat harder to see the growth the departments have gone through, especially for employees who started late 2020. For that reason we also included data on faculty members who joined or left the departments (Tables *tab:appointed omitted* and *tab:departed omitted*). Furthermore, in Table B.4 we provide the number of faculty members at the start of the period (January 2015) as well as at the end (December 2020).

Scientific Staff	2015		2016		2017		2018		2019		2020	
Intelligent Systems	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**
Full Professor	6	2.2	6	2.2	8	2.5	8	2.5	8	2.4	8	2.3
Associate Professor	7	2.8	6	2.5	5	1.9	8	2.9	8	2.8	8	3.0
Assistant Professor	15	6.1	15	6.0	15	6.0	14	5.7	15	5.9	17	6.8
Postdocs	22	15.7	20	13.6	14	9.2	11	7.5	13	9.2	18	13.0
PhD Candidates (standard)	38	30.4	37	29.0	40	31.6	40	31.8	48	37.2	59	46.9
PhD Candidates (ext./scholarship)	22	12.8	17	9.9	20	11.9	20	11.7	20	12.4	20	14.1
Total Research Staff	**111**	**69.9**	**101**	**63.1**	**103**	**63.2**	**102**	**62.1**	**111**	**70.0**	**130**	**86.0**
Support Staff (research)	4	2.9	4	2.8	5	3.4	4	2.8	4	2.7	4	2.5
Visiting Fellows	15	0.0	13	0.0	14	0.0	16	0.0	10	0.0	14	0.0
Total	**129**	**72.8**	**117**	**66.0**	**121**	**66.6**	**122**	**65.0**	**125**	**72.7**	**147**	**88.5**

	2015		2016		2017		2018		2019		2020	
Software Technology	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**
Full Professor	9	3.0	8	2.7	8	2.6	8	2.5	8	2.7	9	3.0
Associate Professor	6	2.2	5	2.0	6	2.3	8	3.3	9	3.6	9	3.5
Assistant Professor	12	4.4	11	4.0	13	5.0	15	5.6	14	5.2	16	6.3
Postdocs	15	11.1	12	9.3	17	13.6	23	17.9	26	19.7	28	20.1
PhD Candidates (standard)	26	19.7	29	22.2	34	25.6	41	31.7	46	35.9	58	46.2
PhD Candidates (ext./scholarship)	12	6.0	12	4.9	14	5.9	14	5.9	16	6.6	16	6.6
Total Research Staff	**79**	**46.3**	**77**	**45.1**	**92**	**55.0**	**109**	**66.9**	**118**	**73.6**	**136**	**85.7**
Support Staff (research)	4	3.1	4	3.0	4	2.4	3	2.3	4	3.3	6	4.2
Visiting Fellows	10	0.0	13	0.0	14	0.0	15	0.0	11	0.0	6	0.0
Total	**93**	**49.4**	**94**	**48.1**	**110**	**57.4**	**127**	**69.2**	**134**	**76.9**	**148**	**89.9**

	2015		2016		2017		2018		2019		2020	
Computer Science (total)	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**	**#**	**FTE**
Full Professor	15	5.2	14	4.8	16	5.1	16	5.0	15	5.1	17	5.3
Associate Professor	13	5.0	11	4.5	11	4.2	16	6.2	17	6.4	17	6.5
Assistant Professor	27	10.5	26	10.0	28	11.1	29	11.3	28	11.1	33	13.1
Postdocs	37	26.7	32	22.9	31	22.8	34	25.4	39	28.8	45	33.0
PhD Candidates (standard)	64	50.1	65	51.2	74	57.2	82	63.5	94	73.1	117	93.1
PhD Candidates (ext./scholarship)	35	18.7	29	14.8	34	17.8	34	17.6	35	19.1	36	20.7
Total Research Staff	**189**	**116.2**	**178**	**108.2**	**195**	**118.2**	**211**	**129.0**	**229**	**143.6**	**266**	**171.8**
Support Staff (research)	8	6.0	8	5.8	8	5.8	7	5.1	8	6.0	10	6.7
Visiting Fellows	25	0.0	25	0.0	28	0.0	31	0.0	21	0.0	20	0.0
Total	**222**	**122.2**	**211**	**114.0**	**231**	**124.0**	**249**	**134.2**	**259**	**149.6**	**295**	**178.5**

Table B.2: Funding and expenditure for the INSY and ST departments, and for Computer Science overall, conform Table E3 of the SEP. The FTE values are *research* FTEs, conform Table B.1. The FTE values in the *Total Funding* row directly correspond to the FTE values in the *Total Research Staff* row in Table B.1. Conform SEP, the percentages below give the relative research funding from the 1st, 2nd, and 3d money tiers. The expenditure column maps the FTE numbers to Euros for personnel costs only, considering no other expenditures were significant during the evaluation period.

Funding	2015		2016		2017		2018		2019		2020	
Intelligent Systems	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**
Funding												
1. Direct Funding	33.1	46%	26.2	40%	31.7	48%	28.8	44%	34.0	47%	42.0	47%
2. Research Grants	10.8	15%	13.3	20%	12.7	19%	14.7	23%	13.2	18%	16.5	19%
3. Contract Research	28.8	40%	26.5	40%	22.2	33%	21.4	33%	25.5	35%	30.0	34%
Total Funding	**72.8**	**100%**	**66.0**	**100%**	**66.6**	**100%**	**65.0**	**100%**	**72.7**	**100%**	**88.5**	**100%**
Expenditure (in Keuro)												
Personnel Costs	3.118	100%	2.927	100%	2.770	100%	2.921	100%	3.235	100%	4.041	100%
Total Expenditure	**3.118**	**100%**	**2.927**	**100%**	**2.770**	**100%**	**2.921**	**100%**	**3.235**	**100%**	**4.041**	**100%**

	2015		2016		2017		2018		2019		2020	
Software Technology	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**
Funding												
1. Direct Funding	29.8	60%	20.4	42%	24.8	43%	28.9	42%	29.4	38%	36.6	41%
2. Research Grants	4.5	9%	12.5	26%	18.2	32%	21.5	31%	20.1	26%	18.1	20%
3. Contract Research	15.1	31%	15.2	32%	14.4	25%	18.8	27%	27.4	36%	35.2	39%
Total Funding	**49.4**	**100%**	**48.1**	**100%**	**57.4**	**100%**	**69.2**	**100%**	**76.9**	**100%**	**89.9**	**100%**
Expenditure (in Keuro)												
Personnel Costs	2.324	100%	2.224	100%	2.518	100%	3.230	100%	3.630	100%	4.396	100%
Total Expenditure	**2.324**	100%	**2.224**	100%	**2.518**	100%	**3.230**	100%	**3.630**	100%	**4.396**	100%

	2015		2016		2017		2018		2019		2020	
Computer Science (total)	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**	**FTE**	**%**
Funding												
1. Direct Funding	63.0	52%	46.6	41%	56.5	46%	57.7	43%	63.4	42%	78.6	44%
2. Research Grants	15.3	13%	25.8	23%	30.9	25%	36.2	27%	33.4	22%	34.7	19%
3. Contract Research	43.9	36%	41.7	37%	36.6	30%	40.2	30%	52.9	35%	65.2	37%
Total Funding	**122.2**	**100%**	**114.0**	**100%**	**124.0**	**100%**	**134.2**	**100%**	**149.6**	**100%**	**178.5**	**100%**
Expenditure (in Keuro)												
Personnel Costs	5.442	100%	5.152	100%	5.288	100%	6.152	100%	6.865	100%	8.437	100%
Total Expenditure	**5.442**	**100%**	**5.152**	**100%**	**5.288**	**100%**	**6.152**	**100%**	**6.865**	**100%**	**8.437**	**100%**

Table B.3: PhD Candidates for the INSY and ST departments, and for Computer Science overall (conform Table E4 of the SEP)

PhD Candidates																
Intelligent Systems				Graduated		Graduated		Graduated		Graduated		Graduated		Not yet		
Starting	Enrolment			<= 4 years		<= 5 years		<= 6 years		<= 7 years		total		finished		Discontinued
Year	M	F	Total	#	%	#	%	#	%	#	%	#	%	#	%	#	%
2012	15	2	17	0	0%	5	29%	8	47%	9	53%	10	59%	0	0%	7	41%
2013	12	2	14	0	0%	8	57%	10	71%	12	86%	12	86%	1	7%	1	7%
2014	10	2	12	0	0%	6	50%	9	75%	9	75%	9	75%	1	8%	2	17%
2015	13	2	15	1	7%	7	47%	10	67%	10	67%	10	67%	1	7%	4	27%
2016	11	4	15	0	0%	3	20%	4	27%	4	27%	4	27%	9	60%	2	13%
Total	**61**	**12**	**73**	**1**	**1%**	**29**	**40%**	**41**	**56%**	**44**	**60%**	**45**	**62%**	**12**	**16%**	**16**	**22%**

Software Technology				Graduated		Graduated		Graduated		Graduated		Graduated		Not yet		
Starting	Enrolment			<= 4 years		<= 5 years		<= 6 years		<= 7 years		total		finished		Discontinued
Year	M	F	Total	#	%	#	%	#	%	#	%	#	%	#	%	#	%
2012	8	1	9	0	0%	4	44%	5	56%	5	56%	7	78%	1	11%	1	11%
2013	6	0	6	0	0%	3	50%	4	67%	5	83%	5	83%	1	17%	0	0%
2014	9	1	10	0	0%	4	40%	5	50%	5	50%	5	50%	3	30%	2	20%
2015	8	1	9	2	22%	7	78%	7	78%	7	78%	7	78%	2	22%	0	0%
2016	11	2	13	0	0%	4	31%	4	31%	4	31%	4	31%	6	46%	3	23%
Total	**42**	**5**	**47**	**2**	**4%**	**22**	**47%**	**25**	**53%**	**26**	**55%**	**28**	**60%**	**13**	**28%**	**6**	**13%**

Computer Science (total)				Graduated		Graduated		Graduated		Graduated		Graduated		Not yet		
Starting	Enrolment			<= 4 years		<= 5 years		<= 6 years		<= 7 years		total		finished		Discontinued
Year	M	F	Total	#	%	#	%	#	%	#	%	#	%	#	%	#	%
2012	23	3	26	0	0%	9	35%	13	50%	14	54%	17	65%	1	4%	8	31%
2013	18	2	20	0	0%	11	55%	14	70%	17	85%	17	85%	2	10%	1	5%
2014	19	3	22	0	0%	10	45%	14	64%	14	64%	14	64%	4	18%	4	18%
2015	21	3	24	3	13%	14	58%	17	71%	17	71%	17	71%	3	13%	4	17%
2016	22	6	28	0	0%	7	25%	8	29%	8	29%	8	29%	15	54%	5	18%
Total	**103**	**17**	**120**	**3**	**3%**	**51**	**43%**	**66**	**55%**	**70**	**58%**	**73**	**61%**	**25**	**21%**	**22**	**18%**

Table B.4: Academic staff by rank and gender per year. For each year we give average headcount per year, rounded to whole numbers, where the "Total" rows directly correpond to the numbers for full, associate, and assistant professors listed in Table B.1. We also provide the starting state in January 2015, and the end state in December 2020, to provide the clearest perspective on the overall growth. Due to the level of aggregation and the calculation of average headcount per year (conform SEP) the totals for each # column might differ from the sum of rows due to rounding differences.

Gender Balance Scientific Staff (full, associate, assistant professors)																
	2015 (jan)		2015 (avg.)		2016		2017		2018		2019		2020 (avg.)		2020 (dec)	
Intelligent Systems	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%
Men	19	68%	19	68%	20	70%	21	73%	22	72%	21	70%	24	71%	25	71%
Women	9	32%	9	32%	8	30%	8	27%	9	28%	9	30%	10	29%	10	29%
Total	**28**	**100%**	**28**	**100%**	**28**	**100%**	**29**	**100%**	**31**	**100%**	**31**	**100%**	**33**	**100%**	**35**	**100%**

	2015 (jan)		2015		2016		2017		2018		2019		2020		2020 (dec)	
Software Technology	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%
Men	24	92%	24	92%	22	92%	24	89%	27	88%	25	83%	30	86%	33	89%
Women	2	8%	2	8%	2	8%	3	11%	4	12%	5	17%	5	14%	4	11%
Total	**26**	**100%**	**26**	**100%**	**24**	**100%**	**27**	**100%**	**31**	**100%**	**30**	**100%**	**35**	**100%**	**37**	**100%**

	2015 (jan)		2015		2016		2017		2018		2019		2020		2020 (dec)	
Computer Science (total)	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%
Men	43	80%	43	80%	42	80%	45	81%	49	80%	47	77%	53	79%	58	81%
Women	11	20%	11	20%	10	20%	11	19%	12	20%	14	23%	14	21%	14	19%
Total	**54**	**100%**	**54**	**100%**	**52**	**100%**	**56**	**100%**	**62**	**100%**	**61**	**100%**	**68**	**100%**	**72**	**100%**

Table B.5: Open Access Publications

Open access publications OA data not available for 2015, CS total may differ due to shared publications	2015		2016		2017		2018		2019		2020	
Intelligent Systems	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**
Refereed article	107	-	85	72%	92	64%	94	77%	102	82%	114	82%
Conference papers	91	-	125	26%	110	35%	110	47%	130	53%	120	73%
Total	**198**	**-**	**210**	**45%**	**202**	**49%**	**204**	**61%**	**232**	**66%**	**234**	**77%**

	2015		2016		2017		2018		2019		2020	
Software Technology	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**
Refereed article	32	-	17	35%	25	80%	70	86%	61	84%	73	86%
Conference papers	99	-	94	54%	116	53%	156	77%	122	72%	121	82%
Total	**131**	**-**	**111**	**51%**	**141**	**58%**	**226**	**80%**	**183**	**76%**	**194**	**84%**

	2015		2016		2017		2018		2019		2020	
Computer Science (total)	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**	**#**	**OA %**
Refereed article	139	-	102	66%	117	68%	164	80%	162	83%	183	84%
Conference papers	188	-	215	38%	220	44%	263	65%	251	62%	236	77%
Total	**327**	**-**	**317**	**47%**	**337**	**52%**	**427**	**71%**	**413**	**70%**	**419**	**80%**

Table B.6: Intake volume of new projects acquired per year

Project Intake	2015		2016		2017		2018		2019		2020	
Intelligent Systems	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**
2. Research Grant	442	27%	459	36%	2.310	70%	748	19%	3.269	82%	2.952	55%
3. Contract Research	1.181	73%	832	64%	986	30%	3.235	81%	722	18%	2.383	45%
Total	**1.623**	**100%**	**1.291**	**100%**	**3.296**	**100%**	**3.984**	**100%**	**3.991**	**100%**	**5.334**	**100%**

	2015		2016		2017		2018		2019		2020	
Software Technology	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**
2. Research Grant	1.132	39%	767	28%	5.376	65%	1.797	21%	4.424	54%	5.459	40%
3. Contract Research	1.737	61%	1.950	72%	2.865	35%	6.861	79%	3.726	46%	8.231	60%
Total	**2.869**	**100%**	**2.717**	**100%**	**8.241**	**100%**	**8.658**	**100%**	**8.150**	**100%**	**13.690**	**100%**

	2015		2016		2017		2018		2019		2020	
Computer Science (total)	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**	**Keuro**	**%**
2. Research Grant	1.574	35%	1.226	31%	7.686	67%	2.545	20%	7.694	63%	8.411	44%
3. Contract Research	2.918	65%	2.782	69%	3.851	33%	10.097	80%	4.447	37%	10.614	56%
Total	**4.492**	**100%**	**4.009**	**100%**	**11.537**	**100%**	**12.642**	**100%**	**12.141**	**100%**	**19.025**	**100%**

Table B.7: Intake of new PhD students per year

PhD intake	2015		2016		2017		2018		2019		2020	
Intelligent Systems	**#**	**%**	**#**	**%**	**#**	**%**	**#**	**%**	**#**	**%**	**#**	**%**
Men	13	87%	11	73%	14	93%	21	75%	14	70%	23	72%
Women	2	13%	4	27%	1	7%	7	25%	6	30%	9	28%
Total	**15**	**100%**	**15**	**100%**	**15**	**100%**	**28**	**100%**	**20**	**100%**	**32**	**100%**

	2015		2016		2017		2018		2019		2020	
Software Technology	**#**	**%**	**#**	**%**	**#**	**%**	**#**	**%**	**#**	**%**	**#**	**%**
Men	8	89%	11	85%	14	88%	12	80%	21	72%	12	57%
Women	1	11%	2	15%	2	13%	3	20%	8	28%	9	43%
Total	**9**	**100%**	**13**	**100%**	**16**	**100%**	**15**	**100%**	**29**	**100%**	**21**	**100%**

	2015		2016		2017		2018		2019		2020	
Computer Science (total)	**#**	**%**	**#**	**%**	**#**	**%**	**#**	**%**	**#**	**%**	**#**	**%**
Men	21	88%	22	79%	28	90%	33	77%	35	71%	35	66%
Women	3	13%	6	21%	3	10%	10	23%	14	29%	18	34%
Total	**24**	**100%**	**28**	**100%**	**31**	**100%**	**43**	**100%**	**49**	**100%**	**53**	**100%**

C | Key Publications

Key publications over the period 2015–2020, as discussed in Chapter 5. Together, the selected papers are an indicator for the breadth of the topics we cover, the depth with which we do this, and the impact our papers can make. The papers are listed in the order in which they are discussed in Chapter 5, which explains how these papers fit in our research themes, how they are connected, and what their impact is in science as well as in society.

All papers are available as (green) open access via https://research.tudelft.nl, the TU Delft repository. For each entry, we offer a link to its entry in the repository, as well as the number of times the paper has been cited, and (when applicable) the rank and h5-index of the venue, all according to Google Scholar.

Decision Support

[1] Pradeep K. Murukannaiah, Nirav Ajmeri, Catholijn M. Jonker, and Munindar P. Singh. "New Foundations of Ethical Multiagent Systems". In: *Proceedings of the 19th Conference on Autonomous Agents and MultiAgent Systems*. AAMAS '20. 2020, pp. 1706–1710. Open access: link; #Cites 11, not ranked, h5:44

[2] Zeynep Akata, Dan Balliet, Maarten De Rijke, Frank Dignum, Virginia Dignum, Guszti Eiben, Antske Fokkens, Davide Grossi, Koen Hindriks, Holger Hoos, Hayley Hung, Catholijn Jonker, Christof Monz, Mark Neerincx, Frans Oliehoek, Henry Prakken, Stefan Schlobach, Linda Van Der Gaag, Frank Van Harmelen, Herke Van Hoof, Birna Van Riemsdijk, Aimee Van Wynsberghe, Rineke Verbrugge, Bart Verheij, Piek Vossen, and Max Welling. "A Research Agenda for Hybrid Intelligence: Augmenting Human Intellect With Collaborative, Adaptive, Responsible, and Explainable Artificial Intelligence". English. In: *Computer* 53.8 (Aug. 2020), pp. 18–28. ISSN: 0018-9162. DOI: 10.1109/MC.2020.2996587. Open access: link; #Cites 10, not ranked

[3] Erwin Walraven and Matthijs T. J. Spaan. "Accelerated Vector Pruning for Optimal POMDP Solvers". In: *Proceedings of the 31st AAAI Conference on Artificial Intelligence*. 2017, pp. 3672–3678. Open access: link; #Cites 27, #59 overall, #4 in Artificial Intelligence, h5:157

[5] Mathijs M De Weerdt, Sebastian Stein, Enrico H Gerding, Valentin Robu, and Nicholas R Jennings. "Intention-aware routing of electric vehicles". In: *IEEE Transactions on Intelligent Transportation Systems* 17.5 (2015), pp. 1472–1482. Open access: link; #Cites 55, #1 in Transportation, h5:107

[6] Marco Loog. "Contrastive Pessimistic Likelihood Estimation for Semi-Supervised Classification". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38.3 (2016), pp. 462–475. DOI: 10.1109/TPAMI.2015.2452921. Open access: link; #Cites 64, #70 overall, #4 in Computer Vision and Pattern Recognition, h5:149

[7] Marco Loog, Tom Viering, and Alexander Mey. "Minimizers of the Empirical Risk and Risk Monotonicity". In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett. Vol. 32. Curran Associates, Inc., 2019. Open access: link; #Cites 13, #12 overall, #2 in Artificial Intelligence, h5:245

[8] Yichao Zhang, Silvia L. Pintea, and Jan C. van Gemert. "Video Acceleration Magnification". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017. Open access: link; #Cites 57, #4 overall, #1 in Engineering and Computer Science, h5:356

[9] Bokun Wang, Yang Yang, Xing Xu, Alan Hanjalic, and Heng Tao Shen. "Adversarial Cross-Modal Retrieval". In: *Proceedings of the ACM Conference on Multimedia, MM2017*. **Best Paper Award**. ACM, 2017, pp. 154–162. DOI: 10.1145/3123266.3123326. Open access: link; #Cites 379, #4 in Multimedia, h5:65

[10] Roger Zhe Li, Julián Urbano, and Alan Hanjalic. "Leave No User Behind: Towards Improving the Utility of Recommender Systems for Non-mainstream Users". In: *WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining*. ACM, 2021, pp. 103–111. DOI: 10.1145/3437963.3441769. Open access: link; #Cites 1, #4 in Data Mining & Analysis, h5:61

[11] Sindunuraga Rikarno Putra, Felipe Moraes, and Claudia Hauff. "SearchX: Empowering Collaborative Search Research". In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018*. ACM, 2018, pp. 1265–1268. DOI: 10.1145/3209978.3210163. Open access: link; #Cites 12, #5 in Databases & Information Systems, h5:66

[12] Dan Davis, Ioana Jivet, Rene Kizilcec, Guanliang Chen, Claudia Hauff, and Geert-Jan Houben. "Follow the successful crowd: raising MOOC completion rates through social comparison at scale". In: *LAK '17: Proceedings of the Seventh International Learning Analytics & Knowledge Conference*. 2017, pp. 454–463. DOI: 10.1145/3027385.3027411. Open access: link; #Cites 112, #8 in Educational Technology, h5:40

Data Management and Analytics

[13] Ujwal Gadiraju, Jie Yang, and Alessandro Bozzon. "Clarity is a Worthwhile Quality: On the Role of Task Clarity in Microtask Crowdsourcing". In: *ACM Hypertext 2017: 28th ACM Conference on Hypertext and Social Media*. 2017, pp. 1–10. DOI: 10.1145/3078714.3078715. Open access: link; #Cites 71, not ranked, h5:21

[14] Jonas Traub, Philipp Grulich, Alejandro Rodriguez Cuellar, Sebastian Bress, Asterios Katsifodimos, Tilmann Rabl, and Volker Markl. "Efficient Window Aggregation with General Stream Slicing". In: *EDBT 2019 - Advances in Database Technology - 22nd International Conference on Extending Database Technology, EDBT 2019*. **Best paper award**. 2019, pp. 97–108. DOI: 10.5441/002/edbt.2019.10. Open access: link; #Cites 39, not ranked, h5:27

[16] Tamim Abdelaal, Lieke C.M. Michielsen, Davy Cats, Dylan Hoogduin, Hailiang Mei, Marcel J.T. Reinders, and Ahmed Mahfouz. "A comparison of automatic cell identification methods for single-cell RNA sequencing data". English. In: *Genome Biology* 20.1 (2019), pp. 1–19. ISSN: 1474-760X. DOI: 10.1186/s13059-019-1795-z. Open access: link; #Cites 143, #2 Genetics & Genomics, h5:124

[17] Abigail L. Manson, Keira A. Cohen, Thomas Abeel, Christopher A. Desjardins, Derek T. Armstrong, Clifton E. Barry, Jeannette Brand, Sinéad B. Chapman, Sang Nae Cho, Andrei Gabrielian, James Gomez, Andreea M. Jodals, Moses Joloba, Pontus Jureen, Jong Seok Lee, Lesibana Malinga, Mamoudou Maiga, Dale Nordenberg, Ecaterina Noroc, Elena Romancenco, Alex Salazar, Willy Ssengooba, A. A. Velayati, Kathryn Winglee, Aksana Zalutskaya, Laura E. Via, Gail H. Cassell, Susan E. Dorman, Jerrold Ellner, Parissa Farnia, James E. Galagan, Alex Rosenthal, Valeriu Crudu, Daniela Homorodean, Po Ren Hsueh, Sujatha Narayanan, Alexander S. Pym, Alena Skrahina, Soumya Swaminathan, Martie Van Der Walt, David Alland, William R. Bishai, Ted Cohen, Sven Hoffner, Bruce W. Birren, and Ashlee M. Earl. "Genomic analysis of globally diverse Mycobacterium tuberculosis strains provides insights into the emergence and spread of multidrug resistance". English (US). In: *Nature Genetics* 49.3 (Mar. 2017), pp. 395–402. ISSN: 1061-4036. DOI: 10.1038/ng.3767. Open access: link; #Cites 178, #1 in Genetics and Genomics, #32 overall, h5:184

[18] Ruben Wiersma, Elmar Eisemann, and Klaus Hildebrandt. "CNNs on Surfaces using Rotation-Equivariant Features". In: *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 39.4 (2020). DOI: 10.1145/3386569.3392437. Open access: link; #Cites 10, #1 in Computer Graphics, h5:85

[19] Vincent van Unen, Thomas Höllt, Nicola Pezzotti, Na Li, Marcel Reinders, Elmar Eisemann, Anna Vilanova, Frits Koning, and Boudewijn P.F. Lelieveldt. "Visual Analysis of Mass Cytometry Data by Hierarchical Stochastic Neighbour Embedding Reveals Rare Cell Types". In: *Nature Communications* 8.1740 (2017), pp. 1–10. Open access: link; #Cites 126, #7 overall, #3 in Life Sciences & Earth Sciences, h5:273

[20] Xiu-Xiu Zhan, Alan Hanjalic, and Huijuan Wang. "Information diffusion backbones in temporal networks". In: *Scientific reports* 9.1 (2019), pp. 1–12. DOI: 10.1038/s41598-019-43029-5. Open access: link; #Cites 17, #28 overall, #5 Life Sciences & Earth Sciences, h5:178

[21] Huijuan Wang, Chuyi Chen, Bo Qu, Daqing Li, and Shlomo Havlin. "Epidemic mitigation via awareness propagation in communication networks: the role of time scales". In: *New Journal of Physics* 19.7 (2017), p. 073039. Open access: link; #Cites 13, #8 Physics & Mathematics (general), h5:76

Software Systems Engineering

[22] Georgios Gousios, Margaret-Anne D. Storey, and Alberto Bacchelli. "Work practices and challenges in pull-based development: the contributor's perspective". In: *Proceedings of the 38th International Conference on Software Engineering, ICSE 2016*. **ACM SIGSOFT Distinguished Paper Award.** ACM, 2016, pp. 285–296. DOI: 10.1145/2884781.2884826. Open access: link; #Cites 142, #1 in Software Systems, h5:74

[23] Moritz Beller, Georgios Gousios, Annibale Panichella, Sebastian Proksch, Sven Amann, and Andy Zaidman. "Developer Testing in the IDE: Patterns, Beliefs, and Behavior". In: *IEEE Transactions Software Engineering* 45.3 (2019), pp. 261–284. DOI: 10.1109/TSE.2017.2776152. Open access: link; #Cites 47, #6 in Software Systems, h5:52

[24] Mozhan Soltani, Annibale Panichella, and Arie van Deursen. "Search-Based Crash Reproduction and Its Impact on Debugging". In: *IEEE Trans. Software Eng.* 46.12 (2020), pp. 1294–1317. DOI: 10.1109/TSE.2018.2877664. Open access: link; #Cites 22, #6 in Software Systems, h5:52

[25] Raja Ben Abdessalem, Annibale Panichella, Shiva Nejati, Lionel C. Briand, and Thomas Stifter. "Testing autonomous cars for feature interaction failures using many-objective search". In: *Proc. 33rd ACM/IEEE Int. Conf. on Automated Software Engineering,(ASE)*. ACM, 2018, pp. 143–154. DOI: 10.1145/3238147.3238192. Open access: link; #Cites 52, #10 in Software Systems, h5:45

[26] Pierre Néron, Andrew P. Tolmach, Eelco Visser, and Guido Wachsmuth. "A Theory of Name Resolution". In: *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015*. Vol. 9032. Lecture Notes in Computer Science. **EAPLS 2015 best paper award.** Springer, 2015, pp. 205–231. DOI: 10.1007/978-3-662-46669-8_9. Open access: link; #Cites 68, not ranked, h5:25

[27] Arjen Rouvoet, Hendrik van Antwerpen, Casper Bach Poulsen, Robbert Krebbers, and Eelco Visser. "Knowing when to ask: sound scheduling of name resolution in type checkers derived from declarative specifications". In: *Proceedings of the ACM on Programming Languages* 4.OOPSLA (2020). DOI: 10.1145/3428248. Open access: link; #Cites 1, #31 in Software Systems, h5:31

[28] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. "Iris from the ground up: A modular foundation for higher-order concurrent separation logic". In: *Journal of Functional Programming* 28 (2018). DOI: 10.1017/S0956796818000151. Open access: link; #Cites 149, not ranked, h5:15

[29] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. "RustBelt: securing the foundations of the Rust programming language". In: *Proceedings of the ACM on Programming Languages* 2.POPL (2018). DOI: 10.1145/3158154. Open access: link; #Cites 177, #7 in Software Systems, h5:48

Distributed and Networked Systems

[30] Pim Otte, Martijn de Vos, and Johan Pouwelse. "TrustChain: A Sybil-resistant Scalable Blockchain". In: *Future Generation Computer Systems* 107 (2020), pp. 770–780. DOI: 10.1016/j.future.2017.08.048. Open access: link; #Cites 122, #2 in Computing Systems, h5:86

[32] Kasim Sinan Yildirim, Amjad Yousef Majid, Dimitris Patoukas, K. Schaper, Przemysław Pawełczak, and Josiah Hester. "InK: Reactive kernel for tiny batteryless sensors". In: *SenSys 2018 - Proceedings*

of the 16th Conference on Embedded Networked Sensor Systems. Ed. by Gowri Sankar Ramachandran and Bhaskar Krishnamachari. United States: Association for Computing Machinery (ACM), 2018, pp. 41–53. DOI: 10.1145/3274783.3274837. Open access: link; #Cites 56, not ranked, h5:33

[33] Xueying Yang, Evgenios Karampatzakis, Christian Doerr, and Fernando Kuipers. "Security Vulnerabilities in LoRaWAN". In: *2018 IEEE/ACM Third International Conference on Internet-of-Things Design and Implementation (IoTDI)*. 2018, pp. 129–140. DOI: 10.1109/IoTDI.2018.00022. Open access: link; #Cites 58, not ranked, h5:22

[34] S. Narayana, R. V. Prasad, V. Rao, L. Mottola, and T. V. Prabhakar. "Hummingbird: Energy efficient GPS receiver for small satellites". In: *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. MobiCom '20. **Best paper award**. ACM, 2020, pp. 110–122. DOI: 10.1145/3372224.3380886. Open access: link; #Cites 2, not ranked, h5:48

[35] Alexey Ilyushkin, Ahmed Ali-Eldin, Nikolas Herbst, André Bauer, Alessandro Vittorio Papadopoulos, Dick H. J. Epema, and Alexandru Iosup. "An Experimental Performance Evaluation of Autoscalers for Complex Workflows". In: *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* 3.2 (2018), 8:1–8:32. DOI: 10.1145/3164537. Open access: link; #Cites 21, not listed

[36] Zilong Zhao, Sophie Cerf, Robert Birke, Bogdan Robu, Sara Bouchenak, Sonia Ben Mokhtar, and Lydia Y. Chen. "Robust Anomaly Detection on Unreliable Data". In: *49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN*. IEEE, 2019, pp. 630–637. DOI: 10.1109/DSN.2019.00068. Open access: link; #Cites 19, not ranked, h5:32

[37] Qing Wang, Marco Zuniga, and Domenico Giustiniano. "Passive Communication with Ambient Light". In: *12th International on Conference on Emerging Networking EXperiments and Technologies*. CoNEXT '16. Irvine, California, USA: ACM, 2016, pp. 97–104. DOI: 10.1145/2999572.2999584. Open access: link; #Cites 28, not ranked, h5:33

[38] Rens Bloom, Marco Zúñiga Zamalloa, and Chaitra Pai. "LuxLink: Creating a Wireless Link from Ambient Light". In: *17th Conference on Embedded Networked Sensor Systems*. SenSys '19. New York, New York: ACM, 2019, pp. 166–178. DOI: 10.1145/3356250.3360021. Open access: link; #Cites 3, not ranked, h5:33

Security and Privacy

[39] Lin, Q., Adepu, S., Verwer, S., and Mathur, A. "TABOR: A graphical model-based approach for anomaly detection in industrial control systems". In: *ASIACCS 2018 - Proceedings of the 2018 ACM Asia Conference on Computer and Communications Security*. Association for Computing Machinery (ACM), 2018, pp. 525–536. DOI: 10.1145/3196494.3196546. Open access: link; #Cites 58, #15 in Computer Security & Cryptography, h5:36

[40] Ghiëtte, V., Griffioen, H., and Doerr, C. "Fingerprinting tooling used for SSH compromisation attempts". In: *22nd International Symposium on Research in Attacks, Intrusions and Defenses (RAID)*. Usenix, 2019, pp. 61–71. Open access: link; #Cites 5, not ranked, h5:22

[41] J. Kim, S. Picek, A Heuser, S. Bhasin, and A. Hanjalic. "Make Some Noise. Unleashing the Power of Convolutional Neural Networks for Profiled Side-channel Analysis". In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* 3 (2019), pp. 148–179. DOI: 10.13154/tches.v2019.i3.148-179. Open access: link; #Cites 98, #14 in Computer Security and Cryptography, h5:42

[42] Lejla Batina, Dirmanto Jap, Shivam Bhasin, and Stjepan Picek. "CSI NN: Reverse engineering of neural network architectures through electromagnetic side channel". In: *Proceedings of the 28th USENIX Security Symposium*. USENIX Association, 2019, pp. 515–532. Open access: link; #Cites 83, #3 Computer Security & Cryptography, h5:80

[43] Chibuike Ugwuoke, Zekeriya Erkin, Marcel Reinders, and Reginald Lagendijk. "PREDICT: Efficient Private Disease Susceptibility Testing in Direct-to-Consumer Model". In: *CODASPY 2020 - Proceedings of the 10th ACM Conference on Data and Application Security and Privacy*. Ed. by Barbara Carminati and Murat Kantarcioglu. Association for Computing Machinery (ACM), 2020, pp. 329–340. DOI: 10.1145/3374664.3375729. Open access: link; #Cites –, not ranked, h5:28

D | Case Studies

D.1 Data Science unraveling Cancer

Cancer is a devastating disease manifested by cells acquiring an accumulation of various biological capabilities that make them grow indefinitely resulting in a tumor that eventually metastases throughout the body. Various genomic instabilities drive the acquisition of these hallmarks, such as genomic alteration, duplication, and epigenetic changes.

The TU Delft Bioinformatics group led by Marcel Reinders applies data science technologies to find causative variations that are hidden in these Big Data. The group has access to various data acquired from tumors, as well as cell lines and mouse models through, amongst others, its long-standing relationship with the Netherlands Cancer Institute (NKI), institutionalized by the part-time appointment of Lodewyk Wessels at the TU Delft. The well-defined NKI datasets are combined with public repositories, such as, for example, The Cancer Genome Atlas (TCGA), to differentiate patients into groups that have different outcomes or react differently on therapies to contribute to personalized treatment plans.

One of the achievements of the TU Delft Bioinformatics group is the development of a statistical tool (RUBIC, Figure D.1, [1]) that can detect frequent recurrence of copy number aberrations across tumor samples which act as a reliable hallmark of certain cancer driver genes. Our breakthrough was that we changed perspective to detecting recurrent copy number *breaks*, rather than recurrently amplified or deleted *regions*, which we combined with a statistic based on the Euler characteristic.

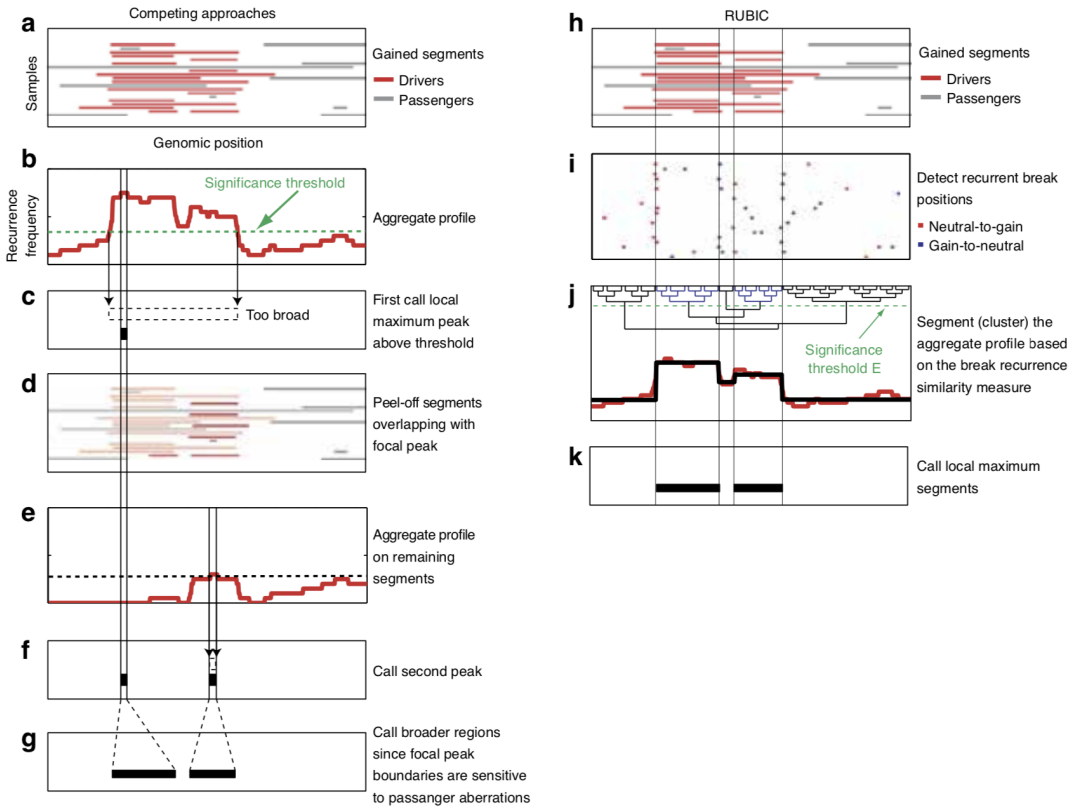

Figure D.1: The steps of RUBIC (Recurrent Unidirectional Break Identification by Clustering) for detection of diver genes (image from [1]).

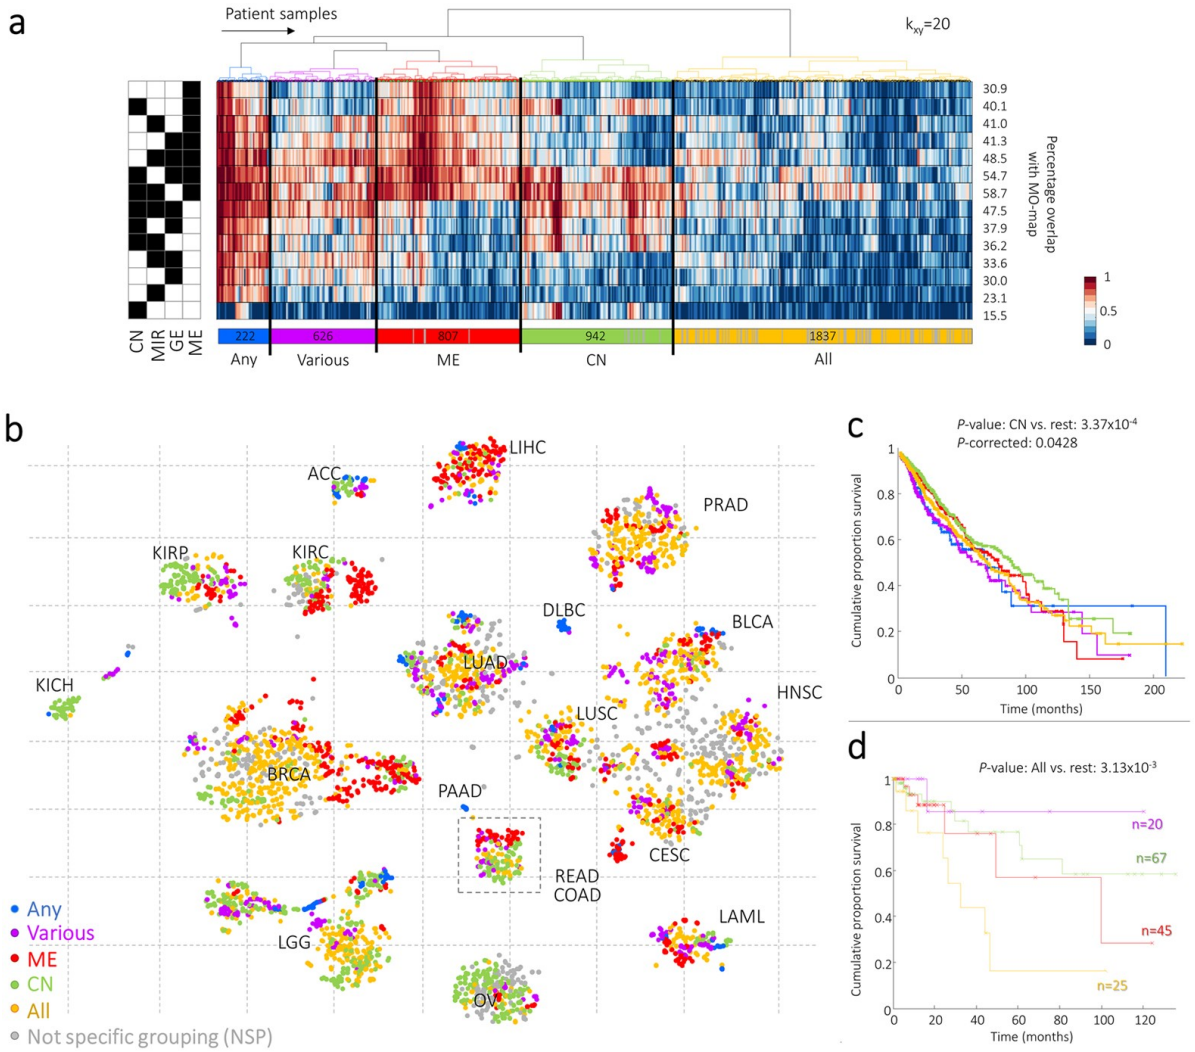

Figure D.2: Key MEREDITH molecular tumor analysis results (image from [2]).

Another innovation by the team of Reinders is the joint analysis of different molecular characteristics of a tumor for which they introduced a novel multi-omic integration approach (MEREDITH, Figure D.2, [2]) supporting visual exploration of the data as well as evaluation of the contribution of the different genome-wide data-types.

The team also proposed a novel computational framework for clinical drug response prediction. This framework employs non-linear manifold learning to capture biological processes active in both pre-clinical models as well as human tumors. Using this non-linear alignment predictors built on cell line response data only can be transferred to human tumors. This framework helps improving personalized therapies as well identification of biomarkers to, for example, resistance of drug therapies.

[1] Ewald van Dyk, Marlous Hoogstraat, Jelle ten Hoeve, Marcel J.T. Reinders, and Lodewijk F.A. Wessels. "RUBIC identifies driver genes by detecting recurrent DNA copy number breaks". In: *Nature Communications* 7.12159 (2016). Open access: link; #Cites 11, #7 overall, #3 in Life Sciences & Earth Sciences, h5:273

[2] Erdogan Taskesen, Sjoerd M.H. Huisman, Ahmed Mahfouz, Jesse H. Krijthe, Jeroen de Ridder, Anja van de Stolpe, Erik van den Akker, Wim Verheagh, and Marcel J.T. Reinders. "Pan-cancer subtyping in a 2D-map shows substructures that are driven by specific combinations of molecular characteristics". In: *Nature Scientific Reports* 6.24949 (2016). Open access: link; #Cites 27, #28 overall, #5 Life Sciences & Earth Sciences, h5:178

D.2 AI for Fintech Research

AI for Fintech Research (AFR)[a] is a collaboration between Dutch bank ING and Delft University of Technology. The mission of AFR is to perform world-class research at the intersection of Artificial Intelligence (AI), Data Analytics, and Software Analytics in the context of FinTech. Funded by ING, AFR was launched in January 2020, and will run for a period of five years. AFR is led by Arie van Deursen (TU Delft, scientific director), Luís Cruz (TU Delft, scientific manager), and Elvan Kula (ING, lab manager).

Advances in AI and data analytics are redefining the financial services sector. With 36 million customers, activities in 42 countries, and a total of 50,000 employees of which 15,000 work in Information Technology, software and data are at the heart of ING's business and operations. AI enables ING to predict customer wants and needs, while becoming a safer and more compliant bank. In this context, AFR seeks to develop new AI-driven theories, methods, and tools in large scale data and software analytics.

The core of AFR consists of 10 research tracks, in which 10 PhD students work topics such as software analytics, data integration, fairness in machine learning, model life cycle management, regulatory compliance, AI-based software engineering, and A/B testing. In each track, both TU Delft faculty members and ING engineers are involved, as well as bachelor, master, and PhD students.

As an example, in the *software analytics track*, we analyze software development data collected from hundreds of teams at ING to both understand and improve the efficiency of development processes. We combine static code analysis data with survey data collected from developers at ING to identify factors and their relationships affecting effort estimation and on-time delivery. Based on these factors, we build models to predict whether software deliveries will suffer delay [2], and formulate guidelines on streamlining software project management. This track is a good illustration of the interactive and participative nature of the collaboration between the students in AFR and developers at ING.

In the *data integration* track, we seek to bring together disparate data sources, which is needed in order to be able to employ machine learning methods to data lakes at the scale as in use at ING. To that end, one of AFR's first results is `Valentine`[b] [1]. Valentine offers a range of open source implementations of state of the art data matching algorithms, as well as synthesized data sets with ground truths that we used to experimentally establish the strengths and weaknesses of these algorithms – Figure D.3 shows screen shots of data fabrication, experimental configuration, and result presentation. Valentine illustrates the essential role open science plays in our work (also with industry), as well as the prominent role students play in AFR (besides PhD student Siachamis two master students, Ionescu and Psarakis, particpated in this research).

AFR is a member lab of the Innovation Center for Artificial Intelligence (ICAI[c]). Launched in 2018, ICAI hosts over 25 labs, each with five or more PhD students, and each a collaboration between one or more Dutch knowledge institutes and industrial or societal partners. AFR benefits from and contributes to this highly successful eco-system of industrial AI labs.

[1] Christos Koutras, George Siachamis, Andra Ionescu, Kyriakos Psarakis, Jerry Brons, Marios Fragkoulis, Christoph Lofi, Angela Bonifati, and Asterios Katsifodimos. "Valentine: Evaluating Matching Techniques for Dataset Discovery". In: *37th IEEE International Conference on Data Engineering, ICDE 2021*. IEEE, 2021, pp. 468–479. DOI: 10.1109/ICDE51399.2021.00047. Open access: link; #Cites 2, #6 in Databases & Information Systems, h5:56

[2] Elvan Kula, Eric Greuter, Arie van Deursen, and Georgios Gousios. "Factors Affecting On-Time Delivery in Large-Scale Agile Software Development". In: *IEEE Transactions on Software Engineering* (2021). Open access: link; #Cites –, #6 in Software Systems, h5:52

a http://se.ewi.tudelft.nl/ai4fintech/

b https://delftdata.github.io/valentine/

c https://icai.ai

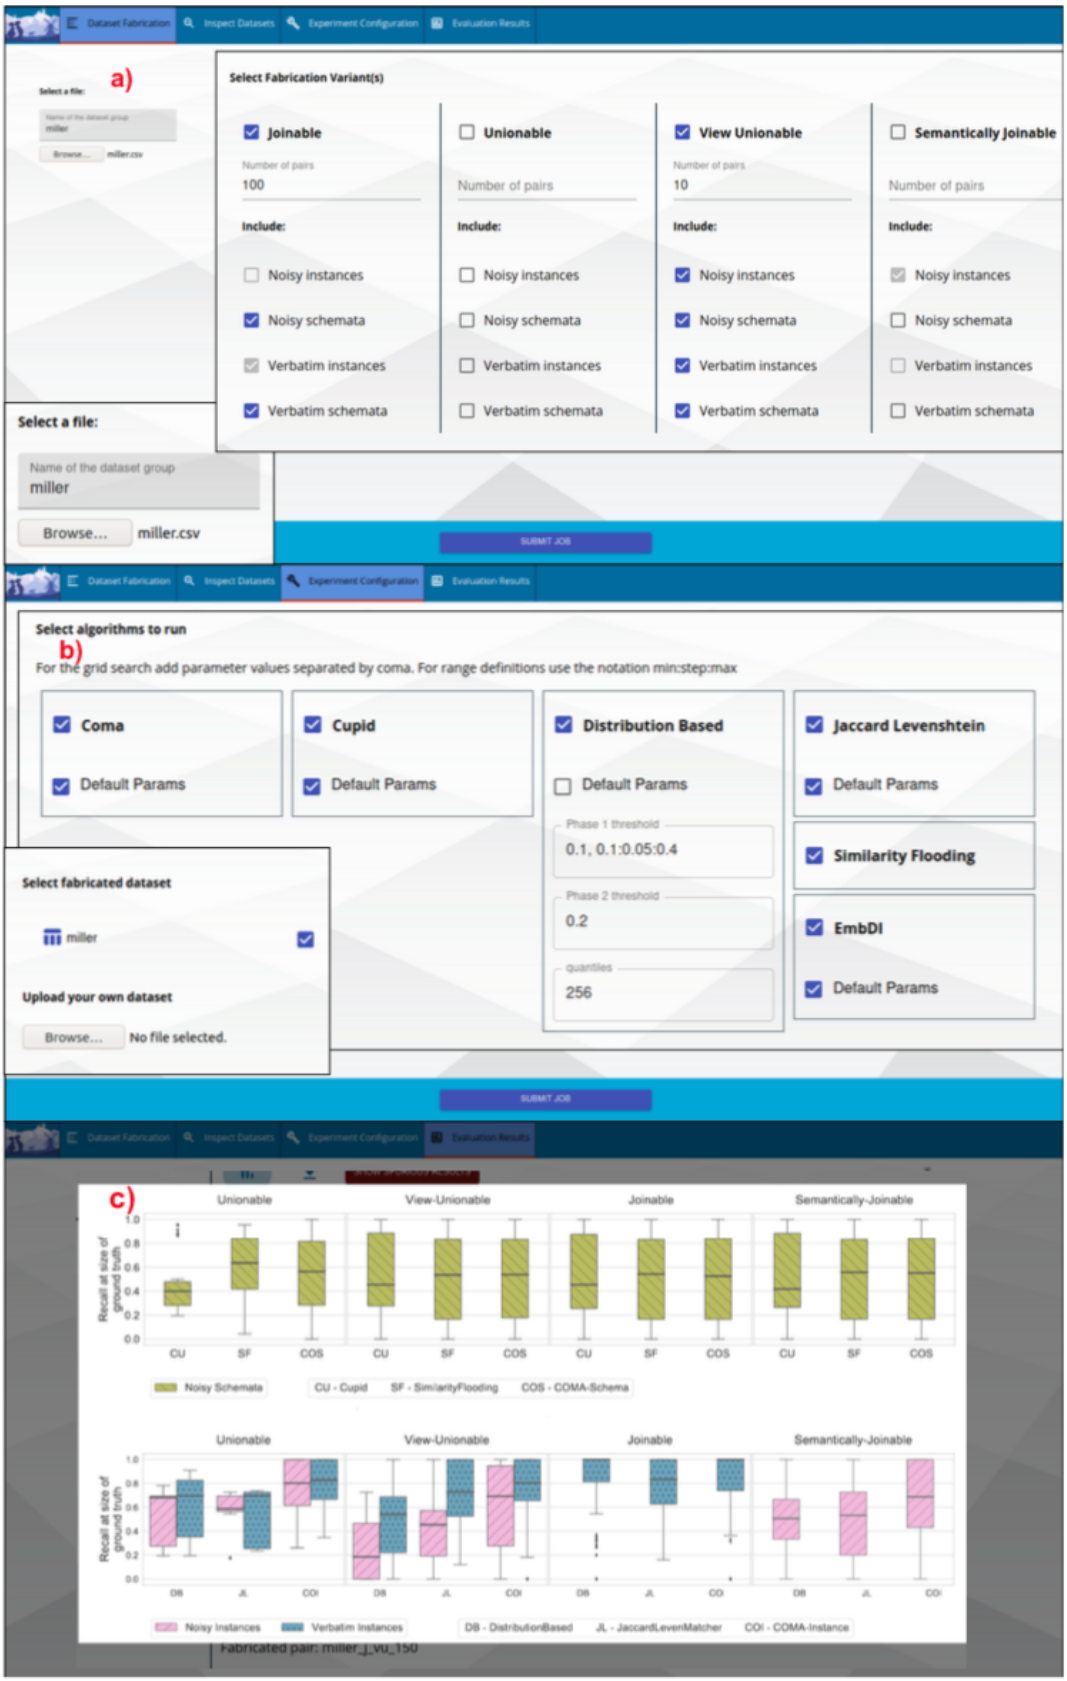

Figure D.3: Using Valentine for conducting tabular data matching experiments: a) data fabrication; b) configuration of experiments; c) presentation of findings.

D.3 Responsible AI in Psychological Assessment

With digitization increasingly being integrated in society, the technologies researched by CS@Delft are increasingly transferred and applied at scale to real-world use cases. While this is happening, questions and concerns are arising on whether this transfer is happening properly and responsibly.

With multimedia data often representing humans or human information, and multimedia search engines and recommender systems affecting automated decision-making based on such information, these developments are of strong interest to the Multimedia Computing (MMC) Group, especially in Cynthia Liem's team. In its academic activities, the team has been striving to explicitly include interdisciplinary perspectives. With Liem's background in music information retrieval, this interdisciplinary expertise already had been built in connection to the humanities, leading to research lines on inclusive information filtering and validity of measurement procedures. During 2015-2020, the latter research line expanded and solidified in connection to the social sciences, regarding computational approaches to psychological assessment.

Because of the MMC group's earlier track record on audiovisual and affective multimedia content analysis, Liem was approached by colleagues in organizational psychology at the Erasmus School of Social and Behavioural Sciences, who were researching the impact of AI on hiring procedures, and fairness considerations with regard to video resumes. Joint interests led to TU Delft becoming the technical partner in the ERASMUS+ 'Big Data in Psychological Assessment' (BDPA) project. With a prime focus on education innovation for psychology curricula, the BDPA project's outcomes needed to be accessible to psychologists who would be confronted with big data, but would not have extensive computer science training. Joint pilot collaborations on this topic included participation in a benchmarking challenge, presented at a prestigious computer vision conference (ChaLearn Looking at People Coopetition, CVPR 2017[a]).

However, the most significant outcome was the realization that psychologists and data scientists consider similar research pipelines, but put their main focus on different parts of these pipelines, also seeking to draw different types of conclusions from them (Figure D.4). This insight, with examples of what realistically can and cannot be researched in data-driven setups, has been documented in a book chapter [2], which is increasingly being integrated in psychology curricula and referenced in courses for HR professionals.

The BDPA project led to both Liem and MMC PhD candidate Andrew Demetriou (a social psychologist by training) gaining considerable experience in educating non-technical audiences about AI fundamentals and applications, and technical audiences about methodological considerations from other disciplines. This expertise has broadly been sought out, e.g. leading to invited talks at the International Convention of Psychological Science, the University of Bucharest and the Dutch-Flemish Network for Selection Research, and participation in a video series[b] by Deloitte. It also led to the European Personnel Selection Office (EPSO) remaining in touch with TU Delft as a trusted collaboration partner, paving the way for further research into responsible (multimedia) AI integration in public-interest applications.

[1] Cornelius J. König, Andrew M. Demetriou, Philipp Glock, Annemarie M. F. Hiemstra, Dragos Iliescu, Camelia Ionescu, Markus Langer, Cynthia C. S. Liem, Anja Linnenbürger, Rudolf Siegel, and Ilias Vartholomaios. "Some Advice for Psychologists Who Want to Work With Computer Scientists on Big Data". In: *Personnel Assessment and Decisions* 6 (1 2020). DOI: https://doi.org/10.25035/pad.2020.01.002. Open access: link; #Cites 3

[2] Cynthia C. S. Liem, Markus Langer, Andrew Demetriou, Annemarie M.F. Hiemstra, Sukma Achmadnoer Sukma Wicaksana, Marise Ph. Born, and Cornelis J. König. "Psychology Meets Machine Learning: Interdisciplinary Perspectives on Algorithmic Job Candidate Screening". In: *Explainable and Interpretable Models in Computer Vision and Machine Learning*. The Springer Series on Challenges in Machine Learning. Springer, 2018, pp. 197–253. ISBN: 978-3-319-98130-7. DOI: 10.1007/978-3-319-98131-4_9. Open access: link; #Cites 21

a https://chalearnlap.cvc.uab.cat/dataset/24/description/

b https://www.youtube.com/watch?v=DUv84dG31TM

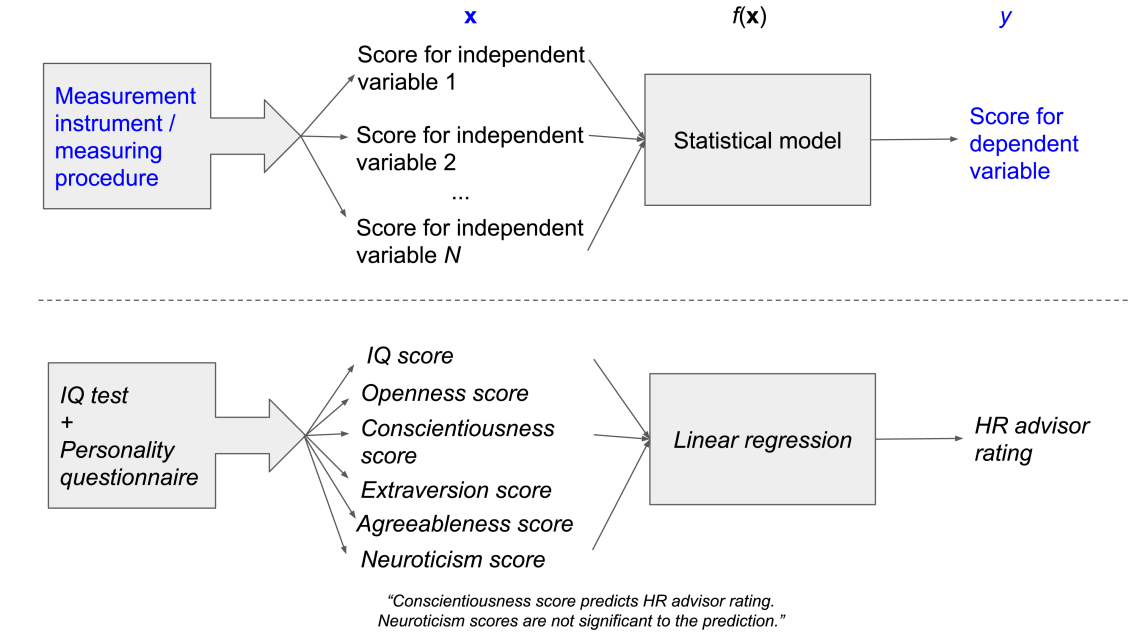

(a) Psychology (in an organizational psychology application).

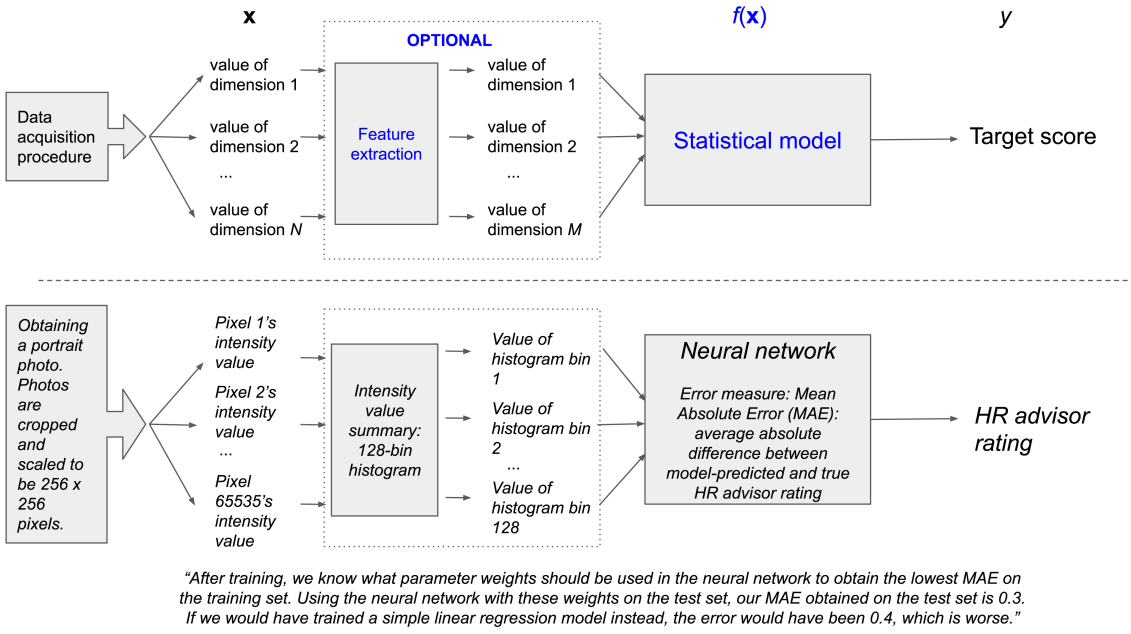

(b) Machine learning (in a computer vision application).

Figure D.4: Prediction pipelines in psychology and machine learning. Abstracted pipelines are given on top, with typical domain-specific research focus areas marked in blue. Simplified examples of how they may be implemented are illustrated at the bottom, together with a typical conclusion as would be drawn in the domain.

D.4 Algorithms for Coordination in a Smart Power System

Caused by the shift to renewable generation, electricity demand needs to be matched to the generation, instead of generating power when it is demanded. For this, the energy system needs to be more smart, scheduling flexible demand depending on weather predictions, taking network capacity into account. This results in algorithmic challenges regarding the autonomous control of components of a large-scale system; components that are adaptable and self-learning. Moreover, there are also significant challenges in supporting strategic decisions regarding such a system under extreme uncertainty, and where different interests play a role. Within the Algorithmics group we contribute to these algorithmic challenges.

A concrete, pioneering example is our work on constrained Markov Decision Processes [2]: De Nijs, now assistant lecturer at Monash University, developed a number of agent-based algorithms that can, for example, be used to control heat pumps such that the capacity of the electricity network is not violated and comfort is optimized. His PhD thesis work, funded by the distribution network company Alliander, concentrates on the fundamental problem of coordinating the use of resources over time with such a capacity limit, while dealing with uncertainty regarding the needed resources (caused by heat loss) as well as the available capacity (caused by other use of the network). The better understanding of this problem has led to a complexity analysis [3], and a number of new algorithms.[a] The most efficient method scales well with both the number of households as well as with the decision horizon (number of hours, see Figure D.5), making it possible to coordinate hundreds of heat pumps to stay within the capacity limit of the electricity network, where a standard optimization approach cannot deal with more than ten (see Figure D.6 for a simulation result of this adaptive planning method against others). Results are not just published as papers, but also include code and benchmark data, such as also done for our work on methods which minimize electricity costs, by scheduling demand using price information from the electricity markets.[b]

Apart from direct funding from industry, funding for this research typically comes from NWO thematic programs, or via RVO from the ministry of Economic Affairs and Climate. At the time of writing there are five PhD students working on this theme in the Algorithmics group. Within the university this research is connected to other energy-related research, among others through the PowerWeb insitute. Nationally, this type of research on the use of Artificial Intelligence for challenges regarding Energy & Sustainability is connected by a working group of the Dutch AI coalition (NL AIC), which is the organization behind the AiNed proposal, funded by the "Groeifonds". This coalition is working on defining new programs for AI research and innovation to strengthen the Dutch economy on the long term.

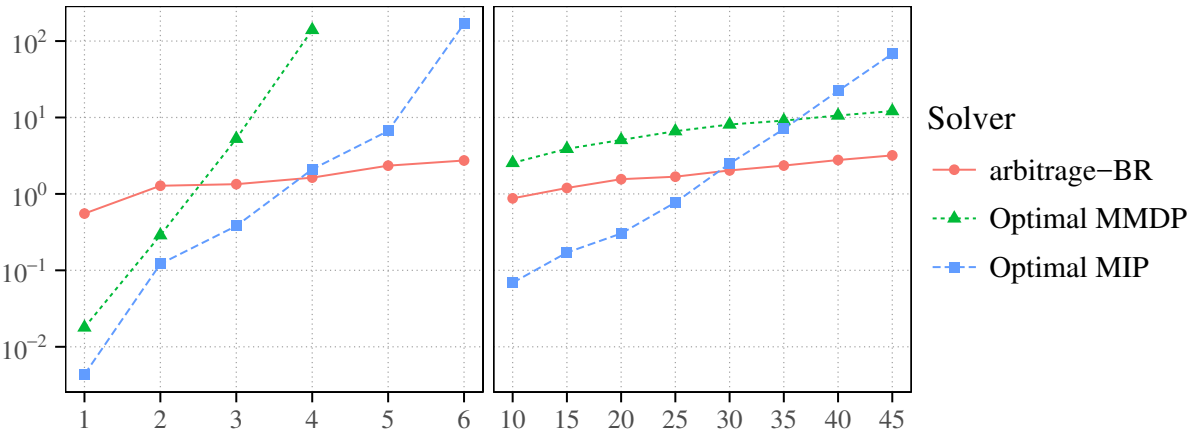

Figure D.5: The runtime of known optimal methods does not scale with the number of agents, while the arbitrage-BR method scales reasonably well in both agents and horizon.

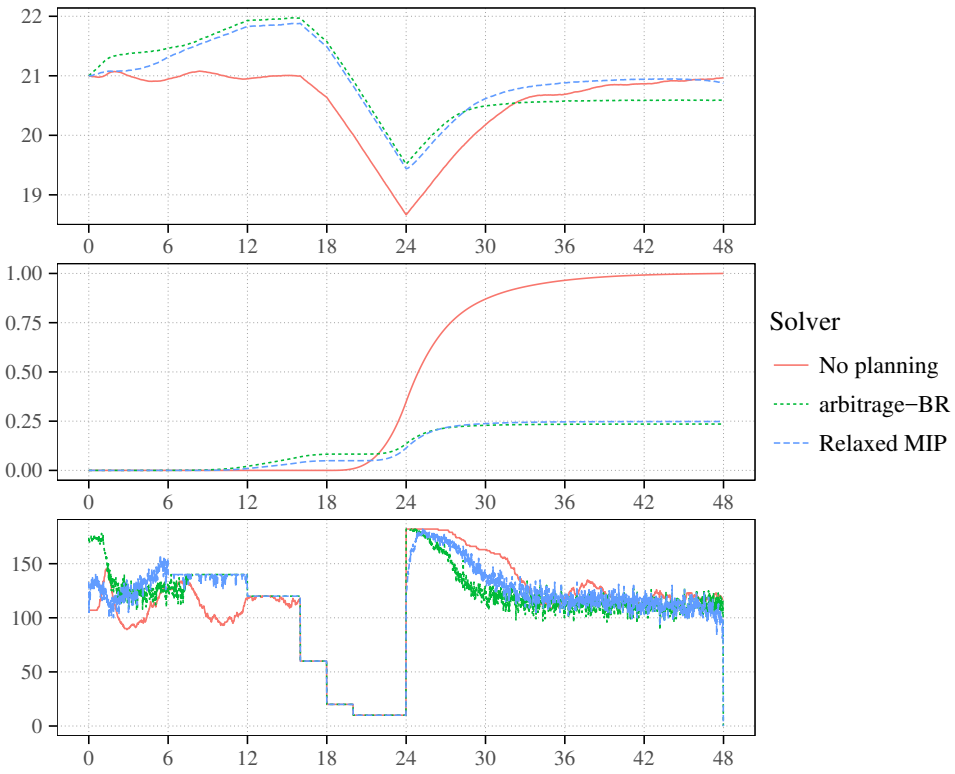

Figure D.6: In this simulation of 182 households with heat pumps the target temperature is 21 degrees for all houses, but a gradual decrease of the available network capacity occurs starting from hour 6, and at hour 20 the minimum capacity is reached and only 10 out of 182 heat pumps are allowed to be switched on. These plots show the average indoor temperature, normalized cumulative penalty, and total load of two different planning methods, compared to not planning at all.

[1] Germán Morales-España, Álvaro Lorca, and Mathijs de Weerdt. "Robust unit commitment with dispatchable wind power". In: *Electric Power Systems Research* 155 (2018), pp. 58–66. Open access: link; #Cites 44, #3 in Power Systems, h5:59

[2] Frits de Nijs, Erwin Walraven, Mathijs de Weerdt, and Matthijs Spaan. "Constrained Multiagent Markov Decision Processes: a Taxonomy of Problems and Algorithms". In: *Journal of artificial intelligence research* 70 (Mar. 2021), pp. 955–1001. Open access: link; #Cites 0, #11 in Artificial Intelligence, h5:64

[3] Mathijs de Weerdt, Michael Albert, Vincent Conitzer, and Koos van der Linden. "Complexity of scheduling charging in the smart grid". In: *Twenty-Seventh International Joint Conference on Artificial Intelligence*. 2018, pp. 4736–4742. Open access: link; #Cites 16, #3 in Artificial Intelligence, h5:95

a https://github.com/AlgTUDelft/ConstrainedPlanningToolbox
b https://github.com/AlgTUDelft/B-FELSA

E | Organization of PhD Training

The TU Delft University Graduate School (UGS) and its local branch, the EEMCS Faculty Graduate School (FGS), of which Dick Epema (DS/ST) is the director, provide a structured Doctoral Program with a PhD Development Cycle, with a clear assessment timeline and a course-based Doctoral Education Program. This appendix presents the process-related organization of the PhD training in CS@Delft supporting the strategic aim of *Getting the most of PhD training* (Section 3.3.5).

PhD Development Cycle In the Netherlands, PhD projects have a standard duration of four years, whether it concerns PhD candidates with a contract (∼70% in the CS departments) or with a scholarship (∼20%). The supervisors are responsible for their recruitment and selection. PhD candidates are registered with the UGS, and are, in cooperation with their supervisors, required to obey the following timeline:

- Intake by the UGS and FGS; following the PhD start-up course; composing their PhD agreement, including a provisional doctoral education programme;
- Initial evaluation meeting after 6 months, and a go/no-go review after 1 year. For this review, the PhD candidate prepares a portfolio with evidence of the research performed (a technical report or a (draft) paper), a plan for the remaining three years, and a self-assessment. The decision to continue or terminate the PhD project is taken by the intended promotor advised by a committee including two independent staff members.
- Formal progress reviews at the end of the 2nd and 3rd year; following the career orientation course in the 3rd year.

At the end of the 4th year, the thesis has to be submitted for a plagiarism check and has to be approved first by the supervisors and then by a Doctoral Committee with at least four independent members and at least one external member.

All information related to the progress of PhD candidates is recorded in the Doctoral Monitoring Application (DMA), which is accessible by PhD candidates and supervisors, but also by the dean, the HR department, the FGS, and the department chairmen.

Supervision and Support Each PhD candidate has at least two supervisors, a promotor (a staff member who has "ius promovendi", that is, the right to graduate PhD candidates) and a second promotor or co-promotor (daily supervisor who does not have this right) to ensure sufficient availability and robustness of supervision and diversity of opinion. Traditionally, in the Netherlands only full professors had this right, but after a change in the national law on higher education, since 2017, at TU Delft associate professors who have successfully acted as co-promotor of two PhD candidates also have this right. Currently, 13 associate professors in the CS departments have the right.

Since the start of the Graduate School (2011), the FGS maintained a PhD mentor system in which a staff member was assigned as a PhD mentor to each new PhD student to discuss any process-related aspects and issues of their PhD project. Around 2017, this system turned out to be ever more unsustainable, and as a consequence, the FGS recently switched to a system in which PhD candidates can get assigned by the FGS an *individual PhD mentor* on their request. In addition, three experienced full professors have the role of *general PhD mentor* and can be accessed in full confidence by PhD students at any time (this has been in place since 2011).

Finally, our PhD students are supported by career counseling and a PhD Council consisting of PhD candidates of all departments in the faculty that represents the interests of the PhD candidates. This council regularly exchanges ideas with the dean and the director of the Faculty Graduate School.

Doctoral Education Programme An essential part of the Doctoral Programme is the mandatory Doctoral Education Programme (DEP), which consists of three skill sets with a minimum requirement of 15 Graduate School points each (1 GS point equals 8 contact hours + 4 hours for preparation/assignment) and which supports the growth of PhD students as scientists:

1. Discipline-related skills: courses to deepen or broaden knowledge selected from national research schools, national of international summer schools, and certified online education;
2. Research-related skills: practical research-related skills such as scientific writing, presenting, and reviewing, and teaching and supervision. PhD candidates typically spend 10% of their time on teaching.
3. Transferable skills: standardised competency courses organised by the UGS.

The supervisors help the PhD students to pick the most suitable program and the appropriate time planning and are responsible for the level and fit of the courses taken. The director of the FGS is mandated by the Board for Doctorates of the university to approve individual doctoral education programs.

The ASCI Research School The PhD students from the INSY and ST departments are enrolled in several national research schools[1], which emerged through collaboration of virtually all Dutch universities that offer CS PhD programs. The schools have different thematic focus, but a joint mission to help train PhD candidates in the context of their doctoral education in their own university's graduate schools.

The Faculty of EEMCS of TU Delft is the national coordinator of ASCI - the *Advanced School for Computing and Imaging*[2]. ASCI operates in the fields of advanced computing and imaging systems. Several sections in the CS departments participate in it, mainly Pattern Recognition and Bio-informatics (PRB/INSY), Computer Graphics and Visualization (CGV/INSY), and Distributed Systems (DS/ST). ASCI offers specialized PhD-level computer-science courses, which PhD candidates can follow to fulfill their obligations for discipline-related courses in the Doctoral Education Program. PhD candidates who complete their ASCI educational plan obtain the formal ASCI certificate. During the evaluation period, 26 PhD candidates of TU Delft obtained this certificate, and currently, 22 are enrolled in ASCI. Community building is an important goal of ASCI, and it sponsors two successful yearly conferences for its two fields, giving PhD candidates low-threshold opportunities to present their work. These are the *International Conference on Computing Systems* (CompSys, started in 2017, around 50 attendees yearly; fully organized by ASCI), and *The Netherlands Conference on Computer Vision* (NCCV), started in 2014 and with up to 150 attendees.

The ASCI research school will be assessed by a different assessment committee around the same time the Computer Science of TU Delft is assessed.

1 Computer Science Research Schools in the Netherlands

2 ASCI Research School

F | Development Track Plan

We include here a verbatim copy of the *Development Track Plan* as it is currently used in an INSY pilot, as well as the TU Delft *Leadership Profile* referred to in Development Track Plan.

About this document

Every individual academic career is unique, and can also in a unique way be valuable to realizing the strategic goals of TU Delft. As the organization, we have together the ambition to excel in many ways and along many dimensions. We can, however, only be successful if all these individual careers complement each other in a productive fashion, by placing each faculty member in the right role and at the right place in the organization.

TU Delft is a large organization and allows for a diversity of academic careers. Therefore, the meaning of 'right' depends largely on the faculty members themselves. While taking the strategic goals of the organization as the reference framework to orient within, it is important for everyone to identify and develop own task portfolio, satisfying one's ambition and bringing sufficient challenge, inspiration and fulfilment, while keeping the frustrations and stress away for the optimal productivity and well-being in the long term. This then also relates to the envisioned career path, bringing each faculty member to the right level of academic hierarchy (assistant, associate or full professor) at the right time.

Needless to say, the meaning of 'right' may also change over time, through the changes in the individual interests and ambition, possibly induced by the increasing knowledge, experience, skills and competences, but also the changing insights, over time. This may open the scope of new task portfolios, for instance by changing the balance between research and education, or by rapidly developing academic leadership, and point to new academic career paths.

This document is designed to help you reflect upon your interests and ambitions at TU Delft at the present time. It serves to give you and your supervisors a first reference framework that will allow us to jointly crystalize the main characteristics of your academic profile, shape up and respond to the expectations regarding your career development, and generate cues for your supervisors and the organization around you to optimally support you. In this sense, please see the action of filling in this document as nothing more than a first attempt to reflect upon your 'academic self'. This document – a living document - will then serve as the main reference for yearly discussions on your career development, subject to change if the circumstances and/or your interests and ambition change.

With this perspective in mind, please think carefully about the questions below, related to the **Education**, **Research** and **Academic leadership**. There are no right or wrong answers, but only sincere answers based on what you believe is best for you from your current perspective and in the organizational context you belong to. Since the preferences you express and the choices you make also affect the organization around you, please also think carefully about the last item About you and the collective which serves to help you place your preferences and choices in your local work context.

Education

You could see your educational activities as a compulsory collegial contribution to the organization's education program, but also as a part of your ambition motivating you to join academia and underlying your academic career planning in the first place. Explain how you see academic education as a part of your academic career.

If you wish to limit your activities in academic education to the shared responsibility, explain your choice by indicating briefly what other priorities you see for your academic career. If you see it as an important part of your academic ambition, possibly also as the main path for your academic career development, explain which reasons and ideals underlie your ambition to make a difference in how academic education at TU Delft and in general is organized and conducted.

Education as contribution:

Estimate and reflect on your target portfolio of education-related contribution: *Sketch your teaching portfolio in terms of the courses developed or given, students (BSc and MSc) supervised, organizational roles and committees taken part in across the bachelor and master, based on the initial discussions with your supervisor. Give a rationale behind this portfolio, reflecting on the interests and preferences for certain parts, like giving massive bachelor courses versus specialized master courses, delivering versus developing versus organizing and managing a course, focusing on knowledge transfer versus developing transferrable skills among the students, supervising master students versus student projects in the bachelor, etc. Furthermore, explain your preference for the type of organizational roles and committees to participate in, for instance, ad-hoc committees, like the one to brainstorm about changing the way we organize master thesis supervision or how we should strengthen thematic profiling of our curriculum, advisory committees, like Board of Education, or administrative committees, like Board of Examiners, or BSA committee. Also reflect on when you think is the best time for you to engage in a committee. Finally, reflect on how your teaching portfolio should be balanced with the portfolios of other colleagues so that we all together can realize the educational goals of our organization and which role you could play in this balancing process.*

Reflect on your performance: *Explain what you consider important to look at in order to evaluate your educational activities. When are you happy yourself with your educational activities, either after having delivered a course or having supervised a student?*

Education as ambition:

If it is your ambition to develop yourself more deeply and broadly in the domain of academic education and increase your impact there beyond your regular education tasks, state and give shape to your ambition below. If you do not have the ambition to further develop the education dimension of your academic profile, you can skip this part.

State and explain your ambition regarding academic education: *State your ambition and explain the underlying ideals. This is essentially about innovating and making a strategic difference in academic education at TU Delft or broader.*

How do you intend to realize your ambition?: *Realizing your ambition requires initiative and leadership in mobilizing the organization. Explain what initiatives are needed from your side and how you intend to show leadership.*

How to reflect on performance?: *Explain how you know you are successful in pursuing your ambition? What do you consider important to look at in order to evaluate your achievements? How should we measure the difference that you intend to make in the organization and in the impact of academic education?*

How can we help?: *Explain how you think TU Delft and your local organization (supervisor, section, department, faculty) can help you, in addition to the UTQ program, to develop skills and competences you need to realize your ambition regarding academic education? State here the skills and competences for which you think that need special attention in the coming years and the mechanisms (e.g. training, coaching/mentoring via group or one-on-one, peer intervision) you think would be most effective for this purpose.*

Research

State and explain your ambition regarding academic research: *State your ambition and explain the underlying ideals. This is essentially about what makes you a scientist, and then in particular in academia, as opposed to an industrial research lab.*

What is the nature of your research?: *In view of your ambition, indicate how do you see your research on the scale from fundamental to applied. Explain the underlying motivation. Your scientific ambition should point naturally to a particular choice. However, please note that your choice also has*

consequences regarding other choices below, like the modus operandi and the type of projects you prefer. It may be therefore productive to address these questions integrally.

What is the best 'modus operandi' for you as a scientist?: *To realize your scientific ambition, you could prefer to focus on (a) generating own ground-breaking contributions, primarily via first- or single-author publications, or (b) establishing yourself by contributing with your key expertise to various collaborative efforts resulting in joint high-impact results, or (c) acting as a senior author, coaching and guiding teams of junior scientists, or on different combinations of these. Each option and combination of options has a particular purpose and effect on the development of a scientific career. Indicate and explain your preference/strategy in this respect and the role each option will play in conducting your own research and in achieving your overall scientific ambition. As already stated above, the nature of research you prefer points naturally to one or another modus operandi, but also to other choices below. It is therefore advised to come back to this question later on to check whether the choices are consistent with the ambition and with each other.*

Research volume: *What is the envisioned size of your team of PhD students and/or postdocs that would help you realize your research ambition? This would be a group of junior/senior researchers around you, under your direct supervision, which should help you generate the research momentum you need, but also which you should maintain over a longer period of time by acquiring sufficient project funds.*

Publication strategy: *Solid publication strategy is critical for achieving impact. Reflect on your envisioned publication strategy and give examples of target venues for the coming years. As a guideline, some research fits strategically better in conferences than in journals, or vice versa. One can consistently go for top-ranked venues in the field, or some venues ranked lower, but better addressing the target community where impact is to be achieved. One can choose to stick to a limited number of venues, or be open to a broader scope, which again depends on the breadth of one's research scope. Finally, one can select venues that strengthen one's scientific profile and visibility in a particular field.*

Research funding strategy: *Acquiring research funding is a critical aspect to foster own career development and achieve impact. In line with the choices and reflection given above, indicate which research funding strategy best matches your ambition. Think of aspects like (a) fundamental vs applied projects, (b) individual grants vs collaborative projects, (c) role of prestigious personal grants (NWO, ERC) in your project portfolio, (d) expected categories of industrial and societal parties as project partners, and (e) expected project acquisition effort with respect to how large your research group should be that needs to be continuously funded (see under Research volume above).*

Key mid- and long-term results: *Explain here what you would like to have achieved in your field 5-8 years from now and in a longer term. What difference would you like to make with respect to the state-of-the-art, or in the way how the technology you contribute to is deployed?*

How does your research link to the collective?: *While developing your own research line, it would be productive to also create synergies with research lines of your colleagues, in order to strengthen the overall scientific visibility and impact of your section, department, faculty and the TU Delft as a whole. Explain how you see such synergies, how your research can contribute to some, and what your role in the process could be.*

Scientific and socio-economic impact: *Explain where you see the biggest impact of your research. Is it in improving science or in helping the society? How do you think the impact of your research should be evaluated? What do you consider important to look at for this purpose?*

How can we help?: *Explain how you think TU Delft and your local organization (supervisor, section, department, faculty) can help you in your development towards becoming a scientific authority in your field and in achieving the impact you wish to achieve? State here the skills and competences, which you think need special attention in the coming years for e.g. developing your research vision, acquiring research funds, attracting junior researchers to your group and inspiring them to join you in pursuing your ambition, communicating about your research, broadening your scientific network, embedding your research line in your organization, etc. What are the instruments (e.g. courses, peer intervision, one-on-one coaching, mentoring) you consider most effective in this respect?*

Academic leadership

Many enter academia in order to develop themselves towards becoming scientific authorities in their field, established through own significant scientific contributions. Such authorities are the pillars of a university and of the international scientific community. Some, however, also wish to develop themselves beyond that, towards becoming academic leaders in a broader sense. Academic leaders are the main catalysts of the change scientific ideas and innovations should induce in the society. However, academic leaders also inspire younger generations of scientists, signal the need for changes in the scientific community itself, and induce such change, by moving the community to reconsider research foci, or by initiating academic debate to redirect the scientific discourse. Finally, academic leaders actively engage in the discussions about the processes in own academic organizations, show initiative towards improving them and make the organization as a whole stronger. To achieve this, academic leaders actively seek opportunities to build up networks and engage into discussions in the organizational, scientific and non- scientific contexts on underlying principles of scientific and/or academic management practice and the role and impact of science in the society. Therefore, academic leadership reaches beyond scientific leadership alone and requires specific motivation, ambition, skills and competences, as well as sufficient room in one's task portfolio.

State and explain your ambition regarding academic leadership: *Explain whether you see yourself grow into an academic leader and give, from the perspective of your own scientific domain and the developments you see there, your vision on academic leadership, for instance whether you see yourself influencing primarily the scientific community, academic organization or a non-academic socio-economic context. Also explain what ideals drive you in your ambition towards academic leadership and in what you aim at achieving.*

Impact: *Elaborate on the type of impact you would like to achieve as an academic leader, what you think you should do in order to achieve impact and how you think the impact of your academic leadership should be evaluated.*

How can we help?: *What skills and competences do you need to become an academic leader, and how can TU Delft help you develop them? To answer this question, you can consult the TU Delft Leadership Profile where there are six leadership components listed: ownership & collaboration, vision & strategy, self-awareness & situational awareness, trust & integrity, attracting & encouraging talent, initiative & courage. These components allow you to take up different leadership roles according to your ambition. Check this document to reflect on own skills and competences, to identify needs for your further development and the instruments you consider most effective to help you.*

About you and the collective

The choices you made above lead to a particular task portfolio, including activities on education, research and academic leadership. You expressed specific ambition for each dimension and gave an indication about the impact to be achieved. This will require a particular distribution of your effort across the three dimensions, which will determine the role you will have in our organization. It might turn out that your ambition points, for instance, primarily towards research or academic leadership, which would mean that extra effort should be invested by others around you to help the organization educate all its students. Explain below how you see the consequences of your specific task portfolio for the colleagues in your section and department, and how you think the discussion about creating a productive collective should be organized, where both the strategic goals of the organization and the ambitions of all individual members of the collective are met. What role do you see for yourself in this discussion?

Insert your narrative here.

Leadership profile

4 leadership roles

1. PERSONAL LEADERSHIP

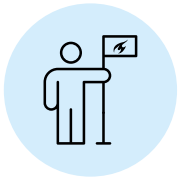

2. LEADING A PROJECT OR TEAM

3. LEADING A PROGRAMME OR MULTIPLE TEAMS

4. LEADING AN ORGANISATION

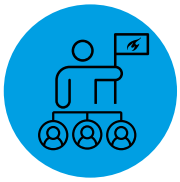

Personal leadership by all

As a staff member of TU Delft, in your own role and on the strength of your expertise, you contribute to science, education, innovation and service provision. As such, you display personal leadership in performing your role as best you can. You possibly also have a role as a leader – leading a team or a project, supervising or helping students or colleagues, or implementing change. This requires good leadership skills, built on professionalism, collaboration and openness. This is the way we build a stimulating and motivating work environment in which everyone can perform at their very best.

This makes life pleasant for you and for those around you. In addition, we need this to be able to contribute to TU Delft's vision: Making a contribution to solving global challenges by educating new generations of socially responsible engineers and by pushing the boundaries of the engineering sciences.
The way in which you can develop your leadership skills is something you discuss with your manager during your development meetings. Read here about the components that form part of the leadership profile, the various leadership roles, and how you can develop your personal leadership skills. See the back of this leaflet for the behaviour that TU Delft expects of you personally and of you as a leader.

The development of your personal leadership skills is built on our core values: **Diversity, Integrity, Respect, Engagement, Courage and Trust (DIRECT).**

The six components of the leadership profile

1 OWNERSHIP AND COLLABORATION

You take responsibility and collaborate with colleagues towards achieving good results.

2 VISION AND STRATEGY

You have vision and are able to get others on board, you set goals and implement them.

3 SELF-AWARENESS AND SITUATIONAL AWARENESS

You consider your actions and are open to the ideas of others. You are aware of developments in your field and in the organisation, and know how decisions are made.

4 TRUST AND INTEGRITY

You have integrity. You behave in a committed and transparent manner. You create an inclusive and diverse work environment which allows everyone to function optimally.

5 ATTRACTING AND ENCOURAGING TALENT

You focus on spotting, attracting and developing talent.

6 INITIATIVE AND COURAGE

You make decisions, show initiative and instigate change.

How can you develop your leadership skills?

The back of this leaflet lists a description of the behaviour required in each leadership role and for each component.

Self-reflection
Read through the leadership profile and consider where you do and do not (yet) exhibit this kind of behaviour.

Feedback
Engage in dialogue with other people, for example your supervisor, a colleague, someone who is receiving your guidance or supervision. Ask these people to provide feedback based on the leadership profile. You could also conduct a team discussion or ask for 360-degree feedback.

Development
Consider what aspects you would like to develop and what you would need for this.
- What experience would you like to gain? Examples might include participating in a project, committee, conference or think tank.
- What training would you like to take in order to develop your competences and increase your effectiveness? It is important for the training to be interactive and to focus on your own situation and learning objectives.
- Who would you like to learn from? Consider an exchange with colleagues and peers in your field and meetings with a mentor or on-the-job coach.

Discuss your leadership profile and your development needs with your manager. You could also ask your HR advisor for advice.

TU Delft

The six components of Leadership	4 Leadership roles			
	1. PERSONAL LEADERSHIP	2. LEADING A PROJECT OR TEAM	3. LEADING A PROGRAMME OR MULTIPLE TEAMS	4. LEADING AN ORGANISATION
				
	Focus on personal leadership in own role and position. Applies to every single staff member.	**Focus on leadership in relation to managing others, projects or changes.**	**Focus on personal leadership when directing groups, programmes, movements and partnerships, within and beyond TU Delft.**	**Focus on administrative leadership. Represents TU Delft in the national and international context.**
1  OWNERSHIP AND COLLABORATION Taking full responsibility, achieving results, working together.	• Where necessary, ensures others are aware of his/her own activities. • Assists colleagues when requested to do so. • Takes an active interest in the specialisation, motivations and interests of others he/she is collaborating with. • Shows respect towards others and their work. • Takes responsibility for the work and achieves results. • Shares successes and involves and acknowledges colleagues who have made a contribution.	• Exchanges information, knowledge and ideas with staff members/colleagues at his/her own initiative. • Identifies problems within the team and applies him/herself to solving these together with others and achieve results. • Is aware of, respects and understands mutual interests.	• Helps colleagues in determining objectives that are not necessarily in their own interest. • Initiates collaboration between different groups in order to achieve results together. • Creates a context in which mutual understanding and respect can be taken for granted.	• Focuses on actions and projects involving productive and constructive collaboration with other divisions or organisations. • Brings together the interests, actions and activities of different parties in order to achieve joint results and create a win-win situation.
2  VISION AND STRATEGY Having vision, implementing it, influencing others and having an impact by taking risks.	• Understands and highlights the broad outlines of his/her own group/team's objectives, within the wider context. • Is able to recognise other people's interests.	• Converts project/department strategy and goals into a course for his/her own research and education and/or own project team. Encourages others to contribute ideas to this. • Sets long-term objectives and develops scenarios to enable the team to achieve them. • Thinks several steps ahead and acts accordingly.	• Has a clear picture of the strategy of the department/faculty/service department, formulates appropriate objectives, works towards them and sets the chosen course in a way that inspires others. • Encourages others to do the same.	• Provides the organisation with direction by translating (societal) developments into new directions in the faculty. • Initiates large-scale movements within and outside the organisation and gets others on board. • Promotes synergy.
3  SELF-AWARENESS AND SITUATIONAL AWARENESS Reflecting on one's own actions and being open to ideas and comments from others, from an independent perspective. Awareness of the specialisation, organisation and the influencing and decision-making processes, and responding to these.	• Reflects on his/her own role and own personal development. • Is aware of his/her own behaviour and how it affects those around them. • Shows that he/she can critically assess and modify his/her own behaviour where necessary. Is open to feedback. • Is aware of his/her immediate surroundings, has insight into how the department is organised, an understanding of other people's interests and points of view and the ability to act accordingly. • Organises feedback on his/her own actions.	• Is aware of his/her own intrinsic motivation and the direct/indirect impact this has on their surroundings. • Shows exemplary behaviour, courage and resilience. • Is aware of decision-making processes and how they are influenced. • Is visible within the organisation to the departmental director, dean, director and colleagues from other (service) departments. • Knows and respects other disciplines and recognises	• Reflects on the development of others and his/her own role in their development. • Shows an understanding of group dynamics, including the impact of his/her own behaviour, and acts accordingly. • Encourages others to reflect on their own behaviour and be aware of their own motivation. • Is a discussion partner on strategic subjects and can generate leverage in this respect.	• Uses forces from within and outside the organisation to achieve strategic goals and create broad support. • Provides administrative support and backing for ideas developed and implemented in the faculty/service department. • Creates team spirit within TU Delft.
4  TRUST AND INTEGRITY Acting with integrity, expertise, commitment and transparency. Creating an inclusive and diverse working environment in which everyone can perform at their best.	• Speaks respectfully to others, also when disagreeing. • Provides accurate information to others, even if it is to his/her disadvantage. • Respects the values, standards, knowledge and skills of others. • Recognises and utilises what others have to offer. • Shares relevant information so that trust can grow. • Acts reliably (do what you say and say what you do).	• Safeguards the integrity, expertise, commitment and transparency of his/her own actions in research, teaching and in the group. • Is transparent about what is and is not part of his/her own area of expertise; encourages others to act in the same way. • Facilitates and creates a diverse group, and encourages everyone to perform to the best of their abilities. Provides a safe social environment. • Makes sure that agreements about duties and powers of authority are clear. • Gives feedback to people from the group and is willing to accept feedback; encourages and motivates others to do the same.	• Raises matters of integrity. • Shows others how to be committed and transparent, as well as being trustworthy and knowledgeable, and how to discuss these issues. • Gives staff members room to question authority and content, and shows that disputes are not settled by power and position. • Communicates the intended goals and results with stakeholders and encourages team members to achieve them together.	• Ensures an environment in which questions of integrity can be discussed. • Acts on violations of integrity. • Is a visible example of acting with integrity to people from both within and outside the university. • Widely promotes the virtues of a diverse organisation, stimulates an inclusive work environment and explicitly raises this subject at various levels.
5  ATTRACTING AND ENCOURAGING TALENT Focusing on spotting, attracting and developing talent.	• Creates a good atmosphere (space, trust and attention), in which others are encouraged to make use of and develop their talents. • Takes care of rapid induction of new colleagues and ensures they feel at home within the team. • Celebrates successes with others and ensures that everyone's contributions are recognised.	• Creates opportunities and encourages colleagues to take full advantage of them. • Puts others in the spotlight. • Focuses on identifying and attracting talent and creating diverse teams.	• Gives talented staff members the chance to grow to a level that may even exceed his/her own level within or outside the organisation. • Encourages colleagues to be visible in the organisation as well as beyond. • Is always actively scouting for talent.	• Always plays an active role in developing and attracting talent for the organisation and creates the right conditions and context for talent recruitment and development.
6  INITIATIVE AND COURAGE Taking responsibility, making decisions and instigating changes.	• Flags up issues and comes up with ideas and proposes strategies at his/her own initiative. • Initiates activities and takes the lead. • Takes responsibility and dares to take risks, understands his/her own impact on decisions.	• Recognises opportunities for improvement, innovates within his/her own specialisation, challenges the status quo. • Organises conferences and network meetings. • invites speakers and is impactful. • Dares to stand out, shows determination and gets things done. • Is visible and active within relevant networks.	• Is visible in the international academic or professional field, brings about changes in the department/faculty/service department. • Takes responsibility in order to increase his/her own impact in a broader context and enable others to do so.	• Explores and utilises new partnership opportunities and applications in teaching, research and valorisation for faculties (and similarly for service departments), both within and beyond TU Delft, beyond the boundaries of faculties and universities. Boosts these opportunities by applying focus and prioritisation. • Takes an active part in improving the university.

G | Joint Programs and Institutes

G.1 Delft Data Science

Delft Data Science[1] (DDS) started as a disciplinary collaboration between the computer science departments ST and INSY, with a focus on three roadmap programs based on TU Delft's unique strengths and opportunities, i.e. *engineering for data science* (Erik Meijer SE/ST, Arie van Deursen. SE/ST, Asterios Katsifodimos, WIS/ST), *social data science* (Alan Hanjalic MMC/INSY, Geert-Jan Houben, WIS/ST) and *health-related data science* (Elmar Eisemann CGV/INSY, Marcel Reinders, PRB/INSY). DDS's disciplinary basis covers the whole stack of engineering for data science, from computer engineering via software engineering to data management, and covers analytics with a focus on visual and social analytics. From this basis, DDS has increasingly engaged into domain-oriented collaborations, which involve researchers from other disciplines and faculties at TU Delft, including urban planning and mobility, online education, and health and sports. In order to maximize quality and impact in these domains, DDS collaborates with key TU Delft initiatives such as Amsterdam Institute for Advanced Metropolitan Solutions (AMS) and Leiden-Delft-Erasmus BOLDCities, TU Delft Extension School and Leiden-Delft-Erasmus Centre for Education & Learning, and Medical Delta and TU Delft Sports Engineering Institute. In this way, DDS has grown into a coordinating program and front-end for data science and big data at TU Delft. In this role, DDS provides an interface for researchers and industrial partners, to connect where focus and interests meet. For example, IBM set up with DDS the TU Delft-IBM-Collaborative Innovation Center on Data Science, to foster collaboration in research, talent exchange and education. In this context, Peter Hofstee (now at the QCE department) has been appointed as part-time professor in the faculty. Another example is the successful launch of several data science MOOCs, for example by Felienne Hermans (SE/ST), Erik Meijer (SE/ST) and Asterios Katsifodimos (WIS/ST). The role of success of DDS was the natural precursor to the development of the *TU Delft AI Initiative* (see Section G.2).

G.2 TU Delft AI Initiative

For the university to excel in implementing the impact of AI, Data & Digitalisation on its core activities, the TU Delft AI Initiative has been founded as a university-wide programme. It functions as a platform and front-end for AI, Data & Digitalisation (-related) research, education and innovation at TU Delft. All eight faculties contribute to the programme on the three core tasks of the university: education, research, and societal impact. It is the central point of contact, coordination, and representation related to AI, Data & Digitalisation activities. The TU Delft AI Initiative is led by the Pro Vice Rector Magnificus for AI, Data & Digitalisation. Geert-Jan Houben (WIS/ST) has in 2020 been appointed in this role, directly reporting to the president of the university. The TU Delft AI Initiative embodies TU Delft's ambition to strengthen its position as a world-renowned university of technology by developing an integrated programme for research, research-driven education and valorisation in and with AI, Data & Digitalisation. The TU Delft AI Initiative shapes this ambition around the university's core activities:

- Education — Employers and society demand engineers that are versed in AI, Data & Digitalisation. This requires a comprehensive approach to the continuing and accelerated introduction of AI, Data & Digitalisation elements into the educational offerings from the university, and thus into the curriculum of virtually all students at TU Delft.
- Research — In virtually every area of research, AI, Data & Digitalisation play an increasingly important or even crucial role. To advance the state of the art in these research areas, and to retain and strengthen TU Delft's leadership, it requires to join forces and combine the fundamental developments 'IN AI' with advancing areas in science, engineering and design by means of domain-specific applications 'WITH AI'. This has led to establishing of a series of Delft Artificial Intelligence (DAI) Labs (see Section G.3), in which CS@Delft plays a major role.

1 Delft Data Science

- Societal impact — Advances in and with AI, Data & Digitalisation occur virtually always 'in context', i.e., within the conditions and challenges from scientific fields or societal sectors. The TU Delft must position itself favourably for fundamental and contextual collaboration with industry and the very substantial future funding opportunities. This requires internal and external coordination.

Members from INSY and ST play pivotal and leading roles in the initiative, for example leading the TU Delft AI Education Taskforce, leading the research and valorisation themes around *AI for Energy and Sustainability* (Mathijs de Weerdt, ALG/ST), *AI for Port and Maritime* (Matthijs Spaan, ALG/ST), *AI for Peace, Justice and Security* (Inald Lagendijk, CYS/INSY), *AI for Health and Care* (Catholijn Jonker, II/INSY), *Machine Learning* (Frans Oliehoek, II/INSY) or *Human-Centered AI Systems* (Alessandro Bozzon, WIS/ST, Geert-Jan Houben, WIS/ST, Catholijn Jonker, II/INSY, Inald Lagendijk, CYS/INSY).

G.3 Delft Artificial Intelligence Labs

As a vehicle to strengthen interdisciplinary AI research at TU Delft, in 2020 so-called Delft Artificial Intelligence Labs (DAI-Labs) have been introduced. Directly funded by the university, these labs provide domain-specific contexts for researchers working on creating and improving AI concepts, theory and algorithms (research *in* AI) and those working in another domain making use of AI (research *with* AI) to join forces, benefit from the developments in AI to give boost to solving challenging problems in a given domain, but also let these challenges inspire further development of AI foundations.

Each DAI-Lab brings together two (tenure-track) assistant professors, one with an *in-AI* and the other with a *with-AI* research signature, and funds four PhD positions. In the two university-wide calls for proposals that took place so far, 16 DAI Labs have been formed, involving different all TU Delft faculties[2]. Among these labs, 6 have been granted to the teams involving CS@Delft faculty members, mainly based on their in-AI expertise. These labs are AIDROLAB (AI for sustainable water management, INSY and Civil Engineering), Design@Scale Lab (combining AI and human intelligence, ST and Industrial Design), DI Lab (humans and AI designing the future in dialogue, INSY and Industrial Design), MACHINA (machine intelligence advances for materials, INSY and Materials Science), HIPPO Lab (AI for fair, efficient and interpretable policy analysis, INSY and Technical Policy and Management) and AI*MAN Lab (transparent & traceable AI in human-AI teamwork, INSY and Aerospace Engineering). The ambition is to have these labs develop, grow (through additional funding to be acquired by the involved faculty members) and keep beyond the initial time period in which the funding for the four PhD students is provided.

G.4 Delft Blockchain Lab

Delft Blockchain Lab[3] (DBL) is TU Delft's initiative in blockchain technology and trust in the internet. DBL started in 2017 as an initiative of the Faculty EEMCS with funding from the Executive Board of the University. In view of the fact that technologically the blockchain technology is anchored in CS, DBL is led by CS@Delft faculty, specifically Dick Epema (DS/ST) as Director and Johan Pouwelse (DS/ST) as Scientific Leader. DBL brings together all blockchain-related research at TU Delft, and has its origin in many years of research in peer-to-peer, reputation, and accounting systems in the ST department. DBL has participants in the INSY and ST departments, and in the faculty of Applied Sciences (science communication) and the faculty of Technology, Policy and Management (ethics). The mission of DBL is to establish a lasting research effort in blockchain and related trust technology at TU Delft, and to perform education and valorisation in this area. Its research addresses the core technology for blockchain and trust in the internet, blockchain applications (targeting, e.g., energy grids and distribution), and societal issues (e.g., ethics, science communication).

DBL participates in the University Blockchain Research Initiative, a program in blockchain technology founded by Ripple. In its valorisation efforts, DBL collaborates closely with Dutch banks and government agencies (e.g., the Agency for Identity Data, RvIG). DBL has created an operational prototype of a new electronic identity infrastructure along with the RvIG, all three main Dutch system banks, Schiphol airport, regulators, and IDEMIA (the manufacturer of the Dutch passport). Through DBL, TU Delft participates in the Dutch Blockchain Coalition, which is an organization for the advancement of the uptake of blockchain technology in Dutch society and the economy with as members many companies and government ministries.

2 Delft Artificial Intelligence (DAI) Labs

3 Delft Blockchain Lab

G.5 LDE Center for Education and Learning

The LDE Center for Education and Learning[4] started up in 2014. In 2018 Marcus Specht (WIS/ST) was appointed as the new scientific director for the strategic alliance. In 2019 a new research program for focusing on Digital Education and Learning and the innovative use of ICT in education was introduced. The research connects to several research lines and focus areas in the departments of INSY and ST. User modelling and Learning Analytics, Intelligent Assessment and Conversational Agents, Hybrid Intelligence for Learning Support and AI for Education are key research topics of the Center linking activities in interdisciplinary projects with Erasmus University and Leiden University. Marcus Specht is also Professor at the Leiden Institute of Advanced Computer Science and the Erasmus School of Social and Behavioural Sciences.

G.6 The ELLIS Delft Unit

The European Lab for Learning and Intelligent Systems (ELLIS) is a European network of excellence focusing on machine learning and closely related disciplines. Our staff not only participate in international activities of ELLIS but have also co-founded the ELLIS Unit Delft. The Unit's goals are to organize international activities, such as Ph.D. student exchanges, summer schools, and workshops, as well as to provide a local platform for machine learning experts across TU Delft to meet and collaborate. Additionally, the ELLIS Unit Delft provides an internationally recognized label of excellence, which is instrumental in exploring new industrial collaborations, and showcasing our machine learning activities. Currently, the ELLIS Delft Unit includes staff from three faculties, with about a half appointed at CS@Delft: Frans Oliehoek (II/INSY), Claudia Hauff (WIS/ST), Marcel Reinders (PRB/INSY), Matthijs Spaan (ALG/ST) and Jan van Gemert (PRB/INSY). The Unit has been organizing monthly meetings and seminars, and started an honors program. The ELLIS Delft Unit has close ties with the other Dutch ELLIS Units, and therefore, it plays an important role in the national coordination of activities related to machine learning and in developing joint proposals.

G.7 CS@Delft and TU Delft Institutes

Scientific staff of CS@Delft has been active in a large number of TU Delft institutes, either as (co-)founders and board members, or as contributors. We mention here only the institutes in which we play a major role:

Robotics Institute Unifies all Delft University of Technology's research in the field of robotics. Its main challenge is to get robots and humans to work together effectively in unstructured environments, and real settings. Within the institute both the 'hard' robot disciplines (mechatronics, embedded systems, control and Artificial Intelligence) and the 'soft' robot sciences (human-machine interaction, user interaction, architecture, ethics, security and design) have a prominent presence. In total, 15 sections from 6 faculties participate in this institute. By joining forces and aligning research, education and valorization, the Robotics Institute takes a leading role in the creation of the next generation robots. The faculty of EEMCS is a strong supporter of this initiative with Catholijn Jonker (II/INSY) as one of the founders and together with Koen Langendoen (ENS/ST) and Marc Neerincx (II/INSY) the leader of the research theme on Intuitive Interaction.

Delft Design for Values Institute This institute focuses on the challenge of how to consider relevant moral and social values already in the design phase of new technologies, services, spaces, businesses, and institutions. There are several reasons for adopting design for values: (1) the avoidance of technology rejection due to a mismatch with the values of users or society, (2) the improvement of technologies/design by embodying these values, and (3) the generation or stimulation of values in users and society through design. The institute is the portal and booster of Design for Values research and education at Delft University of Technology. It aims at intensifying internal and external collaborations in the area and making visible what our university has to offer as a major partner in the debate on / implementation of design for values. The faculty of EEMCS is a strong supporter of this initiative with Catholijn Jonker (II/INSY) as one of the founders and with Myrthe Tielman (II/INSY) as a member of the management team.

4 LDE Center for Education and Learning

Delft Safety and Security Institute This institute brings together TU Delft's wide range of Safety and Security research, education and expertise. Safety relates to unintentional threats, for example caused by natural disasters, unintentional human behavior, technical failures and human error. Security relates to intentional threats caused by intentional human behavior, such as crime, fraud and terrorism. Cybersecurity is one of the domains that receives special attention within DSyS, embodying TU Delft's unique approach to cybersecurity, combining social and technical aspects and focusing not only on the prevention of incidents but especially on detection of, and recovery from, incidents. DSyS carries out research and develops methods and techniques aimed at quantifying safety and security issues and optimizing general and dedicated solutions. It also liaises to the 'real' world where theory and practice can be combined. This enables DSyS to provide high-quality research capacity both nationally and internationally. Topics covered by the institute are, amongst others, safe transport, robots, remote sensing, drones and aerial surveillance, shipping and aviation safety, storage and logistics, forensics, terrorist threat to critical infrastructure and the design of safe cities. Next to scientific contribution from our CYS/INSY section, CS@Delft is actively involved in the operations of this institute through Fernando Kuipers (ENS/ST) and Mauro Conti (CYS/INSY) in their roles as members of the Executive and Advisory Board, respectively.

PowerWeb Institute Founded by the faculties EEMCS 3mE, and TPM, it was set up to pursue challenges induced by the rapid growth of distributed renewable energy sources. Due to this, the structure of our electrical power systems needs to be changed radically and "smartly". This critical energy problem calls for integral solutions that cannot be solved within a single discipline. Therefore, in search of designs for a smart energy system, the institute aims at integrating contributions of several disciplines - from the underlying physical systems (hardware), to the design of smart energy management systems (software), and to the study of smart grids in relation to their societal and economic environment (peopleware) - in the specification and design of a robust and reconfigurable smart grid. Based on this, the institute aims at teaming up for and orchestrating TU Delft's project proposals on smart grids and at educating BSc, MSc students and PhD candidates, and act as a gateway for industry and government on the topic of smart grids. CS@Delft is in particular involved in the development of the smart energy management software, and has been represented in the Board by Pieter Hartel (CYS/INSY), Fernando Kuipers (ENS/ST), and Cees Witteveen (ALG/ST) in the past, and now by Mathijs de Weerdt (ALG/ST) as a member of the executive committee, and specifically as AI/energy liaison.

Delft Bioengineering Institute Founded in 2016 to unite biology-related research at TU Delft, as well as to improve the visibility of Delft in the domain of biological/life sciences. Recent reports, such as McKinsey's *The Bio Revolution*[5] and VNO-NCW's *Toekomstpact Biotechnologie Nederland 2025*[6] 2025 point out that current bioengineering breakthroughs could help us solve major human challenges, but collaboration is needed at all levels. Delft Bioengineering Institute facilitates this much needed collaboration at the TU Delft level. The main goal of the institute is to strengthen internal cohesion of TU Delft principal investigators who work on solutions in, with or for biology. Key stakeholders are early-career PIs, who currently make up about 70% of the collaborative network of close to a hundred PIs, representing six faculties and 22 departments. CS@Delft has gained significant exposure in this research through its research on bioinformatics, resulting in the appointment of Marcel Reinders (INSY/PRB) as institute's scientific director in 2019. An example of the output he created with the institute is the ICAI Artificial Intelligence Lab for Biosciences[7], which has been established with Royal DSM and Planet B.io in 2021. This lab will be the first of its kind in Europe to apply AI to full-scale biomanufacturing, from microbial strain development to process optimization and scheduling, and is mentioned in the VNO-NCW report as an example of strengthening sector specific infrastructure.

TU Delft AgTech Institute TU Delft AgTech Institute is founded in 2020 to boost research and innovation in the field of horticulture and agrifood. Two worldwide challenges feed this opportunity: 1) a lack of resources, e.g., available human labour, water, (chemical) nutrients and energy, which calls for optimization and automation of production (agri/horti) and value chain (retail) processes, and 2) the mega-cities developing worldwide need to be fed with enough, healthy and safe agrifood systems and products. This requires the Dutch agrifood industry (global leader in exports of agricultural products) to shift from producing and

5 The Bio Revolution, McKinsey, 2020

6 Toekomstpact Biotechnologie Nederland 2025

7 AI4B.io Lab

exporting our food products to providing knowledge and technology. The AgTech Institute facilitates this transition and collaboration in key technologies (Robotics, Sensor physics, Computer Vision, Communication Technology, Geothermal engineering, Material Science, Artificial Intelligence, Genetics, Bioinformatics and Crop Modelling) as solution to the global industry reform.

TU Delft AgTech Institute acts both internally and externally as a point of contact for expertise in the AgriFood domain. Key stakeholders are researchers of five faculties (AE, EEMCS, TBM, 3ME and CiTG) and 16 departments that support this institute and a large number of external private and public partners (Ahold, Rabobank, RijkZwaan, Priva, Royal Brinkman, VDL and WUR, 4TU, TKI, PZH, EBZ, IQ). They work together in research programs such as 4TU Plantenna, 4TU Pride & Prejudice, Flexcraft, Synergia) and a large number of research projects. At the moment the institute is working on the final touches of 2 National Growth Fund proposals (NXTGen HighTech, Plant RX) and the approved NWO LTP Plant RX).

With our prominent research on bioinformatics, CS@TUD was, through Marcel Reinders (PRB/INSY), in the lead in creating this institute, which also resulted in the appointment of Roeland van Ham (PRB/INSY) as its scientific director. Furthermore, CS@TUD is strongly involved in the activities of the institute through our research in computer vision and artificial intelligence.

G.8 CS@Delft and the 4TU Federation

The four universities of technology (TU Delft, Eindhoven University of Technology, University of Twente, Wageningen University) are jointly committed to educate a sufficient number of highly qualified engineers and technical designers, to conduct outstanding and socially relevant research of an international standard, and to promote cooperation between research institutes and businesses. The mission of the federation of the universities of technology - called 4TU[8] - is *to contribute to the well-being of the Netherlands by strengthening and pooling knowledge and creativity in the technology sector, and by exploiting it to maximum effect.* 4TU encompasses and partially funds seven joint education programs, the centre for engineering education, the 4TU data center, and nine joint (virtual) research centres. We list below those 4TU entities with a strong involvement of CS@Delft:

4TU NIRICT: Netherlands Institute of Research on ICT Aims to realise the mission of the 4TU federation in the field of information and communication technology. NIRICT was established in 2005 and brings together, inclusively, researchers of all disciplines in ICT research. Between 2014 and 2018, led by Inald Lagendijk (CYS/INSY) as its scientific director, NIRICT focused on building a strong 4TU ICT research community and had one large research program "Big Software on the Run (BSR)". The community building strategy was supported with 150K annually and consisted of (i) stimulation of bottom up community building initiatives through small (10-25K) seed funding, both regarding research and education, (ii) organisation of a yearly community day to discuss joint challenges and direction, and (iii) raising awareness for national ICT leadership and the importance of community building among the new generation academic staff through the ICT Next Generation network. The federation supported BSR with 500k per year, resulting in nine joint PhD positions and several postdocs. This innovative research program used novel techniques in process mining, visualization, security, and analytics for the purpose of taming complexity of software. It was led by Wil van der Aalst (TU/e). In September 2018, Mark van den Brand (TU/e) took over the role of scientific director. Alan Hanjalic (MMC/INSY) joined in 2019 as a member of the Board.

4TU Center for Engineering Education In order to boost engineering education, the Board of the 4TU.Federation has decided to work together in a Centre for Engineering Education. The 4TU.CEE aims to gather, develop and implement up-to-date expertise in engineering education. (New) improvements in education are implemented and the effectiveness of these improvements is monitored and analysed. The expertise and experiences of all four partners is exchanged to benefit all partners. In the new research program for 2022-2025 education innovation with ICT, Datafication and Computational Thinking are explicit topics that link to CS@Delft. Current projects supported include education in Quantum Computing, use of Conversational Agents for education and well-being, as also personalised learning and Learning Analytics.

Since 2020 Marcus Specht (WIS/ST) is leader for TU Delft in the 4TU.CEE. 4TU.CEE facilitates teaching staff to improve their teaching competencies by developing, exploring and demonstrating the newest

8 4TU Federation

educational methods. An explicit task of 4TU.CEE is to contribute to the international community of engineering education and to seek cooperation where appropriate. 4TU.CEE collaborates with educational communities such as SEFI, CDIO, and related organisations. 4TU.CEE organises (international) conferences and workshops on Engineering Education. These events provide a platform for TU teaching staff, sharing their experiences and becoming inspired by (international) colleagues.

4TU Humans and Technology Center realises the mission of 4TU in the field of human-technology interaction for smart social systems and spaces. The center was established in 2014 as a spin-out of NIRICT and brings together humanities and technical sciences. Under the leadership of Dirk Heylen (UT) as the scientific director, the center involves a large number of CS@Delft staff. The objectives are (i) to promote national collaborations in the area of human-technology interaction; (ii) to work on a joint vision embedded in both national and international research programs in consultation with external partners; and (iii) to create prosperous conditions for conducting excellent research in the area of human-technology interaction. In 2014, Humans and Technology received 4TU funding for executing its strategy, encompassing: (i) networking of researchers within and beyond the research centre and with industry and government; (ii) development of joint research agenda's; (iii) training and coaching (young) scientists; and (iv) building a research infrastructure to conduct interdisciplinary research. The center executes the research program "Smart Social Systems and Spaces" that aims to combine multidisciplinary knowledge to explore new collaborations and to come to innovative developments that facilitate and support human values in an increasingly sociotechnical world.

H | Labs and Computational Facilities

Critical for achieving our scientific goals is the availability of the suitable lab facilities and computing infrastructure. Both departments have invested in the past in building up and maintaining these resources. Regarding the lab facilities, the following can be listed:

- INSYGHT Lab (http://www.insyghtlab.tudelft.nl/) is the central lab facility of the INSY department. It hosts experimental installations for showcasing and supporting interdisciplinary research on computer vision, artificial intelligence, and visualization. For example, the lab provides support for research on and the development of robotics technology (RobotLab - 10 NAO robots and 2 Pepper robots), which allows us to strengthen our research and collaboration with other departments within the Delft Robotics Institute. The lab also provides support to conduct user experiments and research on virtual reality, and in particular for virtual reality therapy for mental health issues such as post-traumatic stress, and social phobia, but also skill training, such as negotiation. Finally, the lab contains a sound-proof booth for running speech and listening experiments.
- The Distributed ASCI Supercomputer (DAS, https://www.cs.vu.nl/das5/) is a nation-wide multicluster system purely for computer science research installed by the Research School ASCI with funding from NWO, one cluster of which is located at TU Delft. It is currently in its sixth generation, and is instrumental in much of the research in distributed systems.
- The ES Lab is the main facility for the ST department when it comes to developing custom hardware, which is instrumental for the ENS group in particular. Many PhD and MSc students use the tools (e.g., soldering station, oscilloscopes, 3D-printer) to assemble proprietary PCBs and enclosures to arrive at functional prototypes supporting the experimental evaluation of their designs. To support the research into Visible Light Communication, the ES Lab has been fitted with blackout curtains. To support research into low-level (physical layer) communication protocols (e.g., LoRa) the lab includes Software Defined Radio boards and power monitoring equipment. We plan to upgrade this equipment to ensure that we can also handle 5G (and beyond) networking chipsets.

Regarding the computational resources:

- The TU Delft High Performance Computing Cluster (DHPC) is a heterogeneous computing cluster that will be operational in the Fall 2021. The DHPC consists of 20.000 CPU cores in over 400 compute nodes and incorporates a high-speed parallel storage subsystem, based on a BeeGFS. All the compute nodes and the storage system is interconnected with HDR100 Infiniband technology for high-throughput low-latency inter-node communication. What makes the Delft High Performance Computing Centre special is its flexibility: the facility can be quickly adjusted in line with research and teaching requirements, and when resources are insufficient cloud bursting to SURF or AWS is possible.
- Apart from the large DHPC cluster, both departments maintain local clusters of a total of ca. 250 machines, with tens of GPUs supporting the everyday computational needs of researchers, as well as PhD and master students. More specifically, the INSY cluster works as a shared infrastructure supported by ICT and a research engineer as well as a governing user group. The INSY cluster can be accessed with either a specialized HPC batch system (slurm) or in the form of virtual machines that can be reserved by researchers. Following the need to more GPUs and data-intensive research, the CS@TUD already allocated budget to expand our clusters with more servers and most importantly GPUs. This expansion is planned with 200 CPUs and 30 GPUs per year, for the next five years.
- Finally, CS@TUD also has access to the SURF facilities as part of the national computing infrastructure coalition. Both students and faculty can reserve computational resources in SURF, in terms of CPU and GPUs.

I Computer Science Sections

INTELLIGENT SYSTEMS

Multimedia Computing (MMC)

Head: Prof.dr. Alan Hanjalic (also head of department)

Description: The Multimedia Computing (MMC) group conducts research on new methodological and algorithmic concepts for analyzing, interpreting, enriching, modeling, searching and recommending multimedia data available in stand-alone collections or connected in (complex) networked data structures.

Composition: 1.4 full professors (Hanjalic, Cesar, Larson), 3 associate professors (Liem, Scharenborg, Wang), 1 assistant professor (Urbano), 1 tenure track assistant professor (Isufi), 3 postdocs, 15 PhD students

Interactive Intelligence (II)

Head: Prof.dr. Catholijn Jonker

Description: The Interactive Intelligence (II) group focusses on agent-based reasoning to develop cognitive frameworks for various domains of applications, focusing on robots, human-agent/ robot teamwork, serious gaming, agent-based simulation and negotiation. Another domain of focus is cognitive engineering, aiming to develop methods for creating and applying user-centered designs, with an emphasis on adaptive interfaces, personalization, and behavioral change.

Composition: 1.4 full professors (Jonker, Neerincx), 2 associate professors (Brinkman, Oliehoek), 1 assistant professor (Murukannaiah), 4 tenure track assistant professors (Broz, Oertel, Siebert, Tielman), 4 postdocs, 23 PhD students

INTELLIGENT SYSTEMS (continued)

Pattern Recognition and Bioinformatics (PRB)

Head: Prof.dr.ir. Marcel Reinders

Description: Our core expertise is in pattern recognition; solving societal, industrial and scientific problems by a computer through learning from examples. We aspire to obtain a conceptual understanding of learning methodologies and to exploit them in analyzing complex data, primarily in the (exemplar) domains of computer vision and bioinformatics.

Composition: 1.6 full professors (Reinders, Lelieveldt, Van Ham, Wessels), 3 associate professors (Abeel, Van Gemert, Hung), 3 assistant professors (Goncalves, Loog, Tax), 5 tenure track assistant professors (Baaijens, Krijthe, Pintea, Tömen, Zhang), 5 postdocs, 34 PhD students, 1 data manager

Cybersecurity (CYS)

Head: Prof.dr. Georgios Smaragdakis

Description: Cybersecurity in Delft has a unique sociotechnical approach towards the protection of cyber space. We work on the following topics: Developing a Science of Cybersecurity; Verification and Validation; Secure Hardware; Network Security; Security and Privacy by Design.

Composition: 2.3 full professors (Smaragdakis, Lagendijk, Conti), 2 associate professors (Erkin, Verwer), 1 assistant professor (Picek), 2 tenure track assistant professors (Liang, Zarras), 4 postdocs, 21 PhD students

Computer Graphics and Visualization (CGV)

Head: Prof.dr. Elmar Eisemann

Description: The Computer Graphics and Visualization (CGV) section works on many topics in rendering, visualization of (scientific) information and modelling of 3D objects. Our goal is to develop new algorithms to generate, represent, interpret, display and interact with data. In particular, complex and large data sets are in our focus, as they play an increasingly important role in many scientific, medical and engineering applications.

Composition: 1.2 full professors (Eisemann, Vilanova), 0.7 associate professor (Bidarra), 2 assistant professors (Hildebrandt, Marroquim), 3 tenure track assistant professors (Höllt, Kellnhofer, Weinmann), 6 postdocs, 15 PhD students

SOFTWARE TECHNOLOGY

Software Engineering (SE)

Head: Prof.dr. Arie van Deursen (also head of department)

Description: The Software Engineering Research Group (SERG) addresses the design, construction, and operation of large, complex software systems. Topics of interest include software testing, software analytics and data-driven software engineering, human aspects of software development, software reliability and sustainability, and computational intelligence for software engineering.

Composition: 2.4 full professors (Van Deursen, Zaidman, Spinellis (0.2 fte), Van Solingen (0.2 fte)), 5 tenure track assistant professors (Aniche, Cruz, Kulahcioglu Ozkan, Panichella, Proksch), 4 postdocs, 16 PhD students

Algorithmics (ALG)

Head: Dr. Mathijs de Weerdt

Description: In the Algorithmics (ALG) group we aim to design, and understand fundamental properties of, planning and coordination algorithms for intelligent decision making in real world applications, such as for coordinating electrical loads within network constraints, or logistic processes on a shunting yard or container terminal.

Composition: 0.2 full professors (Bosman), 3 associate professors (De Weerdt, Spaan, Yorke-Smith), 4 tenure track assistant professors (Böhmer, Demirovic, Dumancic, Lukina), 2 postdocs, 21 PhD students

Distributed Systems (DS)

Head: Prof.dr.ir. Dick Epema

Description: The Distributed Systems Group (DS) studies the design, deployment and analysis of distributed systems and algorithms. Topics of interest include trust in the internet, blockchain technology, consensus, anonymity, efficient big-data computing and distributed machine learning.

Composition: 1 full professor (Epema), 2 associate professors (Chen, Pouwelse), 3 tenure track assistant professors (Decouchant, Rellermeyer, Roos), 6 postdocs, 11 PhD students, 5 developers

SOFTWARE TECHNOLOGY (continued)

Embedded and Networked Systems (ENS)

Head: Prof.dr. Koen Langendoen

Description: The Embedded and Networked Systems (ENS) group aims to improve the software development and maintenance process for embedded systems ranging from simple, 8-bit microcontrollers in consumer products to complex, networked controllers in industrial applications, and beyond.

Composition: 2 full professors (Langendoen, Kuipers), 2.7 associate professors (Pawelczak, Venkatesha Prasad (0.7 fte), Zuniga), 3 assistant professors (Iosifidis, Lan, Wang), 3 postdocs, 14 PhD students, 1 developer

Programming Languages (PL)

Head: Prof.dr. Eelco Visser

Description: The Programming Languages (PL) research group aims at improving the effectiveness and reliability of programming languages and systems. Effectiveness such that programmers can express intent at the right level of abstraction and get actionable feedback that is relevant and timely. Reliability such that programmers can trust the execution and analysis of programs. Research topics include language engineering, type systems, software verification, semantics specification, concurrency.

Composition: 1 full professor (Visser), 4 assistant professors (Bach Poulsen, Cockx, Chakraborty, Ahrens), 3 postdocs, 9 PhD students, 3 developers

Web Information Systems (WIS)

Head: Prof.dr.ir. Geert-Jan Houben

Description: The Web Information Systems (WIS) research group aims at making web information systems more effective in retrieving, processing and interpreting data generated by humans and machines. Topics of interest include user modeling, Web science, information retrieval, natural language processing, database systems, Web engineering, Web data management, user interaction, human computing, and human-AI interaction. In addition, via Centre for Education and Learning (CEL), the group looks at digital education innovation.

Composition: 2 full professors (Houben, Specht), 2 part-time guest professors (Bozzon, Tintarev), 1 associate professor (Hauff), 2 assistant professors (Katsifodimos, Lofi), 2 tenure track assistant professors (Gadiraju, Yang), 10 postdocs, 26 PhD students

Contributors & Credits

Main editors and overall responsibility:

- Arie van Deursen
- Alan Hanjalic

Data collection and overall coordination:

- Jasper van Dijck
- Diane Vedder

Full page image credits:

- Cover image: https://www.tudelft.nl/en/stories/articles/looking-over-vermeers-shoulders
- Page v: https://www.hollandptc.nl/wp-content/uploads/2018/07/FvdB-Lucas-van-Vliet-768x1024.jpg
- Page vi: Arie van Deursen (left) and Alan Hanjalic (right). Image credit: Annelies van 't Hul
- Page 4: Odette Scharenborg, MMC, INSY. Picture taken by Frank Auperlé. https://www.tudelft.nl/stories/articles/zo-begrijpen-spraakherkenners-straks-beter-wat-je-zegt
- Page 8: Przemysław Pawełczak, ENS, ST. Picture taken by Marcel Krijger. https://www.tudelft.nl/en/2020/ewi/game-boy-forever-thanks-to-intermittent-computing
- Page 23: PhD student Jaehun Kim winning a best reviewer award at ISMIR 2019, organized in Delft by Cynthia Liem. Image credit: Bart Vastenhouw
- Page 40: The vanishing points and virtual horizon of the floor of De Hooch's painting "Woman with a child in a pantry", as reconstructed by the group of Elmar Eisemann, CGV, INSY. https://www.tudelft.nl/en/stories/articles/so-much-more-than-3d-visualisation
- Page 48: Diomidis Spinellis, SE, ST. Picture taken by Frank Auperlé. https://www.tudelft.nl/stories/articles/de-softwarejungle-temmen-met-evolutionaire-inzichten
- Page 80: TU Delft libary, book sections. Flickr, https://flic.kr/p/9qmZSh.
- Page 81: Image credits: Bart Vastenhouw, Annelies van 't Hul