Perspectief round 2023/2024
Application
Confidential

Foundation for Industry (FIND)

Large AI models for a resilient high-tech industry

P23.016

Utrecht, February 2024
Dutch Research Council (NWO)

Content

1 Overview

1.1 Title

Foundation for Industry (FIND) – Large AI models for a resilient high-tech industry.

1.2 Public Summary

Foundation for Industry (FIND) brings together 5 universities with 10 labs, 11 Dutch companies, ranging from start-ups to multinationals, and 2 knowledge institutes to pave the way for a new wave of AI-based automation that helps the Dutch industry strengthen and keep its international competitive advantage as a leading high-tech nation in the AI-era. FIND researches and develops novel foundation models, i.e., large artificial intelligence (AI) models, and brings their benefits to unexplored applications and data types that are specific to - and requested by - the Dutch industry. FIND uniquely focusses on data types that are underserved by current foundation models and specifically addresses industry-relevant, low-resource, privacy-sensitive edge applications.

1.3 Project leader and co-lead

Project leader: [redacted]
Co-lead (user representative): [redacted]

[redacted] is primarily responsible for the execution of the Industry-domain work packages (WP 1-3) where industry requirements are derived and research outputs are tested in the context of industry-relevant use-cases. The industry co-lead acts as the user representative and as such is also co-responsible for the design of industry and stakeholder workshops.

1.4 Participating organisations

Research institutes and applicants:

Institute	Department - group	Names of (co-)applicants
Leiden University	[redacted]	[redacted]
Leiden University	[redacted]	[redacted]
Leiden University	[redacted]	[redacted]
University of Amsterdam	[redacted]	[redacted]
Vrije Universiteit Amsterdam	[redacted]	[redacted]
Vrije Universiteit Amsterdam	[redacted]	[redacted]
TU Eindhoven	[redacted]	[redacted]
TU Eindhoven	[redacted]	[redacted]
TU Eindhoven	[redacted]	[redacted]
TU Delft	[redacted]	[redacted]
TU Delft	[redacted]	[redacted]
TU Delft	[redacted]	[redacted]
TU Delft	[redacted]	[redacted]

Involved (end)users:

Start-ups and SMEs	*KAIKO - Perciv.AI - KeyGene*
Large enterprises	*AMSL - NXP - Canon - Shell - Technolution - ASMPT - Stryker*
Knowledge institutes	*NKI - TNO*

1.5 Timing and costs

Requested budget from NWO:	Cash contribution from users:	In-kind contribution from users:
[redacted]		
Project duration: 6 years		

Position	Leiden Uni	UvA	TU/e	VU	TU Delft	Total
Number of PhDs	2	2	5	1	2	12
Number of postdocs		1	1		1	3

1.6 Summary

Foundation for Industry (FIND) brings the benefits of Foundation Models (FM), i.e., large AI models, to unexplored industry-relevant applications and data types to realize a Dutch high-tech industry that is competitive in the AI-era. Foundation models are a disruptive AI technology that can simulate complex cognitive tasks and provide a basis for a new wave of AI-based automation.

[redacted]

The relevant stakeholders are involved in FIND through the User Committee, consisting of industry partners, and the Ethical & Societal Impact board. The involvement of all relevant societal, academic, and industrial stakeholders ensures the utilization of the FIND toolbox and its envisioned impact: Ensuring the Netherlands will keep its international position as a leading high-tech nation.

1.7 Keyword(s)

Artificial Intelligence - **Foundation Models** - **Machine Learning** - **Edge Computing** - **Embedded Systems**

2 Problem analysis and Impact

2.1 Problem analysis

2.1.1 Societal problem and causes

FIND contributes to solving the Dutch industry's AI challenge: How to responsibly leverage and harness the international, extremely fast progress in AI such that it strengthens the industry's global competitiveness. The advancements in AI are progressing at a historically unprecedented pace and driven by global, strategic public and private investments [1]. The past years has seen tremendous progress in AI due the usage of, so called, foundation model (FM) technology. Foundation models [2, 3] have been shown to exhibit human-level and even super-human levels of artificial intelligence in focused applications. The current, global trend is to research and develop FM technology to widen the envelope to which FMs can be applied.

Foundation Models and the underlying technology are disruptive in nature. They allow cognitive demanding tasks to be automated and thus allow for more efficient and effective production processes, products, and services, sometimes referred to as the fourth industrial revolution [4]. They have drawn the interest of many public and private investors as they are a solution to the global workforce challenge: fewer people need to do more work to maintain an acceptable cost- and quality of life for an increasing population [5]. This global challenge is driven by demographic trends as well as globalization (pandemics) and climate change. While Foundation Model technology can contribute to solving these challenges, by realizing higher automation levels, they also pose a threat for those who cannot responsibly harness and leverage this disruptive technology.

Foundation models are a potential threat to the Dutch industry and its global competitiveness. If international industries can realize more cost-effective production processes, products, and services by using more advanced AI, then this can have several negative consequences for the Dutch industry and society:

- There is a risk that the production capacity in the Netherlands will diminish in favour of production elsewhere.
- The Dutch industry and society will rely more on products and services that are produced and provided outside of the Netherlands.
- The Netherlands will become less attractive to international talent and investments on which our leading international high-tech industries are based.

The net effect is that economic stability and growth will shift away from the Netherlands in favour of countries that can take a stronger, leading position in AI.

2.2 Societal impact

2.2.1 Societal impact

The application of foundation models, with the high levels of artificial intelligence they embody, to achieve higher levels of automation and automated support in the Dutch production industry, service industry, healthcare, governmental agencies and other sectors, can have a significant positive societal impact [7, 8]. It allows Dutch industries, with the same number of people, to be able to produce higher quantities of goods and services with higher quality while being more robust to internal and external workforce disturbances and while being more resource efficient.

[redacted]

2.2.2 Economic Opportunities

The economic opportunities for the Netherlands are numerous. Foundation model technology can drive economic growth for the Netherlands whereas ignoring foundation model technology poses a huge risk to the global competitiveness of Dutch industries. Looking at ways to leverage FM technology and to realize impact, foundation models can be embedded in existing or new products, services, and related business models. FIND brings together companies that can directly integrate foundation models in their operations to improve the quality, efficiency, and resilience of their business/production process. Other companies in FIND can integrate foundation model technology into their next-generation products and services to increase competitiveness and market share while speeding time-to-market. FIND also supports start-ups and SMEs (such as Kaiko.ai, Perciv.AI, and KeyGene) that can create entirely new products and services based on foundation model technology. Within FIND, these economic opportunities are structurally integrated in a number of use-cases (see Sec. 4.2) across the industry domains below.

HealthTech (KAIKO, STRYKER, NKI):

The healthcare sector is undergoing increasing workload, partly due to an aging population. The landscape within the healthcare and HealthTech sector is already changing. AI is becoming increasingly important and integrated, however due to technical hurdles such as privacy and security concerns, the progress of these advancements has been slow. [redacted]

In FIND, the HealthTech sector is represented by Stryker, The Netherlands Cancer institute and Kaiko.ai. Stryker is a provider of medical technologies, innovative products and services. They aim to support healthcare providers in improving healthcare. The Netherlands Cancer Institute (NKI) is an all-encompassing cancer centre including a hospital (Antoni van Leeuwenhoek) and research laboratories. They treat patients that cannot be helped at other hospitals and carry out research on different types of cancer. They see AI as a key enabler to improve treatments and healthcare. Kaiko.ai is a start-up focused on AI applications in HealthTech sector, providing AI applications to hospitals, trying to bridge the gap between clinical AI research and the application thereof.

Smart Industry (ASML, ASMPT, KeyGene, Canon, Signify, Shell):

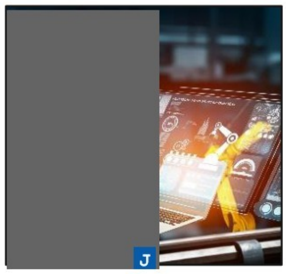

Smart industries are next-generation concepts that utilize AI technology, data, and connectivity to improve flexibility and efficiency of operations and production. It contributes to a better usage of (human) resources. This sector is represented in the consortium by several partners whose turnover varies from a few million to several billion euros per year. KeyGene is a pioneering research company and world leader in agricultural genomics. Signify is the number one company in connected, LED, and conventional lighting supplying solutions to both professional users and consumers. ASML is the leading innovation company and largest supplier in the world of machines to produce computer chips. ASMPT is an international semiconductor company that produces machines for electronics manufacturing. Canon Production Printing creates products and services for markets in printing and workflow management. Lastly, Shell is an energy and petrochemical company that invests in innovative approaches for a sustainable energy future, in which AI will play an important role.

Mobility and Mobile Robotics (NXP, Perciv.AI, Technolution, TNO):

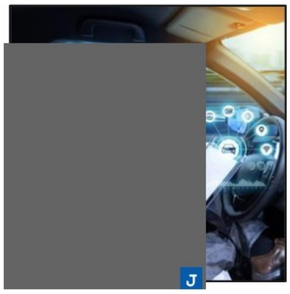

Mobility systems, infrastructure, vehicles, and mobile robots with higher levels of automation are and will be deployed in many areas of modern life, including transportation and logistics. A lot of research on AI-based autonomy solutions is already carried out in the Automotive and Robotics sectors for improving operational efficiency, ensuring precision, and increasing safety. NXP, Perciv.AI, Technolution, and TNO are the representatives of these sectors in the consortium. TNO, the largest independent research organization in the Netherlands, focuses on developing AI software for increased vehicle and robotic autonomy in open world settings. Technolution provides all-encompassing mobility management solutions ranging from smart cameras to measure traffic flows, to traffic management systems that optimize utilization over entire road networks. NXP is a global market leader in automotive semiconductor design and manufacturing with a front-runner position in automotive radar. Lastly, Perciv.AI is a start-up developing AI-driven radar perception solutions for environmental awareness of automated vehicles and mobile robots.

3 Impact pathway

3.1 Outcomes

For the full relationship between outputs, outcomes and impact, see the impact pathway in Ch. 8. To achieve the projected ultimate impact, i.e. maintaining the leading high-tech position of the Netherlands, by means of internally competitive AI-enhanced production, products, and services, and the earlier impacts as stated in Sec.

[redacted]

3.3 Assumptions

When formulating the relationships between the output, outcomes and impact, several key assumptions were made. These key assumptions are:

4 Consortium

4.1 Scientific groups

Industry-relevant research on foundation models requires a team with in-depth expertise in machine learning, data science, as well as embedded and HPC computing. Furthermore, it needs a team with affinity for, and a strong track record in, both fundamental and applied research. Also, non-technical expertise related to ethical, societal, and legal considerations is important for successful and responsible scaling usage of AI technology. Therefore, FIND brings together experts and leading research groups from multiple academic fields, that come with complementary domain knowledge across important data modalities (i.e. vision, radar, time series, text, audio, speech) and different industry verticals (i.e. healthcare, semiconductors, automotive, smart cities), which allows for effective linking academic research to industry. An overview of the expertise and background of FIND's participating research labs/groups is provided in Table 1 and details on each lab/group is provided below. This substantial body of academic know-how is complemented with industry experts from FIND's user-group of which many have PhDs and strong track records in technology R&D. Societal and legal expertise is centred in the FIND Ethics and Societal Impact board and detailed in Sec 4.2.

Table 1: Interdisciplinarity of research groups.

University and lab	Academic expertise	Data modalities
UL-Text Mining and Retrieval	Natural Language Processing, Information Retrieval	Text
UL- Systems and Security	High performance computing and advanced algorithms	Time-series, industrial IoT data
UL- Natural Computing	Evolutionary Computation, Explainable AI, Optimization	Time-series, images, 3D point clouds
UvA - Video & Image Sense	Computer Vision, Machine Learning, Foundation Models	Video, image, text, audio
VU - Quantitative Data Analytics	Machine learning	Time series, medical data
TU/e - Electronic Systems	Low power edge devices, Computer architecture	Streaming real-time data
TU/e - Mobile Perception Systems	Computer Vision, Machine Learning, Robotics	Image, video, radar
TU/e - Generative AI	Generative Modelling, Machine Learning, Deep Learning	Text, time-series, sequence data, images
TUD - Intelligent Vehicles	Computer Vision, Machine Learning, Robotics	Image, lidar, radar
TUD - Computer Vision	Computer Vision, Data-efficient deep learning.	Image, video

The Mobile Perception Systems (MPS) lab (15 FTE), headed by dr. [redacted], focuses on technology by [redacted]

The Video & Image Sense (VIS) Lab (40 fte), headed by [redacted], focusses on computer vision by machine learning and is based at the University of Amsterdam. [redacted] Recent highlights contain efficient self-supervised learning using synthetic rather than real video data [ICCV 23], instilling time-awareness in pre-trained video-language models [CVPR 23A], and the first self-supervised diffusion model [CVPR 23B]. Besides several best papers, the VIS Lab has received prestigious awards, e.g., ACM SIGMM best PhD thesis award, as well as a Veni, Vidi, and ERC Starting Grants. We collaborate closely with industry in our four public-private AI labs, with partners like Qualcomm, TomTom and Elekta Oncology Systems. The lab has spun-off Kepler Vision Technologies and Ellogon.ai. [redacted]

The Electronic Systems (ES) group (80 fte), headed by [redacted]

The Quantitative Data Analytics group (20 fte), headed by [redacted] For the work on efficient deep learning a prestigious QualComm fellowship was awarded. By exploiting information theory), the group has contributed to better generalizable and robust machine learning techniques (see e.g. [CVIU 23, AAAI 23]). The group participates and leads a variety of nationally and internationally funded projects, including high profile grants such as an NWO gravitation grant related to AI and stress understanding and prevention. [redacted]

The Text Mining and Retrieval group (10 FTE), headed by [redacted] The group has an active international and national network, with leading positions in 3 European consortia and coordination of a large NWA consortium on LLMs in low-resource contexts with Dutch industrial partners from the retail and service domains.

The Generative AI group (4FTE), headed by [redacted] The group has made recent contributions to variational auto-encoders [NeurIPS 22-A] and diffusion-based deep generative models [NeurIPS 22]. The research carried out within the Generative AI group is reinforced by multiple applications in Life Sciences (biology, biochemistry), Molecular Sciences (chemistry, physics), and problems ranging from signal processing (e.g., data compression) to self-driving cars, and smart devices, and smart apps (e.g., chatbots, art generation).

The Systems & Security group (32 FTE), headed by [redacted] The group contributes with expertise on high-performance computing aspects of training deep models and AI for multivariate time series, such as anomaly detection [A&A 23, ComSurv. 20].

Intelligent Vehicles group (20 FTE), headed by [redacted]

The group has made key contributions to AI-driven automotive radar perception by producing one of the first public benchmark for imaging radars [RAL 22], and cross-modal supervision of radar models through camera and LiDAR [CVPR 23-A][RadarConf 24]. To tackle the data scarcity problem, its research has explored imputing domain knowledge and structural priors in various perception tasks, e.g. object detection [NeurIPS 22], LiDAR semantic segmentation [TPAMI 22], efficient visual localisation [CVPR 23-B], and uncertainty quantification in visual place recognition [CVPR 24]. The group's research collaborations include the automotive industry, such as Daimler and Toyota, as well the Dutch and European tech/AI industry, including TomTom, NXP, and Bosch.

Computer Vision lab (15 FTE), headed by [redacted] is based at Delft University of Technology (TU Delft). [redacted] Related is our efforts to uphold scientific reproducibility, as demonstrated by our public website for www.reproducedpapers.org [Repro21]. In this project we collaborate with [redacted]

Natural Computing group (35 FTE) headed by [redacted] and represented by [redacted] The NaCo lab has made key contributions in the field of evolutionary computation, particularly in developing and optimizing strategies for genetic algorithms and evolutionary strategies, including adaptive techniques for self-tuning parameters [MIT 96], AI-aided optimization and optimization landscape analysis [GECCO 22] [GECCO 23], and using evolutionary computing for automated machine learning [ICAML 23].

4.2 Users and stakeholders

Co-design: The FIND program was co-designed by industry users, united in our User Committee (UC), and a number of societal stakeholders, united in our Ethics & Societal Impact (ESI) board. As such, FIND is rooted in multiple practical and relevant use-cases; [redacted].

The industry partners have all been deeply involved in the co-design of FIND which took place through project workshops with the entire consortium, specific impact workshops with the UC, and one-on-one contact with all industry partners. [redacted]

[redacted]

[redacted]

Table 2: Industry use-cases.

Technolution	[redacted]
ASML	[redacted]
stryker	[redacted]
kaiko. NETHERLANDS CANCER INSTITUTE ANTONI VAN LEEUWENHOEK	[redacted]
signify	[redacted]
Canon	[redacted]
NXP PERCIVAI	[redacted]
KeyGene	[redacted]
TNO	[redacted]

Shell	[redacted]
ASMPT	K

Ethics and Societal Impact (ESI) board: To create impact, it is important to scale the outcomes of FIND beyond the industry partners in the UC. The FIND program requires an 'interface' to a broader set of stakeholders in industry and society. For this purpose, the FIND Ethics and Societal Impact (ESI) board is established. The ESI board is designed to be highly complementary, bringing together a diverse range of expertise and perspectives. The board comprises professionals from various fields and they will support FIND with: 1) providing advice and guidance on ethical and legal barriers for AI deployment, 2) supporting outreach and dissemination activities to industry and societal stakeholders, and 3) organizing start-up activities for FIND researchers together with other knowledge dissemination activities.

The current list of organizations and persons taking place in the ESI board are provided below. During the FIND program new members will be invited, according to the needs of the FIND program. Budget is allocated (see Sec. 6.1) to cover out-of-pocket costs of ESI board members. The member of the ESI board are provided in Tabel 3.

Table 3: Members of the ESI board.

Ethics & Societal Impact (ESI) board
[redacted] J

Foundation for Industry (FIND) - Large AI Models for a resilient high-tech industry.

[redacted]

[redacted]

[redacted]

5 Project description

5.1 Objectives and technical & scientific challenges

Whereas previous deep AI models were (pre-)trained in a supervised fashion using data that was labelled for a specific unimodal task and output label space, foundation models are pre-trained on massive unlabelled or weakly-labelled (multi-modal) data and can be efficiently adapted or fine-tuned via supervised learning to serve multiple tasks and output spaces [11, 12]. When foundation models are very large, fine-tuning is in some cases even no longer needed; the model can be instructed through the prompt to perform a specific task, potentially with a few (up to 5) examples [13, 14]. The benefits of massive pre-training with unlabelled or weakly-labelled data are:

- The same AI model can be adapted relatively efficiently w.r.t. training time and labelled data to serve multiple (multi-modal) tasks and output spaces and still achieve state-of-the-art accuracy [15, 16].
- It provides the AI model with stronger latent representations that allows the adapted model to generalize better to unseen in- and out-of-distribution data [17, 18].

In addition, the transformer architecture, originally proposed for text data [19], allows for processing data in multiple modalities, and for cross-modal alignment where data in one modality is encoded and data in another modality is decoded (generated). Examples are image-to-text models and text-to-image models. For modalities other than text and images, much fewer developments have been made so far. [redacted]

[redacted]

Foundation for Industry (FIND) - Large AI Models for a resilient high-tech industry.

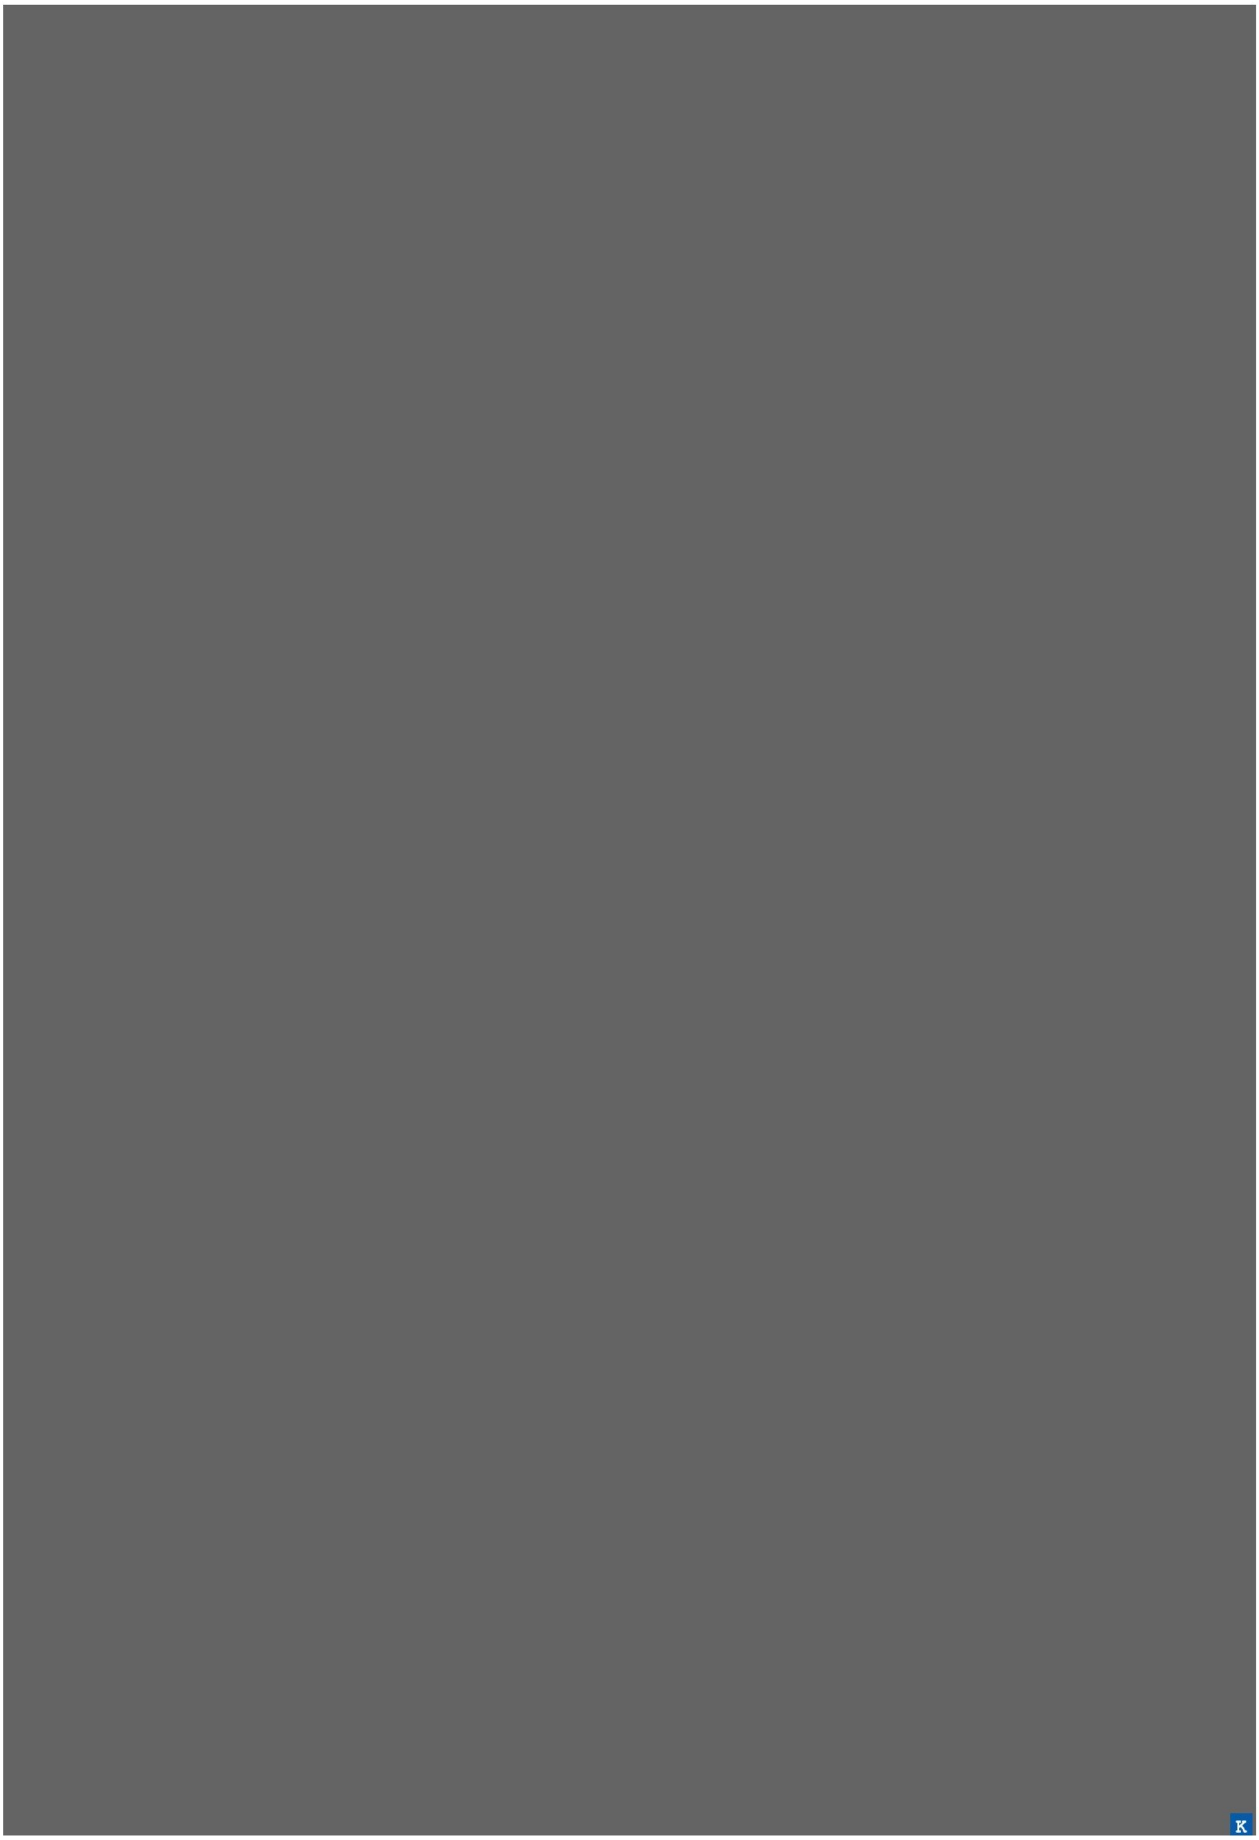

5.2 Project structure

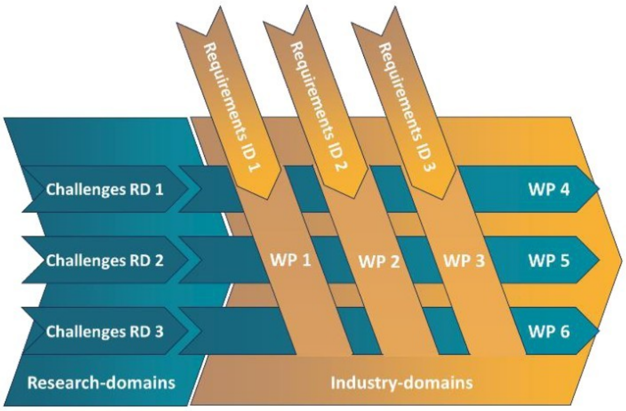

Figure 4: FIND program work package structure driven by Research-domain (RD) challenges and Industry-domain (ID) requirements.

To have an effective program management with the required managerial governance, six work packages are defined together with the work package for Management and Knowledge dissemination that is described in Section 6. In the FIND program three Industry-domain work packages are centred around the use-cases and three Research-domain work packages are centred around the research challenges. Together these work packages form a matrix structure where horizontal enabling AI technologies are integrated in vertical industry use-cases. This concept is illustrated in Figure 4. All researchers are active in one Industry-domain work package and in at least one Research-domain work package. The research position allocation is provided in Table 3. The iterative process by which these work packages are managed, and the governance is described in Sec. 6.1.

Industry-domain work packages: [redacted]

For managerial purposes, use-cases are divided over industry verticals:
WP 1 HealthTech: KAIKO, NKI, Stryker.
WP 2 Smart industry: Signify, ASML, Canon, KeyGene, Shell, ASMPT.
WP 3 Mobility and mobile robotics: NXP, PercivAI, TNO, Technolution.

[redacted]

R[redacted]

[redacted]

Table 4: Schematic overview of the project, including the distribution of PhD/postdoc researchers.

Foundation for Industry												
Industry domains	**Lead** **Co-lead**	**Use-cases and involved users**		**PhD and postdoc allocation (16 total)**								
WP 1: HealthTech	[redacted]	UC 1: KAIKO, NKI										
		UC 2: Stryker										
WP 2 Smart Industry		UC 3: ASML, Canon, ASMPT										
		UC 4: KeyGene										
		UC 5: Signify										
		UC 6: Shell										
WP 3: Mobility and Mobile Robotics		UC 7: NXP, PercivAI										
		UC 8: Technolution										
		UC 9: TNO		1	2	3	4	5	6	7	8	9
Research domains		**Solution lines and involved academia**		**WP 1**		**WP 2**				**WP 3**		
[redacted]		SL 1: TU/e, TUD	WP 4:							2		
		SL 2: UvA, TUD, TU/e				2			1			1
		SL 3: UvA		1								
		SL 4: TU/e	WP 5:							1		
		SL 5: TU/e, UL				1					1	
		SL 6: UvA		1								
		SL 7: TU/e, VU	WP 6:		0.5							1
		SL 8: VU			0.5			1				
		SL 9: TUD					1					

Coherence, added value, synergy, and interdisciplinarity: An important consideration is that the Perspectief funding instrument is primarily meant and used to fund academic PhD and postdoc research, thus at relatively low TRL levels 1-3. At the same time Perspectief requires strong industry involvement which prefer higher TRL levels 6-9. This naturally creates a TRL gap, i.e. the void from TRL 4 to TRL 5 between academic research and industrial R&D. [redacted]

To address the academic and industry challenges, an interdisciplinary team of experts is formed in the FIND program. [redacted]

[redacted]

Best practices from related European programs: The FIND project structure is based on well-established approaches for projects under the Chips Joint Undertaking (Chips-JU), formerly Key Digital Technologies Joint Undertaking (KDT-JU), which is a EU public-private partnership and an overarching objective to support the digital transformation of all economic and societal sectors and the European Green Deal [80]. Chips-JU projects, similar to NWO Perspectief programs, bring together universities, companies, and societal stakeholders in multi-disciplinary research projects aiming to achieve impact in multiple industry verticals. The structures used in Chips projects, build on a decade of experience in knowledge utilization, and is characterized by a matrix of horizontal enabling-technology work packages and vertical industry value-chains with multiple value-chain demonstrators. This creates a clear set of university and industry roles and responsibilities and an effective project management structure. The same proven approach is used in FIND. Three horizontal Research domain work packages, under responsibility of the universities, research and provide enabling FM technology at TRL 1-3. Three vertical Industry domain work packages integrate and test these FM technologies in the context of industry use-cases at TRL 4-5 and are under the responsibility of the companies. Overarching is the joint responsibility of aligning on FM requirements and steering the FM research and integration. TU Eindhoven and Technolution, the lead and co-lead of the FIND program, have decades of experience in Chips-, KDT- and ECSEL-projects and will use their project management knowhow in FIND.

5.3 Utilisation plan

The cornerstone for utilization of the outputs within the UC industry members are their use-cases and the precompetitive FIND FM toolbox. The use-cases, their TRL levels, and their role are described in Sec. 4.2 and detailed further in the Industry-domain work packages WP 1-3 in Sec. 6.2. The FIND FM Toolbox was introduced in Ch. 3 and detailed in Ch. 8. The utilization outside of the UC is also realized via the FIND FM Toolbox, the Ethical & Societal Impact board, and the many AI hubs in the Netherlands, as introduced in Sec. 4.2. The main strategy here is making best use of the many public-private initiatives for AI uptake that already exist in The Netherlands, i.e. not ‘reinvent the wheel’, and to which many of the academic partners are already affiliated to.

The implementation and execution of the utilisation plan in practice during the project will be the responsibility of the project management bodies. The planning and governance structure ensuring the utilisation plan takes place is explained in more detail in Sec. 6.1. The utilisation plan, where the industry-driven use-cases have a pivotal role, consists of three main phases that will ensure the results and outcomes of the project are used:

[redacted]

[redacted]

[redacted]

Table 5: Relevant AI projects.

Name	Focus	Complementarity
Efficient Deep Learning	[redacted]	[redacted]

Hybrid	[redacted]
Lessen	
Megamind	
AI4AI	
National AI-factory	

5.5 Risk management and contingency plan

To ensure the program's successful execution and its potential to societal impact, a comprehensive risk analysis has been conducted. An overview is provided in the table below, which includes the risks identified, likelihood and impact of occurrence and the mitigation measures. These actions are designed to be reactive when necessary and preventive wherever feasible. Many of the latter category have been taken already.

Table 6: Risk management. LI=Likelihood, I=Impact, L=Low, M=Medium, H=High.

Risk	**LI**	**I**	**Mitigation or preventive action**
[redacted]			

[redacted] K

6 Description of work packages

6.1 Project management and knowledge utilisation

WP title:	**Project management and knowledge utilisation**
WP leader: [redacted] J **Project management team:** Project lead & co-lead + project manager + project support officer + WP leaders & co-leaders + NWO program officer.	

6.1.1 Project management

WP Management and Knowledge utilization: In this work package the program coordination takes place, as well as the organisation of the workshops and program events. The FIND workshops are the main drivers to realize the outputs and outcomes that create the desired impact.
The program will be managed through an iterative and agile design which assures knowledge transfer and utilisation both internally and externally. The program has a clear governance structure, which consists out of five main bodies, which are visualised in Figure 6. These are: (1) Project management team, (2) Steering committee, (3) Work package teams, (4) Ethics and Societal Impact Board, (5) User Committee.

Project management team: Management of the program is the main responsibility of the **Project lead** (TU Eindhoven) and the **industry co-lead** (Technolution B.V.) who will be supported by a professional project manager. The (academic) program lead will take responsibility for the research lines (WP1-3) and the industry co-lead of the industry prototypes (WP4-6).
Key tasks also include creating the legal and organizational conditions under which data can be shared between industry and academic partners. Together with the project manager, project support officer an NWO program officer, the project lead and co-lead make up the **project management team:**

- ***Project** (scientific) **lead***: [redacted] J The project lead is ultimately responsible for the program's execution and for it to be successful and scientifically excellent.
- ***Project co-lead** (user representative)*: [redacted] J. The user representative is the spokesperson of the users in UC and is responsible for the program to be practically relevant, state of the art, and impactful.
- ***Project manager***: [*vacancy will be filled by the TU/e team (or external party) after grant*]. The project manager (PM) has the operational lead of the project and will focus on the timely execution of project planning, goals,

scope and deliverables within budget. The PM will be supported by a project support officer for secretarial, administrative, and compliance related tasks (e.g. GDPR, AVG).

- ***Impact officer***: [*vacancy will be filled by the TU/e team (or external party) after grant*]. The role of the Impact officer is to organize, facilitate, and monitor the development of the industry-driven use-cases and the related Co-creation workshops. The Impact officer monitors and steers the development of the precompetitive FIND FM Toolbox and coordinates with the ESI board on knowledge utilization activities external to the members of the UC. The Impact officer is supported by a social-media impact officer to communicate the results of FIND via academic, industrial, and general (social) media channels.
- ***NWO program officer***: [to be assigned by NWO]. The NWO Program Officer monitors the overall progress of the program and provides support and feedback on its execution.

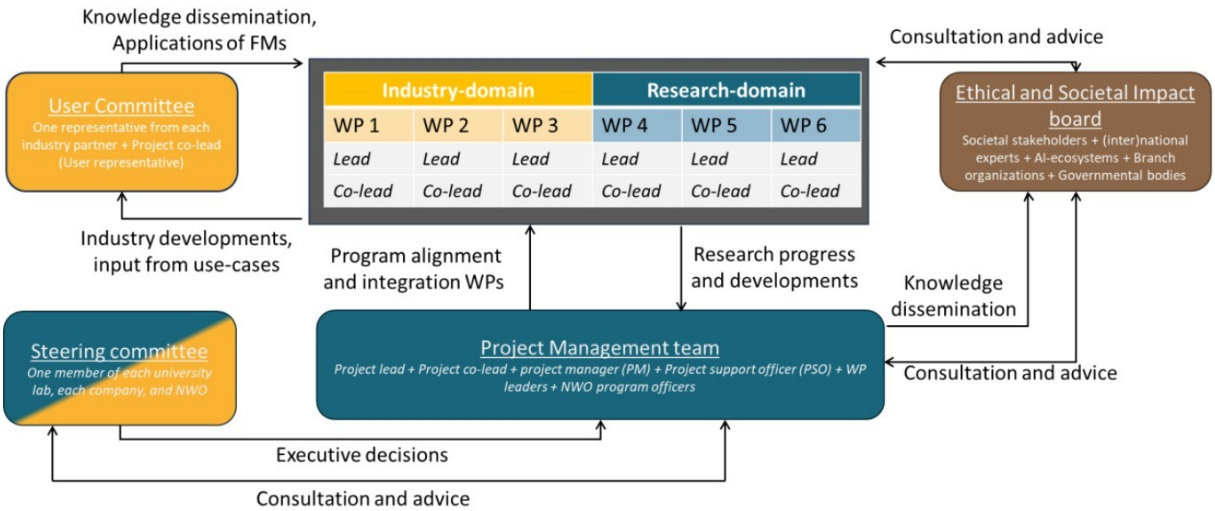

Figure 6: Governance structure and representation of main bodies together with main interactions.

Steering committee: The steering committee is the ultimate decision-making authority within the project. In case a consensus cannot be made in any of the other governance bodies, nor can be extracted from the consortium agreement, the steering committee is convened to vote on a decision. The steering committee is comprised of one representative of every consortium partner, i.e., each university lab and each company participating in FIND.
Work package teams: The researchers and industry partners (co-funders) active in each work package make up the Work package teams. Each Industry WP has a lead and a co-lead, where the Research WP's only have a lead. This setup allows the team to split the work, promotes diversity, fosters co-creation, allows more people to gain management experience and decreases the chance of bias. The co-lead is the acting lead in case the lead is absent. The work package lead and co-lead can escalate to the supervisors of the researchers within the work package firstly or go further up the chain to the project management team if necessary. See the project overview In Table 4 for the leads and co-leads of the work package teams.

Ethical & Societal Impact (ESI) Board: The impact board advises the project lead and co-lead about the direction of the project based on the progress report, with the aim of maximising the chances of scientific and societal impact. The committee is consulted periodically to review and give advice on various levels, such as scientific work (e.g. concept papers), industrial and practical relevance and overall project progress. The Ethical & Societal Impact Board consists of national and international experts from industry, government and science, as well as branch organizations of end-user parties, which are set out in Section 4.2. These experts and organizations also play an important part in upscaling the adoption and utilization of the FIND toolbox through their support for various knowledge utilization and dissemination activities, such as the Impact Workshops, in which the outputs are shared with their respective user bases.
User Committee (UC): This committee consists of a group of experts from the industries involved within the consortium. Every industry partner delegates one representative, who together form the UC. These users (or industry partners) are the envisioned end-users of the FIND toolbox from various sectors. As such, the partners

from this committee will engage in testing and validating and with that provide important user feedback for development and improvement of the toolbox. The testing and validation of the FIND toolbox takes place in the industry partners' respective use-cases (see Sec. 4.2 and 6.2). They also actively participate in the (industry) work packages (WP1-3), and through their feedback help shape and co-create the research.

Meetings: The bodies in the governance structure have frequent governance meetings, as provided in Table 7, in addition to the FIND workshops that are described later in this section. The goal of the governance meetings is to manage the process and monitor the progress. The goal of the workshops is to steer the content of the research, the development of the use-cases, and to realize knowledge utilization. The workshops thus serve a different purpose than the governance meetings. FIND workshops are in-person meetings and the governance meetings will be collocated, as much as possible, with these workshops. In case they cannot be collocated, the governance meetings are held online. The semi-annual and annual meetings will always be collocated with workshops.

Table 7: Governance meeting structure and frequency.

Meeting structure *Planned duration*	**Members**	**Frequency**
WP meeting (one per WP 1-6) *1 hour*	WP (co-)lead, involved researchers, research supervisors.	Monthly
Quarterly PM meeting (WP 0) *1.5 hours*	Project lead & co-lead, WP leads, Project manager, Impact officer.	Quarterly
ESI meeting (WP 0) *3 hours*	Ethics & Societal Impact board members, Project management team.	Semi-annually
Annual PM meeting (WP 0) *3 hours*	Project lead & co-lead, WP leads, Project manager, Impact officer, NWO program Officer.	Annually
Critical decision meeting (WP 0) *variable*	Steering committee.	When needed

Iterative management process: The exponential innovations in AI demand highly agile and iterative program design and management so that external developments and learnings in the field of AI can be adapted quickly by the program. Instead of defining exact deliverables of the program in this proposal, which surely will become outdated with the exponential AI innovation rate, the FIND program details the iterative process under which its deliverables will be determined and realized during the 6-year course of the program. This iterative structure is visualized in Figure 7 and the iterative development of the AI technology and the industry-driven use-cases are detailed in the WP descriptions and contain three main cycles with clear set inputs, outputs, and milestones.

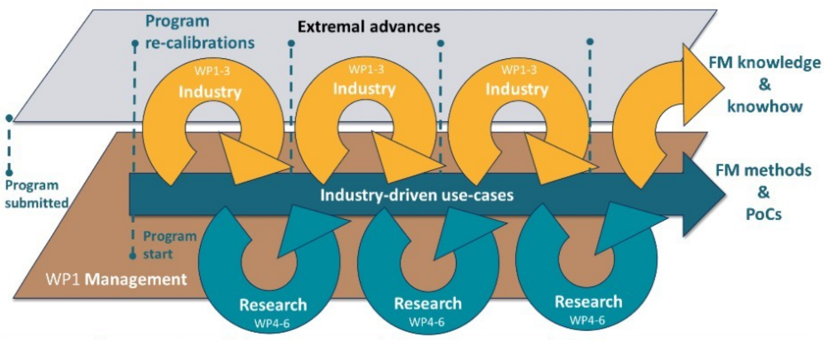

Figure 7: Illustration of the iterative design process of the program structure, to be resilient to external advances in the field of AI and the ability to timely incorporate them into the program.

Workshops: To support and facilitate the iterative process, FIND organizes a series of workshops. The workshops are foundational to steering the research and the use-cases and create impact outside of the FIND user group. The workshops were detailed in Sec. 5.3 but for completeness are again summarized below with more focus on managerial aspects. The FIND workshops are planned semi-annually, see Table 8, in conjunction with each other over the duration of 2 days in June and in December. The are organized by the program lead and co-lead with support of the project manager and the impact officer.

Table 8: Workshop planning.

[redacted]

Communication: Communication to a broad audience outside the FIND UC and ESI board is important to internationally achieve impact with the research outputs. On top of academic and industrial dissemination in the usual channels (international conferences, journals, tradeshows) the role of academic social media is often overlooked. This channel is however becoming increasingly important because it facilitates timely dissemination of AI research results before they are outdated. Putting publications on open platforms as Arxiv is a first step but not enough to attract attention. Academic social media campaigns can help bring the needed attention to research results on Arxiv and significantly boost their impact. To facilitate this, the FIND program will make use of a Social Media Impact (SMI) officer that will support the researchers with obtaining attention in academic social media. The SMI officer will also be responsible for active content creation on the FIND website, LinkedIn account, and other relevant (social) media channels.

[redacted]

6.1.3 Requested budget

[redacted]

	Project management	Stakeholder engagement	Communication	Monitoring and evaluation	Capacity strengthening
Budget	[redacted]				
Description					

6.2 Work packages

Coherent program-level task and milestone planning

We first provide FIND's coherent program-level task planning that transcends work package boundaries. This guarantees effective alignment between activities and prevents siloed activities. All three Industry-domain work packages WP 1-3 and all three Research-domain work packages WP 4-6 follow an identical structure in terms of time plan, division of tasks, milestones, and the role of users.

For efficiency, this structure is provided only once for all Industry domain work packages WP 1-3 and once for all Research domain work packages WP 4-6. Afterwards, only the details of the each of the use-cases (WP 1-3) and the research challenges and solution lines (WP 4-6) are provided per work package. This separation of process and content is a key pillar of effective R&D.

Foundation for Industry (FIND) - Large AI Models for a resilient high-tech industry.

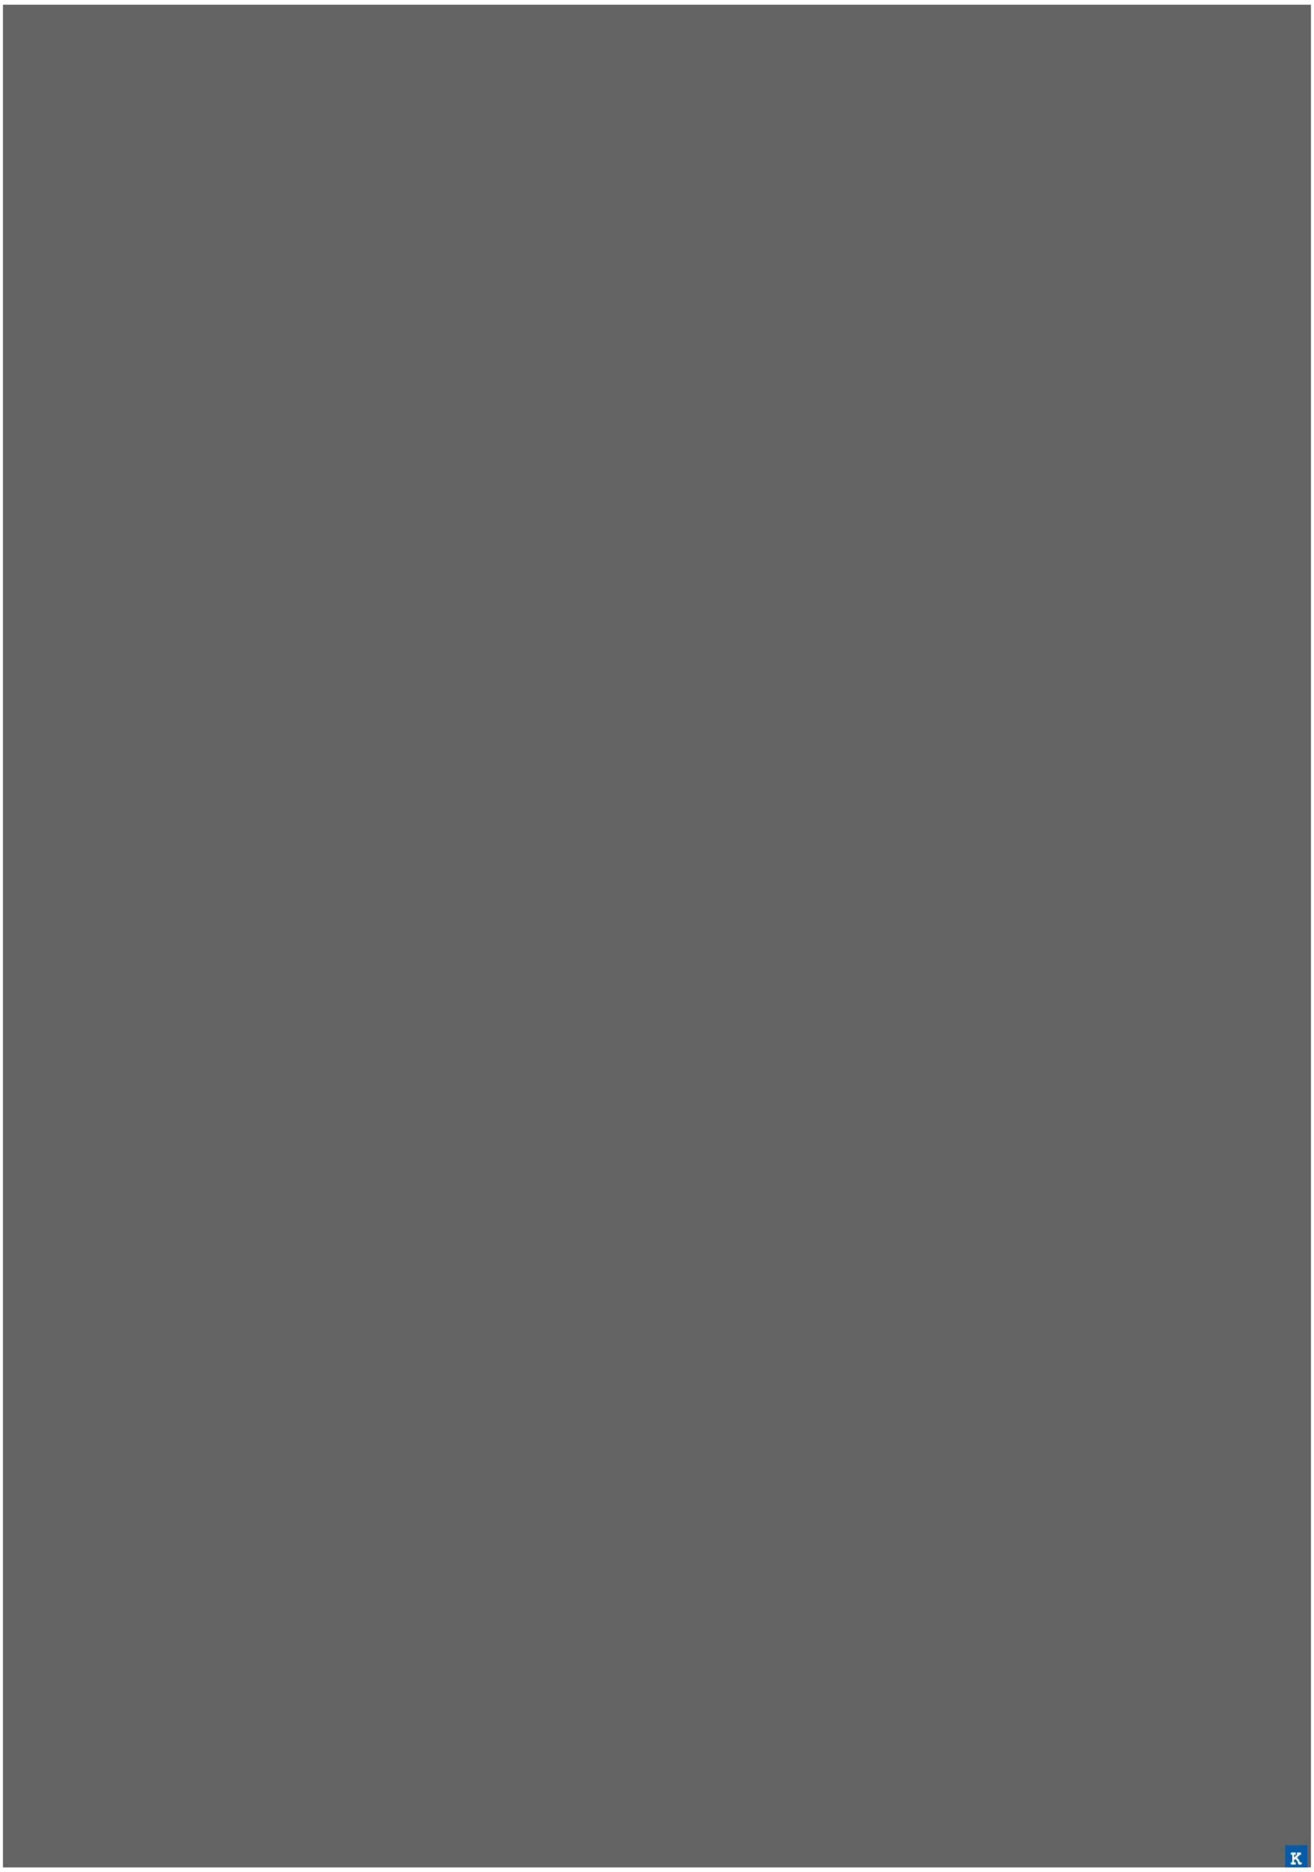

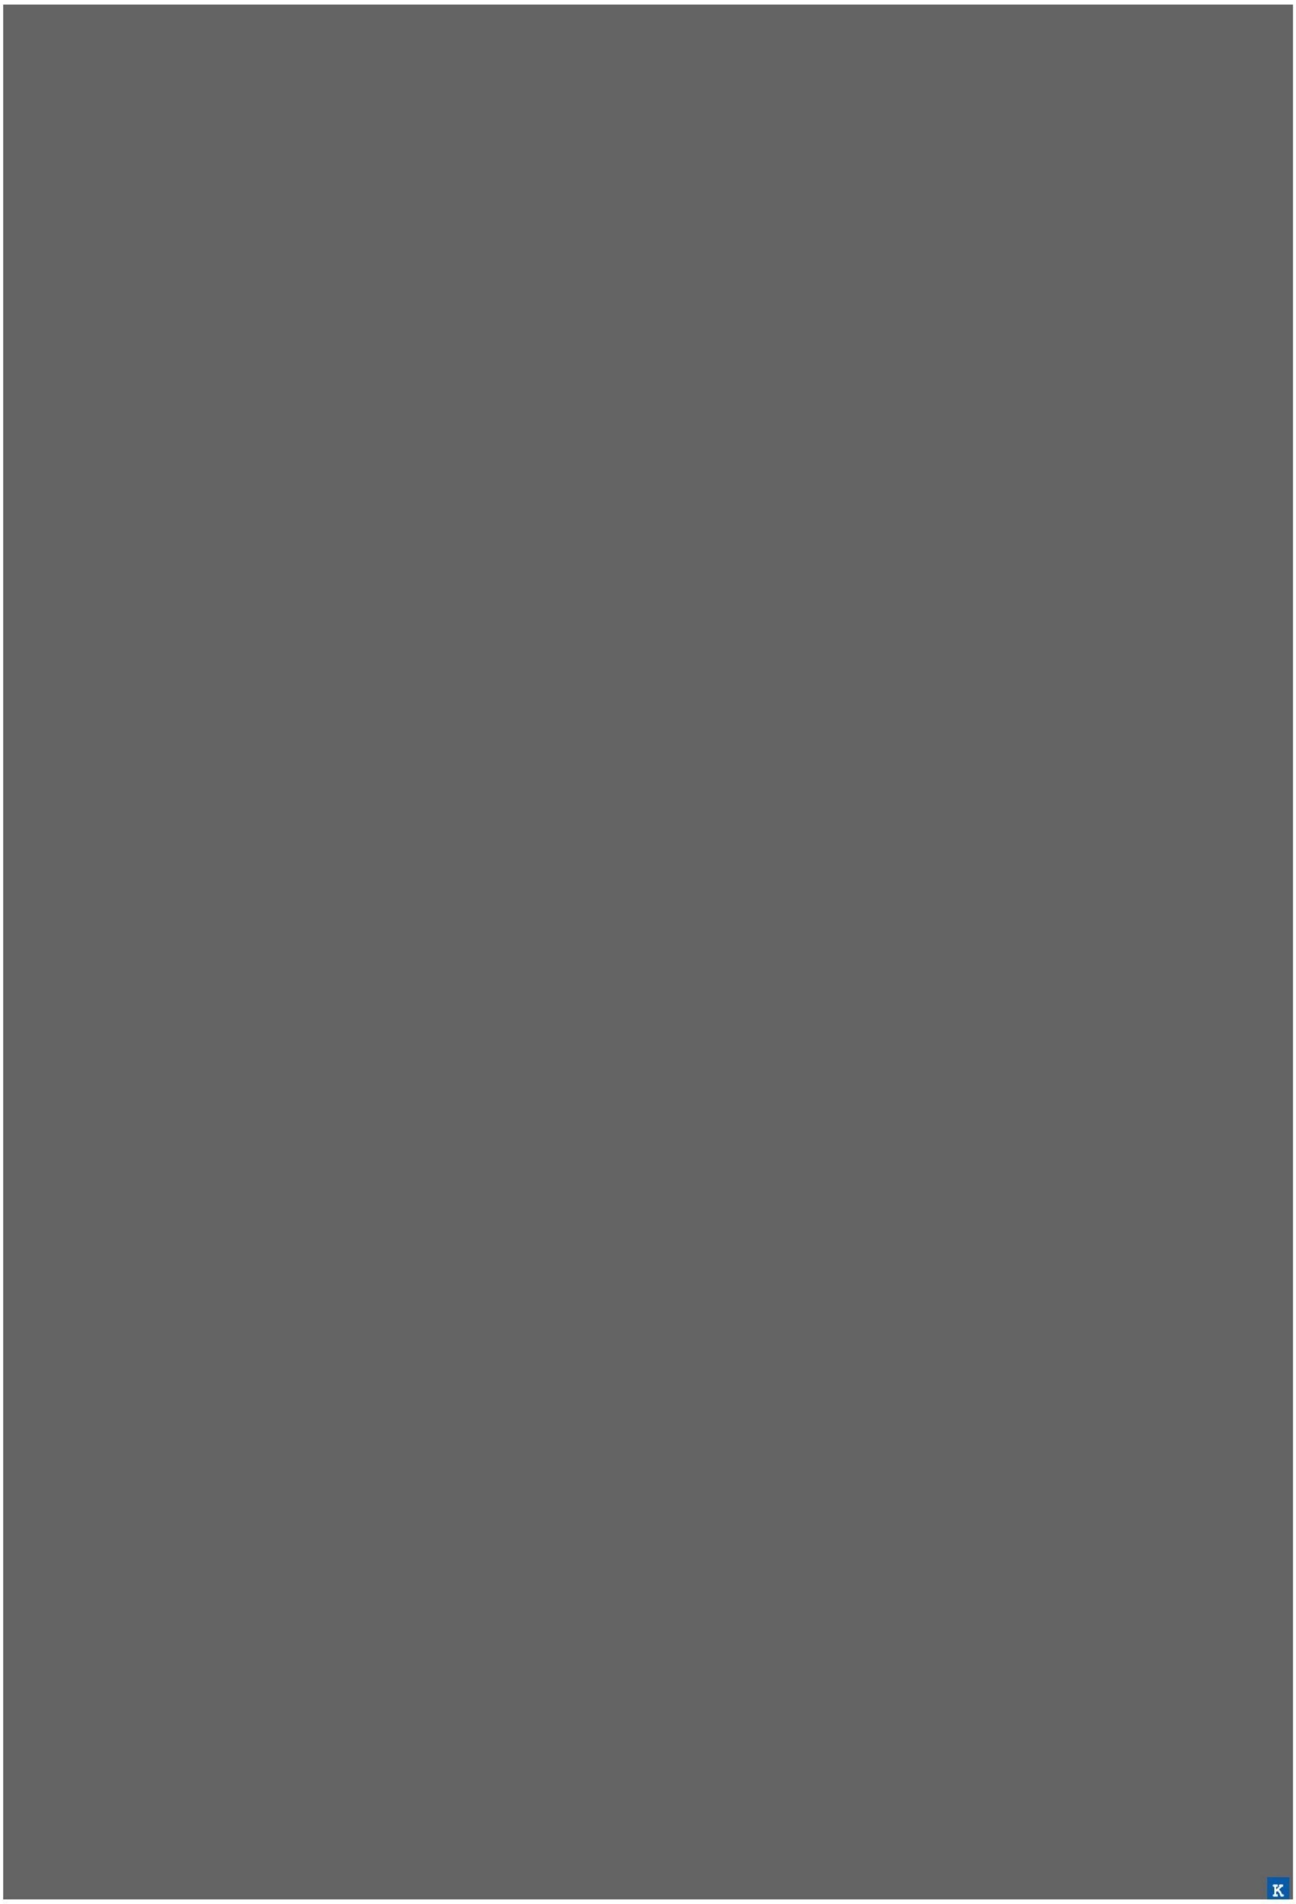

[redacted]

6.2.4 Work package descriptions

WP number:	**1**
WP title:	**HealthTech**

WP leader: [redacted]
Industry co-lead: [redacted]
Involved researchers: [redacted].
Requested research positions: 2 PhDs and 1 postdoc with combined effort of 33 PMs.
Researchers spend on average 25% of their effort in the Industry-domain WPs.
Industry partners: Kaiko, Stryker.
Duration of WP: 6 years

Scientific description of the work package

[redacted]

Osteoarthritis use-case (Stryker): Osteoarthritis (OA) is a severe condition that causes pain, functional impairment, and a decline in life quality, constituting a significant health issue. While joint replacement is an effective remedy for advanced OA, about 15-20% of patients are unsatisfied with their total knee replacement surgery after one year. [redacted]

Oncology use-case (KAIKO): [redacted]

[redacted]

[redacted]

WP number:	**2**
WP title:	**Smart Industry**

WP leader: [redacted]
Industry co-lead: [redacted]
Involved researchers: [redacted]
Requested research positions: 5 PhDs and 1 postdoc with combined effort of 69 PMs.
Researchers spend on average 25% of their effort in the Industry-domain WPs.
Industry partners: Signify, ASML, ASMPT, Canon, KeyGene, Shell.
Duration of WP: 6 years

Scientific description of the work package

[redacted]

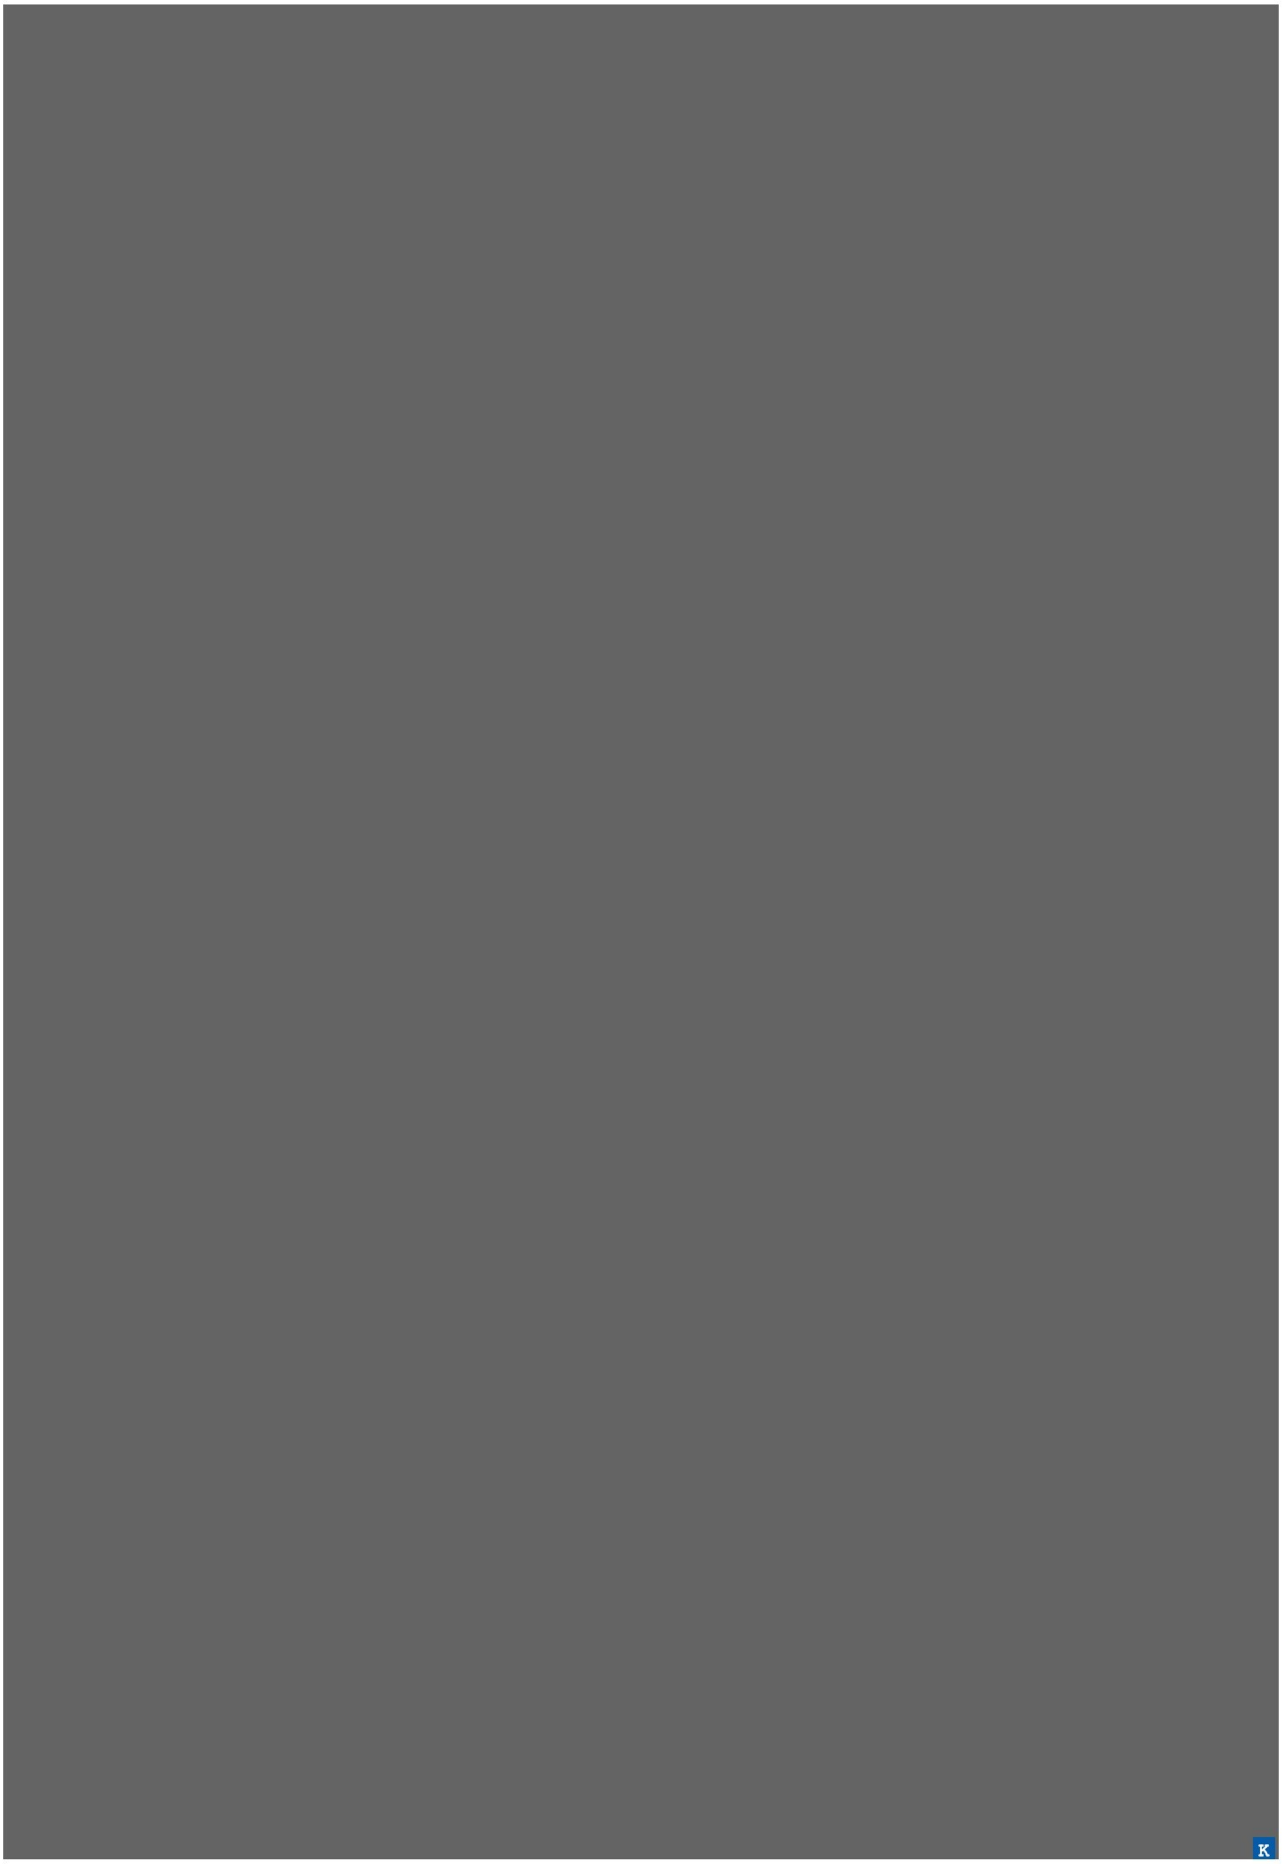

[redacted]

WP number: **3**
WP title: **Mobility and Mobile Robotics**

WP leader: [redacted]
Industry co-lead: [redacted]

Requested research positions: 5 PhDs and 1 postdoc with combined effort of 69 PMs.
Researchers spend on average 25% of their effort in the Industry-domain WPs.
Industry partners: NXP, TNO, Perciv.AI, Technolution.
Duration of WP: 6 years

[redacted]

Imaging radar use-case (NXP-PercivAI): [redacted] With the increase of spatial 3D resolution and processing power, automotive imaging radars allow for deeper semantic understanding of the environment compared to previous generations automotive radar [45]. [redacted]

[redacted]

WP number:	4
WP title:	FM Pre-training and Adaptation

WP leader: [redacted]
Involved researchers (co-leads of WP)**:** [redacted]
Requested research positions: 7 PhDs (2 + 2 x 0.5 at TU/e, 2 at TUD, 1 at UvA) + 1 Postdoc at UvA with combined effort of 244 PMs. *PhDs spend on average 75% of their effort in the Research-domain WPs.*
Related use-case partners: NXP, ASML, Canon, ASMPT, Kaiko, TNO, Perciv.AI, NKI.
Duration of WP: 6 years

[redacted]

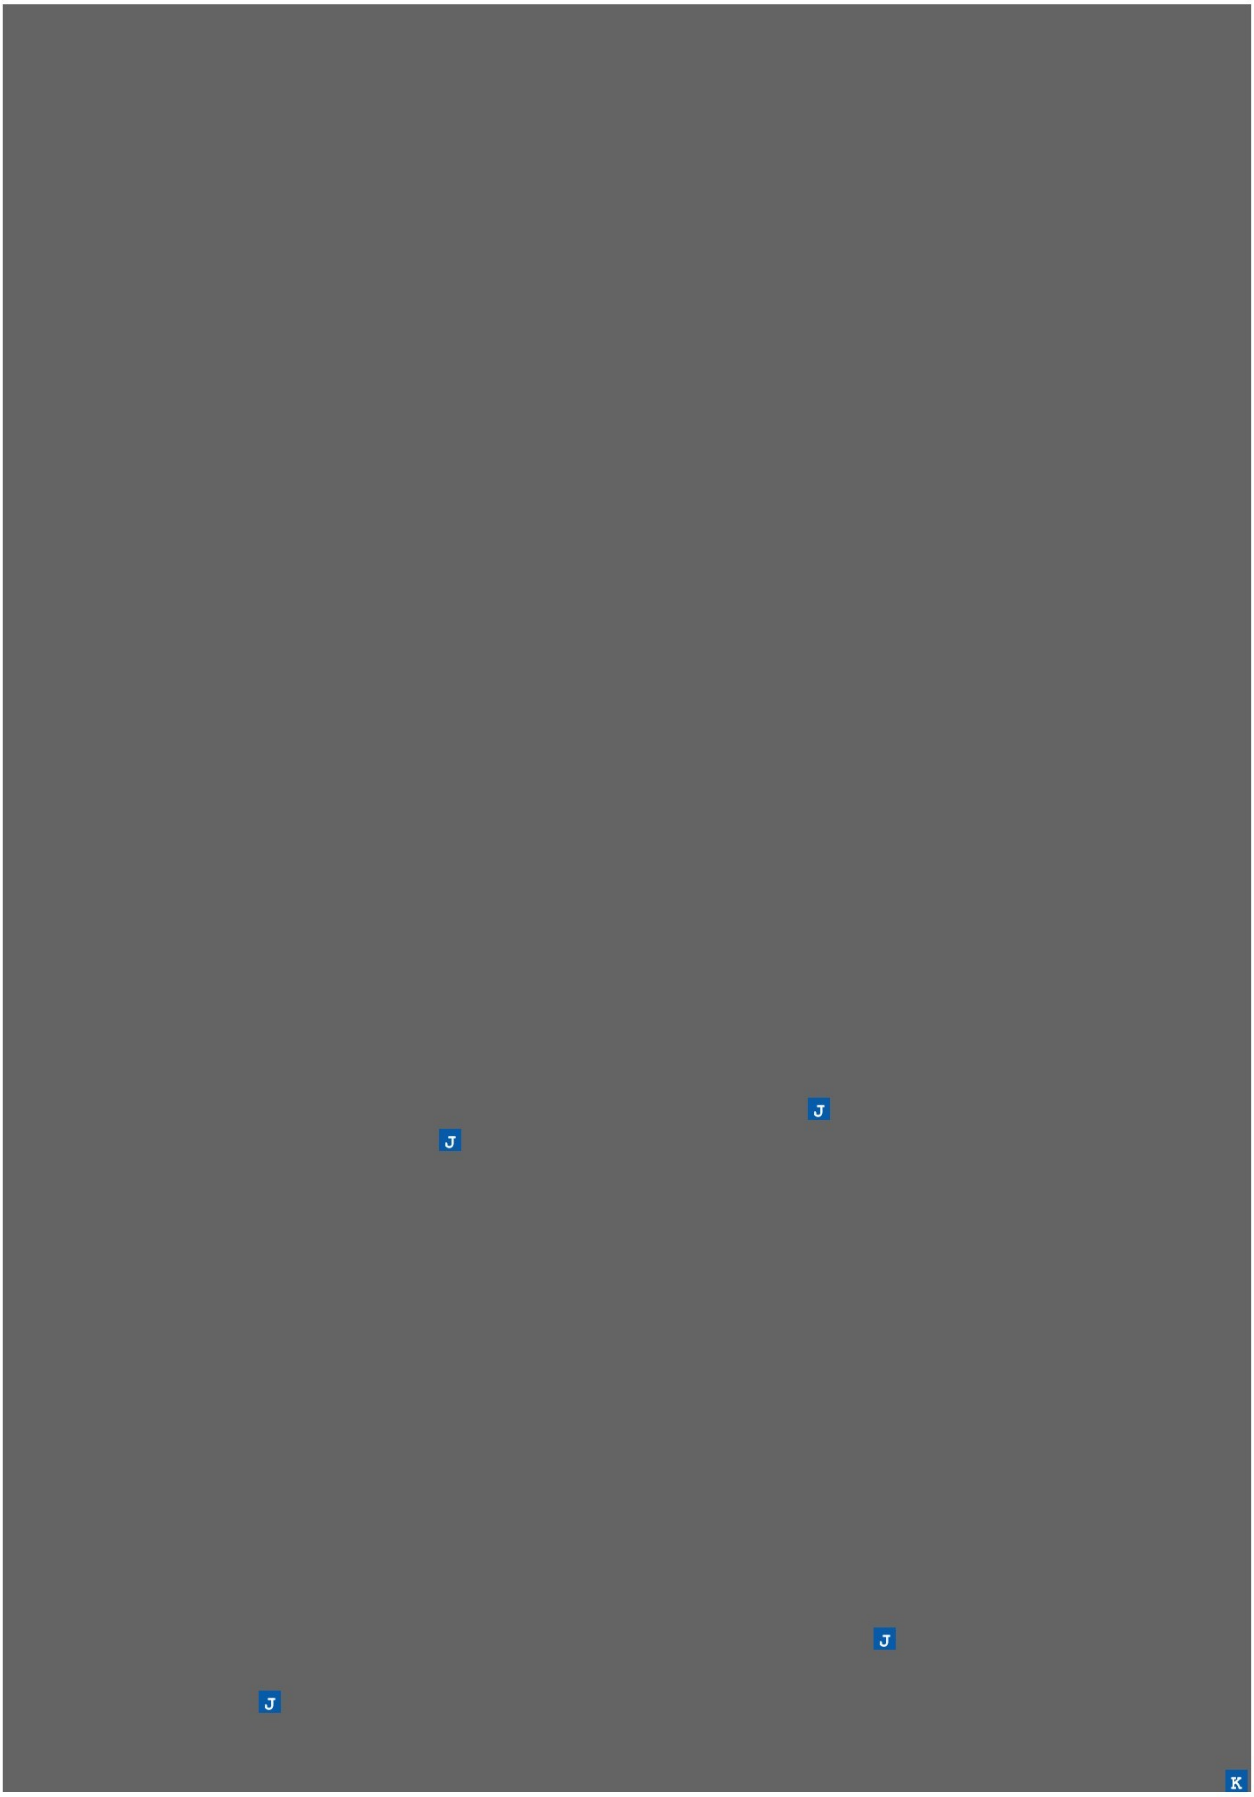

[redacted]

WP number:	**5**
WP title:	**FM edge deployment**

WP leader [redacted]
Involved researchers [redacted]

Requested research positions: 3 PhDs (1 at UvA, 2 x 0.5 at TU/e, 1 at LU) + 1 Postdoc at TU/e with combined effort of 136 PMs. *PhDs spend on av7erage 75% of their effort in the Research-domain WPs.*
Related use-case partners: Kaiko, Technolution, NXP, Canon, ASMPT, NKI.
Duration of WP: 6 years

[redacted]

Foundation for Industry (FIND) - Large AI Models for a resilient high-tech industry.

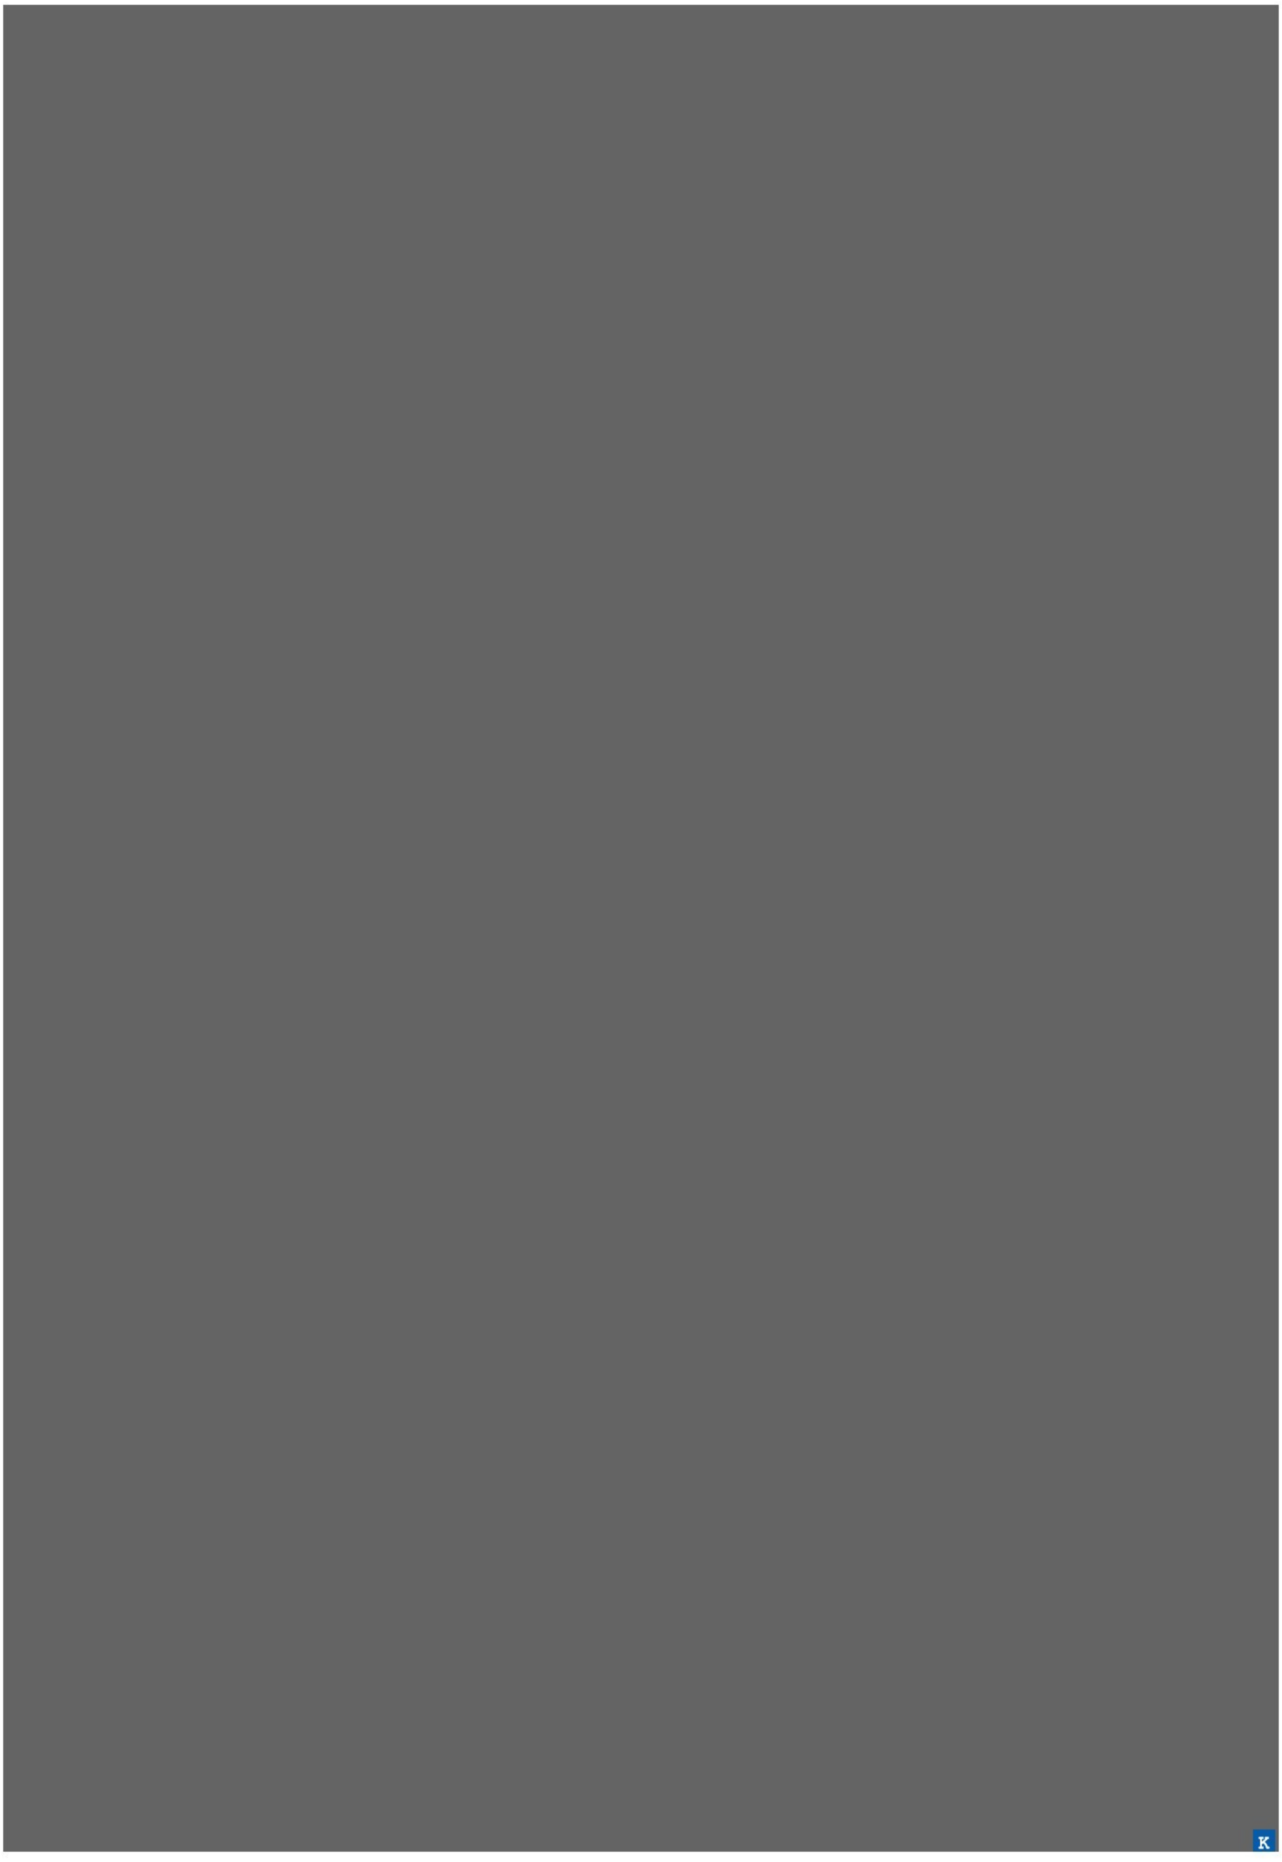

[redacted]

WP number:	**6**
WP title:	**FM Robustness and privacy**

WP leader: [redacted]
Involved researchers (co-leads of WP): [redacted]
Requested research positions: 3 PhD students (1 at VU, 1 at TU/e, 1 at LU) + 1 Postdoc at TUD with combined effort of 136 PMs. *PhDs spend on average 75% of their effort in the Research-domain WPs.*
Related use-case partners: Stryker, TNO, KeyGene, Signify.
Duration of WP: 6 years.

[redacted]

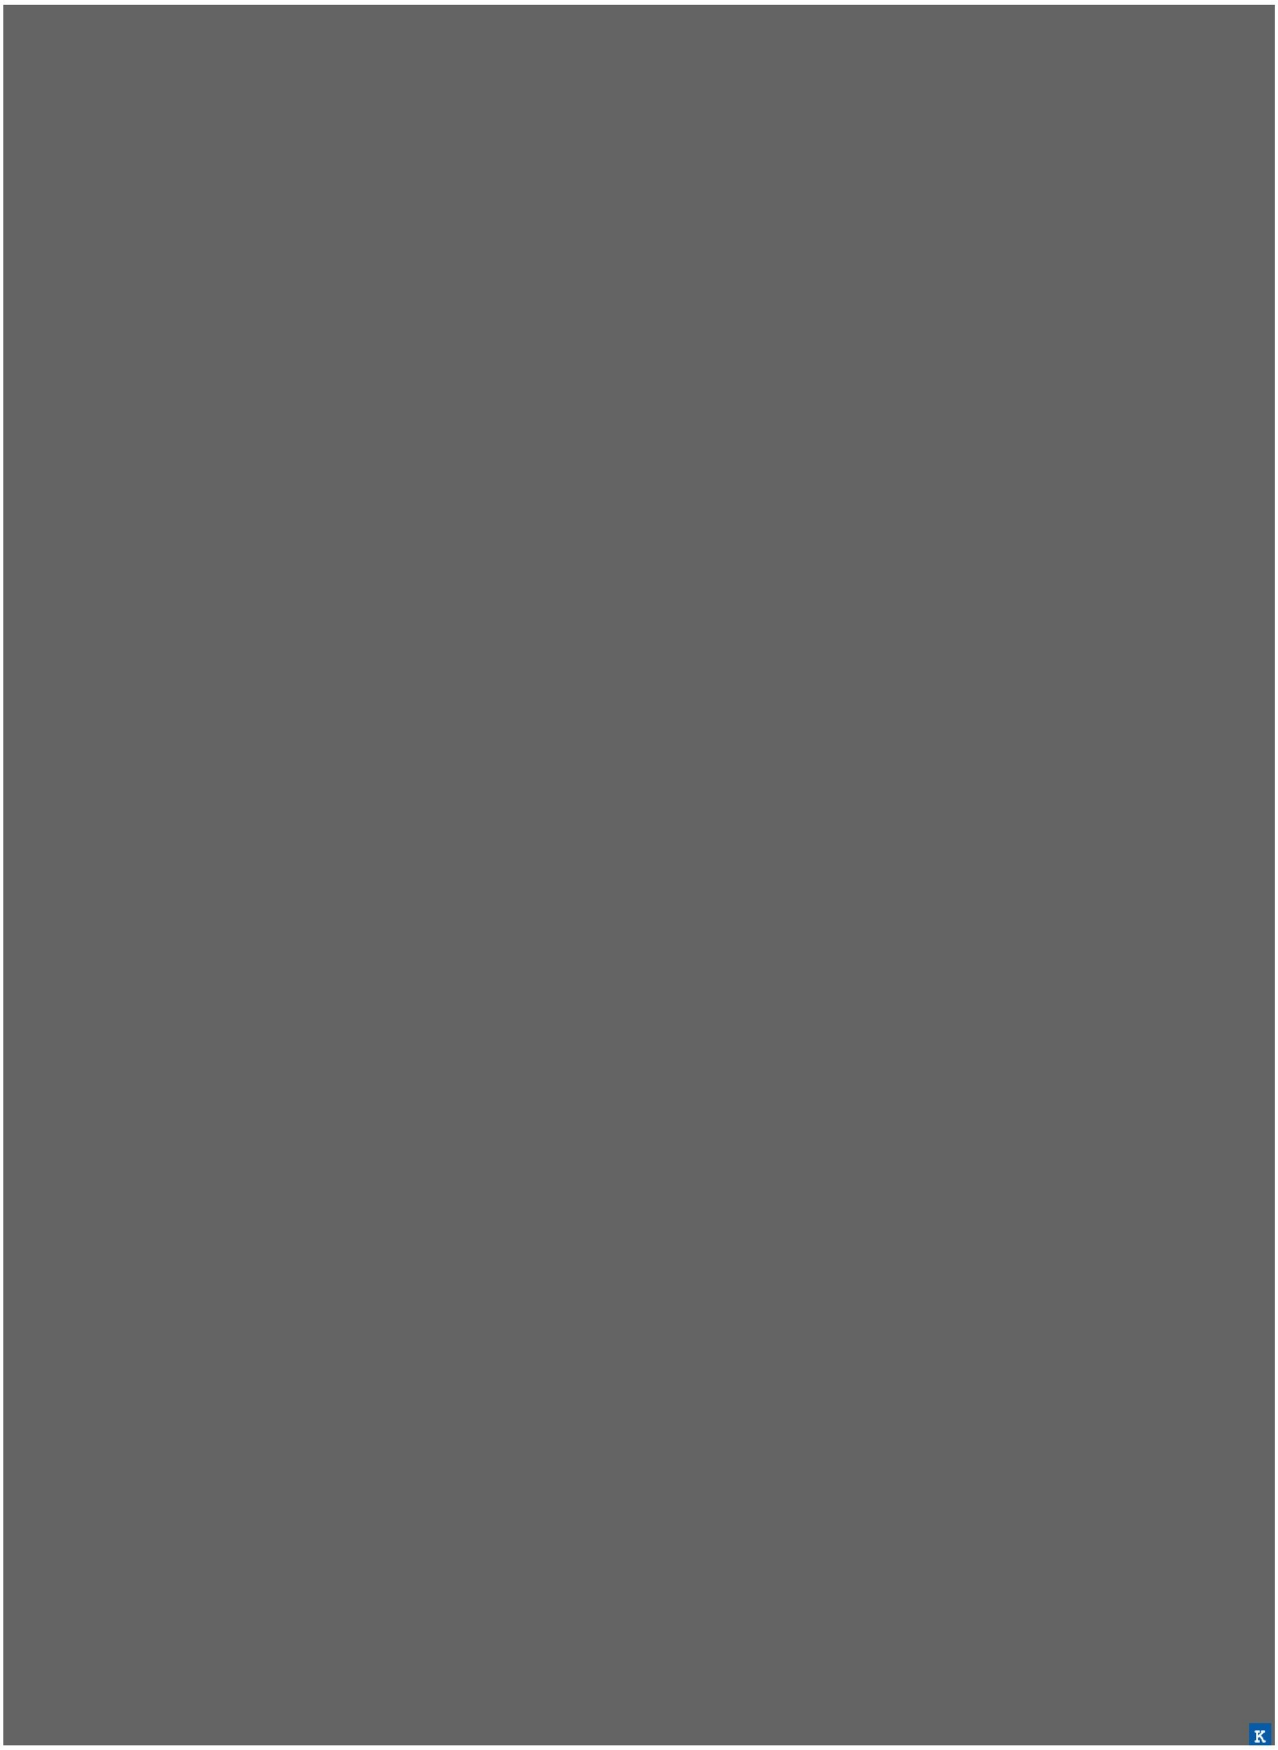

[redacted]

7 Financial planning

7.1 Overview of the budget

[redacted]

[redacted]

8 Diagram Impact pathway and indicators of the research project

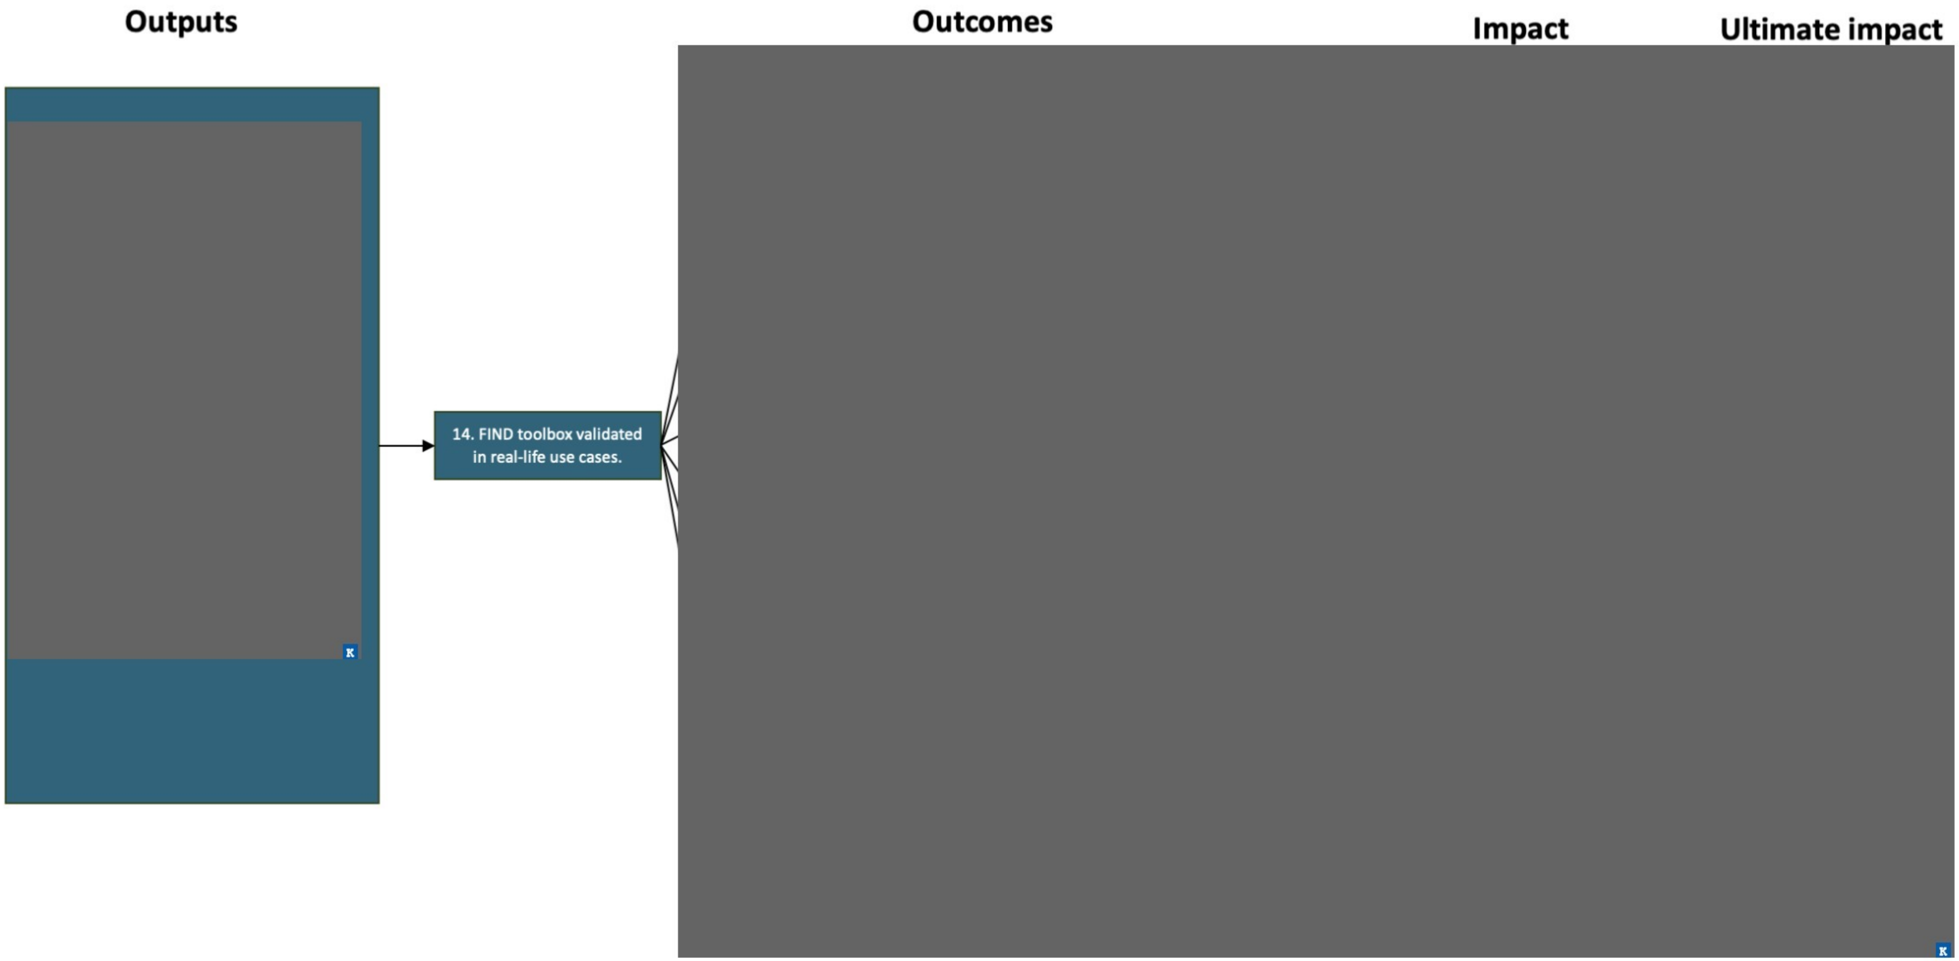

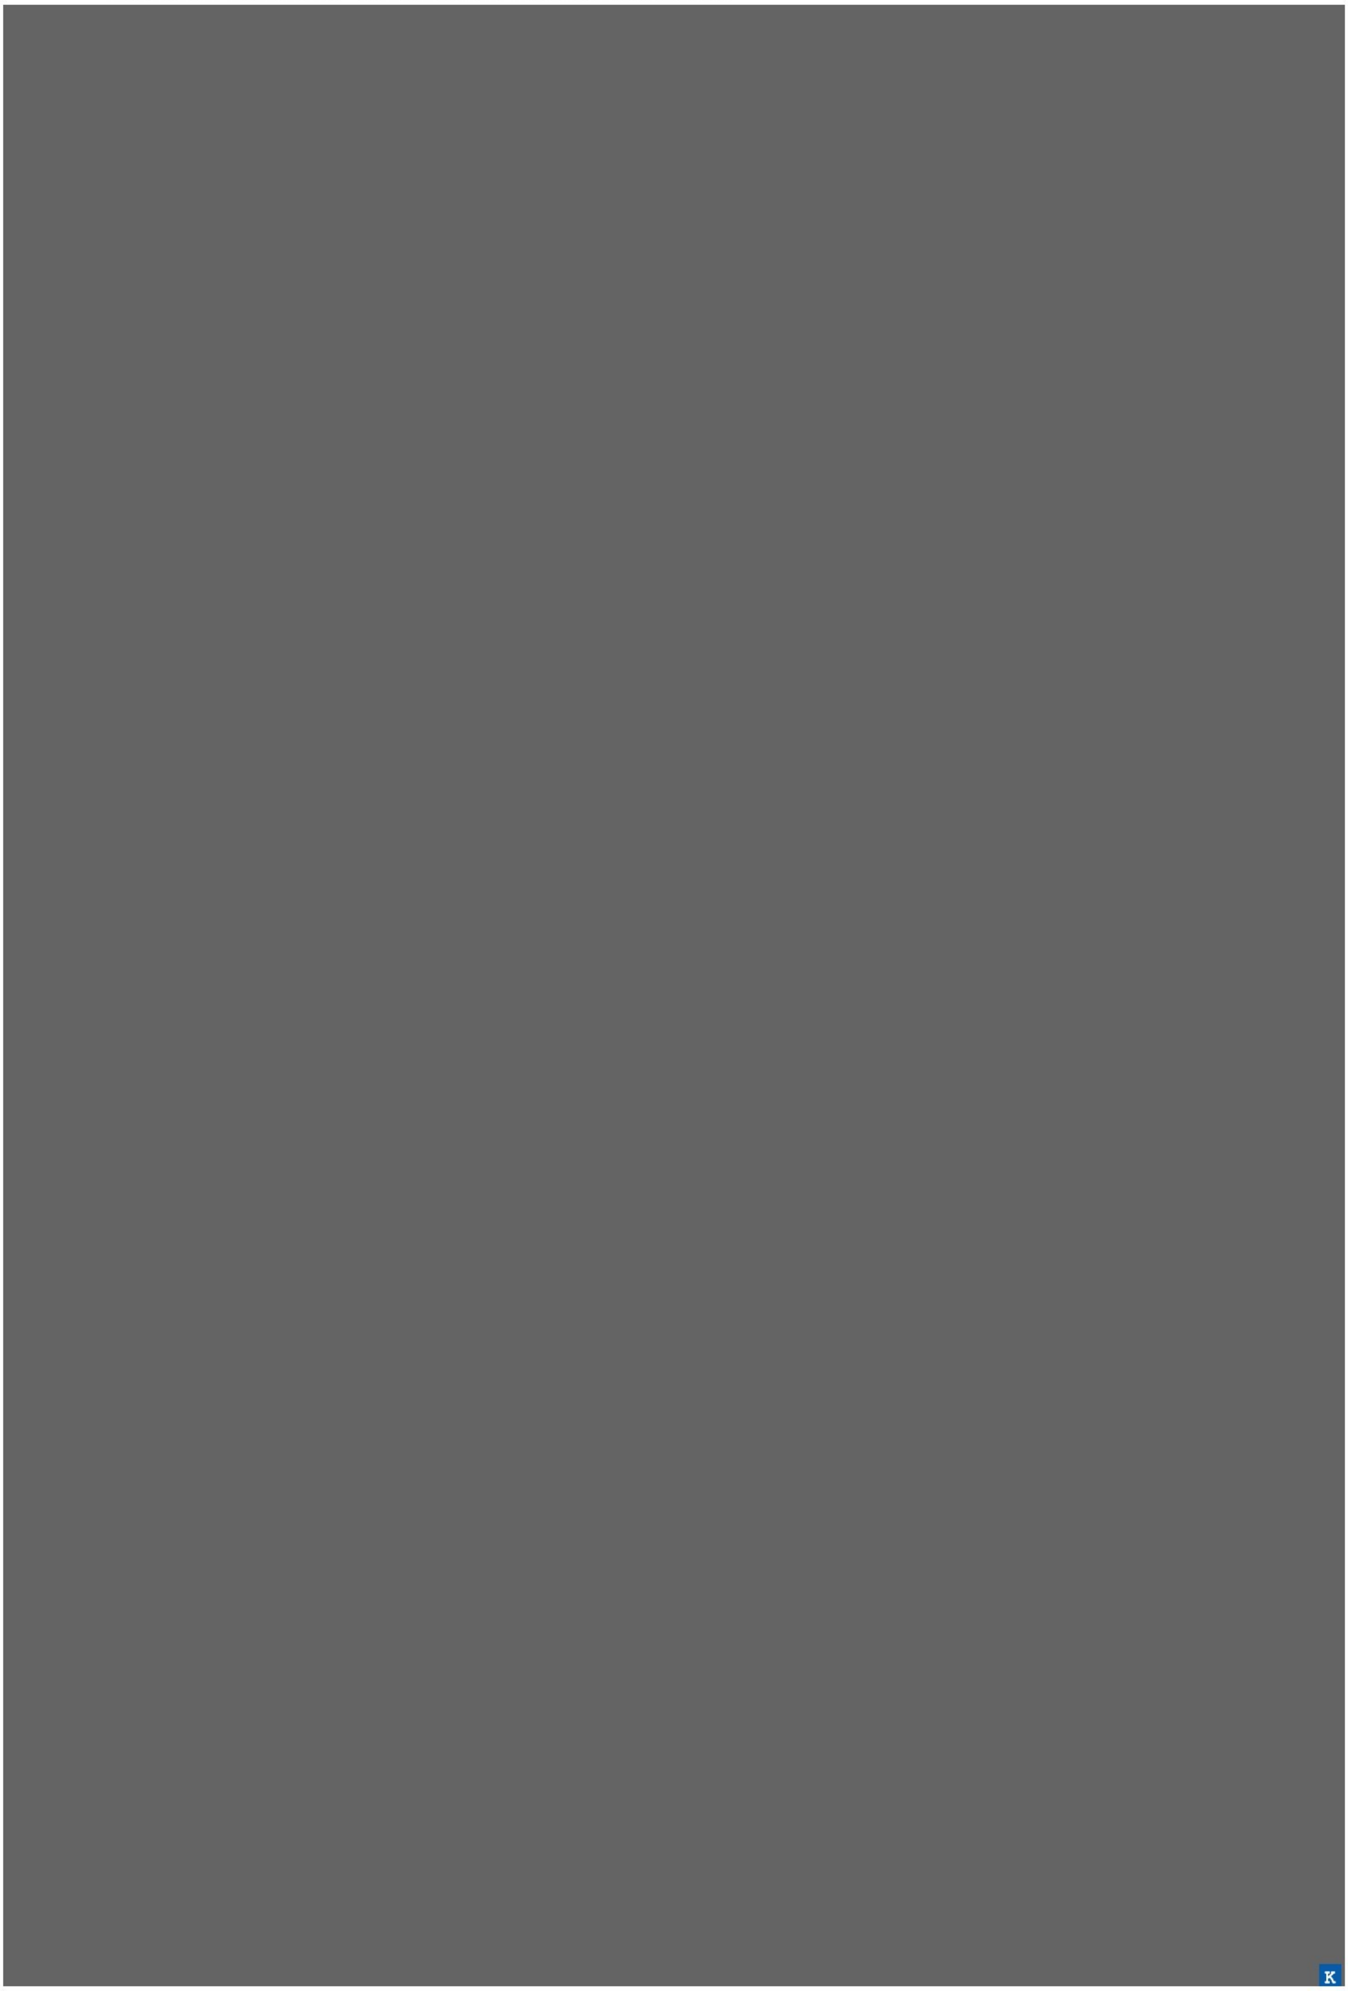

9 Annexes

9.1 Selection of research groups' key publications and patents

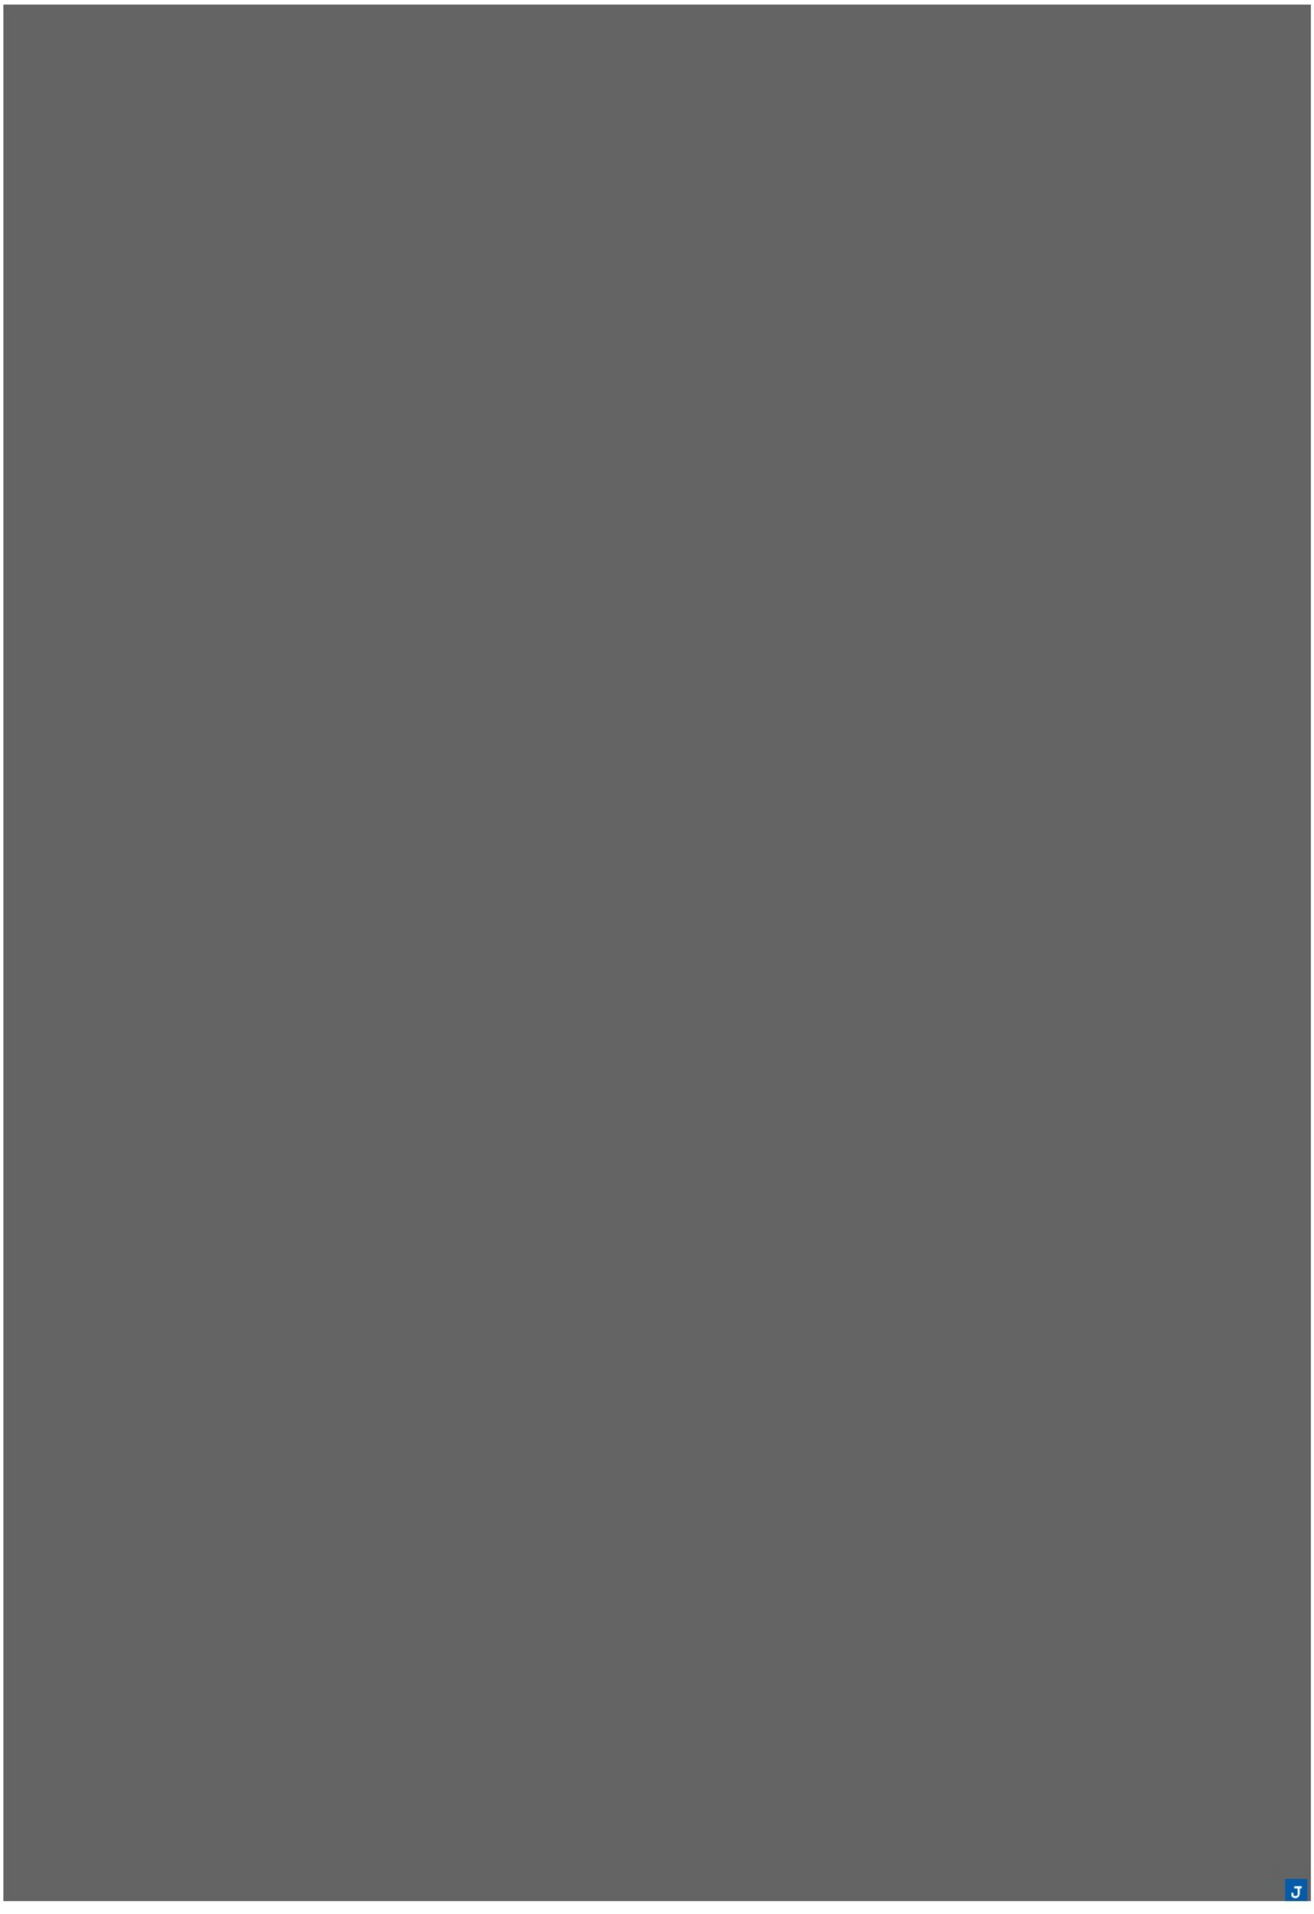

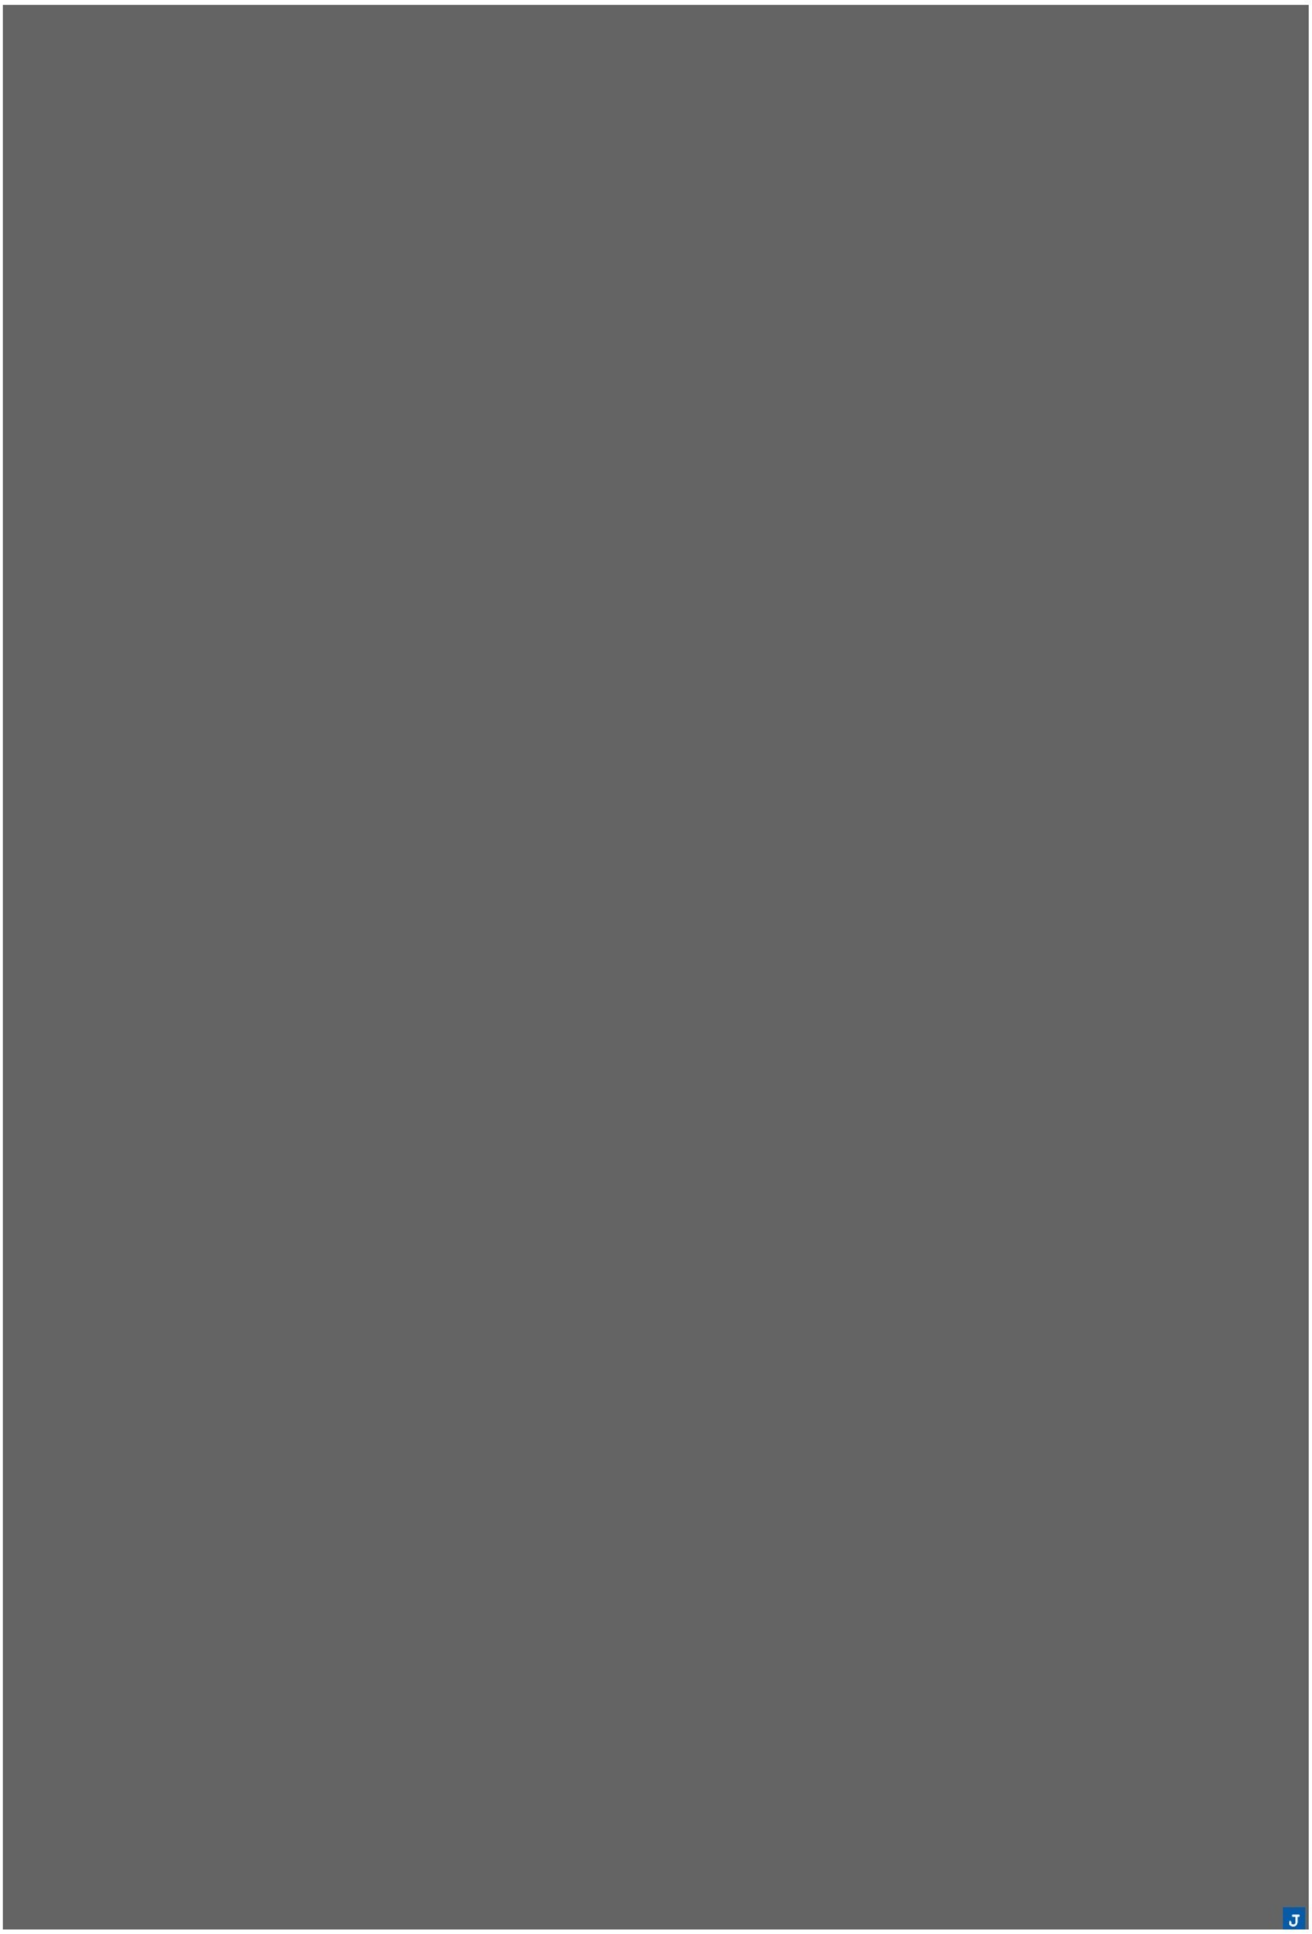

Foundation for Industry (FIND) - Large AI Models for a resilient high-tech industry.

9.2 Abbreviations and acronyms

Artificial Intelligence	AI
High Performance Computing	HPC
Horizon Europe	HE
Natural Language Processing	NLP
Technology Readiness Level	TRL
United Nations	UN
Work Package	WP

9.3 References

[1] Schneider, J., Meske, C., & Kuss, P. (2024). Foundation Models: A New Paradigm for Artificial Intelligence. Business & Information Systems Engineering, 1-11.

[2] “On the Opportunities and Risks of Foundation Models”, R. Bommasani et al., ArXiv, https://crfm.stanford.edu/assets/report.pdf, 2021.

[3] Kolides et al., “Artificial intelligence foundation and pre-trained models: Fundamentals, applications, opportunities, and social impacts”, Simulation Modelling Practice and Theory, Vol: 126, 2023, https://doi.org/10.1016/j.simpat.2023.102754.

[4] Chaka, C. (2023). Fourth industrial revolution—a review of applications, prospects, and challenges for artificial intelligence, robotics and blockchain in higher education. Research and Practice in Technology Enhanced Learning, 18, 002-002.

[5] “Harnessing the economic dividends from demographic change”, Frontier Technology Issues, Department of Economic & Social Affairs, United Nations, July 2023.

[6] https://www.cbs.nl/en-gb/news/2022/19/dutch-digital-skills-at-the-top-in-europe.

[7] “The impact of the technological revolution on labour markets and income distribution”, Frontier Technology Issues, Department of Economic & Social Affairs, United Nations, July 2017.

[8] “Generative AI and jobs: A global analysis of potential effects on job quantity and quality”, ILO Working Paper 96, International Labor Organization, August 2023.

[9] https://efficientdeeplearning.nl/

[10] Katherine Evans, Nelson de Moura, Stéphane Chauvier, Raja Chatila and Ebru Dogan, “Ethical Decision Making in Autonomous Vehicles: The AV Ethics Project”, Science and Engineering Ethics, SpringerLink, Volume 26, pages 3285–3312, 2020.

[11] “A Comprehensive Survey on Pretrained Foundation Models: A History from BERT to ChatGPT”, ArXiv, https://arxiv.org/abs/2302.09419, 2023.

[12] “Emerging trends: A gentle introduction to fine-tuning”, Natural Language Engineering, https://doi.org/10.1017/S1351324921000322, Volume 27, Issue 6, November 2021, pp. 763-778.

[13] T. Brown et al., “Language Models are Few-Shot Learners”, Advances in Neural Information Processing Systems (NeurIPS), 2020.

[14] J. Wei et al.,”Emergent Abilities of Large Language Models”, https://arxiv.org/abs/2206.07682.

[15] “Multimodal Foundation Models: From Specialists to General-Purpose Assistants”, C. Li et al., https://arxiv.org/abs/2309.10020, 2023.

[16] “Few-shot Adaptation of Multi-modal Foundation Models: A Survey”, F. Liu et al., https://doi.org/10.21203/rs.3.rs-3497100/v1, 2023.

[17] “A Study of the Generalizability of Self-Supervised Representations”, A. Tendle et al., Machine Learning Applications, https://doi.org/10.1016/j.mlwa.2021.100124, 2021.

[18] T. Kerrsies et al.,” How to Benchmark Vision Foundation Models for Semantic Segmentation?”, https://arxiv.org/abs/2404.12172.

[19] A. Vaswani et al. “Attention Is All You Need”, Advances in Neural Information Processing Systems (NeurIPS), 2017.

[20] "Training Compute-Optimal Large Language Models", J. Hoffmann et al., Neural Information Processing Systems (NeurIPS), https://arxiv.org/abs/2203.15556, 2022.

[21] “Reproducible Scaling Laws for Contrastive Language-Image Learning”, M. Cherti et al., Computer Vision and Pattern Recognition (CVPR), https://arxiv.org/abs/2212.07143, 2023, pp. 2818-2829.

[22] "data2vec: A General Framework for Self-supervised Learning in Speech, Vision and Language", A. Baevski et al., Proceedings of Machine Learning Research (PMLR), https://arxiv.org/abs/2202.03555, 2022.

[23] L. Zhuang, T. Jiang, J. Wang, Q. An, K. Xiao and A. Wang, "Effective mmWave Radar Object Detection Pretraining Based on Masked Image Modeling," in IEEE Sensors Journal, vol. 24, no. 3, pp. 3999-4010, 1 Feb.1, 2024, doi: 10.1109/JSEN.2023.3339651.

[24] European Commission. (2024). AI factories. Shaping Europe’s Digital Future. https://digital-strategy.ec.europa.eu/en/policies/ai-factories

[25] T. Moutakanni et al., “Advancing human-centric AI for robust X-ray analysis through holistic self-supervised learning”, https://arxiv.org/abs/2405.01469, 2024.

[26] Z Xu, Z Shi, J Wei, F Mu, Y Li, Y Liang, “Towards Few-Shot Adaptation of Foundation Models via Multitask Finetuning”, International Conference on Learning Representations (ICLR), 2024.

[27] Thieme, A., Nori, A., Ghassemi, M., Bommasani, R., Andersen, T. O., & Luger, E. (2023, April). Foundation models in healthcare: Opportunities, risks & strategies forward. In: 2023 CHI Conference on Human Factors in Computing Systems (pp. 1-4).

[28] Li, J., Dada, A., Puladi, B., Kleesiek, J., & Egger, J. (2024). ChatGPT in healthcare: a taxonomy and systematic review. Computer Methods and Programs in Biomedicine, 108013.

[29] US Food and Drug Administration. (2023). Artificial intelligence and machine learning (AI/ML)-enabled medical devices. AI/ML-Enabled Medical Devices.

[30] Ghassemi, M., Naumann, T., Schulam, P., Beam, A. L., Chen, I. Y., & Ranganath, R. (2020). A review of challenges and opportunities in machine learning for health. AMIA Summits on Translational Science Proceedings, 2020.

[31] Cui, C., Yang, H., Wang, Y., Zhao, S., Asad, Z., Coburn, L. A., ... & Huo, Y. (2023). Deep multimodal fusion of image and non-image data in disease diagnosis and prognosis: a review. Progress in Biomedical Engineering, 5(2), 022001.

[32] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE international conference on computer vision (pp. 618-626).

[33] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in neural information processing systems, 30.

[34] Mittermaier, M., Raza, M. M., & Kvedar, J. C. (2023). Bias in AI-based models for medical applications: challenges and mitigation strategies. NPJ Digital Medicine, 6(1), 113.

[35] Dwork, C. (2006). Differential privacy. In International colloquium on automata, languages, and programming (pp. 1-12). Berlin, Heidelberg: Springer Berlin Heidelberg.

[36] Michael Mesarcik, Albert-Jan Boonstra, Marco Iacobelli, Elena Ranguelova, Cees de Laat, and Rob V. van Nieuwpoort: The ROAD to discovery: machine learning-driven anomaly detection in radio astronomy spectrograms. Astronomy & Astrophysics 680, A74, December 2023. DOI: https://doi.org/10.1051/0004-6361/202347182.

[37] Michael Mesarcik, Elena Ranguelova, Albert-Jan Boonstra, Rob V van Nieuwpoort:
Improving novelty detection using the reconstructions of nearest neighbours.
Array journal, Volume 14, July 2022. DOI: https://doi.org/10.1016/j.array.2022.100182.

[38] Nguyen, Duy Minh Ho, et al. "On the Out of Distribution Robustness of Foundation Models in Medical Image Segmentation." NeurIPS 2023 Workshop on Robustness of Few-shot and Zero-shot Learning in Foundation Models.

[39] A. P. Ruiz, M. Flynn, J. Large, M. Middlehurst, and A. Bagnall. The Great Multivariate Time Series Classification Bake off: A Review and Experimental Evaluation of Recent Algorithmic Advances. Data Mining and Knowledge Discovery. vol. 35. no. 2, pp. 401–449. 2021.

[40] L. Pantiskas, K. Verstoep, M. Hoogendoorn, and H. Bal. Reinforcement Learning-Guided Channel Selection Across Time for Multivariate Time Series Classification. In IEEE Symposium Series on Computational Intelligence (SSCI). Dec. 2023.

[41] Ji, Yanrong and Zhou, Zhihan and Liu, Han and Davuluri, Ramana V: DNABERT: pre-trained Bidirectional Encoder Representations from Transformers model for DNA-language in genome, Bioinformatics: vol 37 no 15, pages 2112-2120, https://doi.org/10.1093/bioinformatics/btab083, 2021

[42] Akcay, S., Atapour-Abarghouei, A., Breckon, T.P. (2019). GANomaly: Semi-supervised Anomaly Detection via Adversarial Training. Computer Vision – ACCV 2018. ACCV 2018. Lecture Notes in Computer Science, vol 11363. https://doi.org/10.1007/978-3-030-20893-6_39

[43] Emiel Hoogeboom, Alexey A. Gritsenko, Jasmijn Bastings, Ben Poole, Rianne van den Berg, Tim Salimans, “Autoregressive Diffusion Models”, International Conference on Learning Representation (ICLR), 2022.

[44] Lingkai Kong, Jiaming Cui, Haotian Sun, Yuchen Zhuang, B. Aditya Prakash, Chao Zhang, “Autoregressive Diffusion Model for Graph Generation”, International Conference on Machine Learning (ICML), 2021.

[45] Waldschmidt, C., Hasch, J., & Menzel, W., “Automotive radar—From first efforts to future systems”, IEEE Journal of Microwaves, vol. 1, nr. 1, 135-148, 2021.

[46] Zhou, Y., Liu, L., Zhao, H., López-Benítez, M., Yu, L., & Yue, Y., “Towards deep radar perception for autonomous driving: Datasets, methods, and challenges”, Sensors, vol. 22, nr. 11, 4208, 2024.

[47] Queralta, J. P., Taipalmaa, J., Pullinen, B. C., Sarker, V. K., Gia, T. N., Tenhunen, H., and others, "Collaborative multi-robot search and rescue: Planning, coordination, perception, and active vision", IEEE Access, vol. 8, 191617-191643, 2020.
[48] Peng, S., Genova, K., Jiang, C., Tagliasacchi, A., Pollefeys, M., & Funkhouser, T., "OpenScene: 3d scene understanding with open vocabularies", In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 815-824, 2023.
[49] Ruan, T., Wang, H., Stolkin, R., & Chiou, M., "A taxonomy of semantic information in robot-assisted disaster response", IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), pp. 285-292, 2022.
[50] de Geus, D., Meletis, P., Lu, C., Wen, X., & Dubbelman, G., "Part-aware panoptic segmentation", In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 5485-5494, 2021.
[51] Won, Myounggyu, "Intelligent traffic monitoring systems for vehicle classification: A survey", IEEE Access, vol. 8, 73340-73358, 2022.
[52] Badidi, Elarbi, Karima Moumane, and Firdaous El Ghazi, "Opportunities, applications, and challenges of edge-AI enabled video analytics in smart cities: a systematic review", IEEE Access, 2023.
[53] M. Oquab et al., "DINOv2: Learning Robust Visual Features without Supervision", Transactions on Machine Learning Research, 2024.
[54] H. Touvron et al., "Llama 2: Open Foundation and Fine-Tuned Chat Models", arXiv preprint arXiv:2307.09288, 2023.
[55] Meta AI, "Llama 3 Model Card", https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md (accessed 2024-05-08).
[56] Y. Zhang, H. Doughty, C.G.M. Snoek,"Low-Resource Vision Challenges for Foundation Models", In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024.
[57] J. M. Johnson & T.M. Khoshgoftaar, "Survey on deep learning with class imbalance", Journal of Big Data, 6(1), 1-54, 2019.
[58] A. Redford et al., "Learning Transferable Visual Models From Natural Language Supervision", International Conference on Machine Learning, pp. 8748-8763, 2021.
[59] K. He, et al., "Masked Autoencoders Are Scalable Vision Learners", In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022.
[60] Y. Tian, et al. "Learning Vision from Models Rivals Learning Vision from Data". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024.
[61] Z. Wang,, et al.. Learning to prompt for continual learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022.
[62] M. Wójcik, M. et al. Domain-Agnostic Neural Architecture for Class Incremental Continual Learning in Document Processing Platform. In Proceedings of the Association for Computational Linguistics, 2023
[63] R. Girdhar, et al."ImageBind: One Embedding Space To Bind Them All", In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023.
[64] Y. Zhang, H. Doughty, C.G.M. Snoek: Learning Unseen Modality Interaction. In: NeurIPS, 2023.
[65] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, C. Schmid, "ViViT: A Video Vision Transformer", in Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2021, pp. 6836-6846.
[66] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale", in International Conference on Learning Representations, 2021.
[67] C. Lu, D. de Geus, and G. Dubbelman, "Content-aware Token Sharing for Efficient Semantic Segmentation With Vision Transformers", in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023, pp. 23631-23640.
[68] N. Norouzi, S. Orlova, D. de Geus, and G. Dubbelman, "ALGM: Adaptive Local-then-Global Token Merging for Efficient Semantic Segmentation with Plain Vision Transformers", in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024, to appear.
[69] D. Bolya, C. Fu, X. Dai, P. Zhang, C. Feichtenhofer, and J. Hoffman, "Token Merging: Your ViT but Faster", in International Conference on Learning Representations, 2023.
[70] "Distilling Large Vision-Language Model with Out-of-Distribution Generalizability", X. Li et al., International Conference on Computer Vision (ICCV), https://arxiv.org/abs/2307.03135, 2023.
[71] "Pushing Large Language Models to the 6G Edge: Vision, Challenges, and Opportunities", Z. Lin et al., https://arxiv.org/abs/2309.16739, 2023.

[72] "QLoRA: Efficient Finetuning of Quantized LLMs", T. Dettmers et al., https://arxiv.org/abs/2305.14314, 2023.
[73] "GPTQ: Accurate Post-Training Quantization for Generative Pre-trained Transformers", E. Frantar et al., International Conference on Learning Representations (ICLR), https://arxiv.org/abs/2210.17323, 2023.
[74] Singh, G., Diamantopoulos, D., Hagleitner, C., Gómez-Luna, J., Stuijk, S., Mutlu, O., & Corporaal, H. (2020). NERO: A near High-Bandwidth Memory Stencil Accelerator for Weather Prediction Modeling. In N.
[75] Mentens, L. Sousa, P. Trancoso, M. Pericas, & I. Sourdis (Eds.), 2020 30th International Conference on Field-Programmable Logic and Applications (FPL) (pp. 9-17). Article 9221526 Institute of Electrical and Electronics Engineers. https://doi.org/10.1109/FPL50879.2020.00014
[76] Gu et al., Efficiently modeling long sequences with structured state spaces, arXiv 2023.
[77] M. Beck et al. xLSTM: Extended Long Short-Term Memory, arXiv 2024.
[78] Zhu et al., Vision Mamba: Efficient Visual Representation Learning with Bidirectional State Space Model, arXiv 2024.
[79] Liu et al., Vmamba: Visual state space model, arXiv 2024.
[80] CHIPS JU, https://www.chips-ju.europa.eu/.
[81] Abadi, Martin, et al. "Deep learning with differential privacy." Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, 2016.
[82] Baevski, Alexei, et al. "wav2vec 2.0: A framework for self-supervised learning of speech representations." Advances in Neural Information Processing Systems (NeurIPS), 2020.
[83] Brown, Tom, et al. "Language models are few-shot learners." Advances in Neural Information Processing Systems (NeurIPS), 2020.
[84] Carlini, Nicholas, et al. "Extracting training data from large language models." 30th USENIX Security Symposium, 2021.
[85] Carlini, Nicholas, et al. "Extracting training data from diffusion models." 32nd USENIX Security Symposium, 2023.
[86] Das, Abhimanyu, et al. "A decoder-only foundation model for time-series forecasting." arXiv preprint arXiv:2310.10688, 2023.
[87] Delétang, Grégoire, et al. "Language modeling is compression." arXiv preprint arXiv:2309.10668, 2023
[88] Devlin, Jacob, et al. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2019.
[89] Garza, Azul, and Max Mergenthaler-Canseco. "TimeGPT-1." arXiv preprint arXiv:2310.03589, 2023.
[90] He, Kaiming, et al. "Masked autoencoders are scalable vision learners." Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022.
[91] Jiang, Zhiying, et al. "Inserting Information Bottlenecks for Attribution in Transformers." Findings of the Association for Computational Linguistics: EMNLP, 2020.
[92] Li, Hongming, Shujian Yu, and Jose Principe. "Causal recurrent variational autoencoder for medical time series generation." Proceedings of the AAAI Conference on Artificial Intelligence, 2023.
[93] Li, Qintong, et al. "Explanation Regeneration via Information Bottleneck." Findings of the Association for Computational Linguistics: ACL, 2023.
[94] Mesarcik, Michael, et al. "The ROAD to discovery: Machine-learning-driven anomaly detection in radio astronomy spectrograms." Astronomy & Astrophysics, 2023
[95] Nam, Hyeonseob, and Bohyung Han. "Learning multi-domain convolutional neural networks for visual tracking." Proceedings of the IEEE conference on computer vision and pattern recognition, 2016
[96] Nguyen, Duy Minh Ho, et al. "On the Out of Distribution Robustness of Foundation Models in Medical Image Segmentation." NeurIPS Workshop on Robustness of Few-shot and Zero-shot Learning in Foundation Models. 2023
[97] Singh, Rishabh, and Jose Principe. "Time series analysis using a kernel based multi-modal uncertainty decomposition framework." Conference on Uncertainty in Artificial Intelligence. PMLR, 2020
[98] Sun, Meiqi, et al. "Quantifying uncertainty in foundation models via ensembles." NeurIPS 2022 Workshop on Robustness in Sequence Modeling, 2022.
[99] Triastcyn, Aleksei, and Boi Faltings. "Bayesian differential privacy for machine learning." International Conference on Machine Learning, 2020.
[100] Wei, Jason, et al. "Emergent abilities of large language models." Transactions on Machine Learning Research, 2022.

[101] Yu, Shujian, et al. "Cauchy-Schwarz Divergence Information Bottleneck for Regression." International Conference on Learning Representations (ICLR), 2024

[102] Yu, Shujian, et al. "The Conditional Cauchy-Schwarz Divergence with Applications to Time-Series Data and Sequential Decision Making." arXiv preprint arXiv:2301.08970, 2023

[103] Yu, Yaodong, et al. "ViP: A Differentially Private Foundation Model for Computer Vision." arXiv preprint arXiv:2306.08842, 2023. [24] Zhao, Haiyan, et al. "Explainability for large language models: A survey." ACM Transactions on Intelligent Systems and Technology, 2024

[104] Zheng, Kaizhong, et al. "BPI-GNN: Interpretable brain network-based psychiatric diagnosis and subtyping." NeuroImage, 2024.

[105] Snellius national supercomputer specs, https://servicedesk.surf.nl/wiki/display/WIKI/Snellius+hardware, SURF, 2024.

9.4 Overview budget project

PHASE 3	**Project proposals**														
Project:	**P23.016**														
Title:	**Foundation for Industry (FIND): Large AI models for a resilient high-tech industry**														
		Budget									Income				
	Main applicant	Personnel	Materials*	Investments*	Knowledge utilisation*	Inter-nationalisation & Money follows Cooperation*	Project management*	Total expenditures (applicants)	In kind contribution	Total project costs (expenditures + in kind)	Cash cofinancing	In kind contribution	Total cofinancing (in kind + cash)	% total cofinancing (3)	NWO contribution
TOTAL project cost	[redacted]	[redacted]													
Summary expenditures															
Expenditures universities/UMC/NWO-i/etc.	[redacted]														
Expenditures TO2 instututes															
Expenditures Universities of applied sciences															
Total expenditures															

9.5 Declarations of co-funding

Company / Organization	In-cash co-funding	In-kind co-funding (monetized in-kind hours)
ASML	[redacted]	[redacted]
ASMPT	[redacted]	[redacted]
Brainport Development	[redacted]	[redacted]
Canon	[redacted]	[redacted]
KAIKO	[redacted]	[redacted]
KeyGene	[redacted]	[redacted]
NKI	[redacted]	[redacted]
NXP	[redacted]	[redacted]
PercivAI	[redacted]	[redacted]
Shell	[redacted]	[redacted]
Signify	[redacted]	[redacted]
Stryker	[redacted]	[redacted]
Technolution	[redacted]	[redacted]
TNO	[redacted]	[redacted]

Toelichting grondslagen

In dit document kunt u secties vinden die onleesbaar zijn gemaakt. Deze informatie is achterwege gelaten op basis van de Wet open overheid (Woo). De letter die hierbij is vermeld correspondeert met de bijbehorende grondslag in onderstaand overzicht.

G Art. 5.1 lid 2 sub b

Het belang van de openbaarmaking van deze informatie weegt niet op tegen het belang van de economische of financiële belangen van de Staat, andere publiekrechtelijke lichamen of bestuursorganen

J Art. 5.1 lid 2 sub e

Het belang van de openbaarmaking van deze informatie weegt niet op tegen het belang van de eerbiediging van de persoonlijke levenssfeer van betrokkenen

K Art. 5.1 lid 2 sub f

Het belang van de openbaarmaking van deze informatie weegt niet op tegen het belang van de bescherming van andere dan in art. 5.1 lid 1 sub c genoemde concurrentiegevoelige bedrijfs- en fabricagegegevens