ELISE Guide for Applicants (GfA)

ELISE 2nd Open Call

Open date for proposals: 8 December 2022 at 13:00 CET (Brussels Time).
Deadline: 16 February 2023 at 13:00 CET (Brussels Time).

Version 1/2022

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951847.

1. Basic Info about ELISE **3**

2. What do we offer? **3**

3. Eligibility Criteria **4**

3.1 Eligible applicants 4
3.2 What types of activities can be funded? 4
3.3 Ideal Project 5
3.4 How to apply? 7

4. How will we evaluate your proposal? **9**

4.1 Eligibility Check 9
4.2 Pre-scoring 9
4.3 External Evaluation 10
4.4 Consensus Meeting 12
4.5 Jury Day 13
4.6 What's next? Sub Grant Agreement Preparation and Signature 13

5. Our Support and Payment Arrangements **14**

6. Last but not least - final provisions **16**

Annex 1: Information Clause **17**

Annex 2: ELISE Focus areas **19**

1. Basic Info about ELISE

ELISE [European Network of AI Excellence Centres] is a network of artificial intelligence research hubs. Based on the highest level of research, it spreads its knowledge and methods in academia, industry and society.
ELISE aims to make Europe competitive in AI through a network of excellence. The best European researchers in machine learning and AI work together to attract talent, to foster research through collaboration, and to inspire and be inspired by industry and society.
ELISE aims to support SMEs by:

- Incorporating SMEs and start-ups into their **machine learning [ML] ecosystem,**
- facilitating and accelerating a **broad uptake of machine learning technologies** by SMEs and start-ups,
- ensuring that every SME and start-up has access to ML technologies and expertise within "working distance".

ELISE builds upon the ELLIS (European Laboratory for Learning and Intelligent Systems) network. ELISE Consortium is composed of 23 organisations and coordinated by the Aalto University.

This document summarises the main points of the ELISE 2nd Open Call under the ELISE project, which will be open from **8 December 2022 at 13:00 CET (Brussels Time) with a deadline on 16 February 2023 at 13:00 CET (Brussels Time)**.

Where can you find key information regarding this Open Call?

- ELISE website for general information about the program,

- If you have any technical problem or doubt when filling in the online Application Form at FundingBox, tell us directly at info.elise@fundingbox.com

- Application website, where you can find all documents: https://elise2oc.fundingbox.com/

2. What do we offer?

ELISE will select **16 SMEs that develop novel AI-based services or applications** with the support of a Scientific Advisor in one of the ELISE/ELLIS focus areas (described in Section 3.2).

Selected companies will enter the **6-month** program and receive:

- **Up to € 60.000 (lump sum)** to develop novel AI-based applications in one of ELISE focus areas and
- Visibility through ELISE online channels and events, and dissemination in the ELLIS community.

Apply here: https://elise2oc.fundingbox.com/

3. Eligibility Criteria

3.1 Eligible applicants

We are looking for **SMEs**[1]**.**

Applicants, at the moment of applying, must be officially **registered legal entities** established in:

- The Member States of the European Union and its Overseas Countries and Territories (OCT) or
- Associated Countries to H2020 or
- The United Kingdom.

The ELISE partners, or their affiliates or employees, are NOT considered as eligible applicants and can NOT apply for funding.

3.2 What types of activities can be funded?

ELISE is looking for projects based on AI applications that address the development and implementation of technology and systems in at least one **focus area** listed below. Focus areas are based on ELISE/ELLIS research programs but are not limited

[1] An **SME** will be considered as such if it complies with the European Commission's Recommendation 2003/361/EC. As a summary, the criteria defining an SME are An **SME** will be considered as such if it complies with the European Commission's Recommendation 2003/361/EC. As a summary, the criteria defining an SME are:
• Headcount in Annual Work Unit (AWU) less than 250;
• Annual turnover less or equal to €50 million OR annual balance sheet total less or equal to €43 million.
Note that the figures of partners and linked enterprises should also be considered as stated in the SME user guide. For detailed information check EU recommendation:
https://ec.europa.eu/growth/smes/business-friendly-environment/sme-definition_en

to them as long as the proposals will address high-impact societal and economic challenges using machine learning.

Focus areas:

- Geometric Deep Learning: graph, group and gauge convolutions
- Robust Machine Learning
- Interactive Learning and Interventional Representations
- Machine Learning and Computer Vision
- Robot Learning: Closing the Reality Gap
- Human-centric Machine Learning
- Theory, Algorithms and Computations of Modern Learning Systems
- Quantum and physics based Machine Learning
- Natural Intelligence
- Health
- Machine Learning in Earth and Climate Sciences
- Natural Language Processing (NLP)
- Multimedia / multimodal information
- Information retrieval
- Explainability and Fairness in Data Mining
- Symbolic Machine Learning

Description of focus areas is available in **Annex 2**.
Additional information are also included here: https://ellis.eu/programs

3.3 Ideal Project

The ideal project is the use case of machine learning technologies by SME (startup) that fits in at least one of the focus areas.

The preferable initial TRL (Technology Readiness Level) should be **TRL[2] 6-7** while the initial IRL (Investment) should be **IRL[3] 6-7**, but applications at lower TRL and IRL can be accepted as long as the applicant presents the road map to reach the higher TRL level.

[2] TRL, Technology Readiness Level. Technology Readiness Levels (TRLs) are indicators of the maturity level of particular technologies. This measurement system provides a common understanding of technology status and addresses the entire innovation chain.
There are nine technology readiness levels; TRL 1 being the lowest and TRL 9 the highest. In our project we refer to Annex G of the General Annexes to the Work Programme 2016/17 for a full description of TRLs.
[3] Investment Readiness Level is a scale that runs from 1 to 9 and defines the different stages of a start-ups maturity in relation to value creation and preparedness for commercialization. Investment Readiness can be applied to any type of product, service, technology, and social innovation.
IRL is an evidence based demonstration to investors, that there's a repeatable and scalable business model.

The projects will be supported by the Scientific Advisor, who creates the link between industry and research. A successful project can lead to further collaboration between the SME and the Scientific Advisor. Having a Scientific Advisor is highly recommended at the application stage and it is mandatory for participation in Jury Day.

Who is a Scientific Advisor?

The Scientific Advisor is a **researcher invited by the applicant**, who agrees to advise the SME (startup) in developing the AI application or service. The Advisor must have a PhD in a field that is relevant to the project.

The advisor must be independent and cannot have any formal links with the applicant (NDA, non-disclosure agreement will not be considered as a formal link).

You can invite the Scientific Advisor at any time, before applying to the Open Call or during the evaluation process, but not later than before the Jury Day. We strongly advise having the Scientific Advisor at the moment of applying and **providing the Scientific Advisor statements in your application**, to strengthen your proposal during the evaluation process.

We highly recommend you to have the Scientific Advisor from ELLIS/ELISE community (https://ellis.eu/members), but this is not mandatory.

What is the role of a Scientific Advisor?

The role is strictly advisory, meaning that the advisor cannot participate in the project work, nor receive funding from it or any remuneration. Therefore (s)he cannot have any claims to IPR of the project.

More information about the role and obligations of the Scientific Advisor are in the **Frequently Asked Questions Document.**

To learn about the scope of potential projects, please check below a few examples of projects funded in the 1st Open Call.
More examples can be found here: https://www.elise-ai.eu/news-events.

A French startup developing medical devices for the diagnosis of skin cancers, which incorporate artificial intelligence (AI), wants to drastically improve the quality and timeliness of skin cancer diagnosis by practitioners. The company applied to the Open Call to develop a combination of 2 distinct networks that work together to

improve accuracy of the diagnostic tool through improved reliability and robustness, making the results interpretable and understandable.

A small company from Germany specialising in cyber security wants to develop an early-warning system for cyber risks involving automation that requires reliable, machine-readable information by turning incident and threat reports from text into machine-readable data sets. The company applied to the call to develop the early-warning system using Natural Language Processing and relevant industry standards to automate threat analysis from text documents.

An SME from Spain specialising in remote image capturing wants to develop a project dedicated to inspection of power lines with drones. The goal is to raise the state-of-art automating the camera orientation in long-range drones thus ensuring that all the power line towers are correctly captured. They applied to ELISE to implement ML detection using CNNs to address these issues and consult the process with Scientific Advisor.

iThermAI is a Belgian company dedicated to smart smoke and flame detection. They applied to ELISE to develop a state-of-the-art deep-tech technology to efficiently detect flame and smoke using inexpensive on-the-edge processing based on video surveillance cameras and deep learning models.

3.4 How to apply?

When applying to ELISE open call, please also note that:

- Your project should have a clear European Dimension meaning that you are proposing European Centric Solutions which will benefit the European society and economy.

- We will evaluate **only** proposals submitted through the online form available at https://elise2oc.fundingbox.com/ within the **deadline 16 Feb 2023, 13:00 CET (Brussels Time).** Upon receipt of each proposal, the system will send you a **confirmation** of your submission.

- You have to verify the completeness of the form, as it won't be possible to add any further information after the deadline. However, you will still be able to modify the form after the proposal until the deadline.

- You can submit only one proposal to this ELISE Open Call. If more than one proposal is identified, only the most recent (counting back from deadline) proposal will be evaluated.

- Entities funded in the ELISE 1st Open Call cannot receive support under ELISE 2nd Open Call.

- Your proposal must be written in **English** in all mandatory parts in order to be eligible. Only parts written in English will be evaluated.

- All mandatory sections of your proposal must be completed. The data provided should be actual, true, complete and should allow assessment of the proposal. Additional material, not specifically requested in the online application form, will not be considered for the evaluation. We recommend uploading the Scientific Advisor statements to your online application, but it's not mandatory. We will check the information provided in your application during the external evaluation and Sub Grant Agreement (SGA) preparation phases.

- We will take into consideration the existence of a potential **conflict of interest** between you and one or more of ELISE project's Consortium partners. Consortium partners, their affiliated entities, employees and permanent collaborators cannot take part in the ELISE Open Call. All cases of potential conflict of interest will be assessed case by case.

- Your proposal should address at least one of the focus areas listed in Section 3.2 or the proposal should present the groundbreaking and quality solution in terms of Machine Learning.

- We don't accept entities that are under liquidation or are an enterprise under difficulty according to the Commission Regulation No 651/2014, art. 2.18, or that are excluded from the possibility of obtaining EU funding under the provisions of both national and EU law, or by a decision of both national or EU authority.

- Your project should be based on your original work or your right to use the IPR must be clear. Going forward, any foreseen developments must be free from any third party rights, or those third-party rights must be clearly stated.

After the deadline, all applicants will be requested to fill in an ethics self-assessment form provided by ELISE. Applicants should complete the form within 1 month after the open call deadline (16 March 2022). Additional information about the ethics assessment can be found in the Frequently Asked Question Document.

ELISE plans online webinars about this open call. They will be announced on the ELISE website: https://www.elise-ai.eu/.

4. How will we evaluate your proposal?

Our evaluation process is transparent, fair and equal to all our participants. We will evaluate your project in 5 phases as shown below:

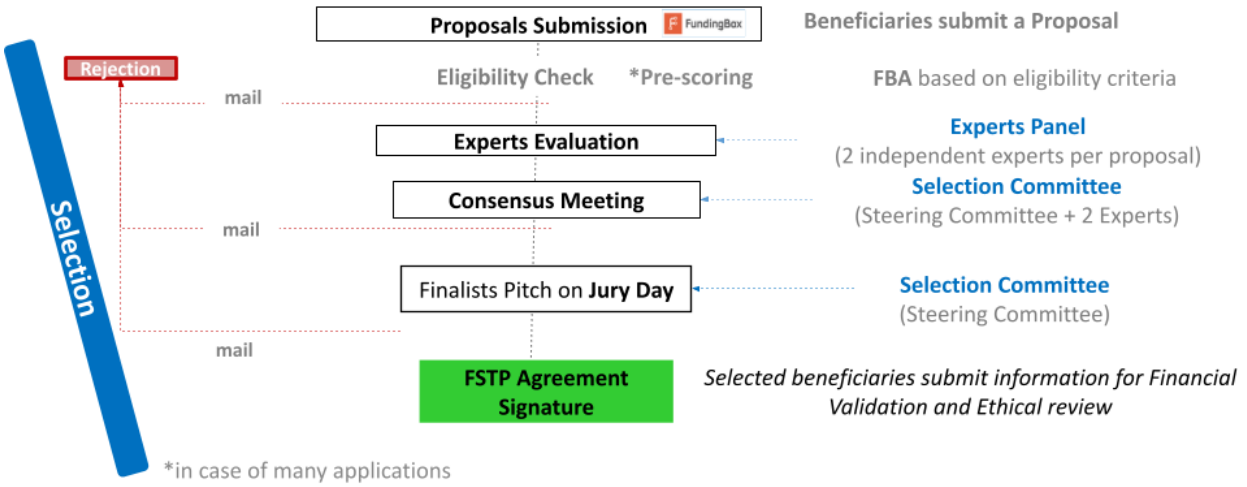

4.1 Eligibility Check

The eligibility check will be done on all proposals submitted before the deadline. All criteria are listed in Section 3 in this Guide for Applicants. The projects that do not comply with those criteria will be excluded and marked as ineligible.

We will inform you about the results of the eligibility check.

4.2 Pre-scoring

As we expect a lot of applications for this call, a pre-scoring system may be included. If we receive more than 100 eligible proposals, the eligible proposals will be automatically scored by the FundingBox System according to the following criteria:

- Dedication to the project (15% weight / up to 15 points),
- Team skills (15% weight / up to 15 points),
- Proved experience in AI /ML (15% weight / up to 15 points),
- Proved experience in the target market (10% weight / up to 10 points),
- Current progress or traction (15% weight / up to 15 points),

- Scalability of the solution (15% weight / up to 15 points),
- Market opportunity (15% weight / up to 15 points)

4.3 External Evaluation

In this phase, all projects will be evaluated by 2 external and independent evaluators with wide expertise in AI. Your project will be evaluated within the following awarding criteria:

EXCELLENCE will evaluate:

- **Ambition:** You have to demonstrate to what extent the proposed Experiment is beyond the State of the Art and describe the innovative approach behind it (e.g. ground-breaking objectives, novel concepts and approaches, new products, services or business and organisational models).

- **Innovation**: You should provide information about the innovation within your market and about the differentiation that your project will bring.

- **Soundness of the approach and credibility** of the proposed methodology.

IMPACT will analyse:

- **Market opportunity**: You have to demonstrate a clear idea of what you want to do and whether the new/improved product has market potential, e.g. because it solves a problem for a specific target customer.

- **Competition:** You have to provide information about the competition for your particular product/service and if the idea is disruptive and breaks the market. i.e. the products/services to be brought to market can be clearly differentiated from the competition.

- **Commercial Strategy and Scalability:** You have to demonstrate the level of scalability of the new/improved product which means that you should explain if your product can solve not only a specific problem but have the potential to be commercialised to solve a structural problem in a specific sector/process/etc.

IMPLEMENTATION will consider:

- **Team:** You have to demonstrate your management and leadership qualities, your ability to take a concept from ideas to market, your capacity to carry through their ideas and understand the dynamics of the market you are trying

to tap into. The team should be balanced and cross-functional, with a strong background and skill base.

- **Resources:** You should demonstrate the quality and effectiveness of the resources assigned in order to get the objectives/deliverables proposed. One important aspect is that you make clear in the proposal that what you propose can be developed using machine learning or AI solutions.

- **Objectives and results validation:** You should properly describe target objectives and expected results of the project and specify the overall structure of the work-plan to achieve them.

The evaluators will score each award criterion on a scale from 0 to 5:

0 = Proposal fails to address the criterion or cannot be assessed due to missing or incomplete information

1 = Poor – criterion is inadequately addressed or there are serious inherent weaknesses

2 = Fair – proposal broadly addresses the criterion, but there are significant weaknesses

3 = Good – proposal addresses the criterion well, but a number of shortcomings are present

4 = Very good – proposal addresses the criterion very well, but a small number of shortcomings are present

5 = Excellent – proposal successfully addresses all relevant aspects of the criterion. Any shortcomings are minor.

Each evaluator will score the application assigning a score from 0 to 5 for each criterion and produce an Individual Evaluation Report.
The threshold for individual criteria will be 3. The overall threshold, applying to the sum of the three individual scores, will be 10.
The final score will be calculated as an average of the individual assessments provided by the Evaluators.
If scores on a project show significant divergence between the two evaluators, a third evaluator will be involved to provide an additional independent assessment of the proposal.

Ties will be solved using the following criteria, in order:

- Impact score,
- Implementation score,
- Excellence score,
- Gender balance[4]: proposals with a higher % of women in the project team go first,
- Environment and low carbon economy contribution
- Social Impact
- Date of submission: earlier submitted proposals go first.

A 'Ranking List' will be elaborated. All proposals obtaining a score above the threshold will pass to the next phase.

Please note that we need time to process all the proposals in this phase, so you probably won't hear back for a while. The detailed timeline you can find in the Frequently Asked Questions Document.

4.4 Consensus Meeting

The 'Open Call Selection Committee' formed by selected consortium partners and two External Experts, will decide by Consensus (or ⅔ majority votes) the 'List of finalists' to pass to the next phase (**from 24 to up to 32 applicants based on the quality of the proposals**) and the 'Reserve List'. The discussion will be based on the ranking obtained as a result of the External Evaluation.
Whilst normally the highest ranked proposals will be selected for funding, the Open Call Selection Committee might have fair reasons for objecting to a specific third party, like the alignment with ELISE goals and scope, the ability to achieve the highest impact possible, commercial competition, as well as the existence of significant ethical concerns or a potential conflict of interest etc. In this case, the choice may pass to the next-ranked proposal.

The exact number of proposals approved will be decided based on the overall **quality** of the proposals.

[4] As beneficiaries, you must take all reasonable measures to promote equal opportunities between men and women in the implementation of the action. You must aim, to the extent possible, for a gender balance at all levels of personnel assigned to the action, including at supervisory and managerial level.

4.5 Jury Day

If your project is among the finalists, you will be invited to an online Jury Day, where you will have the opportunity to pitch your project in front of the 'Open Call Selection Committee'.
Before Jury Day you can receive specific questions formulated by the ELISE 'Open Call Selection Committee' that should be addressed during your pitch. Sending a pitch presentation and attending the Jury Day together with a Scientific Advisor is mandatory, not attending any of these requirements will automatically imply the exclusion of the proposal from the selection process.
During Jury Day your Scientific Advisor should explain why he/she is supporting your project.
Before Jury Day you will be requested to finalise the ethics self-assessment process with support from the Saidot ethics expert (mandatory requirement for Jury Day). More details can be found in the Frequently Asked Questions.
After the Jury Day, the ELISE 'Open Call Selection Committee' will undertake the final evaluation taking into account the following **'Awarding criteria'**:

- Quality of the Team
- Impact on the processes of the company either in terms of cost reduction or increasing productivity or effectiveness
- Business opportunity
- Innovative step
- Scalability

The 'Open Call Selection Committee' will decide by consensus (or majority vote of 2/3 of all members) the 'List of pre-selected grantees' and 'Reserve List'.
Up to 16 proposals will be selected. The exact number of proposals approved will be decided based on the overall **quality** of the proposals.
The Open Call Selection Committee might have fair reasons for objecting to a specific third party, like the alignment with ELISE goals and scope, the ability to achieve the highest impact possible, commercial competition, as well as the existence of significant ethical concerns or a potential conflict of interest etc.

After Jury Day, we will communicate the results to the applicants.

4.6 What's next? Sub Grant Agreement Preparation and Signature

Before the conclusion of the Sub Grant Agreement (hereinafter "Agreement") with the ELISE Consortium, we will ask you to provide documents regarding your formal

status (for the details please check our Frequently Asked Questions Document). Please do it within the deadlines that will be communicated to you.
Please note: If you fail to deliver the requested documents on time, without clear and reasonable justification, we will exclude you from the further formal assessment and you will be replaced with the applicant from the Reserve list.

5. Our Support and Payment Arrangements

The selected grantees will receive a fixed lump sum of up to **€60,000.** The lump sum is a simplified method of settling expenses in projects financed from Horizon 2020 funds. It means that you are not required to present strictly defined accounting documents to prove the cost incurred (e.g. invoices), but you are obliged to demonstrate the implementation of the project in line with the milestones set for it. Simply speaking, it means that we will assess your progress and quality of your work during Final Review, not your accountancy. The milestones (objectives and results) must be defined in the application form and (if necessary) elaborated before signing the Sub Grant Agreement.

The lump sum does not release you from the obligation to collect documentation to confirm the costs under fiscal regulation.
For a more detailed payment schedule please check the Frequently Asked Questions.

The implementation period will last up to 6 months. During this period you will be working on the implementation of your objectives specified in the work plan and the development of the solution with advisory support from the Scientific Advisor (invited at the application stage or assigned later during the selection process by ELISE Consortium).
At the beginning of the support period, you will be invited to participate in a Kick-off Meeting, to learn more about the ELISE Programme and meet other beneficiaries.

The first payment will be released after the Sub Grant Agreement has been signed and the work plan for the project (including objectives and results) has been updated based on the recommendation from the Open Call Selection Committee. The second (last) payment will be **released after the validation of the Final Report** by the 'Open Call Selection Committee'.
Your performance will be reviewed based on the following inputs:
- **Final Report** which includes:
    - Description of the results achieved by your team (max. 5 pages),

- A short summary of the project, which will be published on the ELISE website,
- Video demonstrating results of your project (max. 5 minutes), which will be published on the ELISE website,

- Short statement of the Scientific Advisor on the outputs of the project.

The Final Report has to be submitted within one month after the project is finished. You will demonstrate the results of your project during an **event organised by ELISE**, where all participants will present demo videos showcasing the results of their projects, followed by a Q&A session and a discussion with other participants of the Open Call, members of the ELISE/ELLIS community and the public.

Contact us

How can we help you?

If you have extra questions regarding our Open Call process you can send us a message to info.elise@fundingbox.com
In case of any technical issues or problems, please include the following information in your message:

- your username, telephone number and your email address;
- details of the specific problem (error messages you encountered, bugs descriptions, i.e. if a dropdown list isn't working, etc.); and
- screenshots of the problem.

Complaints

If, after receiving the results of one of the evaluation phases (when foreseen), you consider that a mistake has been made, you can send us your complaint. To do so please send us your complaint in English by email to info.elise@fundingbox.com including the following information:

- your contact details (including email address),
- the subject of the complaint,
- information and evidence regarding the alleged breach.

You have **3 calendar days** to submit your complaint starting from the day of becoming aware of the grounds for the rejection. We will review your complaint within no more than seven calendar days from its reception. If we need more time to assess your complaint, we will inform you by email about the extension.

We will not review anonymous complaints as well as complaints with incomplete information.
Please take into account that the evaluation is run by experts in the given field, and we do not interfere with their assessment, therefore we will not evaluate complaints related to the results of the evaluation other than related to the mistakes in the evaluation of the eligibility criteria.

6. Last but not least - final provisions

Any matters not covered by this Guide will be governed by Polish law and rules related to the H2020 and EU grants.
Please take into account that we make our best effort to keep all provided data confidential; however, for the avoidance of doubt, you are solely responsible to indicate your confidential information as such.
Your IPR will remain your property.
For the selected grantees, the Agreement will include the set of obligations towards the European Commission (for example: promoting the project and giving visibility to the EU funding, maintaining confidentiality, understanding potential controls by the EC/ECA and OLAF) and the ELISE Consortium.
The ELISE Consortium might cancel the call at any time, change its provisions or extend it. In such a case we will inform all applicants about the change. Signature of the agreement is an initial condition to establish any obligations among applicants and any Consortium partners (with respect to the obligation of confidentiality of the application).

Did not find what you were looking for? You may want to check our Frequently Asked Questions Section.

Annex 1: Information Clause

CONTROLLER'S IDENTITY AND CONTACT DETAILS

The data controller is FundingBox Accelerator sp. z o.o. (Al. Jerozolimskie 136, 02-305 Warsaw, Poland).

In all matters regarding personal data, you can contact us using the following email address: **privacy@fundingbox.com**.

PURPOSES, LEGAL BASIS AND PROCESSING PERIOD

The purpose of processing	Legal basis for processing	Period
1) To run an Open Call and collect data necessary to evaluate applications submitted in the Open Call	Legitimate interest of Funding Box (based on Article 6, paragraph 1 (f) of GDPR) which is fulfilling the obligations and our other interests related to this purposes	6 years from the end of the year in which the Project ended
2) To realise the Project goals described in the Grant Agreement (communication, reporting, collaborating with other project partners)		
3) To consider potential complaints		
4) To possibly establish and enforce claims or defend against them		

DATA RECEIVERS

Data controller will transfer personal data only to trusted recipients such as entities belonging to the FundingBox's capital group, IT service providers, accountants, law firms, postal and courier companies (who process personal data on the controller's behalf).

Due to the fact that we use the services of Google LLC, your data may be transferred to the USA. We have concluded an agreement with Google LLC - the so-called Standard Contractual Clauses. This

means that in accordance with the decision of the European Commission No. 2010/87 / EU of February 5, 2010, your personal data may be processed by this company in the USA. More information about the decision at: https://eur-lex.europa.eu/legal-content/PL/TXT/?uri=celex%3A32010D0087

To realise the Project data can be transferred also to Project Partners (complete list of the project partners is available at the email address: privacy@fundingbox.com)and European Commision.

RIGHTS OF DATA SUBJECT

Due to the fact that we process your personal data, you have the right to:

1) request access to your personal data,
2) demand the rectification of your personal data,
3) request to remove or limit the processing of your personal data,
4) complain to the supervisory authority (The President of the Personal Data Protection Office, Warsaw, Poland, https://uodo.gov.pl/en).

You also have a right to object to processing of your personal data (according to the Article 21 of GDPR).

INFORMATION ABOUT VOLUNTARY OR OBLIGATORY DATA PROVISION

Providing data is voluntary, although it is necessary to participate in the Open Call. Without providing your data, it is not possible to contact you and evaluate the application.

Annex 2: ELISE Focus areas

1: Geometric Deep Learning: graph, group and gauge convolutions

A multitude of important real-world data comes together with some form of graph or manifold structure. Examples of graph structured data include the social interactions between users in social networks, protein interactions in biological organisms, bonds between atoms within molecules, transport links within transportation networks, and neuronal connections within the brain connectome. Recent years have seen a surge in deep learning approaches that directly apply neural networks to learn over graph-structured data [Bronstein 2017, Scarselli 2009], including generalisations of convolutional neural networks to graphstructured data |Henaff 2015]. By thinking about symmetries on general manifolds a unified theory of geometric deep learning has emerged that includes group and gauge convolutions [Cohen 2016, 2019].

Europe is currently leading in geometric deep learning research, including its multitude of applications. The programme intends to push the boundaries of the theoretical aspects of geometric deep learning as well as expanding and exploring its applications such as wind patterns on the earth, 3d models of organs, chemical drug design and so on. The programme intends to develop the theoretical aspects of geometric deep learning, such as generalisations to new types of convolutions, uncovering connections to mathematical structures such as fibre bundles and group theory, and proving properties of geometric deep learning architectures.

By collaborating with European industry the tools of geometric deep learning can make European companies leading in important areas such as pharmaceutics (drug design), climate modelling (CNNs on spheres), gaming (rendering on manifolds), knowledge representation (knowledge graphs), and material sciences (deep learning on meshes).

2: Robust Machine Learning

As machine learning technologies are progressively deployed across the sciences and in the real world, it is becoming more important that they can reliably perform well or self-diagnose when not, when applied in settings different to those during training, under adversarial manipulations, or when using unbalanced, messy and heterogeneous data. This programme aims to understand the principles and developing the techniques for such robust and trustworthy ML.

We aim to explore the different facets of robust ML, including,

1) Dealing with domain shifts when testing on new settings or in the real world, 2) Model-based approaches to tackling generalisation to new settings,

3) Robustness to adversarial manipulations and attacks,

4) Quantification and calibration of predictive and inferential uncertainties, and 5) Assessment and verification of robust ML techniques.

Transfer learning, meta learning, continual learning as well as autoML methods can be used to allow ML systems to automatically adapt to new scenarios, while sim2real approaches in robotics involve training and adapting systems to real environments. Model-based approaches, e.g. based on causal models, may allow for more robust generalisation in such situations. Assessment and verification of the robustness of ML systems includes

quantification and calibration of the predictive and inferential uncertainties in ML systems, provable verification of robustness against local perturbation attacks, estimation of the probability of rare failure events, integration of constraints to ensure safe behaviour as well as consideration of the whole data processing pipeline in which an ML system is part of. Application areas include health, environmental sciences, AI-Assisted Design, autonomous vehicles and industrial control. Robust ML also has important links to fairness, interpretability and transparency. Issues of fairness often arise due to biases inherent in datasets, e.g. in terms of lack of coverage of minority groups. Interpretable and transparent systems allow for easier assessment and verification of robustness by users. On the other hand, it can be argued that robustness of learning methods is a prerequisite for explainability of their decisions. Exploration of human-in-the-loop systems is also an interesting avenue towards more robust learning systems for augmenting humans.

3: Interactive Learning and Interventional Representations

The growth of machine learning applications calls for systems with richer sets of capabilities, including the ability to deal with complex feedback loops, cooperate with each other, and adapt to a stream of new tasks. This requires significant conceptual advances that connect different types of interactive learning, including representation learning, online/reinforcement learning, meta-learning, and multi-agent learning.

This programme fundamentally rethinks the principles of interactive models of learning, exploring the role of causal modelling [Peter 2019] in bridging the gap between observational and interventional learning. The ultimate goal is to understand the organising principles underlying robust intelligent behaviour, and to enable reliable learning-based decision systems [Khajehnejad 2019] for high-stakes real-world applications. Interactive learning algorithms [Alon 2017, Cesa-Bianchi 2019], able to intervene in the world and adapt to changing environments, are fundamental in the design of safe and robust AI systems that can cooperate without a centralised data source. The programme seeks to develop the principles of causal learning systems that can interact and intervene in the world, also adapting to situations where the environment evolves in complex ways. Achieving such results is key to applications in IoT, healthcare, renewable energy systems, and other central themes of the European research agenda.

4: Machine Learning and Computer Vision

Computer vision has been revolutionised by machine learning and in particular deep learning. For many problems which have been studied for decades, state-of-the art performance has dramatically improved using artificial neural networks. Does this mean that the knowledge gained in the previous decades is now obsolete?

The goal in this research programme is to make the connection between classical algorithms and modern machine learning more explicit. We wish to bring together researchers working on different aspects of computer vision. Specifically, we envision three subareas:

1) Mid-level vision and image reconstruction [Sun 2010],
2) 3D Geometry from multiple views [Hartley 2003],
3) Object and activity recognition [Lazebnik 2006, Leibe2008].

All three sub areas have recently seen a huge increase in the number of papers that use machine learning, and in all three areas there is considerable domain specific knowledge that has been generated over many years.
While machine learning is a main driver of current computer vision research, computer vision also is a key enabler and driver for a lot of machine learning research ranging from large datasets such as ImageNet, over methods such as Convolutional Neural Networks and Generative Adversarial Networks, to applications such as autonomous driving and tagging of visual data. The key objective of the programme is to identify opportunities and to make foundational contributions to how computer vision and domain specific (top-down) knowledge can be leveraged and combined with (bottom-up) data-driven machine learning methods. Progress in this intersection will have far-reaching implications for both research and applications by enabling to leverage prior knowledge in modern machine learning. Computer vision researchers in Europe are highly recognized and world leaders in the area.

5. Robot Learning: Closing the Reality Gap
Autonomous robots that can assist humans in situations of daily life have been a long-standing vision of robotics, artificial intelligence, and cognitive sciences. Learning techniques have only recently started to enable modelling all the perceptuomotor tasks that a robot should do and have triggered tremendous progress in various areas including self-driving cars, logistics, manufacturing, transportation and home automation. The question of how to effectively train such robot systems becomes a major bottleneck: Unlike in computer vision or natural language processing, it is unlikely that the robot instructor will be able to create gigantic data sets covering all foreseeable circumstances the robots will encounter during their operation. The autonomous systems need to be able adapt to changing conditions and other factors such as aging. Furthermore, robot instructors will not be able to predict all tasks to be carried out by mobile manipulation robots, which immediately raises the question of how such systems can be instructed even by non-expert users. Thus, in this programme, we will address these problems by developing innovative learning solutions specifically for providing future robots with the ability to acquire robust strategies, adapt to new or changing conditions and to be easily teachable even by non-expert users. The goal of this programme is to develop the methodology for effective learning of actions and controls from the right amount of data and instruction for a real-world robot.
The programme focuses on essential research areas, where existing solutions do not suffice for modern service robots, in particular on manipulation, locomotion and collaboration tasks. This programme will collaborate with the NoEs for developing the next generation of intelligent robots

6: Human-centric Machine Learning
Machine learning (ML) models and algorithms have achieved, or even surpassed, human performance in tasks such as recognising images and speech, translating between languages or beating the best players of various games. These recent developments are early evidence of the tremendous potential that ML has to revolutionise industries, public services and society.

ML is a key component of AI and as its use becomes more widespread in human-centric applications and algorithmic decisions are more consequential to individuals and society, the above positive developments have been mirrored by an increasing number of missteps and limitations, from discriminating against minorities, manipulating human decision-making, fuelling the spread of misinformation and increasing polarisation, to being responsible for car crashes, where machine learning has been blamed to play a major role. This research programme focuses on human-centric ML models and algorithms, specifically designed to avoid the above shortcomings. More specifically, human-centric ML will help to ensure widespread benefits to, and acceptance from, the public by guaranteeing transparency, clear accountability, interpretability and fairness [Gummadi 2019, Soares 2019] in the algorithmic decisions they fuel [Oliver 2019, Lepri 2016, 2017]. Moreover, it will be amenable to legal and technical certification, which is key in critical domains such as autonomous driving or health, by guaranteeing accountability and verifiability. Fulfilling these guarantees requires interdisciplinary work and collaboration with experts at the intersection of ML, causality, human-computer interaction (HCI), ethics and the law.
We aim to develop the principles and the framework of the next generation of AI and machine learning that is more human-oriented, human-friendly and human-centred [Angelov 2017, 2018]. The programme plans to connect with companies and policymakers to ensure impact, particularly in the case of highly important domains, such as healthcare, education, autonomous driving, justice, hiring, security and many online social and information systems. In the Human- Centric ML programme we will collaborate with the NoEs for advanced perception or interaction with humans (for human-centred AI) and environments.

7: Theory, Algorithms and Computations of Modern Learning Systems

Even while contemporary machine learning algorithms achieve remarkable results, they are often inefficient and brittle [Hein 2019]. The practice of machine learning is wasteful of resources, both human and technological [Schwartz 2019]. Although the field continues to be primarily focused on models and their applications, the main roadblocks for fundamental progress in machine learning are now arguably computational and theoretical in nature. The programme brings together experts with diverse backgrounds, linked by an interest in methods that are theoretically grounded [Rudi 2015] as well as working in practice [Schneider 2019], to deliver provably efficient and reliable learning systems [Chizat 2019]. Theory is the glue linking all applied areas of machine learning. Much of the recent successes of the field are grounded on the development of widely applicable concepts and algorithms. The presence of this programme within the research aims of this proposal, and the physical presence of its fellows within the Fellowship as a whole, will ensure this crucial aspect is retained going forward. The research aims (provably reliable and efficient methods) are also directly aligned with O5 of the proposal. The programme aims to stimulate theory research within the European AI community and promote mathematical rigour and provable reliability in AI & ML. In fact, theory is an area of ML in which Europe has traditionally punched above its weight, a strength that the programme will aim to sustain.

8: Quantum and physics-based Machine Learning

Conventional CMOS technology is reaching its physical limits and the energy consumption of computing is reaching alarming proportions. The aim of this programme is to use concepts from quantum physics and statistical physics to develop novel machine learning algorithms with the aim to realise breakthroughs in the design of novel future, possibly energy efficient, hardware. Europe has an excellent tradition on this research due to its long tradition of theoretical physicists working in ML and ELISE brings together the top researchers in Europe on this topic. The success of the programme will crucially depend on the understanding of the fundamental principles underlying efficient learning schemes for neural networks devices that are subject to *in situ* physical constraints. This is a fundamental issue which has both conceptual and technological value.
One of the key phenomena that a realistic form of learning in large neural networks at the physical scale has to deal with is noise. Rather than designing reliable bits and use software to generate random numbers, an appealing alternative is to design hardware that exploits the different form of structural noise to implement learning schemes. Such devices would possess one of the key feature to become much more energy efficient, see [Indiveri 2015]. Methods from non-equilibrium statistical physics are well suited to improve our understanding of stochastic systems [Baldassi 2018] and to provide the methodological framework for designing learning schemes which focus on regions on the learning landscape that are accessible to noisy processes and, at the same time, have high learning performance [Baldassi 2019]. Quantum devices are nearing the noisy intermediate scale quantum (NISQ) era [Preskill 2018]. NISQ devices may provide computational advantages over classical supercomputers for various machine learning problems, such as sampling from hard to-simulate probability distributions and linear algebra problems [Kappen 2018]. The aim is to demonstrate that the application of NISQ technology to machine learning may be one of the first instances exhibiting genuine quantum advantages. Better understanding of quantum many body physics will advance our understanding of quantum materials and the design of quantum devices, enabling quantum computers and networks. For this problem, machine learning offers new methods that improve the state of the art [Carleo 2017].

9: Natural Intelligence

The standard paradigm of machine learning is task-centric with an extended training phase and very little change between training and testing. Natural intelligence is agent-centric. While animals might use different tools to solve different tasks, they have only one brain to make decisions and to generate successful behaviour throughout their entire lifetime. The current toolset of AI to describe and understand this agent-centric lifelong learning and decision-making principles of natural intelligence is still very limited. In this programme we focus on important aspects of natural intelligence, especially those that are poorly captured or missing in current machine learning models. We seek to make significant advances in tackling these aspects. The following lists various topics that we will address in the programme:
1) Deep semantics and cross-domain learning,
2) Leveraging large amounts of weak information,
3) Shaping inductive bias via neural network structure and how to meta-learn it (people benefit from a large amount of pre-existing structure, including priors and heuristics),

4) Life-long learning (contrary to current machine learning models, people continually learn new skills, sharpen existing ones, and transfer knowledge across domains, while maintaining unused abilities to a reasonable extent),
5) Adaptive resource deployment (encompassing ways of addressing the internal reinforcement learning problems associated with deploying resources, such as memory and time, adaptively to solve tricky problems),
6) Social reasoning (i.e., we need computationally richer structures to examine and exploit cooperation).

10: Health

The programme seeks to strengthen the connection of AI and ML with the most relevant opportunities in human health. Major drivers in this area include medical data analysis, genomic medicine, biomedical imaging, drug discovery and clinical decision support, which jointly constitute promising and impactful application domains of modern AI/ML. Compared to other fields, the potential of AI/ML has not been fully utilised in health, which is in part due to the domain-specific and complex challenges requiring new solutions. We aim to feed back key challenges of human health applications (e.g., explainability, robustness, uncertainty awareness, privacy, auditability, scale, heterogeneity, causality) into AI/ML methods development. In order to enable rational, risk-aware and informed medical decision making, critical challenges in AI/ML have to be addressed, for instance:
a) how to build accurate models that are explainable and transparent to physicians and affected individuals,
b) how to integrate prior domain knowledge from the life sciences with data-driven approaches,
c) how to quantify uncertainty when information is incomplete, inaccurate and multi-modal,
d) how to turn the vast volumes of primarily observational data into causal insights,
e) how to transfer knowledge across various locations and data sources.
The programme has the ambition to democratise the most relevant and impactful datasets for methods-centric research in health, and to connect ELISE to European initiatives that generate and curate datasets that will build the foundations of this research. Lead application areas of this programme are: Electronic health records, biomedical imaging, human (gen)omics, drug design and identification of targets for therapeutic interventions in personalised medicine. Large volumes of retrospective and prospectively collected data, combined with medical health records pose impactful application domains for AI to identify the most relevant drug targets and biomedical research in the context of personalised medicine.

11: Machine Learning in Earth and Climate Sciences

The Earth is a dynamic networked system, with complex interactions between atmosphere, land, oceans and human activities. The main goal of the programme is to advance in the modelling and understanding of the Earth system with Machine Learning and Process Understanding [Camps-Valls 2019, Reichstein 2019, Runge 2019]. We aim to push frontiers in
1) elucidating and forecasting spatio-temporal dynamics, incl. modelling of greenhouse gases,

2) detection of anomalies and extreme events,
3) classification of climate and land surface patterns,
4) discovery of causal drivers from observational data, and
5) hybrid physics-data driven modelling to achieve robustness, consistency and improved interpretability.
This programme will contribute towards achieving explainable ML and improved extrapolation via with hybrid models and causal methods, within a particularly challenging complex system, the Earth. Our long-term vision is tied to open new frontiers and foster research towards algorithms capable of discovering knowledge from Earth data, a steppingstone before the more ambitious far-end goal of machine reasoning of anthropogenic climate change (O1, O2). The programme fellows have the expertise and knowledge to take this grand challenge, and along with e.g. the European Space Agency, it is a strategic opportunity for Europe to lead also machine learning for Earth and Climate Sciences

12: Natural Language Processing (NLP)
Natural Language Processing (NLP) is a field of computer science and artificial intelligence concerned with enabling computers to analyse and generate natural language text. Recent years have witnessed the rapid evolution of a wide range of NLP systems that translate text [Mallinson 2017], recognize or produce speech, answer questions [Xu 2019], retrieve documents or facts [Mallinson 2017], respond to commands [Dong 2018], summarise articles [Liu 2019], and simplify texts for children or non-native speakers [Zhang 2017]. NLP is transforming the way humans communicate with each other and with machines. Google Translate can instantly translate between any pair of over 100 languages from text; it can also translate 37 languages from images and 32 from speech. Apple's Siri and Amazon's Alexa allow users to accomplish tasks by voice command in many domains of daily life. NLP is embedded in our daily lives, often silently. But there are many applications that are not yet realised, as well as language and speaker populations that current technology does not serve. Improving NLP to address these problems requires a diverse array of expertise including linguistics, data mining, and a broad range of statistical and algorithmic methods such as neural networks and other machine learning models. Importantly, current approaches rely on supervised approaches which are datahungry and costly in cases where human annotations are required. Advances in transfer learning, metalearning and unsupervised methods will address the data acquisition bottleneck leading to innovations in NLP but also related areas such as computer vision, robotics, and human-centric machine learning.

13: Multimedia / multimodal information
The last five years have seen extreme progress in the automatic description of multimedia content, especially visual. Automatic image and video tagging to a high degree of accuracy is now embedded as machine learning applications in web-scale systems like Facebook and on smartphone apps like personal photo management. Image and video captioning is also experiencing huge progress with the best systems able to generate natural language descriptions based on machine learning, comparable to humans, albeit in restricted domains. Multimedia applications, especially the analysis and description of content as tags or as

captions, is based almost entirely on machine learning. Many of the significant developments in recent years in the machine learning field, including deep learning, have arisen from multimedia applications. This is because multimedia applications provide large datasets for training, real users with real needs, and yet the applications as a whole are somewhat forgiving and tolerant of errors when they do occur. Incremental progress has been achieved using annual benchmarking campaigns in the multimedia community (ImageNet, TRECVid, ImageCLEF, MediaEval, the MM Conference Grand Challenges etc.) [Deng 2009, Viitaniemi 2015, Müller 2010, Ionescu 2018]. However, this is not any more enough; a new, more agile, more collaborative and more responsive mechanism for shared progress, is needed. A case in point is the use of machine learning for generative forms of multimedia, from deepfake videos to automatic colourisation of old black-and-white images. This programme will support a greater degree of collaborative benchmarking among researchers across multimedia applications including analysis, indexing, search, summarisation, and the automatic generation of content [Pan 2017]. Topics include:
1) Automatic generation of content descriptions for multimedia in form of natural language-based dialogue, an interaction between human and machine around the topic of multimedia content,
2) Machine learning for generating new content, on demand, based on user needs or requests
3) Machine learning to incrementally analyse multimedia content, on demand, and from a given user's or a given task's perspective.

14: Information retrieval

Machine learning, and in particular deep learning, has transformed information retrieval [Borisov 2016, Joachims 2018]. Highly effective technology to connect people to information, in the form of search engines, recommender systems, and digital assistants, is now part of everyday life across the planet. Future information retrieval systems need to significantly improve their understanding of human information interaction behaviour and act based on this understanding. This requires substantial progress in (1) representation learning of human behaviour, (2) semantic matching methods based on very limited volumes of interaction data, (3) counterfactual learning [Jagerman 2019] to rank methods that are able to learn user preferences, in an unbiased manner, from interaction data collected using a variety of information retrieval environments, and (4) mixed initiative scenarios in which information retrieval systems need to learn to ask questions that are minimally intrusive but maximally informative to the system [Ren 2019]. Advances in this programme will contribute to applications such as (1) transparent search engines with performance guarantees that are minimally invasive in terms of user interaction data required for optimization; (2) recommender systems that are able to infer user's long-term vs. short-term interests in an effective manner; and (3) conversational assistants that are able to engage in natural and informative information seeking dialogues. The information retrieval programme builds on, and contributes to, advances in representation learning, counterfactual learning, reinforcement learning, online learning, and transfer learning.

15: Explainability and Fairness in Data Mining

The past decades have witnessed significant progress in the development of intelligent methods for analysing large volumes of data and performing challenging tasks with success that matches, or in many cases exceeds, the performance of human experts. In their majority, however, the developed methods are opaque and non-interpretable. In addition, they suffer from data biases and produce models that are not fair. The research community has recognized these issues and a lot of recent work has been devoted to design models that are fair, accountable, and transparent [Barocas 2017, Guidotti 2019]. Most of the work, however, focuses on supervised learning. In this ELISE theme we will address the challenges of explainability and fairness in unsupervised learning. In particular, we will develop methods that produce an interpretable clustering along a small number of data dimensions or provide simple and comprehensible explanations for a given clustering produced by a black-box algorithm. We will consider different problem settings, such as overlapping clustering, clustering ensembles, mixture-model clustering, as well as clustering in graph data. As a different research direction, we will investigate the problem of finding interpretable network embeddings [Perozzi
2014]. In the topic of fairness, we will study different notions of fairness in unsupervised learning [Chierichetti 2017], we will formulate relevant problem settings motivated by real-world applications, and we will develop algorithms with quality guarantees. This theme complements the research programme on robust machine learning.

16: Symbolic Machine Learning

Initially AI was dominated by symbolic approaches. In the first AI phase there was an unlimited expectation with respect to the capabilities of computers to “solve tasks for which intelligence is required, if they would be executed by humans”. McCarthy proposed to formalise the complete human knowledge in form of a homogeneous formal representation, first order predicate logic. Predicate logic needs to be fed by background knowledge, so the second phase (ca 1970-1985) was dominated by expert systems which led into some major industrial projects in the 1980s. In the beginning of the 2000er years, the semantic web community used symbolic approaches to describe background knowledge in formal ontologies. In general, there are strong links of symbolic representations to first-order logic, database theory, information retrieval, and NLP. Ontologies are prevalent in medicine and bioinformatics (ICD-10, SNOMED) and document and information and knowledge management. Currently (maybe since 2010) there is a strong focus on knowledge graphs, which are triple oriented forms of knowledge representations [Singhal 2014]. In contrast to previous approaches, knowledge graphs are primarily instance oriented. Knowledge graphs are currently finding widespread applications in industry. Symbolic Machine Learning is often connected with inductive logic programming where the goal is to discover prolog-rule-like constructs in data. In data mining, association rule mining is a popular topic. Probabilistic relational models, Markov logic networks, and Bayesian logic programs are three approaches that are also applicable in probabilistic domains. Representation learning has become the dominant approach in learning with knowledge graphs [Singhal 2014] with industrial applications [Hildebrandt 2019]. Scene graphs in vision are closely related [Nickel 2016]. There are many applications for this work, such as in healthcare (in combination with ontologies and knowledge graphs, machine learning is the basis for modern approaches to medical

decision support [Tresp 2016]), in industry (industrial decision support systems are based on machine learning with knowledge graphs [Hildebrandt 2019], in NLP (various approaches for the extraction of triple statements from text [Dong 2014]), and in vision (a new trend is the extraction of symbolic graph representations from images [Baier 2018]).