

ELLIS

Executive Summary

ELLIS¹ (European Laboratory for Learning and Intelligent Systems) is a European grassroots initiative in AI with a focus on scientific excellence, innovation and societal impact. It aims to create a European AI Laboratory inspired by models such as EMBL (European Molecular Biology Laboratory) and the Vector Institute (Toronto). Like Vector, ELLIS builds upon machine learning as driving modern AI — inspired by a model of human intelligence which is not ‘programmed’ but evolved and learned from data. Virtually all of the dramatic recent progress and impact of AI in today’s world is fueled by data-driven machine learning. Machine intelligence will make further progress as modern forms of deep and reinforcement learning as well as causal inference gain traction, enabled by a new generation of the most talented young students attracted by modern AI. However, Europe is facing major barriers when it comes to retaining top talent in European institutions: lack of competitive salaries, high teaching obligations in universities, rigid environments that do not support a fluid relationship with industry and the creation of startups, and the lack of critical mass due to a fragmented situation with separate islands of excellence. Proposed in 2018, ELLIS tackles these barriers by pursuing a **three-pillar strategy** to foster **European excellence** in this highly competitive field: research programs and fellows, a competitive pan-European PhD program and a network of ELLIS units and ELLIS institutes. Openness for participation throughout Europe is key, subject to strict excellence criteria and mechanisms that ensure the highest quality of research leading to clear economic and societal impact.

Canada established a CIFAR network of Fellows studying Learning in Machines and Brains² before launching the Pan-Canadian AI Strategy. CIFAR Fellows receive CAN\$30k unrestricted³ funding per year and meet at two workshops where almost all Fellows are present, usually co-located with a leading international conference. This network not only significantly advanced the field, but it also gave a voice to machine learning and modern AI, resulting in the pan-Canadian AI strategy and the fast foundation of three major new institutes for modern AI (**Mila** in Montreal, **Vector** in Toronto, and **Amii** in Edmonton), run as non-profit corporations

¹ <https://ellis.eu>

² <https://www.cifar.ca/research/programs/learning-in-machines-brains>

³ <https://www.cifar.ca/ai/pan-canadian-artificial-intelligence-strategy>

with links to partnering universities. Inspired by this model, the **first pillar of ELLIS** is already underway through the establishment of the initial 11 Fellow Programs, dedicated to topics ranging from⁴ basic research in theory and algorithms to applications in health and climate sciences and human-centric elements of AI. The program proposals were reviewed by recent Turing award winner Yoshua Bengio with the help of CIFAR. Currently, there are 14 ELLIS Programs, spanning a wide range of topics, and together, these programs bring together over 100 fellows.

The **second pillar** of ELLIS entails inspiring, identifying, nurturing and connecting the best PhD students in Europe in modern AI through the ELLIS PhD & Postdoc program.

The **third pillar of ELLIS** consists in the creation of a network of ELLIS *units* located at leading existing institutions or created from scratch performing modern AI research across⁵ Europe. According to an open call, an ELLIS unit invests at least 1.5 Mio EUR per year, for at least five years, and requires a commitment to excellence as well as a contribution (at least 300k EUR annually) towards network activities, such as hosting exchange faculty and students, leading research programs and organizing workshops. The response to this call was outstanding, and the set of initial ELLIS units includes many of the strongest European hotspots for modern AI, such as Amsterdam, EPFL, ETH, ISTA, Oxford, and Tübingen, together with new, emerging locations, such as the Valencian region. Jointly, they committed more than 200 Mio EUR of their own funding, a significant part of which is reserved for network activities such as exchange visits and joint projects. They assemble more than 35 ERC grantees, and have already spun out a similar number of startup companies. A second and third set of ELLIS units has been reviewed in 2020 and 2021, including applications from Berlin, Cambridge, Lisbon, , UCL and other sites. At the moment there are 35 ELLIS units in the ELLIS network.

Furthermore, ELLIS aims to create an intergovernmental organization operating a set of fully fledged ELLIS institutes⁶. This is loosely inspired by the EMBL model and should attract outstanding scientists, providing them with the means to generate cutting edge scientific, economic, and societal innovation. All European countries with a strong interest in AI (both inside and outside the EU) should expect economic and societal benefits if they invest in this, and thus not necessarily expect the EU to provide the bulk of the funding. However, EU funding can help in fostering the ELLIS *network* as well as in including scientists in countries that do not have the financial means necessary. It can also ensure ELLIS has a voice at the EU level, where

⁴ <https://ellis.eu/de/news/ellis-programs-launched>

⁵ <https://ellis.eu/de/news/ellis-sites-call-for-proposals>

⁶ <https://ellis.eu/de/news/leading-european-ai-scientists-issue-a-call-to-action>

the “High Level Expert Group” for AI is currently biased towards classical rule-based AI. ELLIS participates in a number of European activities and applications, always dedicated to promote scientific excellence in the field of artificial intelligence and machine learning and to open internationally competitive career opportunities for outstanding young scientists in this field in Europe. In particular, the ELISE proposal (built upon the ELLIS organisation) submitted for the EU call ICT-48-2020 “Towards a vibrant European network of AI excellence centres” received top scores among all proposals and has been funded with ca. 12 million euros⁷. ELISE stands for European Network of AI Excellence Centres, and with its implementation a first set of network activities based on the principles and ideas formulated by ELLIS started in September 2020.^{8 7}

Achieving the ELLIS vision entails attracting scientists from the US in cases where even an ETH full professorship or a Max Planck directorship is not sufficiently attractive. This needs not only attractive salaries, but (at least initially) also co-affiliations with the best academic institutions. ELLIS faculty could split their time with universities or existing research institutes. Their personal position would need to be guaranteed for life, even if their full research budget is not (note that in the US, it hardly ever is). ELLIS co-affiliation should imply partial teaching relief, similar to arrangements for instance in Toronto (between Vector Institute and U of T).

While ELLIS is making significant progress in implementing the first two pillars, the path to the establishment of an intergovernmental organization is more complex. There are multiple paths towards realizing this :

1. Establishment of an intergovernmental working group of a core group of motivated countries.
2. As stepping stones, one can build on existing joint ventures such as the Max-Planck-ETH Center for Learning Systems, or the joint PhD program between MPI for Intelligent Systems and Cambridge, operating since 2014, to create a pan-European ELLIS PhD program.
3. There is significant economic interest in machine learning technology among high tech regions in Europe, and those may want to realize the ELLIS vision faster than national governments. Baden-Württemberg has reserved substantial funds for an institute, and its

7

<https://www.aalto.fi/en/news/the-european-commission-offers-significant-support-to-europes-ai-excellence>

8

https://ec.europa.eu/programmes/horizon2020/sites/horizon2020/files/h2020-leit-ict-2018-2020-05-27_draf_pre-publication.pdf

Minister President Kretschmann has explicitly endorsed the points made in the ELLIS open letter, including flexible governance following the EMBL model, the need for guest houses to host international visitors, etc. Other regions in Europe have made similar public commitments, e.g., Bavaria and the region of Valencia in Spain.

4. There is philanthropic interest to ensure that Europe does not fall behind in this field which is considered crucial for our future.

ELLIS will explore new ways of fostering innovation. It will provide a well-defined interface to industrial partners and support the creation of start-ups. Following the focus on attracting exceptional talent, ELLIS sites will support co-affiliations, encourage leaves-of-absence, and foster startups (done transparently to ensure societal acceptance). This will counteract brain drain, provide economic stimulation, and ensure European technological sovereignty.

ELLIS Vision

An open letter⁹ published in April 2018 describes further details:

- ELLIS should comprise labs with outstanding facilities in all partner countries.
- ELLIS runs programs for visiting researchers (both from academia and industry) as well as workshops and summer schools for students, academics, and industrial participants. Mobility is facilitated by housing, childcare, and international schools at each site.
- ELLIS builds a Pan-European PhD program in cooperation with degree-granting universities. PhDs in this program will be jointly advised by two top AI researchers from two different European countries, and students will spend significant time at both sites.
- ELLIS aims to offer permanent employment to outstanding individuals early on and train them in both academic and non-academic skills. They should be offered a complete career path within ELLIS, paralleling those found in tenure-track programs, from the equivalent of the rank of assistant professor to that of a full professor.
- ELLIS researchers can split their time between ELLIS and local university or industry research labs. Collaboration with industry is encouraged and structured using transparent and simple IP rules that ensure that public funding is used in a way that benefits the

⁹ <https://ellis.eu/letter>

- public. Joint research involving industry and public funding should be published openly.
- ELLIS researchers are encouraged to found startups. ELLIS does not optimize short-term licensing income but aims at sustained impact. To this end, it owns a modest share in those startups and claims no further rights as long as the startup is formed in a partner country, thus generating downstream impact in Europe. ELLIS should support startups in terms of (a) generous leave-of-absence rules, (b) temporary use of infrastructure, and (c) help with administration including legal/financial advice.
 - ELLIS' unique characteristic is outstanding academic quality as measured for instance by publications in the leading competitive conferences of the field. While this sets a high academic bar, it is simultaneously open in that it does not require a broad industrial basis to develop an ELLIS site, as long as there is a commitment to excellence and adherence to mechanisms to select and retain local top talent (which exists everywhere). One such mechanism is that within ELLIS, selection of Faculty and Fellows in one country is always vetted only by ELLIS members from other countries.

Taken together, ELLIS wants to ensure **that Europe will play a major role in the scientific and societal revolution that is underway**. Europe should benefit from it and European values should help shape its impact.

Progress Report

Fellow Programs

The initial 11 programs were admitted, following review by Yoshua Bengio and the CIFAR *Learning in Machines and Brains* program. Currently, there are 14 ELLIS Programs, spanning a wide range of topics, and bringing together over 240 fellows and scholars.

Geometric Deep Learning

- Directors: Michael Bronstein, Taco Cohen, Max Welling
- Fellows and Scholars: Pascal Frossard, Stefanos Zafeiriou, Pietro Liò, Sonia Petrone, Marinka Zitnik, Petar Veličković, Emanuele Rodolà, Dorina Thanou, Cesare Alippi, Frank Nielsen, Jakob Verbeek

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-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, including generalisations of convolutional neural networks to graph-structured data. By thinking about symmetries on general manifolds a unified theory of geometric deep learning has emerged that includes group and gauge convolutions. The program 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.

- Machine learning on non-Euclidean domains
- 4G: geometric, graph, group, gauge convolutions
- Applications: computer vision, graphics, social networks, chemistry, biology, physics, medicine

Robust Machine Learning

- Directors: Chris Holmes, Samuel Kaski, Yee Whye Teh
- Fellows and Scholars: Zoubin Ghahramani, Frank Hutter, Cédric Archambeau, Chris Williams, Sebastian Nowozin, Peter Grünwald, Silvia Chiappa, Yarin Gal, Amir Globerson, Antti Oulasvirta, Richard E. Turner, Isabel Valera, Aki Vehtari, Pushmeet Kohli

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, when applied in settings different to those during training, under adversarial manipulations, or when using unbalanced, messy and heterogeneous data. This program aims at understanding the principles and developing the techniques for such robust ML.

- Principles and methods for Robust ML
- Quantification and verification of Robust ML
- Applications in health, environmental sciences, design, autonomous vehicles, industrial control.

Interactive Learning and Interventional Representations

- Directors: Nicolò Cesa-Bianchi, Andreas Krause, Bernhard Schölkopf
- Fellows and Scholars: Massimiliano Pontil, Barbara Caputo, Christoph H. Lampert, Volkan Cevher, Shie Mannor, Carl Edward Rasmussen, Csaba Szepesvari, Chris Watkins, Shimon Whiteson, Yishay Mansour, Nicolas Heess, Marc Deisenroth, Negar Kiyavash, Gergely Neu, Wouter M. Koolen, John Shawe-Taylor, Doina Precup, Stefanie Jegelka

We rethink the principles of interactive models of learning, exploring the role of causal modelling in bridging the gap between observational and interventional learning. The ultimate goal is to understand the organizing principles underlying robust intelligent behaviour, and to enable reliable learning-based decision systems for high-stakes real-world applications.

- Principles of learning-in-the-loop systems
- Online and reinforcement learning
- Causal inference
- Interacting learning systems (multi-agent learning, games, networks)

Machine Learning and Computer Vision

- Directors: Bernt Schiele, Cordelia Schmid, Yair Weiss
- Fellows and Scholars: Michael J. Black, Tomas Pajdla, Stefan Roth, Andreas Geiger, Thomas Brox, Rita Cucchiara, Zeynep Akata, Luc Van Gool, Andrea Vedaldi, Andrew Zisserman, Michal Irani, Ivan Laptev, Tinne Tuytelaars, Josef Sivic, Jiri Matas, Andrew Blake, Jean Ponce, Andrew Fitzgibbon, Daniel Cremers

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 by using artificial neural networks. Does this mean that the knowledge gained in the previous decades is now obsolete? In fact, many successful neural network approaches to computer vision implicitly employ insight gained during the study of classical algorithms: in the design of the datasets, the cost functions and the architecture. Our goal in this research program 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:

- Mid-level vision and image reconstruction
- 3D Geometry from multiple views
- Object and activity recognition

ELLIS Robot Learning: Closing the Reality Gap!

- Directors: Tamim Asfour, Aude Billard, Jan Peters
- Fellows and Scholars: Ingmar Posner, Oliver Brock, Dario Floreano, Danica Kragic, Manuel Lopes, Gerhard Neumann, Justus Piater, Davide Scaramuzza, Carme Torras, Marc Toussaint, Aleš Ude, Patrick van der Smagt, Sethu Vijayakumar, Wolfram Burgard, Jürgen Schmidhuber, Paul Newman

The goal of this ELLIS program is to develop the methodology for effective learning of actions and controls from the right amount of data and instruction for real world robot. It aims significantly improving robustness and flexibility of robots that interact with the real world. It will enable such robots to cope with changing conditions coming from the environment or potentially the robot itself, while endowing them with the capability of robustly interacting and collaborating with a human user. It focuses on the central questions for closing the real-world gap for any artificial intelligent real-world system:

- How should the robot move?
- How to act?
- How to interact?
- How can sensorimotor behavior be improved by machine learning approaches?

Human-centric Machine Learning

- Directors: Plamen Angelov, Nuria Oliver, Adrian Weller
- Fellows and Scholars: Manuel Gomez Rodriguez, Christian Theobalt, Nozha Boujemaa, Carlos Castillo, Ciro Cattuto, Sergio Escalera, Bruno Lepri, Chris Russell, Karen Yeung, Thomas Hofmann, Novi Quadrianto, Emilia Gómez, Krishna P. Gummadi, Lilian Edwards

As the use of machine learning becomes widespread in human-centric applications and

algorithmic decisions are more consequential to individuals and society, we need to tackle key limitations of today's machine learning systems, including algorithmic discrimination against minorities, manipulation of human decision-making, spread of misinformation and increase of polarization. This ELLIS program focuses on human-centric machine learning models and algorithms specifically designed to avoid the above limitations. Addressing these challenges requires interdisciplinary work and collaboration with experts at the intersection of machine learning, causality, human-computer interface (HCI), ethics and the law. To ensure widespread benefits to, and acceptance from the public the key research areas of the program are:

- transparency, accountability, interpretability and fairness of the algorithmic decisions
- amenability to legal and technical certification, accountability and verifiability

Theory, Algorithms and Computations of Modern Learning Systems

- Directors: Francis Bach, Philipp Hennig, Lorenzo Rosasco
- Fellows and Scholars: Julien Mairal, Thomas Schön, Ulrike von Luxburg, Arthur Gretton, Matthias Hein, Mário A. T. Figueiredo, Asja Fischer, Martin Jaggi, Simo Särkkä, Ingo Steinwart, Michael A. Osborne, Manon Kok, Arno Solin, Florence d'Alché-Buc, Niao He

Although contemporary Machine Learning algorithms achieve remarkable results, they often remain inefficient, unreliable, brittle, and require manual tuning. We will bring together experts with diverse backgrounds, linked by an interest in methods that are theoretically grounded as well as working in practice, to advance the algorithmic foundations of learning systems towards provably efficient and reliable learning systems.

- Many contemporary ML algorithms are still comparably badly understood conceptually. As a result, they require manual tuning, can be inflexible and behave erratically.
- The program connects experts with diverse backgrounds to advance the algorithmic foundations of ML. It will support the development of efficient and reliable learning systems with theoretical guarantees.

Quantum and physics based machine learning

- Directors: Bert Kappen, Riccardo Zecchina
- Fellows and Scholars: Lenka Zdeborová, Florian Marquardt, Giuseppe Carleo, Matthias

Rupp, Gábor Csányi, Giuseppe E. Santoro, Carlo Baldassi, Remi Monasson, Giulio Biroli, Vedran Dunjko, Florent Krzakala, Miguel A. Martin-Delgado, Frank Noé, Hans J. Briegel, Jens Eisert

The aim of the Ellis program Quantum and Physics based machine learning (QPhML) is to use concepts from quantum physics and statistical physics to develop novel machine learning algorithms with the ultimate aim to realize novel future, possibly energy efficient hardware implementations.

- Exploit quantum effects in machine learning
- Accelerate and improve energy efficiency of machine learning algorithms through dedicated physical implementations
- Use machine learning methods to advance understanding of quantum information processing

Natural Intelligence

- Directors: Matthias Bethge, Y-Lan Boureau, Peter Dayan
- Fellows and Scholars: Oriol Vinyals, Timothy Behrens, Matthew Botvinick, Emmanuel Dupoux, Sharon Goldwater, Raia Hadsell, Anne Hsu, Bradley C. Love, Mackenzie Mathis, Jane X. Wang, Nick Chater, Chelsea Finn, Li Zhaoping

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 behavior 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 program we focus on important aspects of natural intelligence, especially those that are poorly captured or missing in current machine learning models.

- Lifelong learning
- Deep semantics and cross-domain learning
- Shaping inductive bias via neural network structure
- Adaptive resource deployment
- Social reasoning

ELLIS Health

- Directors: Gunnar Rätsch, Oliver Stegle, Mihaela van der Schaar
- Fellows and Scholars: Jean-Philippe Vert, Sepp Hochreiter, Fabian Theis, Matthew Blaschko, Lena Maier-Hein, Guido Sanguinetti, Julia A. Schnabel, Christoph Bock, Magnus Rattray, Karsten Borgwardt

The ELLIS Health program seeks to strengthen the connection of the artificial intelligence (AI) and machine learning (ML) community 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 utilized in health, which is in part due to the domain-specific and complex challenges requiring new technical solutions. In addition to strengthening and extending the European community at the interface of AI/ML and health, the program has the ambition to democratize the most relevant and impactful datasets for methods-centric research in health, and to connect ELLIS to flagship European initiatives that generate and curate datasets that will build the foundations of this research. For all these reasons, ELLIS Health represents a central application area for modern AI/ML, a driver of theoretical and methodological innovation, and a visible showcase of what AI/ML can accomplish.

- Connect and promote the ELLIS vision within the broad areas of human health
- Demonstrate the impact of AI/ML on biomedicine and health
- Feed back key challenges of health applications into AI/ML methods development □
Foster training and education of the next generation of interdisciplinary scientists at the interface of health and AI/ML

Machine Learning for Earth and Climate Sciences

- Directors: Gustau Camps-Valls, Markus Reichstein
- Fellows and Scholars: Dino Sejdinovic, Konrad Schindler, Joachim Denzler, Jakob Runge, Sancho Salcedo-Sanz, Devis Tuia, Xiaoxiang Zhu, Veronika Eyring, Maria Piles, Jakob Zscheischler, Jocelyn Chanussot, Miguel D. Mahecha, Nuno Carvalhais, Bjorn Stevens

This program aims to contribute towards achieving explainable ML and improved extrapolation via 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 stepping stone before the more ambitious far-end goal of machine reasoning of anthropogenic climate change.

- Spatio-temporal anomaly and extreme events detection, anticipation and attribution
- Data-driven dynamic modelling and forecasting
- Hybrid modeling: linking physics and machine learning models
- Causal inference, Learning and explaining feature representations
- Earth and Climate model emulation, generative modelling and data-model fusion
- Benchmark synthetic and real datasets

Natural Language Processing

- Directors: André Martins, Iryna Gurevych, Ivan Titov
- Fellows and Scholars: Phil Blunsom, Anna Korhonen, Marie-Francine Moens, Rico Sennrich, Hinrich Schütze, Omri Abend, Isabelle Augenstein, Jonathan Berant, Shay Cohen, Ryan Cotterell, Barbara Plank, Roi Reichart, Lucia Specia, Yoav Goldberg, Mirella Lapata, Ido Dagan, Sebastian Riedel

Natural language processing (NLP) is transforming the way humans communicate with each other and with machines, with applications in multi-document summarization, machine translation, question answering, fact checking, decision support in health domains, or any applications which require making decisions that involve complex reasoning or entail combining different modalities (e.g., vision and text).

The goal of this program is to facilitate collaboration across the leading NLP labs and to encourage closer interactions between the NLP and Machine Learning communities.

Semantic, Symbolic and Interpretable Machine Learning

- Directors: Volker Tresp, Kristian Kersting, Paolo Frasconi

- Fellows and Scholars: Peter Flach, Marco Gori, Marta Kwiatkowska, Nada Lavrac, Jens Lehmann, Eran Yahav, Tim Rocktäschel, Edward Grefenstette, Michela Milano, Thomas Kipf, Stephen Muggleton, Luc De Raedt

This program concerns machine learning (ML) approaches which operate at the human abstraction level, where the world is described by entities, concepts, and their mutual relationships. Among other topics, we cover multi-relational learning, learning with (temporal) knowledge graphs, and the extraction of statistical and logical regularities from data. Methods include, e.g., embedding approaches, graph neural networks, scene graph analysis, neuro-symbolic programming, and inductive logic programming.

The goal is that our program becomes a cumulation point of like-minded researchers and we expect fruitful interactions with closely related programs covering, e.g., NLP, vision, and geometric deep learning.

Multimodal Learning Systems

- Directors: Cees Snoek, Nicu Sebe
- Fellows and Scholars: Alberto Del Bimbo, Björn Schuller, Elisa Ricci, Elisabeth André, Hervé Jégou, Sabine Sússtrunk, Shaogang Gong, Timothy M. Hospedales, Vittorio Ferrari, Vittorio Murino

The goal of this program is to provide a venue to disseminate, discuss and synthesize the latest developments in the area of multimodal learning systems. The program intends to push the boundaries of the foundational aspects of this field, to build bridges between researchers and practitioners currently active in multiple unimodal communities, as well as expanding and exploring the applications of multimodal learning systems.

ELLIS Units

ELLIS announced an [open call for proposals](#) of ELLIS units in September of 2019, with a submission deadline in November 2019. According to the call, an ELLIS unit will invest at least 1.5 Mio EUR per year, for at least five years, and requires a commitment to excellence as well as a contribution (at least 300k EUR annually) towards network activities, such as hosting exchange faculty and students, leading research programs and organizing workshops. Applications were reviewed by an international committee composed of ELLIS Fellows. To avoid conflicts of

interest, all the reviews were done by scientists of countries different from the units being evaluated. The first call resulted in 28 proposals from 13 countries. During the NeurIPS conference in December 2019 in Vancouver, [ELLIS announced](#) the first set of 17 units from 10 different countries that were considered to be ‘ready to be launched’ by the international evaluation committee. Meanwhile another 13 units in 2020, 4 more in 2021 and one unit in 2022 have been accepted. They are briefly described below.

Alicante

The ELLIS unit Alicante will focus on Humanity Centric Artificial Intelligence (HCAI) with three research areas that study different aspects of the intersection between humans and Artificial Intelligence such as (1) Human behavior modeling via machine learning techniques, (2) the development of intelligent, interactive systems, and (3) human-centric challenges in AI, including fairness, transparency, explicability/interpretability, privacy and veracity. The unit will be based at the University of Alicante in Spain. It will collaborate with the other universities and science parks in the Valencian region. It represents a unique opportunity to attract top research talent on AI to the region, to nurture and inspire the next generation of PhD students in AI and to have societal and economic impact. The unit in Alicante is fully aligned and supported by the recently published [AI Valencian Strategy](#).

Directors, Fellows, Scholars: Nuria Oliver

Amsterdam

The ELLIS unit Amsterdam aims at creating AI Technology for People. We strive to create societal and economic impact through fundamental research in deep learning in order to develop the decision making, information retrieval, natural language processing, and computer vision technology to empower people in their roles as citizens, clients, patients, consumers, creators, developers, employees, and entrepreneurs. The unit maintains close links to the University of Amsterdam (UvA) which is an international hotspot for deep learning powered research in AI. The ELLIS unit will further strengthen this research hub and connect it to other AI excellence centers in Europe.

Directors, Fellows, Scholars: Max Welling, Maarten de Rijke, Efstratios Gavves, Ekaterina Shutova, Eric Nalisnick, Vlad Niculae

Berlin

The goal of the ELLIS unit Berlin is to provide the scientific foundations in the fields of ML and, as a result, advance AI applications to yield a substantial benefit and progress for society, economy, and science. The ELLIS unit Berlin would collaborate with existing projects, such as BIFOLD, BZML, BBDC, MATH+ and the Science of Intelligence clusters of excellence, and multiple graduate schools (Graduiertenkolleg) and collaborative research centers (Sonderforschungsbereich), established by the German Research Foundation (DFG). By cooperating with universities in the Greater Berlin Metropolitan Area (e.g., Charite, FU Berlin, HU Berlin, TU Berlin, University Potsdam), scientific associations and societies as well as institutes of applied research (e.g., acatech, BBAW, DFKI, Fraunhofer, Helmholtz, Leibniz, Leopoldina, Max Planck) as well as with companies and startups, the ELLIS unit Berlin will help to spark innovation in a broad spectrum of applications in commerce, services and the sciences. The ELLIS unit Berlin unit will drive forward the following five priority research areas: (1) explainable AI, (2) scalable machine learning and data management, (3) machine learning for the sciences (in particular physics and chemistry, as well as the material sciences), (4) deep learning and (5) learning and inference with structure and priors.

Directors, Fellows, Scholars: Klaus-Robert Müller, Cédric Archambeau, Begüm Demir, Volker Markl, Frank Noé, Manfred Opper, Thomas Wiegand, Grégoire Montavon, Wojciech Samek

Cambridge

The mission of the ELLIS unit Cambridge is to build on the excellent machine learning and AI infrastructure already available within the University of Cambridge and serve as a stepping stone towards creating a center of excellence. Many of the members of the ELLIS unit Cambridge are strong in Bayesian statistics and probabilistic machine learning. These types of methods are expected to play a key role in addressing some of the limitations of existing approaches: lack of robustness, data-efficiency, uncertainty awareness, flexible adaptation and understanding causality. Other unit members work on specific application areas: language modeling, healthcare, computer systems and molecular modeling. The significant strength in probabilistic and Bayesian machine learning makes the ELLIS unit Cambridge unique in Europe and a valuable asset for the ELLIS society. Furthermore, the joint methodology used by the members of the unit, together with a diverse portfolio of expertise across different application areas, creates an excellent foundation for broad impact.

Director, Fellows, Scholars: Carl Edward Rasmussen, José Miguel Hernández-Lobato, Gabor Csanyi, Zoubin Ghahramani, Mark Girolami, Anna Korhonen, Neil D. Lawrence, Mihaela van der Schaar, Richard E. Turner, Adrian Weller

Copenhagen

The ELLIS Unit Copenhagen consists of machine learning faculty from Technical University of Denmark (DTU), University of Copenhagen (UCPH) and the IT University of Copenhagen (ITU). Its research agenda is focused on both machine learning methods and applications such as computer vision, health, earth and climate sciences that link to the corresponding ELLIS fellowship programs. Additional application areas are natural language processing (NLP) and material science. The overall mission of the unit is to strengthen machine learning research and innovation, increase the presence and visibility within the European research community and make Europe more competitive internationally.

Directors, Fellows, Scholars: Ole Winther, Christian Igel, Lars Kai Hansen, Serge Belongie, Isabelle Augenstein, Søren Hauberg, Jonas Martin Peters, Barbara Plank

Darmstadt

The ELLIS unit Darmstadt will lay the *foundations for intelligent systems that understand context and meaning, and are able to adapt accordingly*. The aim is to make computers learn as much about the world, so rapidly and flexibly, as humans. To this end, the research activities of the unit comprise problem areas such as deep (probabilistic) models, hybrid models, physics-aware machine and robot learning, Bayesian (nonparametric) models, (deep and/or inverse) reinforcement learning, probabilistic programming, vision/NLP/robotics, explainable and interactive AI, and human-centric AI. The unit is part of a strategic high-level initiative of TU Darmstadt called Intelligent Systems at TU Darmstadt (IS•DA) which broadly connects the local activities in the areas of artificial intelligence (AI).

Directors, Fellows, Scholars: Stefan Roth, Iryna Gurevych, Kristian Kersting, Jan Peters

Delft

The ELLIS unit Delft will offer unique research and education strengths of 13 selected,

high-profile faculty members in the following fields: (1) sequential decision making with applications in domains like robotics, vision, self-driving cars, smart energy systems, transportation, smart cities, and cybersecurity, (2) interaction, with applications such as in human-AI interaction and collaboration, crowdsourcing, and web technologies, (3) machine learning techniques for bioinformatics in the context of healthy living, disease diagnosis and prevention, and microbiological industrial processes. A unique characteristic of Delft University of Technology is the human aware context: The Delft engineer designs solutions with inherent focus on humans as part of the overall systems and human values such as autonomy, privacy and responsibility. The unit aims to create economic impact and create jobs in the Netherlands and Europe at large, maintaining an extensive project portfolio with societal and industrial partners.

Directors, Fellows, Scholars: Robert Babuska, Claudia Hauff, Wiro Niessen, Jens Kober, Frans A. Oliehoek

Edinburgh

The ELLIS unit Edinburgh will mainly focus on machine learning (ML) and natural language processing (NLP) research. The ML group studies computational processes that find patterns and structure in data. This forms the largest academic ML group in the UK and includes former programme chairs for three top international conferences in ML (NIPS, ICLR, ICML). They develop new ML and AI methodology often based on generalising the demands of cutting-edge application areas, including astronomy, systems biology, neuroscience, econometrics and healthcare. In Natural Language Processing, the University of Edinburgh has the highest concentration of academic NLP and speech technology in Europe. It is the world's most productive NLP research group, outranking Stanford, CMU, and Johns Hopkins.

Directors, Fellows, Scholars: Ivan Titov, Sharon Goldwater, Frank Keller, Alex Lascarides, Iain Murray, Mirella Lapata, Sotirios Tsafaris, Stephen Renals, Guido Sanguinetti, Amos Storkey, Sethu Vijayakumar, Chris Williams, Alexandra Birch, Shay Cohen, Oisín Mac Aodha, Siddharth Narayanaswamy,

Freiburg

The mission of the ELLIS unit Freiburg is to act as one of the best places in Europe to work in the field of robot learning and to foster European collaborations in the intersection of robotics, computer vision and machine learning. The unit has a unique set of outstanding robotics

researchers, accompanied by ELLIS fellows in the areas of computer vision and machine learning, and close connections to the industry (Toyota, Amazon, and Bosch). The unit will facilitate collaborations on the intersection of these fields within the ELLIS network, building on highly successful interactions between the respective groups within Freiburg itself. With the recent convergence of much of computer vision, robotics and machine learning under the common theme of deep learning, there are great opportunities for exploiting synergies between these fields.

Directors, Fellows, Scholars: Frank Hutter, Wolfram Burgard, Thomas Brox, Abhinav Valada

Heidelberg

The ELLIS unit Heidelberg will foster innovations at the interface of artificial intelligence (AI), machine learning (ML) and the biological and medical sciences. The mission of the unit is to facilitate breakthrough applications of AI/ML, delivering leading-edge analytics to fully exploit the rapidly growing volumes of biomedical data across Europe. The unit will conduct foundational research to address key challenges and obstacles for deploying AI in biomedicine.

This includes methods to cope with the heterogeneous and often noisy nature of “omics” data and the scarcity of labeled data in medical imaging, algorithms and infrastructures to deal with ethical and privacy constraints of data access, algorithms to infer causal relationships, as well as novel modelling strategies to deliver interpretable, auditable decisions. The research will be carried out with a particular focus on (1) un-/semi-/self-supervised learning to tackle lack of training data, (2) explainability and uncertainty quantification, (3) privacy aware and federated learning, (4) sparse predictions from dense inputs and (5) interpretable low-dimensional representations and metric learning. Research activities of the unit will be complemented with a pan ELLIS service mission, democratizing access to curated reference datasets, as well as organizing and certifying benchmarks of novel AI-based analytics.

Directors, Fellows, Scholars: Oliver Stegle, Fred A. Hamprecht, Wolfgang Huber, Carsten Rother, Klaus Maier-Hein, Lena Maier-Hein, Anna Kreshuk

Helsinki

The ELLIS unit Helsinki builds on the long tradition and track record of pioneering machine learning research in Finland and seeks to contribute to a concerted European effort in basic research in machine learning. In particular, the unit focuses on (1) Probabilistic modeling and Bayesian inference, (2) Simulator-based inference, (3) Data-efficient deep learning, (4)

Privacy-preserving machine learning and (5) Interactive artificial intelligence. The faculty and the operations of the ELLIS unit Helsinki has close links to the Finnish Center for Artificial Intelligence (FCAI) which is a nation-wide center for AI, combining fundamental AI research with a broad range of applied AI research. The ELLIS unit Helsinki will support the FCAI mission to create a new type of AI, which is able to operate with humans in the complex world - and to renew industry.

Directors, Fellows, Scholars: Samuel Kaski, Aapo Hyvärinen, Aki Vehtari, Antti Oulasvirta, Jaakko Lehtinen, Jukka Corander, Petri Myllymäki, Simo Särkkä,.

IIT & Genoa

The overarching goal of the ELLIS unit IIT & Genoa is to facilitate synergies between machine learning, robotics and the study of natural intelligence. The included research areas are (1) Machine Learning: from DATA to Artificial Intelligence, (2) From Natural to Artificial Intelligence and (3) Robotics: From Intelligence to Action. The mission of the ELLIS unit IIT & Genoa is first to develop foundational research in ML/AI and to strengthen ties with applied research area at IIT and University of Genoa (including Robotics, Natural Intelligence, and ML for Health), cultivating an intellectually stimulating and engaging environment for faculty (PIs), junior researchers, and PhD students working across these areas. The unit also plans to strengthen and expand the research activity in AI Genoa by hiring in key areas such as Human-centric ML, NLP, Interactive and Reinforcement Learning, and Online Learning, Mathematics of Machine Learning. Last but not least, the ELLIS unit IIT & Genoa commits to promote research excellence in ML and modern AI in Italy via dissemination and training activities, ultimately making the unit the reference point for machine learning in the country and a key node for machine learning in Europe.

Directors, Fellows, Scholars: Massimiliano Pontil, Lorenzo Rosasco, Alessandra Sciutti, Arash Ajoudani, Lorenzo Natale, Stefano Panzeri, Vittorio Murino

Jena

The ELLIS unit Jena is involving two research institutions and the university at the interface between climate/environmental science and machine learning with an emphasis on scientific knowledge integration in and knowledge generation from machine learning approaches. The goal of the ELLIS Unit Jena is thus to combine fundamental development in machine learning with challenges concerning spatio-temporal environmental and climate dynamics for a better

understanding of the Earth system and its components. An important aspect here is the integration of knowledge into machine learning methods as appropriate assumptions - this can be qualitative knowledge about causal relationships ("causal modeling") or quantitative knowledge about functional relationships, which can be "cast" into physical, chemical, biological or geo-ecological formulas ("hybrid modeling"). The ELLIS unit Jena will also strive to help AI to best serve sustainable development via links to politics, business and civil society.

Directors, Fellows, Scholars: Joachim Denzler, Markus Reichstein, Jakob Runge, Nuno Carvalhais

Linz

The ELLIS unit Linz will contribute to coordinating machine learning excellence in Europe and to establish a local sustainable ecosystem of machine learning stakeholders covering the entire value network to facilitate and accelerate a broad uptake and integration of Machine Learning technologies. The unit will conduct basic machine learning research at the highest levels in coordination with other ELLIS sites and thereby advance theories, algorithms, and applications of machine learning. The unit will be established on the premises of the LIT AI Lab located at the Johannes Kepler University Linz (JKU), and will have financial support from University and industrial partners with a budget of ~25.0 million € across five years.

Directors, Fellows, Scholars: Sepp Hochreiter, Gerhard Widmer, Günter Klambauer, Johannes Fürnkranz

Lausanne (EPFL)

The ELLIS unit Lausanne brings together researchers working on various aspects of creating intelligent systems. The fundamental research and its applications cross the boundaries of schools at EPFL and create opportunities for direct interactions with Swiss industry. It is organized around 4 main pillars (Math & theory, software & hardware, technology and ethics) and reaches out to various application domains such as autonomous systems, wearable technologies, robotic co-workers, smart devices, drone technology, augmented reality, intelligent houses, digital twins, intelligent manufacturing, and many more. The unit is aligned with one of the strategic domains of EPFL, and of the Swiss Confederation, and is closely linked to the newly created Center for Intelligent Systems (CIS) at EPFL, a joint initiative of the schools IC, SB and STI, which seeks to advance research and practice in the strategic field of intelligent systems.

Directors, Fellows, Scholars: Pascal Frossard, Alexandre Alahi, Amir Zamir, Aude Billard, Dario Floreano, David Atienza, Devis Tuia, Florent Krzakala, Giuseppe Carleo, Lenka Zdeborová, Mackenzie Mathis, Martin Jaggi, Negar Kiyavash, Nicolas Flammarion, Sabine Süsstrunk, Sofia Olhede, Volkan Cevher, Wulfram Gerstner

Leuven

The ELLIS unit Leuven comprises the research activities of five highly active faculty selected for academic excellence and complementarity of research areas. The focus of the unit is to develop a fast conduit between key application areas, core machine learning technologies, and real-time implementations on edge devices. The research focuses on advancing machine learning methods such as representation learning, continual learning, neural network compression and discrete optimization methods suitable for optimized edge implementations. In addition, it focuses on applications such as computer vision, self driving cars, and cultural heritage, natural language processing and multimodal data modeling, as well as chip design, resource efficient machine learning processing, and health applications in close collaboration with the KU Leuven Medical Imaging Research Center.

Directors, Fellows, Scholars: Matthew Blaschko, Johan Suykens, Luc De Raedt, Luc van Gool, Marian Verhelst, Marie-Francine Moens, Tinne Tuytelaars

Lisbon

The mission of the ELLIS unit Lisbon is (1) boosting collaborative research and higher education in artificial intelligence (AI) and machine learning (ML) in Portugal and Europe, and (2) empowering AI researchers to become active agents in maximizing the social and economic impacts of ML&AI in Europe and the world. The ELLIS unit Lisbon will conduct cutting-edge research in the following AI-related areas: Natural Language Processing, Machine Learning and Optimization, Reinforcement Learning & Robotics, Computer Vision & Cognitive Robotics, Networks and Infrastructure, and Computational Biology. The ELLIS unit Lisbon will bring together researchers in these fields with the common goal of designing human-interacting explainable AI systems: this involves a strong bet on human language technologies, social and cognitive robotics, and computational biology, well supported on the development of efficient and reliable ML systems with theoretical guarantees. The unit will exploit synergies between these areas towards the unified goal of designing AI systems that interact seamlessly with humans, learning from them and for them.

Directors, Fellows, Scholars: Mário A. T. Figueiredo, André Martins, Manuel Lopes

London (UCL)

The proposed unit will integrate research in artificial intelligence in UCL, which takes place across multiple departments. The UCL centre for artificial intelligence in the Computer Science Department is one of the world's leading AI research organizations. It comprises the Machine Reading Group, whose goal is to build machines that can read and "understand" unstructured textual information; the Computer Vision Group, which aims to extract useful information from images and movies; and the Computational Statistics and Machine Learning Group, which aims to make methodological progress in foundational AI using techniques from statistics, mathematics and computer science. The core aim of the AI Centre is to create new AI technologies and advise on the use of AI in science, industry and society, as AI becomes more prevalent across all aspects of modern life. The AI Centre takes inspiration from the vast array of applications across UCL and acts as the engine of methodological progress. It has close research and training links to other local AI communities, including the Gatsby Computational Neuroscience Unit and the Alan Turing Institute.

Directors, Fellows, Scholars: Arthur Gretton, David Barber, Emine Yilmaz, John Shawe-Taylor, Maneesh Sahani, Marc Deisenroth, Matt J. Kusner, Pontus Stenetorp

Madrid

The ELLIS Unit Madrid is a partnership made up of professors and researchers in machine learning from the six public universities in Madrid: Universidad Autónoma de Madrid, Universidad Carlos III de Madrid, Universidad Complutense de Madrid, Universidad de Alcalá, Universidad Politécnica de Madrid and Universidad Rey Juan Carlos. The focus of the unit is to develop ground-breaking interpretable probability-based causal machine learning methods for dynamic scenarios and cross-cutting quantum technologies for intelligent systems. Successful application areas include biometrics, computer vision, healthcare, renewable energy, climate, robotics and intelligent vehicles.

Directors, Fellows, Scholars: Miguel A. Martin-Delgado, Sancho Salcedo-Sanz

Manchester

The ELLIS Unit Manchester offers (i) a link to a large, strong community across a broad range of data science and AI, which aims to considerably strengthen its machine learning action through the ELLIS Unit and new recruitments to it, (ii) outstanding translational opportunities in health; it is affiliated with the Pankhurst Institute having its main mission in translation, and (iii) a regional hub. Through the unit, it is intended to significantly contribute to a concerted European effort in basic research in machine learning.

Directors, Fellows, Scholars: Magnus Rattray, Samuel Kaski

Milan

The ELLIS unit in Milan brings together excellent researchers from four institutions: Bocconi University, Politecnico di Milano, University of Milan, and University of Milan-Bicocca. The unit is active in several research areas: interactive learning and game theory, statistical learning and non-convex optimization, health and computational biology, natural language processing, computational social sciences, and neural networks in connection with classical AI and neuroscience.

Directors, Fellows, Scholars: Nicolò Cesa-Bianchi, Riccardo Zecchina, Sonia Petrone, Carlo Baldassi,

Munich

The ELLIS unit Munich focuses on developing novel machine learning methods and deploying them in the three key applications areas Biomedicine, Computer Vision and Earth Observation. It integrates expertise from the Technical University of Munich, Helmholtz Munich and collaborating partners and it will represent the machine learning core for the Munich Data Science Institute (MDSI). The members of the ELLIS unit Munich define four areas they build on their unique expertise and thriving collaborations: (1) foundations of machine learning on graphs and time series, optimal control, knowledge graphs. novel training methods, (2) biomedicine (with applications in human functional genomics, biomedical imaging and bioengineering, drug research and electronic health records), (3) Earth observations (elaboration of remote sensing data, in particular those from current and the next generation of Earth

observation missions such as TerraSAR-X, TanDEM-X, TerraSAR-X follow-on, Tandem-L and EnMAP, big data for 4D global mapping) and (4) Computer Vision (with applications such as image restoration, multiple-view reconstruction, visual SLAM, dynamic scene understanding, multiple object tracking). Directors, Fellows, Scholars: Fabian Theis, Daniel Rückert, Eleftheria Zeggini, Hinrich Schütze, Julia A. Schnabel, Laura Leal-Taixé, Massimo Fornasier, Mathias Drton, Matthias Nießner, Patrick van der Smagt, Volker Tresp, Xiaoxiang Zhu

Oxford

The ELLIS Unit Oxford will serve as a focal point, bringing together the disparate departments and institutes under one roof. While initially ELLIS@Oxford will be a virtual unit with bases across the 3 core departments, our intention is for it to crystallise into a research institute conducting world leading AI/ML research, educating the next generation of scientists and technologists, incubating startups and high impact applications, and guiding thinking on the societal impacts on new AI/ML technologies. Establishing an ELLIS Unit is a crucial first step towards this goal. It allows a single point of contact with industries and government bodies to raise funds and investments. It creates a critical mass for the development of university structures enabling impact, including incubators and teaching and admin buyouts. And it enables dissemination of AI/ML technical knowhow across the university, which will be important for the AI/ML revolution to take hold across Oxford.

Directors, Fellows, Scholars: Michael Wooldridge, Stephen Roberts, Yee Whye Teh, Alison Noble, Andrea Vedaldi, Andrew Zisserman, Chris Holmes, Dino Sejdinovic, Edith Elkind, Ingmar Posner, Marta Kwiatkowska, Michael A. Osborne, Phil Blunsom, Philip H. S. Torr, Shimon Whiteson, Timothy Behrens, Varun Kanade, Yarin Gal

Paris

The main goal of the ELLIS unit Paris is to offer a common platform for Parisian AI researchers and students to collaborate within the ELLIS network. The unit will foster exchanges and collaborations both within the Paris area and across Europe. The unit will create a bridge on topics related to AI between the two main geographical locations (Paris center, Paris Saclay) and academic entities (PSL Université, Sorbonne Université, Université Paris-Saclay, Institut Polytechnique de Paris, Université de Paris) and is supported by the two main research agencies CNRS and Inria. The research topics of the ELLIS Unit Paris are focused on core research topics of AI (ranging from theory to applications), with strong interdisciplinary nature. It will 1) advance fundamental research in AI, in particular in core machine learning and related fields

(vision, robotics, NLP), 2) support interdisciplinary research in AI, in particular in physics, health, biology and humanities, and 3) promote open source software and reproducible research.

Directors, Fellows, Scholars: Gabriel Peyré, Bertrand Thirion, Christian P. Robert, Cordelia Schmid, Emmanuel Dupoux, Florence d'Alché-Buc, Francis Bach, Giulio Biroli, Isabelle Guyon, Ivan Laptev, Jean Ponce, Michèle Sebag, Nicolas Vayatis, Rémi Flamary, Stephane Mallat

Prague

The ELLIS unit Prague is committed to making ELLIS the leading open science AI organization in the world by: (1) outstanding foundational research in AI and related disciplines, (2) supporting the mobility of researchers within ELLIS and elsewhere, (3) building a European brand of PhD and Postdoc program, and (4) transferring research results to Czech as well as European industry to boost economic and societal innovation in Europe. The unit is hosted at the Czech Institute of Informatics, Robotics and Cybernetics (CIIRC) of the Czech Technical University and brings together five internationally recognized researchers and their teams that cover several key research areas necessary for building intelligent autonomous systems. In five years the target is to reach at least 10 ELLIS faculty with their research teams, commitment from funding sources for 15 years, and an innovation engineering team with an industry coordination office to ensure impact on economy and society in the Czech Republic and Europe via technology transfer and support for the creation of startups.

Directors, Fellows, Scholars: Josef Sivic, Tomas Pajdla, Josef Urban, Tomas Mikolov, Torsten Sattler

Nijmegen

The ELLIS unit Radboud promotes fundamental research in machine learning and their application in life sciences, by focusing primarily on statistical approaches. The machine learning research at the ELLIS unit Radboud focuses on elucidating the basic mechanism of information processing in biological systems as well as improving healthcare. This research is supported by several institutes and programs as (1) the neural computation theme of the Donders Institute for Brain, Cognition and Behaviour, which focuses on elucidating the computational mechanism underlying neural information processing, (2) the ICAI AI for Health lab, which aims to implement AI techniques for improving healthcare, (3) the bits and brains program which focuses on neuromorphic computing, (4) coordination of a recently established European network on brain-inspired computing (iNavigate, H2020-MSCA-RISE-2019 Program), (5) coordinating the AI component of a recently awarded 14.3M euro CrossOver grant on Neurotechnology and (6)

the recent establishment of a 65M euro IMEC OnePlanet Research Centre at Radboud University, which focuses on health and sustainability.

Directors, Fellows, Scholars: Marcel van Gerven, Bram van Ginneken, Elena Marchiori, Hilbert Johan Kappen, Tom Heskes, Umut Güçlü

Saarbrücken

The ELLIS Unit Saarbrücken brings together six PIs from four research institutions of Saarland Informatics Campus (SIC). It comprises the Max Planck Institute of Informatics (MPI-INF), the Max Planck Institute of Software Systems (MPI-SWS), Saarland University (UdS), and the CISPA Helmholtz Center for Information Security (CISPA). The PIs have agreed to jointly work on both the foundations for enhanced functionalities of Artificial Intelligence and Machine Learning (AIML) systems and the pressing needs for security, privacy, and trustworthiness that arise from the widespread use of AIML systems. In the future, these systems will capture reality through a multitude of sensors, interact with humans, derive knowledge, and influence our lives. They will make autonomous decisions and enable enhanced functionalities, e.g., for personalized assistants, immersive multi-person platforms, and autonomous systems. This ELLIS Unit will research the foundational algorithmic solutions needed to enable systems capable of such new functionalities and design new algorithms for high trustworthiness, controllability, and explainability.

Directors, Fellows, Scholars: Bernt Schiele, Manuel Gomez Rodriguez, Christian Theobalt, Isabel Valera, Mario Fritz

Stuttgart

The Stuttgart ELLIS Unit brings together an interdisciplinary team of PIs at the University of Stuttgart and the Stuttgart site of the Max Planck Institute for Intelligent Systems (MPI-IS). The PIs have joined forces to advance research in learning and intelligent systems from four synergistic perspectives: Interactive Intelligent Systems, Natural and Programming Language Processing, Learning Theory, and Robot Learning.

Directors, Fellows, Scholars: Ingo Steinwart

Haifa

The ELLIS unit Haifa is developed as part of the Technion's newly devised "Interdisciplinary

Program for Research in Machine Learning and Intelligent Systems” that is formed by the Technion’s president to deepen the Technion’s commitment to machine learning. The program was recently approved and a budget of \$1M was already established. This ELLIS unit Technion serves three main purposes: (i) creates a platform to engage and collaborate in cross faculty projects in machine learning; (ii) coordinates research efforts with the industry to benefit both the Technion and the industry in advancing the reach of machine learning (iii) establishes an international network, with Europe, that will increase the impact of machine learning, via student exchange program, visiting faculty program and holding ELLIS workshops.

Directors, Fellows, Scholars: Shie Mannor, Alex Bronstein, Yonatan Belinkov, Roi Reichart

Tel Aviv

The ELLIS unit Tel Aviv covers broad aspects of the field including machine learning theory, natural language processing, machine vision, reinforcement learning and others. The unit will make a concerted effort to advance specific topics such as (1) Vision and Language, (2) Theory of Deep Learning, (3) Privacy and Fairness, (4) Generative Models, and (5) Common Sense Knowledge. The unit further aims at addressing societal challenges that are a new and important aspect of machine learning, as the predictions of algorithms impact humans in a significant way. The unit’s research on privacy and fairness has high potential for enabling new technologies that can lead to broader and safer usage of AI in improving a broad range of aspects of society (e.g., healthcare). In addition, the unit is embedded in Israel’s flourishing startup community, and many startups have been founded by and are populated by its machine learning graduates.

Directors, Fellows, Scholars: Amir Globerson, Yishay Mansour, Lior Wolf, Jonathan Berant, Nadav Cohen, Tomer Koren

Tübingen

The goal of the ELLIS Unit Tübingen is to build a novel public research institution to attract the best scientists to advance AI, train top international students, and generate positive impact in science and society. The research agenda of the unit aims at building learning systems that approach the versatility and robustness exhibited by natural intelligent systems. Machine learning (ML) is at the heart of a technological and societal revolution, yet today’s learning systems do not generalize well to new situations, cannot learn from few examples, and do not infer causal relationships. Addressing these deficits and developing robust AI systems will be key to efficient robot teaching and explainable AI and thus help ensure technological leadership and deploying AI systems responsibly and to the benefits of society. Furthermore, the ELLIS Unit actively

supports its members to tackle highly visible research challenges and to engage in start-ups, part-time employment, and other activities that help society benefit from AI research.

Directors, Fellows, Scholars: Matthias Bethge, Bernhard Schölkopf, Andreas Geiger, Bob Williamson, Gerard Pons-Moll, Jakob Macke, Matthias Hein, Michael J. Black, Moritz Hardt, Peter Dayan, Philipp Berens, Philipp Hennig, Ulrike von Luxburg, Zeynep Akata

Vienna (IST Austria)

Building on IST Austria's mission statement the ELLIS unit Vienna will conduct basic research in machine learning and related areas, foster interdisciplinary interaction between scientists and scientific disciplines, and provide a world-class environment for science and an attractive destination for doctoral students, postdocs, and professors from all countries. The research focus comprises (1) *Core Machine Learning* such as transfer learning, trustworthy learning and theory of deep learning, (2) *Optimization* covering both continuous as well as discrete optimization, large-scale distributed optimization, and inference in graphical models, (3) *Computer Vision and Image Processing* with emphasis on generative image models, natural image statistics, and automatic scene understanding, and (4) *Statistical Models for the Life Sciences* to unravel the information processing capabilities of biological networks, build predictive models for large scale biological data and develop Bayesian methods for personalized medicine.

Directors, Fellows, Scholars: Christoph Lampert, Dan Alistarh, Matthew Robinson, Vladimir Kolmogorov

Zürich (ETH)

The ELLIS unit Zürich will be established primarily around a core group of ETH faculty members working on methodological aspects of machine learning and emphasize the use of ML in health, life sciences, environmental sciences, and human-machine interaction. The unit will be developed in close coordination and integration with a new initiative supported by the ETH leadership towards establishing an AI Center/Institute within ETH. It will primarily focus on the following research areas: (1) Theoretical and methodological foundation for reliable and trustworthy ML (e.g. inductive bias of deep networks, fairness/ transparency/ robustness, causality, reinforcement learning), (2) Machine learning for Personalized Health, (3) Interactional intelligence and computational pragmatics and (4) Machine Learning for Remote Sensing and Environmental Modeling.

Directors, Fellows, Scholars: Andreas Krause, Gunnar Rätsch, Joachim M. Buhmann, Konrad Schindler, Niao He, Niko Beerenwinkel, Otmar Hilliges, Siyu Tang, Thomas Hofmann, Valentina Boeva, Karsten Borgwardt

ELLIS Society and Governance

The ELLIS society was established at the NeurIPS conference in December 2018 in Montreal. Membership in the ELLIS society is limited to scientists that (i) have a PhD or equivalent; (ii) regularly publish at top tier venues in machine learning driven fields¹⁰, and (iii) work in Europe, are European (working elsewhere) or have some other relation to Europe. Fellows and Scholars of ELLIS Programs automatically become members; others require endorsements from two existing members. ELLIS initially was governed by the founding board composed of Barbara Caputo (Italy), Nuria Oliver (Spain), Bernhard Schölkopf (Germany, chair), Max Welling (Netherlands) and deputies Matthias Bethge (Germany), Andreas Geiger (Germany), Sepp Hochreiter (Austria) and Josef Sivic (France and Czech Republic).

In July 2022 the General Assembly voted on a new board of directors. This board is composed of Serge Belongie (University of Copenhagen & Cornell University), Nicolò Cesa-Bianchi (Università degli Studi di Milano), Florence d'Alché-Buc (Télécom Paris), Nada Lavrac (Jožef Stefan Institute), Neil D. Lawrence (University of Cambridge), Nuria Oliver (DataPop Alliance, Vodafone Institute & The Spanish Royal Academy of Engineering), Bernhard Schölkopf (Max Planck Institute for Intelligent Systems) and Josef Sivic (Czech Technical University, École Normale Supérieure & INRIA). Sepp Hochleitner (Johannes Kepler University Linz) will provide support as permanent guest of the board.

The ELLIS Society e. V., an association under German law, is legally represented by Bernhard Schölkopf as President, Nuria Oliver as Vice President and Daniela Diaconu as Treasurer.

In August 2022 the ELLIS Society e.V. has over 900 members.

¹⁰ Examples include NeurIPS, ICML, ICLR, JMLR, CVPR, ICCV, ECCV, IJCV, PAMI, RSS, CoRL, ACL, EMNLP, KDD, AAAI or IJCAI.

PhD & Postdoc Program

The pan-European ELLIS PhD & Postdoc program aims to attract, keep and educate the best AI talent in Europe, and to prepare the next generation of talent for European research and industry. The highly selective ELLIS PhD & Postdoc program creates a recognized world-class brand and coordinated selection process that will be the reference entry point for PhD & postdoc candidates in machine-learning driven fields in Europe and world-wide. This program also disseminates knowledge by mixing academic and industrial sites. The main mechanism for selecting the participants is as follows: candidates apply to a joint application portal where they indicate their preferences for fields of study, locations and advisors. Candidates are then matched to and interviewed by potential advisors. Candidates retained by individual advisors then go through the admission process in the degree granting program of the advisor's home institution.

The PhD & Postdoc program relies on the ELLIS Fellows' network described above that attracts top talent to Europe. Each ELLIS PhD student will have a primary advisor from one of the degree granting participating institutions and a secondary advisor from another site (in another country) in the network. This secondary advisor can be from an industrial laboratory. One of the advisors must be an ELLIS Fellow or Faculty, the other one must be an ELLIS Member. The PhD & postdoc programme encourages internships at companies or startups, in particular companies and startups from within Europe to foster the local ecosystem. Towards this goal, yearly joint academic/industry workshops will be organized where invited faculty/advisors will hold tutorials/keynotes, PhD candidates and postdocs present their research, and companies present their problems and research programs. Similar smaller workshops will also be organized at local sites to encourage interaction with local industry and the local AI ecosystem. To spread knowledge across sites and to industry we will organize regular summer/winter schools. The primary example is the regular Machine Learning Summer School of which the next edition will be held in 2020 in Tübingen.

The ELLIS PhD program started in two phases:

- In 2020, all ELLIS fellows were invited to nominate PhD students and postdocs as ELLIS PhD students and postdocs. The conditions for becoming an ELLIS PhD student or postdoc in this first phase of the program is that the PhD student or postdoc is supervised by an ELLIS fellow or faculty and conducts a research exchange at an institution located in another country than the primary institution, hosted by an ELLIS member and lasting for at least 3 months.
- Starting from 2021, ELLIS PhD students and postdocs are also recruited via the central

recruitment as described above. The deadline is once a year, usually scheduled for November, where applicants submit their materials and advisor preferences, followed by remote interviews and matchmaking. Newly recruited ELLIS PhD students are co-supervised by (1) an ELLIS Fellow/Faculty and (2) an ELLIS member from another country, and conduct a research stay at the exchange institution that lasts between 6 months and 1 year which will be supported by central ELLIS travel funds with a lightweight application procedure. The research exchange period may also be split into several shorter periods.

ELLIS PhD students (and postdocs) will be invited to participate in ELLIS workshops of suitable ELLIS programs. This gives students access to the strongest European researchers in their field and further adds to the attraction of the PhD program. A number of industry researchers are directly involved in ELLIS research as Fellows and Members and are thus eligible to co-supervise ELLIS PhD students (and postdocs).

Donations and Sponsorship Program

In order to fund ELLIS activities, such as the PhD & postdoc programs, research workshops and exchanges, summer schools, online courses, etc. we are requesting donations and sponsorship contributions from institutions, companies and individuals at three possible levels: Gold, Silver and Bronze.

The ELLIS Society e.V. is a non-profit association and can accept both tax-deductible material and monetary donations for the implementation of its activities

We offer four options to contribute::

1. Becoming a Supporting Member of the ELLIS Society e.V. or
2. donating directly to the ELLIS Society or
3. sponsoring an ELLIS unit and donating an additional reduced amount to the ELLIS society or
4. sponsoring PhD students and/or postdocs supervised by ELLIS Fellows and/or ELLIS faculty and donating an additional reduced amount to the ELLIS society.

Contribution **option 1:** Becoming a Supporting Member of the ELLIS Society

There are several levels of Supporting Memberships defined:

- Platinum: min. €50k per year
- Gold: min. €2k per year
- Silver: min. €10k per year
- Bronze: min. €5k per year
- Standard: min. €1k per year

Contribution option 2: Become a donor of the ELLIS Society e. V.

You can support the work of the ELLIS Society at any time by making a donation of any amount.

Contribution option 3: Sponsoring the ELLIS Society

For the first, standard sponsorship program, three levels of sponsorship are defined:

- Gold: €100k per year
- Silver: €50k per year
- Bronze: €25k per year

Contribution option 4: Sponsoring an ELLIS Unit.

For the second sponsorship program the company can donate an unrestricted gift to a local ELLIS unit. The sponsorship levels are the same as for the first model above, but with an extra 20% fee to support the ELLIS society:

- Gold: €100k per year for the unit + €20k per year for the society
- Silver: €50k per year for the unit + €10k per year for the society
- Bronze: €25k per year for the unit + €5k per year for the society

Sponsoring (local) ELLIS Faculty

In the third sponsorship program, a company can sponsor PhD students and/or postdocs associated with an ELLIS Fellow or an ELLIS Faculty (the latter associated with an ELLIS unit).

Being a sponsor of ELLIS or an ELLIS unit has the following benefits:

- Your company's name will appear on the ELLIS sponsorship website indicating your level of sponsorship (Gold, Silver or Bronze)
- In case your company supports a PhD student at an ELLIS unit, you can recruit the student from a central pool of ELLIS PhD applicants.

- Preferential access to ELLIS conferences and workshops, summer-schools, and other ELLIS-wide events with benefits and representation proportional to the level of sponsorship (depending on the event)
- Subscription to the ELLIS newsletter
- The option to advertise internships and job openings directly to the ELLIS PhD-student and postdoc population through the newsletter and the ELLIS website.

The supported “ELLIS industrial PhD students and/or postdocs” under this program are co-advised by an ELLIS fellow or ELLIS faculty at the university and an ELLIS member at the company, and is subject to the requirements for standard ELLIS PhD students and postdocs, including mandatory exchanges. However, ELLIS industrial students and postdocs have the opportunity to spend their exchanges at the sponsoring company. The expectation is that all research results will be published at top tier conferences and journals, including the release of software for the purpose of scientific reproducibility (as is required by many of the top publication venues). IP agreements should be negotiated between the company and the university.

To become a Gold sponsor, the company should sponsor at least 4 PhD students at the local rate of the employing university (or the equivalent amount on a mix of PhD students and postdocs), plus an additional donation of €20k/yr for the ELLIS society. (Note that this amount may be different for different countries depending on the cost of a PhD student.) To become a Silver sponsor, the equivalent of 2 PhD students need to be invested plus an additional donation of €10k/yr for the ELLIS society, and for Bronze the equivalent of 1 PhD student plus an additional donation of €5k/yr for the ELLIS society must be invested.

The details of the sponsorship agreement will be valid for one year (option 1, 2 and 3) or the duration of the PhD/postdoc appointment (option 4) and will be regularly reevaluated by the ELLIS board. Any questions about this sponsorship program should be directed to sponsorship@ellis.eu.

Relationship to other initiatives

ELLIS wants to foster European research excellence in machine learning and related fields by offering European researchers and students outstanding opportunities to carry out their research in Europe. Towards this end, ELLIS established (I) research programs connecting Fellows across Europe, (II) the ELLIS PhD program, (III) the ELLIS sponsorship program to link European enterprises of all sizes with top researchers and graduates, and (IV) the first ELLIS units. ELLIS

is open to work together with other initiatives that share our goals and are willing to help us realize the above measures which we consider effective instruments to have a positive impact on European society. After ELLIS published its open letter in April 2018, another initiative was created, called CLAIRE, with a stronger emphasis on classical AI. We early on agreed with CLAIRE that *we will endorse each other's proposals, acknowledge their complementary nature, and coordinate efforts while preserving the unique characteristics of the two proposals*. We welcome that both initiatives were successful in the ICT-48 call (see also below), providing them with the means to work towards strengthening European AI. For ELLIS, this allows us to retain our focus on learning-based AI and scientific excellence in this highly competitive field. Research in democracies is characterized by plurality and competition of ideas, and it would thus be counterproductive if Europe were to put one group alone in control of the overall research ecosystem.

Partner Networks

ELISE

In 2019, the EU issued the call ICT-48-2020: *Towards a vibrant European network of AI excellence centres*. The objective of this call is to develop networks of excellence centres aiming at boosting the research capacity in Europe and the status of Europe as a research powerhouse for AI, and making it attractive for scientists and new talent. This initiative is also expected to contribute to the development of ethical and trustworthy Artificial Intelligence, the trademark for AI “made in Europe”.

The proposal “**ELISE, for European Learning and Intelligent Systems Excellence**” initiated by ELLIS has been selected to enter Grant Agreement Preparation (to be funded with 12 Mio EUR) with the top scores awarded among all proposals for this round, which is testament to the fact that machine learning is at the core of modern AI research, and a main driver in this thriving research field.

The proposal was formed by 202 core contributors from 105 organizations. The proposal is a combination of excellent research in ELLIS and related fields, and a set of actions, such as PhD and Fellows programs, designed to make Europe even stronger according to the call. Two sample facts illustrate the scientific excellence and innovative economic impact of ELISE: its members have founded a large number of startups (one of every two members of ELISE did start a

company), and in total 60 ERCs have been awarded to members of ELISE.

With the ELISE proposal, ELLIS and partners volunteer to take responsibility for supporting research excellence, and of becoming a reference in AI foundations, providing a balanced view of AI in the EU and integrating the EU vision into AI research.

An integral element of the proposal is the collaboration between different research fields and AI approaches and with industry, which has an urgent need for modern AI approaches, and well-informed crossovers between industrial/societal needs and AI expertise. Subject to the common goal of making Europe attractive to new and established talent in AI, the proposed network of excellence centers is very open for new collaboration complementing the already strong network, as excellent AI researchers are the motor of digital innovation also in Europe.

ELSA - European Lighthouse on Secure and Safe AI

ELSA¹¹ is a virtual center of excellence that will spearhead efforts in foundational safe and secure artificial intelligence (AI) methodology research. A large and growing network of top European experts in AI and machine learning is to promote the development and deployment of cutting-edge AI solutions in the future and make Europe the world's lighthouse of AI. ELSA builds on and extends the existing internationally recognized and excellently positioned ELLIS network of excellence.

ELSA rests like ELLIS on three pillars: selected research programs, local research units, and a PhD and postdoc program. ELSA will establish a fluid two-way communication between academic research and real-life application scenarios through the creation of the ELSA Innovation Lab. These dynamics will be strengthened and facilitated by new instruments (open call, joint PhD supervision) that foster academia-industry collaboration. The ELSA innovation lab will provide a platform for driving and utilizing socially-beneficial innovation. This European project is particularly important because AI solutions will be applied in such sensitive areas as the healthcare sector or autonomous driving. Artificial intelligence is also used in robotics, cybersecurity, media and document security. All of these application areas are the focus of the network. To achieve its goals, the network is addressing three major challenges: The development of robustness guarantees and certificates, data-secure and robust collaborative learning, and the development of human control mechanisms for the ethical and secure use of AI.

¹¹ <https://www.elsa-ai.eu/>

ELSA is funded under Horizon Europe, one of the largest research and innovation funding programs in the world.

ELSA is based on the understanding that safe and secure AI can only be achieved by building on foundational research. Relying on heuristics alone is susceptible to arms races and will not yield appreciable, reliable or sustainable levels of safety or security and is prone to fail overcoming the key obstacles in shaping AI in compliance with our European values. ELSA will overcome this by combining rigorous approaches such as robustness certification and differential privacy with machine learning and deep learning that are ubiquitous in applications. Hence, ELSA is based on a research agenda that encompasses the key research directions and fosters vital research by bringing together researchers across communities as well as different stakeholders. The research community is unified and directed towards addressing grand challenges that reach across use-cases by benchmarks that will drive innovation and provide measurable progress. The use-cases call for different, domain-specific requirements and risk analysis that require solutions, tools, and software that overcome obstacles in a trustworthy approach that conforms with legal and ethical principles grounded on European values. Our investigations emphasize state of the art methodology and technology (with a particular focus on Deep Learning techniques) in order to ensure practical relevance of the solutions.

The ELSA consortium consists of the following research institutions and companies:

- CISA Helmholtz Center for Information Security
- University of Helsinki
- KTH Royal Institute of Technology
- Computer Vision Center
- Italian Institute of Technology
- Valeo.ai
- University of Modena and Reggio Emilia
- National Interuniversity Consortium for Informatics
- University of Cagliari
- University of Genoa
- University of Milano
- Polytechnic University of Turin
- Max Planck Society
- National Institute for Research in Computer Science and Automation
- European Molecular Biology Laboratory
- Pluribus One
- Yooz

- Leonardo
- PAL Robotics
- École polytechnique fédérale de Lausanne
- ETH Zurich
- NVIDIA
- University of Oxford
- Lancaster University
- University of Birmingham
- The Alan Turing Institute