

The sample topics of doctoral dissertations are listed below.

If you have not found topics that interest you most, you can always suggest your own, we like open-calls very much.


as of 03.07.2023

-	Research topics offered by IDEAS NCBR	Abstract	Group/team leader in IDEAS NCBR	E-mail
1	Stochastic online optimization algorithms	Often when solving optimization problems, we are given some apriori information about the data, online requests, or other players taking part in the game. In this research challenge we aim to develop new algorithms that would be able to solve such problems when stochastic information about online requests is given up front.	Piotr Sankowski	piotr.sankowski@ideas-ncbr.pl
2	Learned algorithms for real-world incremental data	Data that we need to handle in real-world is never static and keeps on changing, e.g., vertices are added to social networks, or new ties appear. Hence, typically, we need to update the solution to our problem constantly. This not only poses efficiency issues, but requires that we do not change the solution too much each time. In this research challenge we want to face these problems from new perspective and create algorithms that can learn and adapt to changes.	Piotr Sankowski	piotr.sankowski@ideas-ncbr.pl
3	Explainable Algorithmic Tools	In this research project we aim to propose tools that would provide explanations for the different basic optimization problems, e.g., assignment problem, shortest paths, minimum cuts, or basic graphical neural networks. This research is motivated by the fact that even when faced with problems that can be solved exactly, we still would like to understand why this solution was computed.	Piotr Sankowski	piotr.sankowski@ideas-ncbr.pl

4	Learned Data-Structures	<p>ML tools enter as interior components into basic data structures or state-of-the-art approximation algorithms resulting in solutions that have better practical properties, e.g., indices. These new hybrid constructions are called learned data-structures. As the work on these ideas has just started we miss the right framework and tools for implementing state-of-art solutions and thus the research on new tools and models is hampered. This research aims to continue research on this problem and create new algorithms and data structures together with their implementations. This could prove tools to bridge the gap between theory and practice in algorithms and show that new theoretical advances can have practical implications.</p>	Piotr Sankowski	piotr.sankowski@ideas-ncbr.pl
5	Algorithmic Tools for Data Science and ML	<p>Although, different parallel computation models have been studied for years already. A new model that describes real-world systems has been proposed recently - the Massively Parallel Computation (MPC) frameworks includes systems such as MapReduce, Hadoop, Spark, or Flume. It comes with a completely new possibilities as well as requirements. MPC computations are executed in synchronous rounds, but implementing these rounds on real-world systems takes considerable time. One round takes orders of magnitude longer than on classical Cray type system. Thus we would like to solve problems, in particular graph problems, in as few rounds as possible. With this challenge in mind, this project aims to design methods to break barriers that were impossible to overcome using classical techniques and models. More specifically, we are going to work on new algorithmic tools that would improve efficiency of both parallel and non-parallel algorithms used in data science.</p>	Piotr Sankowski	piotr.sankowski@ideas-ncbr.pl




6	Deep NLP Models for Polish Language	In recent years we observe a huge progress in development of deep NLP models. In many applications these models can effectively compete with humans, and there usage is growths. However, the main works on these models are limited to mayor languages, and recent developments are not directly available for Polish language. The aim of this project is twofold: develop cutting edge NLP models for Polish language; use the experience gained this was to extend and improve models for other languages.	Piotr Sankowski	piotr.sankowski@ideas-ncbr.pl
7	Universal and Multi-modal Neural Networks	In this project we aim to work on multi-purpose and multi-modal neural networks. The tasks we aim to cope with will different problems where we aim to integrate different kind of information and aim to deliver joint representation that would allow for example: translate text to images and vice-versa for general and medical usage; transform natural language to animations, or approach no-code programming challenges.	Piotr Sankowski	piotr.sankowski@ideas-ncbr.pl
8	Continual learning of neural networks	Despite the recent successes in the fields of image, text, and sound processing, based on neural networks, adapting the models to changing data conditions still poses a significant challenge. Continual learning is a discipline that deals with the problem of changing the characteristics of the data used to train a model over time. The most important challenge is catastrophic forgetting, which causes the model learned sequentially on two datasets to lose its accuracy on the former with training on the latter. The project will develop methods for training deep neural networks that can address the problem of forgetfulness and create new application possibilities for continual learning.	Tomasz Trzcinski	tomasz.trzcinski@ideas-ncbr.pl
9	Learning data representations for computer vision and machine learning	Various data representations are crucial for solving multiple real-life applications, including autonomous driving, robot manipulations and language processing. In this project, we plan to develop novel methods for learning data representations leveraging neural network architectures. We will focus specifically on visual and multimodal representations and investigate methods	Tomasz Trzcinski	tomasz.trzcinski@ideas-ncbr.pl



		using supervised and unsupervised (e.g. generative) models to that end.		
10	Zero-waste machine learning	The computations run by contemporary machine learning models to process the increasing amount of data come at an enormous price of long processing time, high energy consumption and large carbon footprint generated by the computational infrastructure. Moreover, neural networks become increasingly complex, which leads to high monetary costs of their training and hinders the accessibility of research to less privileged communities. Existing approaches to reduce this burden are either focused on constraining the optimization with a limited budget of computational resources or they attempt to compress models. In this project, we plan to look holistically at the efficiency of machine learning models and draw inspiration to address their main challenges from the green sustainable economy principles. Instead of limiting training of machine learning models, we want to ask a different question: how can we make the best out of the information, resources and computations that we already have access to? Instead of constraining the amount of computations or memory used by the models, we focus on reusing what is available to them: computations done in the previous processing steps, partial information accessible at run-time or knowledge gained by the model during previous training sessions in continually learned models.	Tomasz Trzcński	tomasz.trzcinski@ideas-ncbr.pl
11	Partial information in self-supervised learning	Self-Supervised Learning (SSL) was introduced as a remedy for massive amounts of labeled data required by supervised approaches to building intelligent generalized models. It exploits the freely available data to generate supervisory signals which act as labels. For this purpose, in the case of image classification, SSL uses different image distortions, also referred to as	Tomasz Trzcński	tomasz.trzcinski@ideas-ncbr.pl

		<p>augmentations. While self-supervised approaches provide on par or superior results to their fully supervised competitors, they are computationally demanding, requiring large batches or momentum encoders.</p> <p>This project aims to leverage partial information into self-supervised strategies to increase their efficiency and reduce computational costs. Partial information assumes that a set of labels corresponding to a given image is known during inference, and it can be used to improve the performance of the model. This corresponds to a real-life application, where, for instance, we know that the image was captured in a forest or in a cave.</p> <p>The proposed research topic will leverage partial information into SSL, among others, by developing augmentation methods that use contextual information as a distortion source and utilize it as supervision in self-supervised learning.</p>		
12	Partial information in attention-based models	<p>Ever since the transformer was introduced in 2017, there has been a huge success in the field of Natural Language Processing (NLP). The main reason for the effectiveness of the transformer is its ability to handle long-term dependencies compared to RNNs and LSTMs. After its success in NLP, there have been various approaches to its usage for Computer Vision tasks. However, while transformers provide state-of-the-art results, they require large-scale training to trump an inductive bias. This project aims to leverage partial information into attention-based models to increase their efficiency and reduce computational costs. Partial information assumes that a set of labels corresponding to a given image is known during inference, and it can be used to improve the performance of the model. This corresponds to a real-life application, where, for instance, we know that the image was captured in a forest or</p>	Tomasz Trzciński	tomasz.trzcinski@ideas-ncbr.pl



		in a cave. The proposed research topic will leverage partial information into attention-based models, among others, by incorporating partial evidence to model sparsity in attention layers.		
13	Is self-learning really enough? Unsupervised representation learning.	In this project, we will study the problem of building efficient representations for downstream tasks in a continual learning scenario. Recently, novel self-supervised approaches showed promising results when they are properly regularized. We will investigate how internal network representation can be prepared for best re-use in the downstream tasks when trained continuously without supervision. Such an approach can be applied later to many downstream tasks in a cost-efficient way, i.e. with only a simple fine-tuning of a small and task-dedicated part of the model.	Tomasz Trzciński	tomasz.trzcinski@ideas-ncbr.pl
14	Efficient experience usage for life-long learning. Are exemplars all you need?	Storing exemplars directly in an additional memory buffer is the most common way to get acceptable performance in continual learning tasks. This allows you to easily learn cross-task features. Exemplar-based methods for class incremental learning or experience replay methods for online continual learning are focused on the efficient use of a given memory buffer by appropriate selection and retention of exemplars. Different methods directly optimize stored exemplars or use given memory to store models that allow generating samples or features -- the so-called pseudo rehearsal. The research question we ask in this project is the following: Is there a way to store previous knowledge more efficiently? Can we prompt saved representation in memory better, i.e. learn to prompt or query it?	Tomasz Trzciński	tomasz.trzcinski@ideas-ncbr.pl

15	<p>Representation alignment - What should not be regularized in continual learning?</p>	<p>Regularization-based methods are one of the easiest to apply and most common techniques for incremental learning, where we cannot store exemplars. There are two main types of regularization techniques: based on the weights like EWC or on the network outputs, such as LwF. Both try to alleviate the problem of catastrophic forgetting by keeping the network regularized -- enforcing the current network to remain similar to the old model and be able to solve similar tasks to the old ones by that. This increases the stability of the model, and can therefore hurt plasticity. In this research topic, we attempt to answer the following research questions: Is there a good trade-off for that? Maybe some aspect of the network should not be regularized at all, or regularized in a completely different way?</p>	Tomasz Trzciński	tomasz.trzcinski@ideas-ncbr.pl
16	<p>Cooperate to learn continually</p>	<p>Most class incremental learning methods assume one network, one backbone encoder to solve all the tasks, previously seen and the new ones. The signal goes through all the networks to solve any task. In living organisms, this is not exactly the case. Sensory information goes through different compartments that focus on various aspects of input signals. In addition, they are coordinated by a more global signal, e.g. gated by dopamine. In this line of research, we would like to focus on the cooperation of many learners - usually smaller, more energy-efficient, and weaker in comparison to the one-big model. Continual learning of them needs additional coordination.</p>	Tomasz Trzciński	tomasz.trzcinski@ideas-ncbr.pl



17	Continual Federated Learning	<p>In Federated Learning (FL) we have a central server node and many peers - clients that learn on their own data. We exchange only the model gradients to and from the server. Clients do not share their data or any information that can break privacy. Usually, a differential privacy model is applied to enforce that. This is an attractive way to train models in many domains, e.g. healthcare, advertisements, and mobile applications, just to name a few. Most of the use cases are based on the static data, split for clients, and then the FL training process proceeds. Neither tasks nor data are changing along the way. New concepts are not emerging at the client's level. Simply, they will be averaged (FedAvg) and lost. Most of the methods do not consider learning anything and how to integrate and propagate this knowledge from the server to other clients. In this work, we address the problem of incremental learning on clients' devices, usually edge devices, low-energy, and memorable ones. Learning new concepts in such an environment is challenging and not well explored so far.</p>	Tomasz Trzcinski	tomasz.trzcinski@ideas-ncbr.pl
18	Foundational models in continual learning	<p>An enormous amount of training data used recently in language modelling (e.g. GPT) led to emerging properties (for example the models can handle task, which they never encountered via prompting). We propose to study such models in the area control (e.g. for robotic tasks). We speculate that by gathering in one model, a large number of skills can lead to more efficient learning on new tasks. The key questions to be studied during the project are: a) how to train a model capable of storing many tasks, b) how to query such a model efficiently, in order to learn new task faster, c) how to update such a model with new task, while not forgetting the previous tasks.</p>	Piotr Miłoś	piotr.milos@ideas-ncbr.pl



19	Mechanism of transfer in CLR	<p>Knowledge transfer is key for obtaining a good performance for complex tasks. Intuitively, it is much more effective to pre-train a model on related (and perhaps easier/cheaper tasks) and later 'just to fine-tune' it to a new task. Such approaches have been widespread in practice. However, they lack a proper understanding in the case of neural networks. It is not clear what is really transfer, if these are useful features, or perhaps good weight initialization. In the project, we plan to evaluate existing hypotheses explaining transfer in the case of control tasks (e.g. robotic manipulation).</p>	Piotr Miłoś	piotr.milos@ideas-ncbr.pl
20	Recalling not to forget	<p>Experience replay has proven to be one of the most powerful technique mitigating forgetting in long sequences of tasks. Its main drawback is large usage of memory, which prevents in scalability for long sequences.</p> <p>This project aims for a systematic study of this experience replay techniques with the goal of making them more efficient. To this end, we conceptualize two major tasks:</p> <ul style="list-style-type: none"> - what are the quantitative and qualitative properties of the experience replay samples, which are needed for successful mitigation of forgetting - what are mechanism of experience replay <p>For the second questions, we speculate that the experience replay loss gradients are sparse and can be distilled into much more compact form. Perhaps, also they could be factorized with respect to network weights and, therefore, expressed as a sum of simple per-weight losses.</p>	Piotr Miłoś	piotr.milos@ideas-ncbr.pl
21	Obtaining efficient representations for C	<p>Learning a long sequence of tasks might be facilitated by active representation learning. In this project, we aim to study two high level questions. The first one is, how much the structure of the space can facilitate efficient learning. For example, one can introduce a learning bias such that representations related to various tasks can be easily disentangled, for example expressed linearly. The second</p>	Piotr Miłoś	piotr.milos@ideas-ncbr.pl

		questions, is how much data augmentations can facilitate forming better representations.		
22	Planning with subgoals	When dealing with problems that require long-term planning, the search depth often needs to be reduced due to a large branching factor (for example, solving the Rubik's cube). One promising solution to this problem is the use of subgoals, which are intermediate milestones towards the final solution. Some previous implementations of this concept have already demonstrated impressive results by allowing for deeper search and solving problems with much lower computational costs. The project aims to explore the design and testing of new methods related to subgoals for a diverse range of problems	Piotr Miłoś	piotr.milos@ideas-ncbr.pl
23	Efficient neural planning	Neural networks has brought spectacular progress in solving many problems. In some cases, however, we expect that they have natural limitations and cannot solve each problem also. An archtypical example concerns combinatorial puzzles, like Rubik's cube, but has much broader applicability in discrete optimization. The project will explore how to use neural network efficiently with other computational mechanisms (e.g. classical search techniques). The core question is understanding situations, in which neural networks make errors and in which can be trusted. A proper analysis should lead to more efficient planning methods.	Piotr Miłoś	piotr.milos@ideas-ncbr.pl
24	Transformers with external memory, towards neural knowledge systems	Transformers have been extremely successful architectures in sequential modeling, however, they have a practical limitation of a relatively short context span due to the quadratic cost of the attention mechanism. The project aims to explore practical solutions to mitigate this problem by providing access to an external memory, which can be thought of as a external knowledge system. The aim is to factorise the reasoning capabilities, which could be stored in the	Piotr Miłoś	piotr.milos@ideas-ncbr.pl


		weights of transformer, from trivia facts, which can be stored in memory.		
25	Language modelling for scientific reasonings	Large language models like GPTs have revolutionized the field of machine learning by introducing new ways of learning, such as in-context learning, chain-of-thoughts, and scratchpads. Interestingly, they also appear to possess rudimentary reasoning capabilities. This project aims to investigate how we can improve and utilize these capabilities to achieve better results.	Piotr Miłoś	piotr.milos@ideas-ncbr.pl
26	Massive multi-agent neural simulators	Classical reinforcement learning operates under the assumption of perfect knowledge of the environment, which is only applicable in limited, idealized scenarios. In more typical situations, an agent only has access to a subset of information about the environment, particularly in multi-agent systems like traffic control, where an agent's understanding of other agents' intentions may be limited. Our project aims to explore this by scaling the number of agents and observing patterns that emerge, with the goal of designing improved control mechanisms.	Piotr Miłoś	piotr.milos@ideas-ncbr.pl
27	Tree species detection using close-range remote sensing data	The problem of tree species recognition based on different remote sensing technologies is represented by a large number of different publications. In this respect, especially works related to aerial and satellite data acquisition systems have a long publication history. Remote sensing at close range is developing strongly in this field and it is necessary to carry out scientific work in this area in order to keep up with the intense technological developments that are taking place. The aim of the doctoral thesis will be the recognition of tree species with AI in order to automatically inventory them in forest management. Depending on the competence and commitment of the PhD candidate, the work may involve several different	Krzysztof Stereńczak	krzysztof.sterenczak@ideas-ncbr.pl

		remote sensing technologies and forest and/or urban environments.		
28	Tree quality assessment using close-range remote sensing data	The size of trees and whether there are defects on the side of the trunk affect the economic value of individual trees. However, the sides of the trunks may also contain various parasites or the effects of various biotic and abiotic factors, which in turn tell us about the current or future health status of the trees. The detection of such lateral objects is important for the protection of forests or the management of urban greenery management. Close-range remote sensing provides data that is highly likely to help visualise various artefacts on trees. The use of artificial intelligence algorithms can further increase the probability of detecting these artefacts. The aim of the PhD is to use AI to recognise the size and quality of trees in order to automatically inventory them in forest management. Depending on the expertise and commitment of the PhD candidate, the work may involve several different remote sensing technologies and forest and/or urban environments.	Krzysztof Stereńczak	krzysztof.sterenczak@ideas-ncbr.pl
29	Determination of selected biometric single tree characteristics using close-range remote sensing data	An inventory of trees, whether in the city or in the forest, always involves measuring at least some of their characteristics such as diameter at the breast height, tree height, crown diameter and crown base height. These measurements are made under different environmental conditions, with different tools and by people with different training and experience. These measurements are labour-intensive and difficult to verify. However, they are the basis for most decisions related to forest management, urban greening or tree protection.	Krzysztof Stereńczak	krzysztof.sterenczak@ideas-ncbr.pl


		<p>The aim of the PhD is to use AI to determine selected individual biometric characteristics of trees in order to automatically inventory them during forest management. Depending on the expertise and commitment of the PhD candidate, the work may involve several different remote sensing technologies and forest and/or urban environments.</p>		
30	<p>Measuring lying trees snags using using close-range remote sensing data</p>	<p>The effects of catastrophic winds or snowfalls sometimes result in many thousands of hectares of forest being overturned and destroyed. The damaged areas are very dangerous, so it is difficult to inventory them in order to determine the economic damage associated with the event or to plan future activities. In these areas, trees are lying on top of each other, often with varying degrees of damage, making it virtually impossible to move around the area on the ground. Another example of an area with lying trees is a situation where foresters plan to harvest raw wood for the timber industry. This involves cutting down trees that are then lying on the ground, and the forester has to measure each one, which is often labour-intensive and sometimes dangerous.</p> <p>The aim of this project is to use AI to detect and measure fallen trees in order to automatically inventory them on site. The development of automatic recognition methods for measuring lying trees is therefore of great practical and cognitive importance. On the one hand, research in this area is quite limited, but on the other hand, the development of such tools will improve the quality and safety of the work of many people involved in forest management and protection.</p>	<p>Krzysztof Stereńczak</p>	<p>krzysztof.sterenczak@ideas-ncbr.pl</p>



31	<p>Game-Theoretic Aspects of Blockchain Technologies</p>	<p>While there are countless potential application of blockchain, almost all of them share a common feature: the parties that use it are assumed to be, in principle, self-interested utility maximizing individuals. Given this, many aspects related to the blockchain technology should be analysed using the apparatus of game theory. These include such issues like: selfish mining, majority attacks and Denial of Service attacks, computational power allocation, reward allocation, and pool selection, and energy trading. While the literature that analyses game-theoretic aspects of blockchain is growing, there are many interesting open questions that have not yet been answered in a satisfactory way. For instance: how to design rules that lead to the development of payment channel networks that are secure, reliable and efficient.</p>	Tomasz Michalak	<p>tomasz.michalak@ideas-ncbr.pl</p>
32	<p>Adversarial Social Network Analysis</p>	<p>How can individuals and communities protect their privacy against social network analysis techniques, algorithms and other tools? How do criminals or terrorists organizations evade detection by such tools? Under which conditions can these tools be made strategy proof? These fundamental questions have attracted little attention in the literature to date, as most tools are built around the assumption that individuals or groups in a network do not act strategically to evade social network analysis. To address this issue, a recently novel paradigm is social network analysis explicitly models strategic behaviour of network actors using the apparatus of game theory. Addressing this research challenge has various implications. For instance, it may allow two individuals to keep their relationship secret or private. It may also allow members of an activist group to conceal their membership, or even conceal the existence of their group from authoritarian regimes. Furthermore, it may assist security agencies and counter terrorism units in understanding the strategies that covert organizations use to escape detection,</p>	Tomasz Michalak	<p>tomasz.michalak@ideas-ncbr.pl</p>




		<p>and give rise to new strategy-proof countermeasures.</p>		
33	<p>Adversarial Social Detection based on Graph Neural Networks</p>	<p>Social Networks have become a primary media for cybercrimes. For instance, attackers may compromise accounts to diffuse misinformation (e.g., fake news, rumors, hate speeches, etc.) through a social network. Fraudsters may also trick innocent customers into conducting fraudulent transactions over online trading platforms. Meanwhile, on the defense side, defenders (e.g., network administrators) are increasingly employing machine-learning-based tools to detect malicious behaviors. Graph Neural Networks (GNNs) have become the <i>de facto</i> choice of social detection tools due to their superior performance over a wide spectrum of tasks.</p> <p>In this project, the overall goal is to develop robust and effective GNN-based social detection tools in an adversarial environment. This goal is decomposed into three coherent objectives. First, design more effective GNN-based tools to detect crimes in social networks that could achieve a better detection accuracy as well as a lower false positive rate. Second, from the standpoint of an attacker, investigate effective evasion techniques to bypass the detection of the GNN-based tools. Third, as a defender, enhance the robustness of the GNN-based detection tools to mitigate evasion</p>	<p>Tomasz Michalak (jointly with Kai Zhou, PolyU, Hong Kong)</p>	<p>tomasz.michalak@ideas-ncbr.pl</p>



		attacks. Overall, the expected outcomes significantly advance our knowledge in developing trustworthy AI systems in a real-world adversarial environment.		
34	Robust Deep Learning Systems through Explainable AI (XAI)	Machine learning, especially deep learning, has transformed the way how data is processed. Recent studies have revealed that deep learning systems lack transparency and are also vulnerable to adversarial attacks. The fundamental reason is that deep learning systems rely on a large amount of data possibly collected from the wild, which gives the opportunity for attacks to inject \textit{adversarial noise} to mislead the systems. Meanwhile, an active line of research, termed Explainable AI (XAI), aims to interpret the decisions made by AI systems, which essentially identify a subset of data that is important for the decision. In this project, we will investigate how to use XAI to build robust deep learning systems against attacks. This goal is decomposed into two major objectives. First, enhance existing or develop new XAI techniques to effectively identify adversarial noises from data. That is, we employ more advanced XAI to sanitize the data for deep learning systems. Second, provided with the sanitation results, develop new algorithms to train robust deep learning systems from the \textit{noisy} data. Overall, the expected outcomes will make	Tomasz Michalak (jointly with Kai Zhou, PolyU, Hong Kong)	tomasz.michalak@ideas-ncbr.pl

		significant contributions toward developing more transparent and robust deep learning systems.		
35	Fraud Detection in Bitcoin Transaction Networks based on Unsupervised Machine Learning	<p>Cryptocurrency based on blockchain technology has significantly reduced our dependence on the central authority. Meanwhile, due to its anonymity nature, cryptocurrency trading platforms have also become the perfect media for financial crimes. For example, many known studies have revealed that criminals are increasingly using Bitcoin transaction networks for money laundering. Thus, a very significant while underexplored problem is how to effectively detect fraudsters in bitcoin transaction networks utilizing machine learning techniques. The major objectives of this project are as follows. First, design unsupervised machine learning algorithms (e.g., clustering, contrastive learning, etc) to effectively identify fraudulent transactions and malicious accounts in a transaction network. Essentially, this objective calls for new approaches to detect anomalies at the node level, edge level, and sub-graph level within a graph. Second, investigate the vulnerabilities of prior detection methods by designing more practical evasion techniques. Especially, besides considering the evasion objective, the design of evasion attacks should simultaneously consider the need for stealthiness and preserving malicious utilities. Third, faced with strategical</p>	<p>Tomasz Michalak (jointly with Kai Zhou, PolyU, Hong Kong)</p>	<p>tomasz.michalak@ideas-ncbr.pl</p>

		<p>evaders, further improve the robustness of the detection methods. Successfully achieving these objectives will contribute to applying unsupervised machine learning in anomaly detection from a technical perspective, and enhancing the security of the trading environment of cryptocurrencies.</p>		
36	<p>Federated Learning over Distributed Graph Data</p>	<p>Federated learning is a computation paradigm for training machine learning models from distributed data while preserving data privacy. Most of the existing research has been devoted to investigating federated learning algorithms over well-structured data such as tabular data. Since graphs are widely used to represent various kinds of relational data (e.g., social networks, recommendation systems, communication networks, etc.), there is an urgent need to investigate and design new federated learning algorithms for graphs. Especially, graphs have some unique features which make previous algorithms not suitable. For example, the features of nodes in graphs are highly heterogeneous, which makes federated training algorithms hard to converge. Also, graphs distributed into different subgraphs will inevitably miss those interconnected edges, which represent a kind of information loss for learning. Thus, in this project, the primary goal is to design new federated learning algorithms for graph learning models (e.g., graph neural networks) over</p>	<p>Tomasz Michalak (jointly with Kai Zhou, PolyU, Hong Kong)</p>	<p>tomasz.michalak@ideas-ncbr.pl</p>



		distributed graphs. In expectation, these algorithms will mitigate a series of issues of learning over graph data, including heterogeneity, information loss, and so on.		
37	Logical Interpretable Learning	The objective is to enable a paradigm shift from correlation-driven to scaled-up causality-driven machine learning. The project deals with the unresolved learning challenge in logical engineering, the scaling challenge of probabilistic causal models, and the correlation-reliance of deep learning using explainable AI, advanced time series analysis, and multimodal deep learning.	Tomasz Michalak (jointly with Youcef Djenouri, Norce Norwegian Research Center)	tomasz.michalak@ideas-ncbr.pl
38	Developing advanced market mechanisms for Local Energy Markets	Renewable energy sources have to be integrated with the whole electricity grid in a way that satisfy all the market players and make the whole system sustainable in the long run. To this end, various market design concepts have been studied. However, no comprehensive model that takes into account all the key aspects of the problem has been developed so far. In particular, no tractable market mechanisms has been developed that simultaneously address uncertainty, strategic behavior, non-convexity of market participants' cost/utility function, and network constraints. The objective of this ambitious project is the development of such a mechanism.	Tomasz Michalak	tomasz.michalak@ideas-ncbr.pl



39	Collaborative forecasting of renewable energy resources	<p>Recent works on forecasting renewable energy production demonstrate that using data from neighboring locations improves the accuracy of prediction. Given this, there is a need to develop methods to share data that, on the one hand, respect privacy and confidentiality constraints, and, on the other hand, are based on market mechanisms that incentivize data owners to participate in the whole system. To this end, in this research project, we will develop a forecasting system that combine statistics, machine and deep learning with cryptography, blockchain, mechanism design, and sociology.</p>	Tomasz Michalak	tomasz.michalak@ideas-ncbr.pl
40	Proofs-of-Space in practice	<p>The most popular blockchain platforms use consensus based on the so-called Proofs-of-Work, where the participants are incentivized to constantly solve many computational puzzles (this process is also called mining). This leads to massive electricity consumption. Several alternatives to Bitcoin mining have been proposed in the past. Stefan Dziembowski (who leads this research at IDEAS) is one of the authors of another approach to this problem, called the Proofs-of-Space. In this solution, the computational puzzles are replaced with proofs that a given party contributed some disk space to the system. Several ongoing blockchain projects are based on these ideas. This student will work on improvements to these protocols.</p>	Stefan Dziembowski	stefan.dziembowski@ideas-ncbr.pl
41	Real-life side-channel security of blockchain wallets	<p>Another critical weakness in the vision of decentralizing internet services is that interacting with blockchains is more complicated than in the case of centralized solutions. Moreover, the decentralization makes it impossible to revert the transactions that were posted by mistake or as a result of an attack. Due to this, the users often rely on the so-called hardware wallets, which are dedicated devices protected against cyber-attacks. This student will work on analyzing the security of the existing hardware wallets. In particular, we will be interested in their side-channel security, i.e., security against attacks based on information such as power</p>	Stefan Dziembowski	stefan.dziembowski@ideas-ncbr.pl

		consumption or electromagnetic radiation.		
42	Privacy in machine learning	Several machine learning applications involve issues where privacy plays a special role. This includes cases in which secrecy applies to the training data (e.g., when it contains medical information) and those in which the algorithm itself is subject to protection because, for example, it reveals specific information about the training data. The student will work on addressing these problems using methods such as multiparty computation protocols, differential privacy, and trusted execution environments.	Stefan Dziembowski	stefan.dziembowski@ideas-ncbr.pl
43	Formal analysis of cryptographic protocols using machine learning	One of the main problems in the blockchain space is that decentralized solutions are typically more complex and error-prone than centralized ones. In particular, errors in smart contracts can lead to considerable financial losses. Furthermore, some blockchain algorithms in the past had serious mistakes that could be exploited to steal large amounts of money. This student will work on addressing these problems using tools from formal methods from machine learning, in combination with proof assistants and formal theorem provers such as Coq, Easycrypt, Why3, and others.	Stefan Dziembowski	stefan.dziembowski@ideas-ncbr.pl
44	Learning to teach	Classically, reinforcement learning agents optimize the sum of discounted rewards, where the reward structure is assumed as given. This is a bottleneck if we want the agent to generalize to other tasks or when the reward structure is unknown and de facto is a part of the solution (e.g., as is the case for large language models). The goal is to formalize a meta-learning algorithm where we allow agents to autonomously discover interesting tasks, skills, or generate interesting data. This goal is to structure this problem as a game where we train both a pupil and a teacher, allow them to cooperate or compete with one another and improve in a closed feedback loop. Other objectives	Łukasz Kuciński	lukasz.kucinski@ideas-ncbr.pl

		include applying these ideas e.g., to seamlessly train a subgoal generator and a low-level policy in subgoal search, discover skills to improve transfer in continual learning, learn to improve optimization algorithms.		
45	Truthfulness and uncertainty-aware agents	Large language models (LLMs) have proven to be very strong general-purpose architectures. Recent successes of systems like chatGPT highlighted several important areas that the research community needs to address. These include truthfulness, alignment, or uncertainty awareness. The lack of truthfulness results in model making stuff up, or hallucinating. We want strong models to be aligned with human values and act in accordance with human intentions. Lastly, models that are not aware of their uncertainty may either unnecessarily withhold information, or hallucinate; whereas if the contrary happens the model can e.g., delegate queries to external APIs. The goal of this research stream is to formulate the aforementioned problems as a solution to a multi-player cooperative game and via agents autonomous interaction.	Łukasz Kuciński	lukasz.kucinski@ideas-ncbr.pl
46	Instruction-based RL	Natural language is a very exciting modality, which has opened up to other parts of Machine Learning mostly due to the power of large language models (LLMs). In particular it allows constructing reinforcement learning (RL) agents that can interact with the world via instructions or communicate their internal state in natural language. Furthermore, we can use LLMs as AI-generated environments to be used in RL, which opens up new possibilities such as performing interventions or asking counterfactual questions in a natural way. The goal is to study the capacity of RL agents to learn in this regime and simultaneously ask questions about LLMs consistency, truthfulness, or their knowledge graph.	Łukasz Kuciński	lukasz.kucinski@ideas-ncbr.pl

47	Alternative Learning Objectives in RL	<p>Classically, RL algorithms use some approximation of future rewards as a learning signal to improve a policy. Recent research has shown that RL can be viewed through the lens of representation learning. Here the premise is that the policy guided by the similarity between representation of the current state and the goal state can be a valid alternative. The goal of this research stream is to investigate old and recent ideas from statistics and self-supervised learning to propose new RL algorithms. In particular, study the impact of the methods in subgoal search, for subgoal generation, learning the latent, or guiding the search.</p>	<p>Łukasz Kuciński</p>	<p>lukasz.kucinski@ideas-ncbr.pl</p>
48	Parallel graph algorithms: theory and practice	<p>The research objective is to develop new algorithms for classical graph problems - related to reachability, shortest paths, and maximum flow - on parallel machines with shared memory. First of all, we would like to obtain some new theoretical work-depth trade-offs. We are also interested in engineering fast implementations of graph algorithms on modern multicore machines.</p>	<p>Adam Karczmarz</p>	<p>adam.karczmarz@ideas-ncbr.pl</p>
49	Semantic Networks for Experience-Oriented Computational Psychiatry	<p>Mental health diagnosis remains prone to systemic and predictable cognitive errors resulting from not-fully-conscious simplistic inference schemes called heuristics, which may lead to epistemic injustice affecting medical assessments. To reduce such epistemic injustice, this topic aims to develop tools for computer-assisted recognition of experience patterns in selected mental conditions by incorporating a semantic network of lived experience founded upon a hybrid database combining both third-person medical data and deep first-person reports. The advantage of semantic networks is that they can be used to represent meanings (including metaphors) of natural language in a way easily interpretable by both AI and humans, which may facilitate overcoming distrust in computer-assisted diagnostic tools. NLP algorithms can analyze language patterns and features in text data to uncover</p>	<p>Marcin Moskalewicz</p>	<p>marcin.moskalewicz@ideas-ncbr.pl</p>

		meaningful insights concerning mental health.		
50	Navigating Ethical AI in Mental Health	<p>Human ethical biases (conscious and unconscious) can significantly skew data and influence the behavior of AI systems; the way data is collected, labelled, and processed affects training datasets that mould AI systems. The aim of this topic is to identify and overcome such biases regarding mental health. When these biases seep into AI algorithms, they perpetuate harmful stereotypes and misinformation, potentially leading to misdiagnosis and improper treatment. One of the most pervasive demographic biases in AI research is the overrepresentation of WEIRD (White, Educated, Industrialized, Rich, Democratic) populations. If data labels categorize certain behaviors as indicative of mental illness, the AI systems will likely perpetuate these misconceptions. When an algorithm is optimized for a certain goal that reflects a human bias (such as prioritizing efficiency over fairness), an AI system such as a therapeutic bot can advise solutions that are ethically skewed. An AI system might e.g. maximize overall utility leading to treatments or resources being disproportionately directed towards groups that are easier or less costly to treat, potentially neglecting those with more complex or costly needs, thus disregarding the bioethical principle of autonomy or individualized care. The broader aim of this topic is to stimulate conversation around the ethical use of AI in mental health, inspire the development of more robust, bias-aware data collection methods, and promote transparency in AI algorithm design.</p>	<p>Marcin Moskalewicz</p>	<p>marcin.moskalewicz@ideas-ncbr.pl</p>

51	Reducing Cognitive Errors in Autism Diagnosis	<p>The main goal of this topic is the design of a digital decision-making tool for mental health professionals working with people on the autism spectrum. In contrast to prevailing perceptions, extensive research underscores the vulnerability of expert decision-making to myriad of systematic cognitive biases, such as representativeness, anchoring or blind spot bias. The project will analyze (and implement in the form of appropriate modules) factors influencing expert decisions, such as: prevalent social stereotypes regarding autism, cognitive humility, cognitive fatigue, etc. to equip experts with support enhancing the precision and reliability of their judgment as well as mitigating the adverse repercussions stemming from cognitive errors that may arise during the evaluation or diagnosis of individuals with autism. The larger context for this topic is prioritizing patient-centered approaches to autism.</p>	<p>Marcin Moskalewicz</p>	<p>marcin.moskalewicz@ideas-ncbr.pl</p>
52	Advancing Synthetic Phenomenology Through Algorithmic Tools	<p>The aim of this topic is the algorithmization of some key analytic tools used in applied (qualitative) phenomenology, such as reduction and imaginative variation. Synthetic phenomenology in this context refers to the creation of algorithms specifically designed to analyze the “essential” facets of phenomenal states via linguistic/conceptual representations. The goal is to merge the realms of phenomenology and machine learning by leveraging a neural network architecture trained on carefully preselected corpora of lived experiences. The core challenge is the meticulous selection of first-person training data encapsulating both regular and pathological lived experiences. Another critical aspect involves the application of the phenomenological principle of epoché signifying the suspension of judgment via the implementation of debiasing strategies reflecting various levels of phenomenological reduction. Blending phenomenological methodologies with advanced machine learning tools offers an unprecedented level of insight into</p>	<p>Marcin Moskalewicz</p>	<p>marcin.moskalewicz@ideas-ncbr.pl</p>

		subjectivity, unattainable through conventional phenomenological tools reliant solely on the human mind.		
--	--	--	--	--

